Army Research Laboratory

ROBOTICS COLLABORATIVE TECHNOLOGY ALLIANCE (RCTA)

FY 2012 Annual Program Plan

Submitted March 2012

Members
Boston Dynamics
Carnegie Mellon University
Florida State University
General Dynamics Robotic Systems (Integration Lead Organization)
Jet Propulsion Laboratory
QinetiQ North America
University of Central Florida
University of Pennsylvania

All contents are public releasable.

For any purpose other than to evaluate the *Robotics CTA FY 2012 Annual Program Plan*, the technical and financial data contained herein shall not be disclosed outside the Government in whole or in part without prior written permission from a duly authorized representative of General Dynamics Robotic Systems. The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Cooperative Research Agreement.

1. INTRODUCTION 1
Background of Robotics Collaborative Technology Alliance Program 1
Vision of the Robotics CTA 2
Technical Barriers to the Vision 3
Fundamental Research Thrusts for Overcoming Technical Barriers 6
Overview of Technical Research Thrusts 11
Assessment of Integrated Research 22
Progress During 2011 27
Collaboration: Both Within and Outside the Alliance 33
Relationship to Military Needs 40
Conclusion 44
2. INTELLIGENCE OVERVIEW 47
Intelligence Tasks 56
I1-2012 – Framework for Intelligence 56
I2-2012 – World Model and Applied Intelligence Architecture 64
I3-2012 – Combining Cognitive and Probabilistic Reasoning 87
I4-2012 – Generating Adaptive Tactical Behaviors 97
I5-2012 – Learning through Experience 108
I6-2012 – Life-long Learning 118
I9-2012 – Distributed Intelligence for Human/Robot Teams 126
3. PERCEPTION OVERVIEW 136
Perception Tasks 146
P1-2012 – Exploiting Novel Sensor Phenomenology 146
P2-2012 – Compact, High Performance Sensors 157
P3-2012 – Static Scene Understanding 171
P4-2012 – Perception for Missions in Complex Environments 193
P5-2012 – Dynamic Scene Understanding 212
P6-2012 – Perception for Missions in Dynamic Environments 223
4. HUMAN-ROBOT INTERACTION OVERVIEW 239
Human-Robot Interaction Tasks 245
H1-2012 – Shared Mental Models for Soldier-Robot (SR) Teaming 245
H2-2012 – Situation Awareness in Human-Robot Teams 260
H4-2012 – Communication Protocol and Language Processing 271
H5-2012 – Evaluating Tactical Command and Coordination Vocabulary and Protocols 282
H8-2012 – Social Dynamics Modeling and Simulation 292
H9-2012 – Social Dynamics Experimentation 305
5. DEXTEROUS MANIPULATION AND UNIQUE MOBILITY OVERVIEW 327
Dexterous Manipulation and Unique Mobility Tasks 333
M1-2012 – Theory and Principles of Mobile Manipulation 333
M2-2012 – Principles of Generalized Grasp Mechanics 348
M3-2012 – Sensor-based Dexterous Manipulation 357
M4-2012 – High Degree-of-Freedom Dynamic Manipulation 365

M5-2012 – Geometric Mechanics for Multi-modal Highly Articulated Systems 370
M6-2012 - Principles of Locomotion Mechanics 383
M7-2012 – Learning Terrain Interactions..... 392
M8-2012 - Dynamic Multi-modal and Reconfigurable Mechanisms..... 398
M9-2010 – Next Generation Actuators and Materials..... 412
6. INTEGRATED RESEARCH OVERVIEW 426
Integrated Research Philosophy..... 426
Integration and Assessment Process 427
Integration and Assessment Team 430
Task Breakdown of Integrated Research 430
Integrated Research Tasks 432
IR1-2012 – Integration, Experimental Design, and Assessment 432
IR2-2012 – Integration, Experimental Design, and Assessment 435
IR3-2011 – Platforms and Testbeds for Integrated Research 456
IR4-2012 – Modeling and Simulation Environment 461
IR5-2012 – Collaboration Software Integration 469
APPENDIX 1. COMMON LANAGUAGE TABLE 477

1. INTRODUCTION

Background of Robotics Collaborative Technology Alliance Program

The United States Army Research Laboratory (ARL) issued a Program Announcement (PA)[1] on February 2, 2009 which solicited proposals for a new fundamental research program entitled the Robotics Collaborative Technology Alliance (RCTA) in order to help fulfill the research and development goals of the U.S. Department of the Army. The stated purpose of this Alliance is to "bring together Government, industrial, and academic institutions to address research and development required to enable the deployment of future military unmanned ground vehicle systems ranging in size from man-portables to ground combat vehicles." The PA identified four key technology areas expected to be critical to the development of future autonomous systems, namely:

- Perception
- Intelligence
- Human-Robot Interaction (HRI)
- Dexterous Manipulation and Unique Mobility (DMUM)

It further emphasized the overlap and interplay among these technologies and the need to better understand their interactions through relevant integration and assessment activities.

To accomplish this mission, the PA called for the formation of a "consortium of a small number of industrial and academic institutions acting as equal partners in a research enterprise." The Consortium is to conduct both basic research[2] and applied research[3]. It is expected to partner with ARL and other Government agencies to advance technology by formulating and executing a number of individual, but coordinated, research tasks.

The PA calls for the preparation of a proposed Annual Program Plan (APP) for the research in each year of the RCTA program. The APP provides a detailed plan of research activities, down to the task and even subtask levels. Each year's APP is to be presented to the Research Management Board (RMB) for comments and suggestions.

This document is the proposed 2012 APP, which consists of six sections. The first section is an introduction which presents the vision of the RCTA, the barriers to achieving that vision, and the primary technical thrusts we are undertaking to overcome the barriers. The introduction also provides a brief overview of each research thrust as well as shows how the thrusts map to the four technology areas identified above, summarizes the process of integrating research outcomes and assessing progress, and shows the relationship of the research to military needs. Section 2

[1] http://www.arl.army.mil/www/DownloadedInternetPages/CurrentPages/CTA/Documents/ROBCTAFINALPA11FEB09.pdf
[2] funded by Project H09 of Program Element (PE) 0601104A
[3] funded by Project TS2 of PE 0602120A

through Section 5 of the APP describes in detail the research by technology area, down to the task and subtask levels. These sections specifically identify the state of the art related to each research task, describe how the present research moves beyond the state of the art, and identify metrics and goals that quantify the progress toward achieving the research goals. Finally, Section 6 describes the detailed plans for integrating research outcomes in order to achieve needed capabilities that overcome the barriers to our vision for autonomous ground systems.

Vision of the Robotics CTA

Unmanned systems have begun to have a significant impact on warfare; for example, unmanned drones providing sustained surveillance, swift precise attacks on high value targets, and small robots being used for counter-IED missions. While unmanned and highly complex, these systems are still generally remotely piloted systems, reliant upon near-continuous control by a human operator and vulnerable to break-downs of communications links. The future for unmanned systems lies in the development of highly capable systems, which have a set of intelligence-based capabilities sufficient to enable the teaming of autonomous systems with Soldiers. To act as teammates, robotic systems will need to reason about their missions, move through the world in a tactically correct way, observe salient events in the world around them, communicate efficiently with Soldiers and other autonomous systems, and effectively perform a variety of mission tasks. These capabilities certainly do not need to be at a human level, but they do need to be at a level that moves well beyond the current state of tele-operation or closely supervised autonomy.

More specifically, our vision is one where robotic systems have greatly enhanced capabilities in the following five problem domains:

Adaptive Tactical Reasoning. In our vision, robots understand the concept of a mission or task, including stages of progress and measures of success. They work with Soldiers, using the shared concept of METT-TC – mission, enemy, troops, terrain, time, and civilian considerations. They generate tasks to accomplish the mission at hand, reacting appropriately to unforeseen events. They understand their teammates, human or otherwise, and what they need to know during the mission. They make clear distinctions among teammates, adversaries, and non-combatants. They are able to learn from experience, including their own mistakes, generalizing appropriately from specific examples.

Focused Situational Awareness. Future autonomous ground systems maintain situational awareness (SA) that is relevant to the current task and the larger mission. They monitor friendly forces and neutrals and look for threats. They contribute to the general SA of the unit, looking for any salient unexpected events. They continuously predict the future situation so they can better detect anomalies and learn from experience.

Efficient Proactive Interaction with Humans. In our vision, robots interact with each other and especially with Soldiers in an efficient and proactive way relevant to the evolving situation. They receive, understand, and acknowledge orders, asking for clarification if needed. They send

relevant messages to their teammates about salient events, using whatever communication means are available. They also receive and understand messages about unfolding events from others and, thus, are able to take any needed actions. They have little or no need for operator control units (OCUs), instead working with Soldiers as team members utilizing the same modes of command and control that Soldiers use among themselves.

Safe, Secure, and Adaptive Movement. We envision robots that move on orders or their own initiative from one tactical position to the next with little or no reliance on metric inputs such as GPS. They can move, as Soldiers do, to semantically described locations (e.g., "third building on the left after the next intersection"). They also move safely in the presence of people, vehicles, and even animals. They move securely in the context of the current METT-TC, moving with speed or stealth as appropriate. Additionally, they move in a manner that is adaptive to both natural and cultural terrain, including hills, rocks, mud, ice, walls, vehicles, tunnels, and other features.

Interaction with the Physical World. Finally, in our vision, robots are able to observe objects at close quarters to enable 3D interaction with them. They pick-up and move objects, either upon semantic direction or their own initiative. They use tools as necessary for digging, cutting, drilling, etc. They also manipulate doors, windows, hoods of vehicles, etc. as needed to gain access to buildings, vehicles, or confined spaces to execute their missions. They have the dexterity to manipulate a small wire, the strength to pick-up heavy objects, and the range of motion to reach around obstacles. While interacting with the physical world, they can learn, for example, that an object is deformable.

We use a convenient anthropomorphic shorthand – "Think," "Look," "Talk," "Move," "Work" – to encapsulate these five capability building blocks.

Technical Barriers to the Vision

The above vision is an appealing one that promises very great capabilities for future autonomous systems. However, there are significant technical barriers to each envisioned capability. Here, we examine them in turn.

Think. Adaptive tactical reasoning requires both declarative and procedural knowledge with which to reason. Neither exists in current systems, which generally have no data structures for mission level information. Tactical reasoning also requires some kind of model of the other members of the team, both human and robot, so that reasonable predictions of expected behavior can be made. Present systems do not take into account uncertainties in the observed world and the very large decision space in which reasoning occurs. They are forced to reason in a simplified world that does not match reality. Finally, their "thinking" is programmed rather than allowing for adaptation through learning.

Look. The second capability, focused SA, requires a semantic/cognitive description of the robot's environment that current systems do not have. At best, current systems have a map of static and

some dynamic obstacles to support navigation. SA also requires a sense of salience, what is important based on a shared understanding among teammates. This sense of salience is missing in the prevailing bottom-up approaches to autonomous perception that are not guided by context. Another critical missing element for effective SA is the ability to learn at a "deep" level, both offline and during operations. Better learning is needed to develop a more human-like hierarchical understanding of object categories in the first place as well as to refine perception capabilities in the field.

Talk. Existing robotic systems are notoriously opaque and distrusted. For example, they will change course or simply stop during a mission for no apparent reason. They cannot explain what they are doing, primarily because they do not have meta-cognition; in other words, they do not have a model of their own behavior. Current systems also lack the ability to understand human (i.e., semantic) communication of orders or other information. They correspondingly lack the ability to formulate semantic communication to Soldiers to explain what they are doing or ask for guidance.

Move. Safe, secure, and adaptive movement through a complex world is hampered by many technical barriers. First, current systems have insufficient descriptions, or models, of the world in which the robot is moving. They typically have a "green-yellow-red" map of mobility surfaces and possibly a kinematic list of movers. Existing systems struggle to distinguish a stationary person from a barrel or mailbox which represents very different challenges to safe and secure movement. Useful movement is also hampered by the lack of task or mission context so that a robot may persist in trying to reach a particular location that is not needed for the mission. Robots also need to be able to move in crowded and unpredictable environments, where existing algorithmic approaches are probably intractable but new learning approaches may work. They cannot yet adapt to mobility challenges from terrain, weather, etc. by adjusting their gait or form of locomotion.

Work. The above four capabilities (think-look-move-talk) largely enable the performance of the main goal of the mission – the "work" the robot is to do. The work most often involves direct physical interaction with the world: entering and searching a building or vehicle, loading and delivering supplies, inspecting a suspected IED, etc. This direct interaction with the physical world raises several important barriers. First, there is generally great uncertainty about the objects with which the robot is attempting to interact; for instance, exactly what and where are they? An object may be slippery or deformable. Also, the number of objects and the number of degrees-of-freedom of a mobile manipulator create a state space that is intractably large. Consequently, current approaches are almost entirely tele-operation based, with some attempts at supervised autonomy.

The barriers described above tend to impact multiple desired capabilities of future ground robotic systems. For example, the lack of effective semantic perception affects not only situation awareness but also the abilities to move safely and securely, to communicate about the world, and to interact with objects in the world. Similarly, the other barriers, while daunting, cut across multiple capabilities; therefore, overcoming each barrier results in multiple benefits.

Based upon the above discussion, we identify five primary cross-cutting technical barriers to achieving our vision:

Simplistic/Shallow World Model. Existing autonomous systems fall into two categories: either they have a world model that is at only a metric level, thus precluding any cognitive reasoning, or they have a model that exists at only a cognitive level without physical grounding in the metric world. Neither approach is sufficient for our vision where robots must behave cognitively while interacting in the physical world.

Lack of Semantic Understanding. In existing systems, objects in the world are perceived primarily or only as mobility regions, not as discrete objects of semantic and cognitive importance. Thus, one cannot tell a robot, "Go block the back door of this building" and expect it to do anything useful.

Scripted and Brittle Planning. Robots are almost always tele-operated or, at best, only perform simple scripted behaviors. Scripting all needed behaviors is not tractable and does not allow for learning new or alternative behaviors. Planning algorithms in robots work well only when the planning space is both small and certain enough, but the real world is fraught with uncertainty and high dimensionality. The inability to reason in complex and uncertain environments means that users must intervene frequently in robot operations and are trapped at a close level of "supervised autonomy."

No Shared Understanding of Missions and Roles. Robots now are opaque and distrusted and cannot explain what they are doing. Not only do they not know what they are doing, but also they do not understand what their teammates are doing or what the expectations for roles and communication are. Consequently, current systems must use tedious OCUs to bridge the enormous cognitive gap between humans and robots.

Missing or Shallow Learning Capability. Robots now must be explicitly programmed to do tasks, so producing the needed scope of behavior is intractable. Existing learning capability is shallow and lacks generalization. Thus, we cannot retrain robots without bringing engineers to the field or sending the robots back to the developer.

Figure 1-1 summarizes these barriers as columns and relates them to desired capabilities that are listed in rows. The fact that most of the table entries are filled-in demonstrates how the technical barriers impact many capability gaps.

	Barriers to Achieving our Vision -->	Simplistic and Shallow World Model	Mobilty-Focused Perception	Tele-operation or (at best) Scripted Planning	No Shared Understanding of Missions and Roles	Missing or Shallow Learning Capabilities
		World Model is either at only a metric level, precluding reasoning, or at only a cognitive level without physical grounding	Objects in the world are perceived primarily/only as mobility regions, not as discrete objects of semantic and cognitive importance	Bots are almost always tele-operated or at best only perform simple scripted behaviors -- and scripting all needed behaviors is not tractable	Bots are opaque and distrusted, and cannot explain what they are doing -- primarily because they don't know	Bots must be explicitly programed to do tasks, so it is intractable to produce the needed scope of behavior. Any learning capability is shallow and lacks generalization
"Think"	**Adaptive Tactical Reasoning**					
	Understand tasks, missions (METT-TC)	World model needs to represent concepts such as missions, tasks and generally METT-TC		Robots need to generate behaviors pertinent to achieving the mission, adapt to changing situation	Robots need to be able to follow instructions given at a semantic or cognitive level, not just "go to (x,y)"	
	Follow semantic instructions					
	Generate behaviors to achieve mission, adapting to changing situation					
	Understand teammates and what they need to know					
"Look"	**Focused Situational Awareness**					
	Maintain SA relevant to current task/mission	World model needs to represent, maintain, monitor and correct all info needed for SA	Robot needs to contribute to the general SA of the unit, noting salient observations		Robot needs to report on salient observations as needed to other elements of its unit	Robot should learn by comparing its observations and actions to those of its human counterparts
	Contribute to general SA of unit					
	Look for salient unforseen events					
	Observe and report on salient activity					
"Move"	**Safe, Secure and Adaptive Movement**					
	Move cognitively in relation to salient entities in the world (as people/dogs do) w/o GPS or other metric crutches	World model needs to store and operate upon all entities needed to relate movement to tactical constraints	Robot must perceive all entities in its environment relevant to safe, secure and adaptive movement	Robots must move in a tactically correct manner and react to changes in mission or circumstances		Robot needs to learn from its movement experience, whether from mobility challenges or tactical behavior
	Move in tactically and contextually relevant manner					
	Adjust to mobility challenges such as terrain, weather, barriers					
"Talk"	**Efficient Interactive Communication**					
	Receive and acknowledge semantic instructions	World model needs to have shared mental models as a basis for human-robot interaction	Robot needs to send and information relevant based on shared perception (common ground)		Roobot needs to receive and acknowledge cognitive-level instructions and similarly explain its own behavior	Robot needs to be able to learn through cognitive-level interaction with human teammates
	Explain own behavior					
	Report information relevant to mission					
	Seek guidance as needed					
"Work"	**Interaction With Physical World**					
	Inspect and manipulate objects	World model needs to represent wide variety of objects to be manipulated	Robot needs to perceive well enough to interact effectively with objects in a 3D world	Robot needs to figure out how and when to manipulate or transport objects as needed		Robot needs to learn from interaction with the physical world, e.g., when door is locked
	Transport objects as needed					
	Open doors, windows, hoods, trunks, etc					
	Use tools as needed					

Figure 1-1: Five primary barriers limit the capabilities of autonomous ground systems.

Fundamental Research Thrusts for Overcoming Technical Barriers

Derivation of Research Thrusts

Each of the technical barriers described in the preceding section has spawned one or two technical thrusts to address and overcome it.

1) To replace the existing shallow and simplistic world models, we are developing a hybrid cognitive/metric world model that is the foundation for much of the work of the RCTA program. This is the first thrust in the Intelligence technical area. The world model must simultaneously handle both cognitive constructs, such as missions on the one hand and the details of vehicle traction on the other. Thus, our architecture weds a top-down cognitive/deliberative framework to a bottom-up algorithmic/reactive framework and joins them in the middle via statistical reasoning to manage uncertainty. The world model must ultimately include most or all of the following elements:

- Hierarchical data store combining memory resident and traditional relational database management system (RDBMS) techniques to form a Hybrid Database.
- Bi-directional linking of metric and cognitive data.

- Long-term declarative memory which integrates, fuses, and infers from multi-sourced data (perception, semantic relations, *a priori*, training data, teammate knowledge) across time and space.
- End-to-end support for multi-hypothesis classification and reasoning.
- Prediction: how does the world change due to action X? Inclusion of time in planning cycles.
- Internal data structures to optimize memory and CPU usage to support task-based focus of attention.
- Merging of *a priori* GIS and sensed metric data.
- Resolution independent storage.
- Support and enable learning.
- Using Shared Mental Models to maintain context and state within a team while minimizing bandwidth.

2) To provide a much more complete description of the world in which a robot moves, we are developing a Perception approach that labels the environment semantically. This semantic labeling can then populate the world model at a cognitive level suitable to support higher-level reasoning. We have identified two related thrusts in semantic perception – one focusing on static scene understanding and the other on dynamic understanding and prediction.

3) To move beyond tele-operation and scripted behavior, we have three thrusts in adaptive behavior generation. All of these thrusts have a large focus on dealing with uncertainty. In the Intelligence technical area, the first thrust uses cognitive reasoning approaches to generate adaptive tactical behaviors like searching for an alternative method to enter a building if the first one fails. In the DMUM area, the second thrust enables adaptation to challenges encountered while interacting with the physical world; for example, changing gait in response to slippery conditions. A third thrust supports both high- and low-level planning: it seeks to overcome the barriers posed by a world that is fraught with uncertainty and complexity. When a high-DOF manipulator needs to grasp an object of uncertain size, shape, and position, the state space quickly becomes unreasonably large. This thrust seeks to bind the state space through better reasoning and sensing. Related work in world modeling tries to bind the problem using new representations of the state space.

4) To make robots more trusted partners, we have two thrusts in the area of transparency and meta-cognition. The intra-team cognition thrust, in the HRI area, develops shared mental models to provide common ground and a basis of higher trust. In the Intelligence area, the transparency thrust gives the robot self-knowledge and the ability for a two-way semantic communication with Soldiers.

5) To overcome the intractable problem of trying to program all needed behaviors, our learning thrust is aimed at designing in the ability to learn rather than attempting ad hoc learning after the fact. Key to our approach is the notion of deep learning – we want robots to learn at a conceptual/cognitive level, as humans do, rather than shallow, imitative learning where the true

lesson may be missed. This deep learning enables better generalization so the robot can more readily learn in novel situations.

Figure 1-2 summarizes the barriers we have identified as well as the new approaches we have taken to overcome them. The table also lists the research thrusts we have defined within the technical areas of Intelligence, Perception, Human-Robot Interaction, and Dexterous Manipulation and Unique Mobility.

Barrier	Research Thrust	Research Elements
1) Simplistic and Shallow World Model	Cognitive-to-Metric World Model	New Intelligence Framework (Intelligence)
2) Perception Limited to Generic Obstacle Detection	Semantic Perception using Bottom-up Context and Top-down Guidance	2a. Terrain and Object Classification, Identification and Reasoning (Perception)
		2b. Activity Detection and Understanding (Perception)
3) Scripted, Brittle Planning	Adaptive Behavior Generation From Reasoning with Uncertainty in Very High-Dimensional Cognitive and Metric State Spaces	3a. Cognitive Reasoning and Behavior Generation for Tactical Missions (Intelligence)
		3b. Behavior Generation for Manipulation (DMUM)
		3c. Adaptive and Unique Mobility Behavior Generation (DMUM)
4) Missing or Shallow Learning Capabilities	Deep Learning Coupled to Hierarchical World Model	Learning and Adaptation (Intelligence)
5) No Shared Understanding of Mission and Roles	Shared Mental Models Based on Cognitive World Model	5a. Intra-Team Cognition (HRI)
		5b. Common Ground for Shared SA (Perception)

Figure 1-2: We have defined new approaches for the five main barriers to autonomy, which lead to nine inter-related research thrusts.

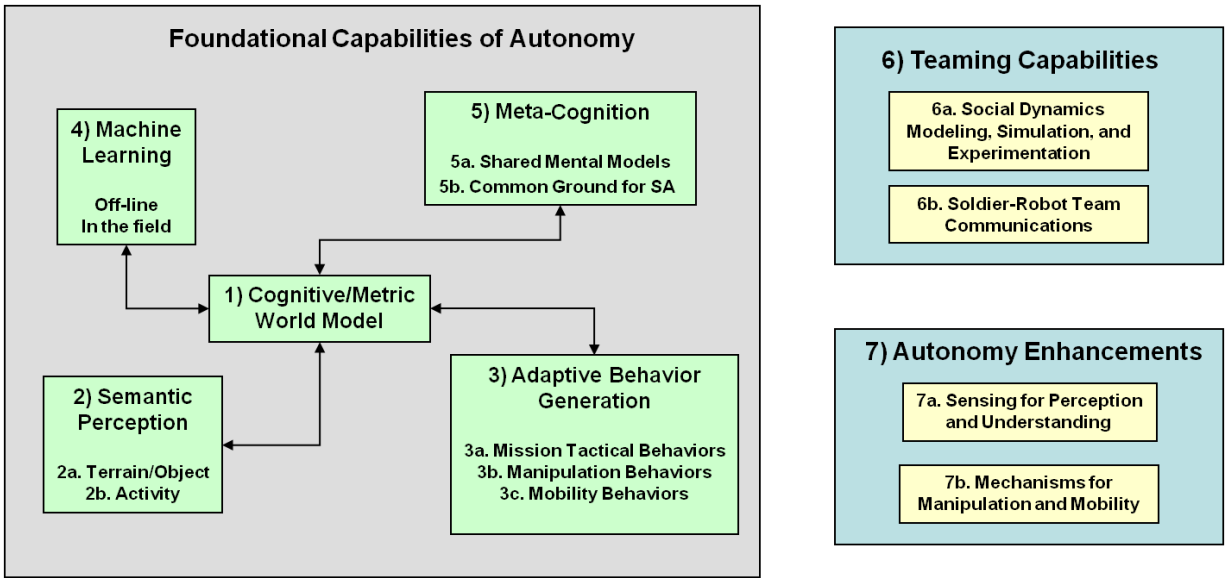

Figure 1-3: Five interrelated thrust areas (green boxes) represent the foundation of our autonomy vision. They correspond to eight technical thrusts in our Annual Program Plan. Four additional thrusts in the areas of teaming and enhancements for autonomy build on that foundation.

Overcoming the five barriers listed in Figure 1-2 will result in an entirely new level of autonomy consistent with our vision. To overcome these barriers, we have taken the technical thrusts corresponding to the green boxes in Figure 1-3. All of these approaches represent foundational capabilities – cognitive-to-metric world model, semantic perception, adaptive behavior generation, meta-cognition, and machine learning – which are essential to intelligent autonomy. Without semantic perception, for example, the world model cannot be populated with information to generate adaptive tactical behaviors. Without meta-cognition, the robot cannot understand and communicate its role in a mission. Moreover, without learning, its mistakes can only be corrected through intractable reprogramming.

To implement the five thrusts, we have defined ten research elements as shown in Figure 1-2. For semantic perception, there are two elements: one in terrain and object classification, identification, and reasoning and another in activity detection and understanding. There are three technical components of adaptive behavior generation: one each for mission-level tactical behaviors, for dexterous manipulation control, and for unique mobility planning. Meta-cognition includes both team cognition associated with shared mental models and transparency that arises from introspection.

The above five technical thrusts provide the foundational capabilities for individual autonomy and, thus, merit the highest priority in our program plan. However, we can more fully advance the state of the art in ground autonomy by pursuing two additional research areas – first in the area of Soldier-robot teaming and second in hardware enhancements for improved sensing, mobility and manipulation. Through the five thrusts, we have a solid foundation for autonomy. The additional areas build upon that foundation. Figure 1-4 illustrates the relative investment in

the five foundational autonomy areas, in teaming research, and in hardware research for autonomy enhancements.

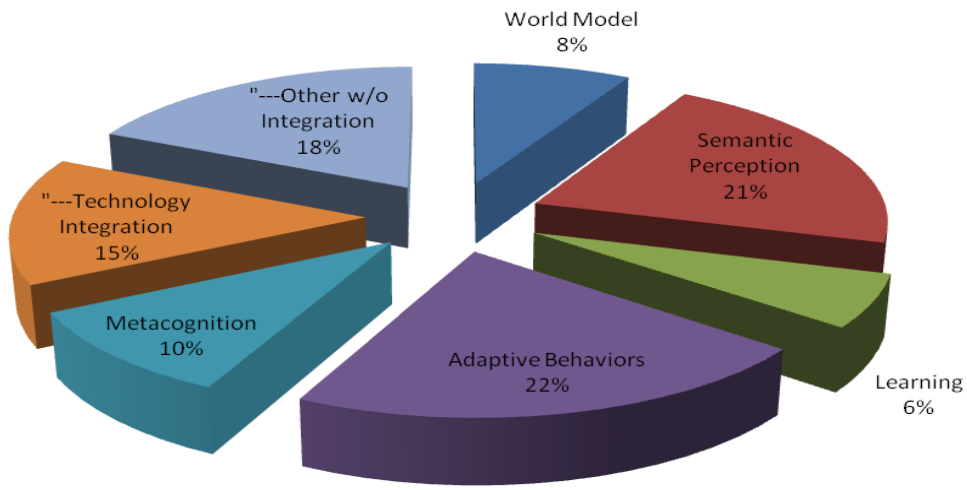

Figure 1-4: Our Annual Program Plan for 2012 invests primarily in the foundational building blocks for autonomy (67%).

All members of the Consortium are also making cost sharing contributions to RCTA. Through this cost share, they are contributing platforms, simulation software, in- kind research, workshops, seminars, and short courses for the benefit of the entire Alliance.

Through the teaming research area, we will capitalize on the benefits of autonomy as well as test and refine it. To achieve this teaming capability, we have identified two components as indicated by the yellow boxes along the right side of Figure 1-3:

- Social dynamics modeling, simulation, and experimentation
- Soldier-robot team communication

The second additional area enhances the foundational capabilities previously described in two specific ways. First, we pursue an effort in sensing focused on the specific needs of robotic perception to enhance semantic perception. Second, to enhance the algorithm-focused thrust in adaptive behavior generation for manipulation and mobility, we pursue applied research in the

associated mechanisms as well as basic research in a new generation of actuation materials and approaches.

Capability Area	Technical Approach	Research Elements
6) Soldier-Robot Teaming	Computational Models of Trust	6a. Social Dynamics Modeling, Simulation, and Experimentation (HRI)
	Communication Building on Shared Mental Models	6b. Soldier-Robot Team Communication (HRI)
7) Hardware Enhancements to Autonomy	All-weather, multi-spectral and proprioceptive sensing	7a. Sensing for Perception and Understanding (Perception)
	New forms of actuation based on smart materials	7b. Mechanisms for Manipulation and Mobility (DMUM)

Figure 1-5: We have identified key supporting research in collaboration, sensing, and actuation, which leads to four key supporting research thrusts.

Thus, our primary effort is represented by the five foundational capabilities (green boxes) in Figure 1-3 which correspond to the ten research elements of Figure 1-2. There are four supporting, but very important, research elements as listed in Figure 1-5. All of these thrusts are briefly summarized below. Those detailed discussions of technical research thrusts are organized according to the four technical domains that were previously identified: Intelligence, Perception, Human-Robot Interaction, and Dexterous Manipulation and Unique Mobility.

The third column of Figure 1-2 and Figure 1-5 identifies which of these technical domains corresponds to each of the technical research thrusts discussed below and in the following sections of this document.

Overview of Technical Research Thrusts

This section provides a brief overview of all technical research thrusts in our 2012 Robotics CTA Annual Program Plan. Detailed descriptions of all research, to the task and subtask levels, is given in Section 2 through Section 5 of this document. Figure 1-6 relates all of the tasks described in those sections to the research thrusts previously described. Figure 1-6 also lists the integration and assessment tasks, which are described in detail in Section 6. These tasks combine research outcomes into integrated autonomous capabilities; they also assess the integrated capabilities in order to measure our progress toward the goal of robotic autonomy.

These research thrusts are occurring at the basic (6.1) and applied (6.2) research levels. Technology is first being developed through the development of basic concepts (the Basic Research Element), then through initial implementation including laboratory experimentation in

simulation and live environments (the Applied Research Element). This is followed by integration onto testbed platforms for assessment in initial technology experiments conducted using relevant environments to examine the interplay between different technology elements and gain quantitative performance measures.

Mega Thrust	Research Area	Technical Thrust	Task	Page Reference
Cognitive/Metric World Model	Intelligence	Intelligence World Model	I1: Framework for Intelligence (6.1)	56
			I2: World Model and Applied Intelligence Architecture (6.2)	64
Semantic Perception	Perception	Static Understanding	P3: Static Scene Understanding (6.1)	178
			P4: Perception for Missions in Complex Environments (6.2)	201
		Dynamic Understanding	P5: Dynamic Scene Understanding (6.1)	221
			P6: Perception for Missions in Dynamic Environments (6.2)	232
Adaptive Behavior Generation	Intelligence	Adaptive Tactical Behaviors	I3: Combining Cognitive and Probabilistic Reasoning (6.1)	88
			I4: Generating Adaptable Tactical Behaviors (6.2)	98
	DMUM	Dexterous Manipulation (Planning)	M1: Theory and Prinicples of Mobile Manipulation (6.1)	348
			M2: Principles of Generalized Grasp Mechanics (6.1)	364
			M4: High Degree-of-Freedom Dynamic Manipulation (6.2)	382
		Unique Mobillity (Planning)	M5: Geometric Mechanics for Multi-modal Highly Articulated Systems (6.1)	388
			M6: Principles of Locomotion Mechanics (6.1)	401
Deep Learning	Intelligence	Learning	I5: Learning through Experience (6.1)	110
			I6: Life-long Learning (6.2)	121
Meta-Cognition	HRI	Intra-Team Cognition	H1: Shared Mental Models for Soldier-Robot (SR) Teaming (6.1)	255
			H2: Situation Awareness in Human-Robot Teams (6.2)	271
Teaming Capabilities	Intelligence	Collaborative Behaviors	I9: Distributed Intelligence for Human/Robot Teams (6.1)	130
	HRI	Multi-Modal Communications	H4: Communication Protocol and Language Processing (6.2)	283
			H5: Evaluating Tactical Command and Coordination Vocabulary and Protocols (6.1)	295
		Organizational Collaboration	H8: Social Dynamics Modeling and Simulation (6.2)	306
			H9: Social Dynamics Experimentation (6.1)	320
Autonomy Enhancements	Perception	Sensing	P1: Exploiting Novel Sensor Phenomenology (6.1)	151
			P2: Compact, High Performance Sensors (6.2)	163
	DMUM	Dexterous Manipulation	M3: Sensor-based Dexterous Manipulation (6.2)	373
		Unique Mobillity	M7: Learning Terrain Interactions (6.2)	410
			M8: Dynamic Multi-modal and Reconfigurable Mechanisms (6.2)	417
		Next Generation Actuation	M9: Next Generation Actuators and Materials (6.1)	432
N/A	Integrated Research	N/A	IR1: Integration, Experimental Design, and Assessment (6.1)	453
			IR2: Integration, Experimental Design, and Assessment (6.2)	456
			IR3: Platforms and Testbeds for Integrated Research (6.2)	479
			IR4: Modeling and Simulation Environment (6.2)	485
			IR5: Collaboration Software Integration (6.2)	494

Figure 1-6: The research summarized below is described in detail in Section 2 through Section 6 of this document.

Cognitive-to-Metric World Model (1)

The new cognitive-to-metric world model is at the heart of the RCTA program. As Figure 1-3 depicts, all other elements of autonomy interact closely with the world model and indeed must operate through it. To construct this new world model, we pursue a "best in class" approach to architecture that weds a top-down cognitive/deliberative framework to a bottom-up algorithmic/reactive framework and joins them in the middle via statistical reasoning to manage uncertainty. In the current state of the art, cognitive architectures have been used to implement sophisticated behaviors in simulation, but these simulators do not capture the difficulties of a real robot interacting with the real world. We investigate the use of a cognitive architecture to recognize and compensate for robot failures using a cognitive model and associated processes for the robot and the task. The idea is to produce a robot system that does not break as soon as the first unexpected event occurs; instead, it is able to recover and plan a workaround.

Our world model must also deal with the complexities introduced by the addition of legged mobility and whole-body manipulation, which add multiple degrees-of-freedom beyond traditional wheeled or tracked based systems, and with a focus on smaller platforms. We will investigate the appropriate algebra to simplify the computations required for these systems to interact with the environment. One potential method is to inject traditionally cognitive capabilities into the lower level dynamic planner to allow navigation using a context-based comprehension of the physical environment. This reflects a common theme through intelligence – while we believe in an overall hierarchical architecture to wed cognitive/metric/physical levels, we also feel gains can be made by applying combined techniques at the module level.

For robots to deeply understand environments, terrains, situations, and activities, they require a broad range of data types. We employ a world model strongly tied to the robot's architectural framework. The world model is more than a data repository, rather an adaptive process that answers questions about its aggregated information. It is a substrate for data from widely distributed sources and inferring greater SA.

The use of hierarchal world models within robotics is well proven, yet implementations are constrained to the needs of sub-specialties of robotics or for the support of specific algorithms. Instead, we seek a cross-discipline world model capable of storing and indexing metric, semantic, and cognitive information. The world model should support multiple capabilities for the robot from navigation through manipulation and adaptive behaviors supplied by semantic knowledge, learning, and cognitive control. We will investigate applying state of the art concepts developed in the database and machine learning communities to enable a query-based world model, feeding historically separate algorithms from a common probabilistic data store. Assessment of this concept involves the applicability of this mechanism throughout the intelligence architecture: providing focus of attention to perception, feeding world changes to cognitive agents, and enabling algorithms to move beyond traditional 2D raster map-style representations into volume and depth-based 3D.

Semantic Perception (2)

The world model described above must be populated with information about the physical world in which the robot resides. In current systems, this information is very high in data content but very low in information content – the robot is often given massive sets of data such as pixels, point clouds, and radar returns, while it actually needs much more compact information about what *objects* are in the world around it and what they are *doing*. It is the role of Semantic Perception, aided by the cognitive level of the world model, to provide this needed information at the appropriate level of abstraction so other elements of the architecture can reason about them. Cognitive guidance via the world model is critical since unaided state of the art perception algorithms appear to be approaching their limit of performance. We divide the semantic perception effort into components, one for the understanding of static entities and the other for understanding activity. Supporting both components, we pursue a third, smaller focused effort to improve robotic sensing.

Terrain and Object Classification, Identification, and Reasoning (2a)
In order to reason beyond safe driving, robots must have a detailed understanding of the world, including a description of objects, material, and other salient features in their environment. In addition to "naming" the entities in the scene, perception should also derive qualifiers (e.g., parked car, occluded wall) and relations between scene parts (e.g., car in front of door) from sensor data. This level of scene understanding remains a challenging problem that has focused the attention of the computer vision community in the past decade. In particular, much progress has been made in object recognition. But since these techniques are still brittle, we bring into play both learning and the use of context.

One major objective of this thrust is to design efficient learning and recognition algorithms through the use of techniques like deep inference that replace intractable global optimization problems commonly used in state of the art vision systems by approximations that are considerably more efficient and which, when integrated, achieve similar or better recognition accuracy.

Another approach is to use the constraints induced by operation in a particular type of environment. For example, we can use the known context of an urban setting to generate accurate representations of the environment from sensor data. Here, we will incorporate facade detection algorithms and combine the urban scene analysis tools with motion and visibility cues. Similar context constraints can be used in other settings such as wooded terrain or indoor settings. Major advances in scene parsing, scene surface layout analysis, and 3D reconstruction will be combined and leveraged in a uniform framework to advance the state of the art in overall scene understanding.

Activity Detection and Understanding (2b)
A major objective of the second semantic perception thrust is to investigate robust approaches to detecting, tracking, and identifying objects in general configuration, to use the resulting intermediate description to identify behaviors of individuals and groups, and to predict

distributions of likely behaviors based on learned models. As part of this thrust, we also develop algorithms for mid- and long-range perception, situational awareness, and recognition of posture, body parts and gestures.

While vast literature exists in the current state of the art for short-term prediction based on classical temporal filtering techniques, longer-term prediction of motions and behaviors remains open, and recognition of behavioral cues has been confined to well-structured environments.

The other key factor contributing to degraded performance of current state of the art detection and tracking algorithms is that typically environments are densely cluttered, thus causing the perception system to lose track of the objects as they move through extended occlusions. We plan to pursue two approaches to address this difficult problem. The first approach is an application of our general "purposive prediction" model; the second approach is based on forward-simulating the motion of tracked people.

Another area of research of this thrust is in pose estimation for behavior understanding. Recent work in human pose estimation has focused on two different domains. Historically, researchers have worked in laboratory environments, which make a variety of simplifying assumptions to sidestep the difficulty of detecting human body parts in natural images: using multiple cameras or active sensors, a known background, a limited known range of poses, or even motion capture markers. In these simplified settings, researchers have enjoyed success in estimating and tracking 3D locations of joints and classifying basic actions with relatively high accuracy. However, many of the techniques do not translate to our setting for mobile robots in complex and dynamic environments. We propose to handle these computational barriers by using a cascade of tractable models which successfully filter-out more and more unlikely pose configurations, allowing focus of computation resources on the most likely models.

Also, as part of this thrust, a subtask will be recognizing a small set of actions from "cooperative" subjects. The output of this task will be used in two ways: 1) action recognition for the purpose of communicating with the robot, and 2) to provide data for generalizing our prediction models to sequences of actions.

Adaptive Behavior Generation (3)

Based on the first two foundational approaches, cognitive-to-metric world model and semantic perception, we posit that we can produce a robot that has the ability to perceive and understand its environment and a world model in which to save and analyze that information over time. However, in order to be useful, the robot must be able to *do* something. Thus, we need a foundational capability for adaptive behavior generation. There are three fundamental types of behavior we wish the robot to possess:

- *Tactical mission behaviors* – this is the macro platform level of planning or behavior generation where the robot decides "what to do next" based on the current circumstances, i.e., METT-TC. Such decisions include changing sensor modes and platform location to

provide better overwatch, adjusting speed and course to avoid a pedestrian, and sending a message with important surveillance results.

- *Dexterous manipulation behaviors* – this is the level of behavior where the robot interacts directly with the physical world in order to perform tasks such as picking-up and carrying objects, digging a hole, using tools, or opening a door.
- *Unique mobility behaviors* – this is the level of behavior relating to how the platform interacts directly with the physical world in order to achieve the higher-level mission goals. The robot may, for example, need to change its locomotion in order to climb a wall, navigate rough terrain, or simply get through the mud.

These three behaviors share planning paradigms at an abstract level. They all involve continuously deciding "what to do next" but in differing problem domains, timescales, and scale sizes. Therefore, we group them together conceptually; we anticipate much synergy among the efforts but also recognize the differences. Our fundamental approach is to focus upon the algorithmic aspects of behavior generation while using available approaches to realizing those behaviors through conventional actuation mechanisms.

Each behavior generation area is described in more detail below:

<u>Tactical Mission Behaviors (3a)</u>
To build robust tactical behaviors, we combine knowledge-intensive approaches from cognitive architectures with algorithmic approaches from traditional robotics. To bring the two disciplines together, we quantify, manage, and reduce inherent model uncertainty that arises in the interface between them. Typically, in modern robotic systems, planning is always done under the most likely hypothesis as uncertainty in perception estimates is massive, while decision-theoretic planning with uncertainty is computationally hard. We will develop principal approaches to planning under multiple hypotheses given by perception to support planning with uncertainty at the algorithmic and cognitive levels. Other methods to reduce uncertainty come from defining how the robot interacts with its environment. We devise *policies* for a robot to interact with an unpredictable human behavior to result in a desired outcome. Progress will be assessed via theoretical analysis of the performance guarantees we can provide as well as experimentation both in simulation and on a physical platform. We will test how our approaches to handling, reducing, and reasoning over uncertainty allow for a more robust behavior that focuses perception efforts and produces actions in such a way as to provide a higher level of robustness.

<u>Dexterous Manipulation Behaviors (3b)</u>
We seek a formal understanding and framework that will enable mobile systems to perform highly dexterous manipulation operations, including planning, perception, control, and user interaction. A goal of this research is to provide a foundation for whole-body manipulation, where manipulation moves beyond traditional arms to include the abilities of the platform itself to assist with the task. For example, with the BigDog platform, we achieve added "lift capability" by using the momentum of the arm to increase the impulse delivered to the object being moved, much like the wind-up of a pitcher increases the ultimate velocity of a baseball thrown at a batter. The state of the art provides examples such as the manipulator twisting a door

knob, and the platform reverses to pull open the door. To begin moving toward this goal, we investigate abstractions and algorithms for reduced dimensional planning for coordinated mobile manipulation and workspace exploration.

Unique Mobility Behaviors (3c)
The next generation of mobile robots must be able to maneuver through complex three-dimensional environments, including urban, mountain, jungle, and riverine terrain. This requires platforms to go beyond reliance on a single, fixed mobility configuration by exploiting multiple, innovative methods of locomotion. While some platforms of this nature exist in the state of the art, they fall short of their true potential due to a lack of intelligent motion planning and control. We apply learning techniques to current gait models to improve and discover new gaits during online execution of "fixed topology" robots. We also investigate the opportunities and challenges of reconfigurable robots, such as snakes, whose topology may be changed at runtime to meet task goals. The current state of the art focuses upon solitary gaits, allowing us to move beyond by creating principles for adapting gaits to variable non-steady terrain as well as to transition between horizontal and vertical gaits.

Machine Learning (4)

The technical approaches described thus far provide the basis for a highly capable robot to perceive the world, reason about it, and take useful actions. However, the environment in which the robot operates is full of surprises, and the robot must be able to learn from them rather than repeating its mistakes. To plan explicitly for all possibilities is intractable, so we plan to program what is well understood and learn that which is not. Thus, we have defined a research thrust in learning. The thrust emphasizes deep learning in order to deal with the difficulties of the assignment problem often encountered in conventional learning.

We hypothesize that we can efficiently grow intelligence through intensive training with domain experts (e.g., Soldiers rather than researchers or engineers) and learning from other "smarter" robots, creating an intelligence that continues to expand and grow through new and diverse experiences. To enable this, we develop methods to map gestures and vocal instructions into actions that can be understood by the robot. The techniques that lead to good performance in deep systems of learned modules remain poorly understood, and we will investigate methods such as "no-regret learning" and the use of boosting to leverage existing supervised learning algorithms to handle more complex problems. We investigate generalizations of data sharing between different robots and sensor systems to enable robots to automatically tune themselves.

We perform extensive comparisons on a range of complex tasks that compare performance using fully specified, human programmed implementation of tasks within our system; pervasive learning using only local signals; end-to-end training of the system including using imitation learning to improve performance; and full end-to-end training leveraging both supervised training and self-supervision. We test the speed-up of planning on tasks from complex motion control to cognitive decision making to establish the benefits of enabling adaptiveness in planning.

Meta-Cognition (5)

With learning added to perceiving, reasoning, and acting, the set of foundational capabilities for autonomy is nearly complete. However, we must still lay the groundwork for teaming through technical thrusts in team cognition, transparency, and common ground. We can build additional teaming capabilities with this foundation.

Shared Mental Models (SMMs) (5a)
In this thrust, we leverage expertise in team research to expand the science of teamwork to the realm of Soldier-Robot Teams. The tasks within this thrust provide data and procedures to give robots the basic cognitive functionality required for effective dynamic collaboration with humans. Our approach involves examining the creation, elicitation, and combination of mental models for both humans and robotic agents to understand the complexity of knowledge required to combine teammates into one cohesive unit. This research will be conducted in close collaboration with Intelligence in determining the best representation for shared mental models of team structure, populating the SMMs, and dynamically maintaining these models during mission execution. To facilitate this collaboration, key members of research staff from CMU and UCF are members of both the Intelligence and HRI research teams. In the first six months of the project, we have developed a close cooperation between these, specifically with respect to our task H1: Team Shared Mental Models.

Common Ground for Shared Situation Awareness (5b)
The primary objective of this thrust is to create representations of space that will enable communication between humans and robots. Metric maps, such as occupancy grids or coordinates of low-level features, are commonly used. However, more intuitive and more compact representations are required for communication between human and robot team members. This thrust includes collaboration work with ARL. To enable communication, we need to determine a network of traversable space (similar to road networks) plus semantic names for the landmarks to be labeled. Such landmarks are necessary both for a change of mobility (like stairs) and manipulation (elevator door) and for verification of route instructions ("turn right when you pass the restroom").

Teaming Capabilities (6)

The current generation of robots essentially considers the external world, including humans, vehicles, and other robots, as navigational issues rather than as team members, opponents, or part of the ambient culture. In contrast, the RCTA vision for future robotic systems calls for highly effective Soldier-Robot Teams where each part of the team understands the roles, responsibilities, and required actions of the others and has the capability to provide the communication necessary to make the team successful. Accomplishing this within a mission context, accepted military doctrine, and social norms of the society in which the Soldier-Robot Teams operate will be a major technical challenge but will provide a quantum leap in effectiveness and capability.

Our HRI research also seeks to enable the future robotic team members to communicate with their human teammates using multiple simultaneous communication modalities. This approach allows for (a) integrated situation awareness (SA), world model development, and command interpretation; (b) more natural interaction and communications redundancy; and (c) awareness of human team member affect. More revolutionary, though, is our approach to integrating robots into the Soldier team structure, into social structures, and into societies. We propose innovations to achieve this futuristic vision by drawing on successful research in human team behavior, human factors, live/virtual/constructive (LVC) simulation, computer science, and neuroscience.

Social Dynamics Modeling, Simulation and Experimentation (6a)
In this thrust, we are examining multi-level collaboration issues with research aimed at social, organizational, and cultural factors that are required so that robots can be collaborating partners within Soldier-Robot Teams that operate in the real world. This research is fundamental to making it possible for robots to function effectively within human social situations. This thrust addresses three levels of social interaction in separate tasks: within a team, within a social environment, and within a culture. The best form for representing this information is very much a research issue and, as in thrust 1, will require close collaboration with related tasks within Intelligence to be effective.

Soldier-Robot Team Communications (6b)
In the second thrust, we will apply our understanding of how communication unfolds (explicitly and implicitly) in dynamic team contexts. Soldier-Robot Team communication in the dynamic team context poses unique challenges; our research seeks to overcome these challenges by taking advantage of all available modalities to both scaffold and augment communication. Our approach includes focused efforts in each prominent modality and a dedicated 6.2 effort into methods for test and integration of multi-modal communication in live and virtual environments. The ultimate goal is to facilitate collaboration between humans and robots at multiple levels. We are addressing the distinctive and complex issues created in socially, organizationally, and culturally charged situations across a series of inter-related tasks. The tasks within this thrust examine both explicit and implicit communication modalities expected to provide effective team communication in ways that are intuitive or intrinsic to humans. The focus on multi-modal and redundant modalities is expected to help overcome interaction problems intrinsic to acoustically noisy and visually challenging situations and to ensure the dynamic bi-directional communication necessary to realize shared team awareness. We include in this research an emphasis on extending modalities that traditionally have been considered line-of-sight (LOS) to be useful even when out of visual contact.

Autonomy Enhancements (7)

In addition to the Foundational Capabilities of Autonomy (1 through 5 in Figure 1-3) and the Teaming Capabilities (6), we have identified a set of needed enhancements to autonomy. While these are not foundational, they extend unmanned system capabilities in three areas:

- Sensing for perception and understanding
- Mechanisms for manipulation and mobility
- Scaling issues for autonomy

Each of these enhancements to autonomy is discussed below.

Sensing for Perception and Understanding (7a)
Closely coupled to the above two tasks of understanding the world through the processing of sensor data is the generation of that data through the sensing process. Research into new sensing technologies has consumed major resources for many decades, and therefore, this program primarily uses the outcomes of previous and ongoing work in sensor research. However, we believe a small amount of focused research in sensing specifically aimed at robotic autonomy is warranted. Thus, we have targeted efforts to improve LADAR (active), stereo (passive), and contact sensing.

In active sensing, we will extend the development of Spectral LADAR, which allows object recognition using spectral signatures in addition to 3D spatial information. In this effort, we build upon a large body of prior work in multi-spectral sensing which is able to identify material types represented in single pixels and, thus, provide shape-independent object classification.

Another area of this thrust is in the use of smart materials systems to incorporate microarrays of pressure sensors into the “feet” of legged robots and in robot grippers as well as micro-electromechanical systems to design embedded temperature and moisture sensors for terrain classification and to integrate these with vision sensors to classify terrain for a variety of legged and wheeled vehicles.

Also, as part of this thrust, we are collaborating with ARL researchers in the areas of MEMS LADAR, acoustics, and very small radars. The MEMS LADAR is capable of real-time 3D images with high fidelity at frame rates and ranges suitable for SUGV application. The acoustic work includes audio classification.

Mechanisms for Dexterous Manipulation and Unique Mobility (7b)
We have previously described foundational thrusts in low-level control and behavior generation for both dexterous manipulation and unique mobility. In order to better realize those behaviors, we have defined a mechanisms thrust that encompasses three efforts: in mechanisms for manipulation, in mechanisms for mobility, and in next generation actuation approaches for both mobility and manipulation.

We explore manipulation systems that take advantage of increased range of motion afforded by a highly articulated mobility base in pursuit of whole-body manipulation. Using an existing and readily available experimental platform, the DARPA BigDog, we apply the research in basic manipulation behavior in order to seek a control system that can perform useful manipulation tasks with a high degree-of-freedom system. These tasks, such as pushing, pulling, lifting, and throwing, are beyond the state of the art. We follow two paths to accomplish this. First, we

pursue "legipulation," a bio-inspired approach where a legged system uses one of its legs to manipulate the environment. Second, we will integrate a preliminary manipulator arm onto a BigDog to investigate behaviors with full-body articulation while maintaining balance. We assess these capabilities using success rate and speed of task completion.

To investigate mechanisms for unique mobility, we take advantage of the Canid platform under development though a collaboration between UPenn and ARL. Building upon these topics, we investigate a hybrid framework for selecting, mixing, and transitioning between gaits at runtime. These topics are then assessed using metrics of locomotion power efficiency, velocity, efficiency per unit weight, and grasp versus release force.

Finally, we investigate potentially revolutionary approaches for next generation actuation of manipulation and mobility mechanism. Actuators based on smart materials have great potential to transform robotic systems by improving the strength-to-weight ratio, speed, range of motion, compactness, efficiency, controllability, and reliability of manipulators, legs, sensors, and other robotic components. The state of the art technologies with respect to robotics include ferroelectric materials, magnetostrictive compounds, shape memory alloys, and dielectric elastomers. These materials have constraints that require special attention during the design process; therefore, synergistic research on material characterization, model development, and development for robotic platforms is critical. To overcome these challenges, we develop active materials for legged robotic platforms. With electro-active elastomeric materials, we expect to enable robotic limbs that can change their shape, stiffness, and potentially viscoelasticity with an applied electric field. We will investigate how electrically and thermally activated materials can be utilized to create passive mechanical joints with variable stiffness and damping. Smart structures utilize shape-changing actuation, simultaneous sensing, and real-time material property control for dynamic adaptation and superior maneuverability. For each material investigated, we will quantify the relevant materials properties and their response to stimulation and compare their capabilities to standard actuation techniques. We will also assess difficulties in and progress toward integration into robotic structures

Assessment of Integrated Research

Our plan for measuring progress toward autonomy over time is built around our integration and assessment (I&A) plan. As a collaborative fundamental research effort, the Robotics CTA's assessment process differs from that of traditional system development efforts. Instead of building a system to meet a particular performance specification, we deliberately undertake high-risk basic and applied research that may ultimately result in breakthrough technologies. Instead of managing development to meet pre-defined goals, we assess our research against performance benchmarks to evaluate how well that research stands to enhance or even revolutionize robotics and related disciplines.

Our integration and assessment is performed by a team consisting of research integrators and support staff, an assessment team, and I&A management. The I&A management team consists of the four Technical Area Leads, the Integration Lead, and the ARL I&A Lead. The I&A

management team plans integration events, defines assessment methodologies, and conducts quantitative assessments. The assessment team prides quantitative and objective results based on accepted experimental practices.

Our assessment protocol is based on a two-stage approach. The first stage consists of task- and/or subtask-level assessments of stand-alone research outcomes. These assessments are typically conducted by the researchers, but they are reported to and monitored by the I&A team. Task and subtask assessments are provided in quarterly reports and at other times as requested. An example of an individual assessment is the measurement of precision/recall performance for object or activity detection on a given dataset. Results of such individual task assessments will help determine which research outcomes are ready for integration. Task-level research that is producing demonstrated results beyond the state of the art creates a “push” to be included in the integrated research described below. As described in Section 6, we also define integrated capability goals each year that reflect expected outcomes. Thus, we also create a “pull” to set an expectation for research outcomes. The I&A management team considers both push and pull to decide which research outcomes are suitable for integration and assessment.

The second stage of assessments focuses on integrated capabilities that result from bringing together results from multiple research tasks. At this second stage of assessment, we conduct a series of experiments which we call integrated research assessments (IRAs). Each IRA combines two or three outcomes from research thrusts to achieve a capability from the think-look-talk-move-work spectrum. Similar assessments will be repeated and extended over time to provide regression testing and integrate improving technologies. The assessments involve formal experimental design in collaboration with the Government with reported results. They include elements of both modeling and simulation and laboratory experiments appropriate for the assessment of basic and applied research outcomes. Two IRAs are planned for each of the years 2012 - 2014 and are described in more detail in Section 6. The exact timing of the IRAs will depend upon the status of research outcomes.

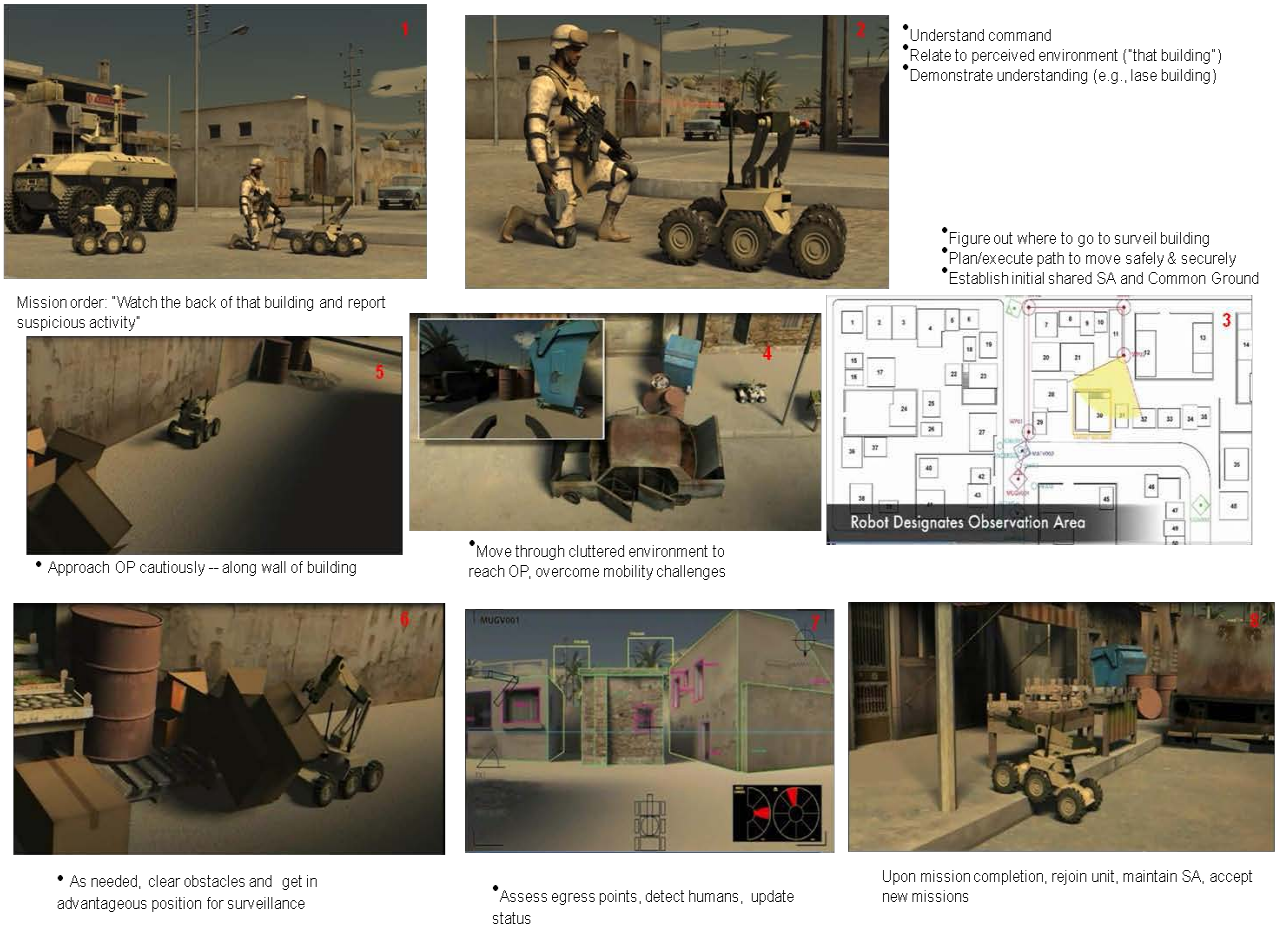

Figure 1-7: Our planned Capstone assessment will test key enabling capabilities of autonomy – adaptive tactical reasoning, situation awareness, efficient communication, robust movement, and interaction with the physical world.

In order to better focus RCTA research and experimentation, we have developed a vision for a Capstone Experiment to be conducted toward the end of calendar year 2014. The narrative components of the experiment are exemplary of capabilities needed across a wide range of missions. This experiment is illustrated graphically in Figure 1-7 and described in some detail in Section 6. Briefly, it is centered around a notional cordon and search operation: during urban transit by a small (4-5 Soldier) unit, a fugitive is reported to have entered a building the unit is approaching. A man-transportable robot is instructed to "cover the back door" of the building by the unit commander since he cannot safely split up his limited resources (see Figure 1-7, panel [1]. The robot must understand and acknowledge the order [2], associate the order with its perceived environment [3], move safely and securely [4] to an appropriate vantage point [5], observe activity behind the building and report any salient events to the unit commander [6]. As needed, it enters the building and negotiates stairs or other mobility obstacles [7]. It then returns to its unit, maintaining SA, and is ready for another assignment [8]. Again, while this narrative

occurs in the context of a cordon and search operation, its underlying capabilities support a broad range of potential operational missions.

As an illustrative example of an integrated research assessment, we plan to combine elements from three research thrusts into a basic "look-think-talk" capability that enables the building surveillance portion of the capstone narrative. In this "Autonomous ISR" assessment, a robot observes a series of actions. Its semantic perception capability populates the cognitive level of the world model with information about observed activity. This perception uses state of the art "bottom-up" techniques but is guided by contextual information from the world model's long- and short-term memories. In this case, the system may observe a person exiting the back of the building. Based on a shared mental model of what the current mission is, the system identifies one or more activities as salient events. In this case, the person may just be a bystander or may represent a threat, depending on context and prior knowledge as well as the immediate visual evidence. If an event is important enough, it triggers a number of possible actions, such as making a report, moving to gain a better vantage point, taking evasive action, etc. This integrated assessment combines semantic perception, shared mental models, and adaptive behavior generation, all of which are mediated by the new world model. Of course, the first integration of these capabilities is likely to reveal shortcomings. Thus, we will subsequently add improved capabilities as well as assess more complex situations.

The IRA example just described is one of a set of planned assessments leading up to the Capstone Experiment in late 2014, as shown in Figure 1-8. The integration and assessment events build capabilities across the "think-look-move-talk-work" spectrum. Each assessment focuses on an integrated capability, such as intelligent navigation, autonomous ISR, or physical interaction (e.g., manipulation of objects). The columns of Figure 1-8 highlight the buildup in each of the capabilities, with a growth toward higher levels of behavior and more challenging environments.

IRA ID	Integrated Research Assessment Name	Timeframe	Think	Look	Move	Talk	Work
0	2011 Baseline: Small Robot Autonomous Navigation	Aug. 2011	Legacy Waypoint Planning Algorithms	Hokoyu LADAR	Stationary & Moving Obstacle Avoidance		
1	2011A: Recognition of Salient Behavior	Dec. 2011	Cognitive Architecture (ACT-R) on Small Robotic Platform	Stereo Tracking on Small Robot Feeds Cognitive Reasoning			
2	2011B: Physical Interaction Baseline	Apr. 2012	Include WM Entities/Descriptors to Enable Interaction	Close-in 3D Sensing for Manipulation	Autonomous Stair/Hill Climbing, Crawling		Simple Autonomous Trenching, Caging & Manipulation
3	2012A: Intelligent Navigation	Oct. 2012	World Model with Short/Long-Term Memory Linked to ACT-R	Semantic Labels and Tracked Movers to WM, Using Stereo & LADAR	Intelligent Movement to Described Locations	Accept & Understand Movement Orders, Provide Status	
4	2012B: Physical Interaction	Apr. 2013	Reason about Moveable Objects	Perceive Wide Range of Mobility Challenges	Adaptive Locomotion, Quadruped & Hexapod Mobility, Snakes	Supervisory Communication for Manipulation	Complex Autonomous Trenching, Caging & Manipulation
5	2013A: Autonomous ISR	Oct. 2013	Reason about Adversary Actions to Provide Prediction & Context	Perceive Apertures, Transit Routes for Fugitives Maintain	Move to Good Vantage Point(s) for Observation	Update Teammate(s) wrt Ongoing Mission	Recognize & Report Salient Events
6	2013B: Tactical Team Movement	Apr. 2014	Reason about Troop, Adversary and Neutral Movement Reason	SA and Common Ground with Soldiers Monitor	Move with Soldiers in Correct Tactical Positions	Communicate only as Needed, Accept New Orders	Move with Team, Accept New Mission on the Fly ("Follow that Guy!")
7	2014A: Capstone Assessment	Oct. 2014	and Adapt through Stages of Mission	Egress Points and Recognize Mission Relevant Activity	Move Adaptively to Evolving Mission and Mobility Challenges	Update Status & Communicate Salient Information	Report Events, Move Obstacles, Track Fugitives, Seek Guidance

Figure 1-8: A series of Integrated Research Assessments (IRAs) systematically tests autonomous capabilities and culminates with a Capstone Experiment in late 2014.

In our assessments, we apply established principles of scientific experimentation. The assessment plan includes the platforms, sensors, human participants, and simulation tools used, along with a specification of datasets to be collected. The datasets will include data to be sequestered as well as data for possible posting to the broader research community.

For the "Autonomous ISR" example discussed previously, the overall experimental hypothesis is that an autonomous sensing system can observe a wide range of human activity, recognize a subset of that activity as salient to the mission and conditions at hand, and then report those salient observations in a human-understandable form such as, "A man in a green coat just left the safe house." If that activity actually occurs and the system reports it correctly, that constitutes a true positive. If the system does not report salient activity, we have a false negative. In this example, the human activity which is sensed constitutes the independent variable, while the reported message is the dependent variable.

In addition to the end-to-end performance, we will also examine the individual components of the experimental system. In this case, the components are:

- Activity detection
- Salience assessment
- Activity reporting

For the integrated end-to-end capability to work correctly, the activity of a man departing a building must be detected, that activity must be correctly assessed as salient to the mission, and a

suitable report must be constructed. If the overall system fails in a given case, we need to understand what the cause of failure was. Thus, we will separately assess each of these capabilities, which can be tested both in an integrated form and individually. For example, we can assess the activity reporting module by providing many examples of output that could come from the prior stages and evaluating how well it constructs messages that are both accurate and readily understood by humans. Similarly, the salience assessment module can be fed a variety of activity detections across a variety of mission contexts. A given activity may be salient in one context but not another.

This type of integrated assessment is quite challenging because we are conducting fundamental research: we seek to create capabilities that do not exist, rather than simply making incremental improvements to existing capabilities.

Progress During 2011

The RCTA program is integrating technologies to produce basic skills of autonomy, and then determining progress through a series of assessments as described previously. The technologies we are integrating are being developed by researchers organized into four functional groups – intelligence, perception, human-robot interaction (HRI), and dexterous manipulation and unique mobility (DMUM). These following sections of this document describe the research within each of these technical areas in great detail. Briefly:

- Intelligence is building the cognitive-to-metric world model, and providing the basis for adaptive behavior generation and learning
- Perception is developing the algorithms that process sensor data to build a semantic understanding of the robot's environment, including both static and dynamic entities
- HRI harnesses both Intelligence and Perception to enable a common ground between the robot and Soldiers
- DMUM is developing the capabilities for robots to interact effectively with their physical environment through manipulation of objects and traversal of challenging terrain

Substantial progress has been made in all areas during the past year. That progress is summarized below by technical area.

Intelligence

During the past year we have made significant progress in the three major thrust areas of Intelligence that are critical to achieving the capstone vision– the new world model, adaptive behavior generation, and learning. In the world model area, we developed a revision to the architectural framework that is suitable for handling the combination of navigation with manipulation for a variety of platforms, able to reason about and manage uncertainty, benefit from human training and experience, and support communication with human teammates. We also achieved the first integration of a cognitive module built around ACT-R with a traditional

perception model for human recognition and tracking to perform a surveillance task. And we developed a first version of shared mental model and updated it automatically based on robot observations of human actions.

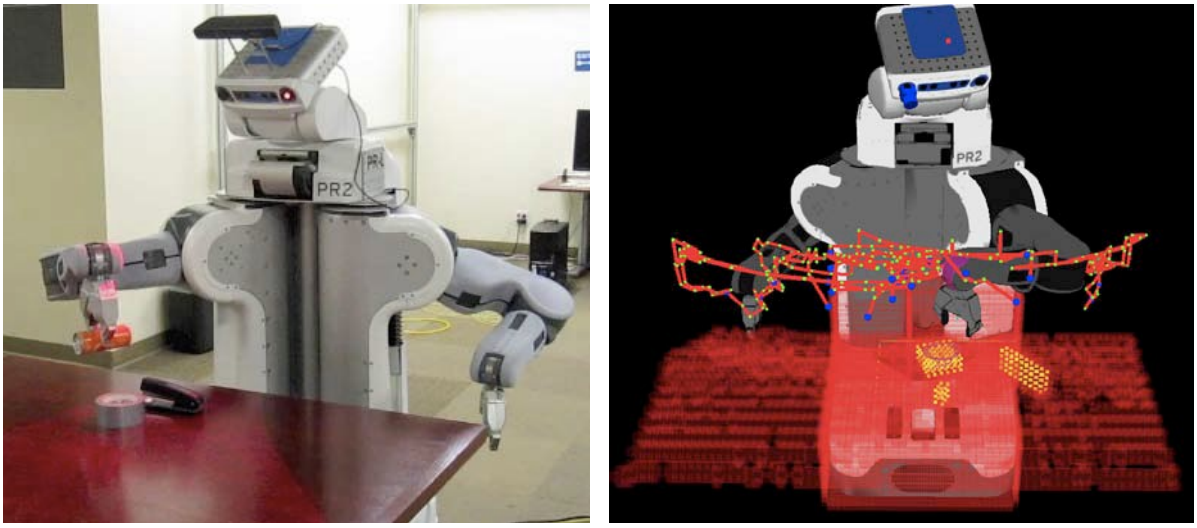

Figure 1-9: Learning to Plan on PR2 (our ongoing efforts in I4-2).

In the adaptive behavior generation thrust, we developed two planning approaches for belief spaces, one that manages uncertainty and the other that reduces it, for maximizing the probability of plan success. We also developed two planning approaches for dealing with high-dimensional spaces, one that learns to plan faster over time by reusing old plans, and the other that reduces the problem to a low-dimensional space. And we developed mapping algorithms for single and multiple robots that enable distributed search, exploration, and mapping using entropy minimization techniques. Figure 1-9 shows an example of behavior generation research on the PR2 platform which has many degrees-of-freedom and thus poses a very high-dimensional planning challenge.

The Intelligence learning thrust is developing a toolbox that will be employed throughout the program. To date it has been used heavily in semantic perception to segment and label static scenes. We have also developed learning technique to parse from natural language directions to a formal robot control language. The parser enables a robot to interpret human commands and generate corresponding control programs that contain complex statements involving counting and while loops. And we have developed online learning techniques for self-calibrating stochastic vehicle models for predicting tire/soil interaction.

Other key Intelligence accomplishments include the development of algorithms for constructing geometrically consistent colorized range maps for use in the world model, and for semantic labeling of objects/terrains using acoustic data. Finally, we developed metacognitive techniques that provide robust symbol grounding through deep integration of cognitive and perceptual process.

Perception

As an important step toward assessment and deployment of state of the art techniques in the area of understanding of static environments, we have developed algorithms for classification from both images and 3D data, showing competitive performance on extended datasets in urban areas using different 3D sensors. In addition to environment interpretation at long range for outdoor operation, we also developed algorithms for parsing RGBD images of indoor scenes, as input to a higher level interpretation scheme which reasons about the overall floor plan of an indoor scene. To enable localization in human created semantic maps we designed and tested a new approach for object based localization which does not rely on accurate decisions on object locations. These achievements are important contributions to the elements of the capstone vision concerned with building, sharing, and navigating a semantic map from sensor data and human input. Figure 1-10 shows recent results from our semantic perception research, where imagery is accurately labeled with up to 30 classes of objects and terrain.

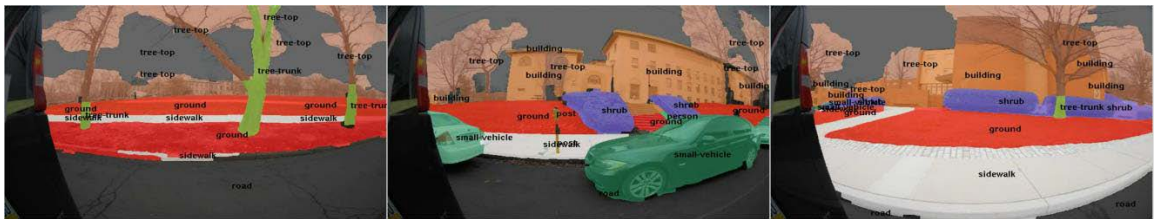

Figure 1-10: Examples of classification from images using 30 classes of objects and terrain.

In the area of understanding dynamic environments, we made key progress in developing new techniques for tracking in cluttered environments and for generating behavior prediction from partial observations. Specifically, we developed a novel approach for visual object "tracking by planning" in crowded scenes. The approach is unique because it combines visual tracking with recent methods from robot motion planning. In the area of 3D representations in dynamic environments, we developed a spatio-temporal representation to address the key barrier of real-time temporal accumulation and segmentation of 3D point-cloud data (from stereo) in dynamic scenes. The representation results in a 3D voxel map of the static elements of the scene and a segmentation of all dynamic elements in the scene. In the area of higher level semantic descriptions of dynamic scenes, we developed behavior prediction algorithms which, for the first time, were capable of combining semantic information from the environment with motion information, and to reasons about dynamic goals, while on-the-move. These algorithms were implemented on several platforms and sensors with initial testing at FTIG, showing the feasibility of this approach for future integration.

In addition to being tested on data collections from crowded urban environments, the 3D tracking work was fielded as part of IRA1 at the FTIG MOUT site to segment and classify pedestrians, which were then fed via the world model to a cognitive reasoning algorithm to recognize pedestrian behaviors. This integration enabled the first experiment integrating robust tracking with behavior classification in a complete architecture. These achievements are

important contributions to the elements of the capstone vision concerned navigating through a populated area and detecting specific behavior patterns.

To support our objectives in semantic perception, we have made progress in developing high performance sensing solution. A key accomplishment in this area is the completed implementation of stereo vision and visual odometry on the OMAP3530 system-on-a-chip (SoC), demonstrating stereo rate of 46 fps at 320x240 and visual odometry from up to 200 features. Overall, implementation shows a 3.75x improvement in energy efficiency of stereo over prior art with the added odometry functionality. This part of the work is critical for the operation of the capstone vision elements on small platforms.

HRI

During 2011 we have made significant progress in defining how the three most important characteristics of effective teams, Shared mental models, shared situation awareness, and trust, can be embedded within the control structure of advanced robots. HRI provided important design information to colleagues within HRI and in Intelligence and Perception in building the software that can use these higher-level concepts for more autonomy and effective collaboration with Soldiers.

We examined the best current mental model elicitation and shared mental model assessment techniques, then used the results to build a mental model measurement methodology suitable for tactical HRI. In collaboration with our colleagues in Intelligence, we are building computational models using the ACT-R framework. These models will allow experimentation with robot decision-making abilities, and permit validation against human counterparts. This work enables moving beyond isolated, simple scenarios to the ability to train computerized mental models to comprehend and act on novel situations.

During 2011 HRI provided guidance for using anthropomorphic and zoomorphic designs in robots (physical and cognitive characteristics). Results included understanding how design decisions (both physical and cognitive) will impact a human's mental model of a robot. These are very important factors that influence human estimation of robot ability and perceptions of trustworthiness. We produced design guidance for a continuum of tool-like to teammate-like relationships including tactical capabilities enabled by each to address mission primitive requirements.

Based on previous research in situational awareness in human teams, we empirically compared candidate SA metrics in the laboratory. We are using the best of existing SA measurement as the groundwork for development of a non-invasive, field-ready situation awareness measures. These results will guide what specific information a robotic teammate should provide to human teammates, facilitating team SA within the mixed human-robot team.

We conducted a meta-analysis of trust within robotics literature (and outside of it) to reveal key factors that influence a human's trust in robots: human, environment, and robot characteristics It

was found that factors related to the robot itself (as opposed to environmental or individual factors) most heavily influenced trust, emphasizing the benefits of robot reliability and transparency with respect to action selection and execution.

HRI developed representation models capable of interpreting and internally representing multi-modal communication signals. This is foundational work to support processing of raw communication data to internal formats for use in Intelligence and Perception. Prototype communication interfaces supporting bi-directional communication between humans and robots were developed under this thrust. This enables a near term collaborative experiment with DMUM using the BigDog platform to provide a technology feasibility demonstration of integrating multi-modal communication capabilities within an operationally relevant platform.

We investigated ways to layer context over content through empirical studies. These results set the stage for the 2012 development of an HRI Tactical Communication Protocol to be used by Consortium members. Related to this, we developed a novel specification language for virtual human interactions using parameterized behavior trees. In addition, we created an open-source framework for authoring and simulating narrative driven interactive virtual environments, populated with goal and event-driven autonomous agents. These will be important tools for both implementing effective communication building the simulations needed for development and evaluation.

To support shared understanding of social interactions suitable for use within robots we developed an interactive tutorial representing multiple modalities of social signals relevant to basic robot functionality in inhabited spaces. This understanding is fundamental to instantiating computational models in ACT-R, designing experiments to establish the parameters and dynamics of social intelligence, and ultimately provide design guidelines and implementation assistance within Perception and Intelligence.

DMUM

To date the DMUM mobile manipulation task has created tools to rapidly develop dynamic manipulation behaviors, generated several dynamic full-body behaviors and tested those behaviors using a simulation tool (Digital Biomechanics). During the final quarter of this project year we will be demonstrating those behaviors on the BigDog platform. As part of this task and the associated Integration task BigDog has been modified with the addition of a custom designed hydraulic 7-degree-of-freedom manipulator that includes a gripper. By the end of the 2011 project year, we expect to demonstrate Dynamic Heavy Lifting behaviors as well as Throw behaviors in a laboratory, using a 17.5kg cinder block.

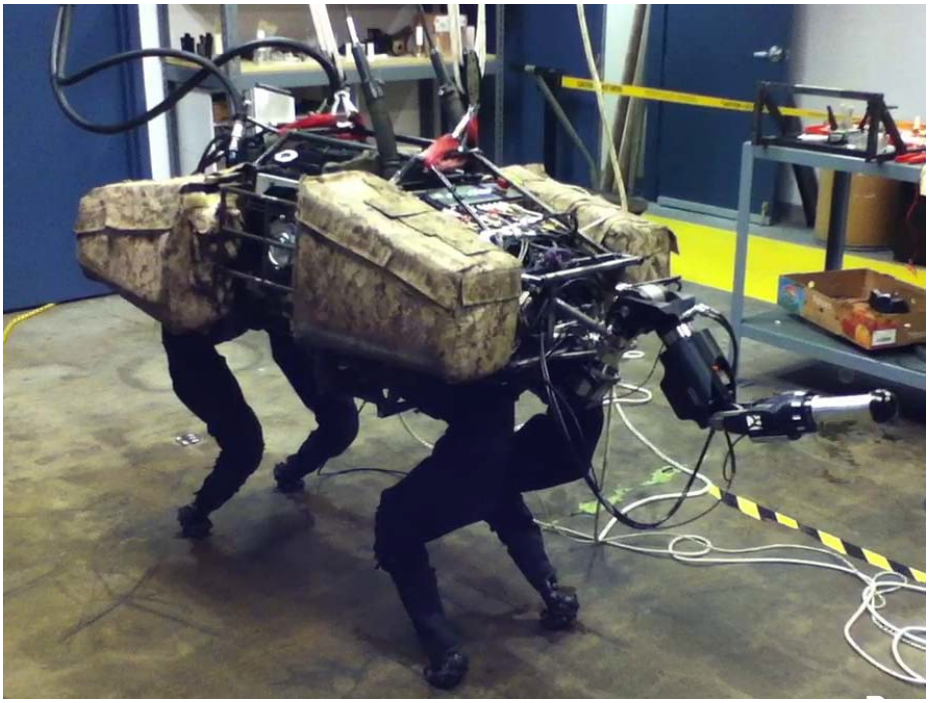

During the past year DMUM demonstrated near-autonomous negotiation of synthetic and natural climbing terrain by a rugged legged robot, achieved through sequential composition of appropriate perceptually triggered locomotion primitives. The first, simple composition achieves autonomous uphill climbs in unstructured outdoor terrain while avoiding surrounding obstacles such as trees and bushes (see Figure 1-11). The second, slightly more complex composition achieves autonomous stairwell climbing in a variety of different buildings (see Figure 1-12). In both cases, the intrinsic motor competence of the legged platform requires only small amounts of sensory information to yield near-complete autonomy. Both of these behaviors were developed using X-RHex, a new revision of RHex that is a laboratory on legs, allowing a style of rapid development of sensorimotor tasks with a convenience near to that of conducting experiments on a lab bench. Applications of this work include urban search and rescue as well as reconnaissance operations in which robust yet simple-to-implement autonomy allows a robot access to difficult environments with little burden to a human operator.

Figure 1-11: The X-RHex robot on a forested hill.

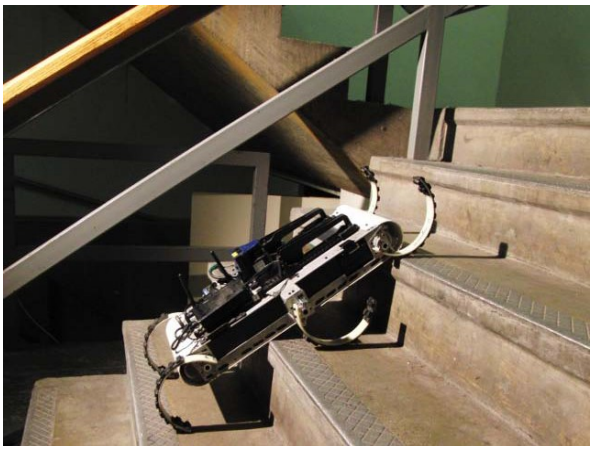

Figure 1-12: The X-RHex robot on a set of stairs with laser scanner, IMU, wireless repeater, and handle payloads.

Collaboration: Both Within and Outside the Alliance

Clearly, collaboration is at the heart of the RCTA program. A fundamental underlying premise of RCTA, as for the other CTA programs, is that great progress in fundamental research can be achieved by bringing together researchers who have not previously worked together – as well as by fostering further collaboration among those who have. Enabling far greater autonomy for ground robotics is an inherently interdisciplinary undertaking. As Morley Stone, Chief Scientist with the Air Force Research Laboratory's 711th Human Performance Wing and Lead of the DoD Autonomy Priority Steering Council, stated in a February 2012 Armed with Science article entitled "Self-Sufficient Robots" (http://science.dodlive.mil/2012/02/21/self-sufficient-robots/):

> *"… one of the key obstacles is…to try to get different communities at work on this problem. The key communities we need to bring together are those working on things like machine learning, working together with folks like human factors engineers and those who do cognitive modeling, the group trying to understand human cognition from a top-down perspective. Those are three communities that typically do not work together. But if we're going to make progress to get machines that can reason on par with the human, we're going to need to make progress on getting those communities together."*

The RCTA program is, in fact, for the first time in one program, bringing together researchers from formerly diverse disciplines such as cognitive architectures, robotic navigation, machine learning, human-machine interaction, legged locomotion, and semantic perception. The interactions among RCTA researchers occur at three levels: within the Consortium, within the Alliance, and between the Alliance and the broader research community.

One challenging aspect that can arise from all three forms of interaction is the communication problem among researchers from diverse technical backgrounds. We are deliberately bringing together these different "technical cultures" to achieve major progress as we have said, but the

various technical cultures sometimes bring with them differing terminology. As a step toward bridging those cultural gaps, we have begun construction of a table of "common language." This table is a living document that will evolve as the program evolves. Appendix 1 provides the current version of that document.

Within the Consortium – The Consortium consists of eight member organizations as well as a changing group of sub-awardees. The Consortium researchers come from a wide variety of technical backgrounds, and many of them are collaborating for the first time in the RCTA program. These collaborations are described throughout the APP in the corresponding task descriptions. For convenience and clarity, these collaborations are also summarized below in Table 1-1.

Collaborating Tasks	Objectives	Participants
P3 – Static Scene Understanding P4 – Perception for Missions in Complex Environments I2 – Data Mapping for Inference and Focus	Define interactions between semantic perception and WM for static environments	Tony Stentz (CMU), Martial Hebert (CMU), Alonzo Kelly (CMU), Bob Dean (GDRS)
P3 – Static Scene Understanding P4 – Perception for Missions in Complex Environments I5 – Learning through Experience I6 – Life-long Learning	Learning techniques for semantic classification	Drew Bagnell (CMU), Martial Hebert (CMU), Bob Dean (GDRS), Michael Turmon (JPL)
P5 – Dynamic Scene Understanding P6 – Perception For Dynamic Environments I2 – Data Mapping for Inference and Focus	Define interactions between semantic perception and WM for dynamic environments	Drew Bagnell (CMU), Martial Hebert (CMU), Bob Dean (GDRS)
P5 – Dynamic Scene Understanding I5 – Learning through Experience	Develop imitation learning techniques for motion prediction and activity understanding	Drew Bagnell (CMU), Martial Hebert (CMU)
P4 – Perception for Missions in Complex Environments P6 – Perception For Dynamic Environments I2 – Data Mapping for Inference and Focus I1 – Framework for Intelligence IRA	Develop and integrate perception components for IRA4. Specify architecture	Tony Stentz (CMU), Drew Bagnell (CMU), Martial Hebert (CMU), Bob Dean (GDRS), Larry Mianzo (GDRS), Max Bajracharya (JPL), Larry Matthies (JPL), Kostas Daniilidis (UPenn), Jianbo Shi (UPenn), Dave Duggins (QNA)
P6 – Perception For Dynamic Environments I1 – Framework for Intelligence IRA	Integration of perception capabilities in cognitive architecture	Tony Stentz (CMU), Martial Hebert (CMU), Christian Lebiere (CMU), Max Bajracharya (JPL), Bob Dean (GDRS), Brad Stuart (GDRS)

Collaborating Tasks	Objectives	Participants
P2 – Compact, High Performance Sensors M3 – Sensor-Based Dexterous Manipulation IRA	Precise short range sensing for manipulation; sensing on difficult materials	Christoph Mertz (CMU), Dave Duggins (QNA), Juan Pablo Gonzalez (GDRS)
P3 – Static Scene Understanding M3 – Sensor-Based Dexterous Manipulation IRA	Object recognition for manipulation/grasping tasks	Kostas Daniilidis (UPenn), Dave Duggins (QNA), Juan Pablo Gonzalez (GDRS)
P1 – Exploiting Novel Sensor Phenomenology P2 – Compact, High Performance Sensors M3 – Sensor-Based Dexterous Manipulation IRA	Stereo sensing for manipulation tasks such as trenching	Max Bajracharya (JPL), Dave Duggins (QNA), Juan Pablo Gonzalez (GDRS)
I1 – Framework for Intelligence H1 -- Shared Mental Models for Soldier-Robot Teaming	Align underlying SMM definitions and needs with representational and computational aspects of WM	Tony Stentz (CMU), Florian Jentsch (UCF), Christian Lebiere (CMU)
H1 – Shared Mental Models for Soldier-Robot Teaming I6 – Lifelong Learning	Specify format and content of external mental model representations that are helpful for incorporation into robots and have reasonable memory requirements	Drew Bagnell (CMU), Florian Jentsch (UCF)
H2 – Situation Awareness in Soldier-Robot Teams I3 – Combining Cognitive and Probabilistic Reasoning	Specify critical situation elements, identify uncertainties, and study when and how to disambiguate “important” uncertainties for Situation Awareness	Florian Jentsch (UCF), Maxim Likhachev (CMU)
H2 – Situation Awareness in Soldier-Robot Teams I6 – Lifelong Learning	Specify critical situation elements critical to life-long learning that need to be transmitted among robots and between humans and robots	Drew Bagnell (CMU), Florian Jentsch (UCF)
H2 – Situation Awareness in Soldier-Robot Teams P5 – Dynamic Scene Understanding	Provide situation-critical input to new perception tools for classifying interactions between robot and human agents from sensor data	Martial Hebert (CMU), Florian Jentsch (UCF)

Collaborating Tasks	Objectives	Participants
I1 – Framework for Intelligence H1 -- Shared Mental Models for Soldier-Robot Teaming H8 – Social Dynamics: Modeling, Simulation	Take advantage of overlapping work by collaborating to formalize and instantiate abstract socio-cognitive models (for SMM and social intelligence) into computational models in ACT-R, including methods for ACT-R coding & learning and experimental scenario design	Florian Jentsch (UCF) Christian Lebiere (CMU) Stephen M. Fiore (UCF) Jonathan Streater (UCF) Scott Ososky (UCF) Lotzi Boloni (UCF)
I2 – Learning to Understand Human Instructions and Annotations H2-4 – Dialogue Management for Robust HRI	Connect Fox's learned human input parsers to Roy's and Teller's graphical models for perceptual grounding	Dieter Fox (UoW), Nick Roy (MIT), and Seth Teller (MIT)
I3 - Combining Cognitive and Probabilistic Reasoning P6 – Perception for Missions in Dynamic Environments	Integrate multiple hypotheses from perception with planning techniques that can handle them. For now this is done in the context of navigating in dynamic environments under uncertainty in the intentions of people	Maxim Likhachev (CMU), Jianbo Shi (UPenn)
M1 – Motion Planning and Control for Mobile Manipulation I4 - Life-long Improvement of the Robustness of Tactical Behaviors	Integrate domain-independent techniques for repeated search-based planning in high-dimensional graphs with methods for constructing and searching graphs that exploit the specifics of mobile manipulation problems	Maxim Likhachev (CMU), Vijay Kumar (UPenn)

Table 1-1: Selected interdisciplinary collaborations within the RCTA Consortium.

Within the Alliance – The Alliance consists of the Consortium plus researchers at ARL who are engaged in collaborative work. Again, such collaborations are described throughout the APP in the corresponding task descriptions and are summarized below in Table 1-2.

Subtask	Subtask Title	Principal (Additional) Investigators	Organization
I1-3	Robotics CTA Architectural Development and Integration	Troy Kelley Eric Avery	ARL-HRED
I2-8	Meta-cognitive Development for Cognitive Architectures	Troy Kelley	ARL-HRED
I5-4	Anytime Learning for Robust Navigation and Control	Jonathan Fink Nicholas Fung	ARL-CISD

Subtask	Subtask Title	Principal (Additional) Investigators	Organization
I5-5	Integrate Visual Texture Recognition	Douglas Summers-Stay	ARL-CISD
I9-2	Cooperative Task Allocation for Surveillance	Ethan Stump	ARL-CISD
P2-5	MEMS-Scanned LADAR for Ground Robot Integration and Data Collection	Barry Stann Mark Giza William Lawler	ARL-SEDD
P3-6	Developing Metric/Topological Maps to Promote Common Ground	Jason Owens MaryAnne Fields	ARL-VTD
P4-5	Integration of Human Detection Capability for Ground Robots and Data Collection	Alex Chan Shuowen Hu Prudhvi Gurram	ARL-SEDD
H2-4	Investigate the Role of Trust and System Transparency in Shared Mental Model (SMM) Development during Tactical HRI	Jessie Chen	ARL-HRED
H4-4	Multi-modal Controls and Displays for Soldier-Robot Interaction	Ellen Haas Chris Stachowiak	ARL-HRED
H9-6	Social and Cultural Impact on HRI	Susan Hill	ARL-HRED
M8-7	Tactile Sensing for Control of Hyper-redundant Mechanisms	Raymond VonWahlde Geoffrey Slipher Justin Shumaker	ARL-VTD
M8-8	Dynamic Modeling of the CANID Platform	Raymond VonWahlde Jason Pusey	ARL-VTD
M9-7	Engineering Large Displacement Actuators for Enabling Biologically Inspired Modes of Mobility on Small Robotic Platforms	Geoffrey Slipher	ARL-VTD
IR1-2	Integrated Ressearch Assessments	Marshal Childers Barry Bodt	ARL-VTD
IR4-2	RIVET/ROS Interface	Ralph Brewer	ARL-VTD

Table 1-2: ARL Contributions to Robotics CTA Research.

Between the Alliance and the broader research community – Although the Alliance already includes a large body of researchers at the forefront of their respective disciplines, we constantly

seek more progress through interaction with other researchers who are not presently in the Alliance. These interactions are summarized below in Table 1-3.

Alliance Researcher(s) and Technical Area	External Organizations/ Researchers/Sponsor	Description of Interaction
Kelly (CMU) *Intelligence*	ARO Information Sciences Directorate, Randy Zachary	Unify the mathematics of on-line vehicle model interaction
Lebiere (CMU) *Intelligence*	ONR/NSWC, SWRI, Battelle, Soar Technology	Explore different methods to integrate cognitive architecture into robotic control stack in context of SUMET program (upcoming)
Likhachev (CMU) *Intelligence*	Willow Garage/Sachin Chitta	Improving the performance of planning over time, in particular in the context of mobile manipulation
Bagnell (CMU), Lee (UPenn) *Intelligence*	Willow Garage/Sachin Chitta	PR2 platform to obtain training data
Stentz, Bagnell (CMU) *Intelligence*	ONR MURI/Nick Roy	Imitation learning for uncertain environments
Bagnell, Stentz (CMU) *Intelligence*	Intel Science and Technology Center/Sidd Srinivasa and Mei Chen	HERB platform, co-development of life-long learning technologies
Bagnell, Stentz (CMU) *Intelligence*	DARPA ARM-S/Nancy Pollard	Use of low-level perception, control and planning primitives for manipulation
Bagnell, Hebert (CMU) *Intelligence*	ONR BIRD/Goshawk MURI/Rus Tedrake (MIT), Emilio Frazzoli (MIT), Yann LeCun (NYU)	Advances in optimal control libraries for UAVs and ground vehicles
Hebert (CMU), Shi, Taskar (UPenn) *Perception*	UIUC (Forsyth, Hoiem), UMD (Chellappa, Davis), ONR	Research on object recognition and scene understanding as part of ONR MURI
Bagnell, Hebert (CMU), Mianzo (GRDS), Bajracharya (JPL) *Perception*	DARPA, ARL, Mind's Eye teams	Research on video interpretation and real-time feature computation and event detection as part of DARPA Mind's Eye program
Mertz (CMU) *Perception*	Princeton Vision	Sensing and algorithms for collision mitigation (STTR)
Daniilidis (UPenn) *Perception*	SRI Sarnoff, RTC	Collaboration as part of R-MASTIF program: Robotic Mobile Autonomous System for Threat Interrogation and Object Fetch
Daniilidis, Kumar (UPenn), Bagnell, Hebert (CMU) *Perception*	iRobot, DARPA, CMU-NREC, USC	Perception, grasping, planning for autonomous manipulation tasks
Hebert, Mertz (CMU), Daniilidis, Lee, Shi (UPenn) *Perception*	University Transportation Research Center (DoT RITA)	Perception, planning, sensing for Intelligent Transportation Systems

Alliance Researcher(s) and Technical Area	External Organizations/ Researchers/Sponsor	Description of Interaction
Daniiliidis, Kumar (UPenn), Matthies (JPL) *Perception*	ARL, MAST consortium	Autonomous multifunctional mobile microsystems
Shi (UPenn) *Perception*	Kitware	Labeling buildings by video activities (STTR)
Matthies (JPL) *Perception*	UCLA (Soatto, Tsotsos)	collaboration on 3-D perception of non-Lambertian scenes
Matthies, Ansar (JPL) *Perception*	UCLA (Terzopoulos, Chan)	Exploratory research on self-calibration of stereo zoom cameras
Matthies, Turmon, Bajracharya (JPL) *Perception*	Robotic Research (Kluge)	Collaboration on work related to near-to-far learning of terrain type (Army STTR)
Bajracharya (JPL) *Perception*	CMU NREC, Boston Dynamics, DARPA	Research on perception for off-road navigation in dynamic scenes under DARPA LS3 program
Bajracharya (JPL) *Perception*	CMU NREC and USC; DARPA	Research on perception for manipulation under DARPA ARM-S program
Matthies, Turmon, Bajracharya (JPL) *Perception*	SPAWAR (Bruch), ONR	Collaboration on sensor fusion for off-road terrain understanding for ONR
Jentsch (UCF) *HRI*	Army/ARL: SOURCE (Barnes and Fields)	Project under SOURCE investigating HRI at different levels of autonomy
Jentsch, Fiore et al. (UCF) *HRI – H1/H2*	Consortium, ARL Researchers, and others outside the Alliance (Barnes)	Held joint area workshop on Shared Cognition and Shared Mental Models in Human-Robot Teams
Shumaker, Jentsch, Fiore (UCF) *HRI*	Human Factors and Ergonomics Society	Held special session on HRI at HFES 2011
Billings (UCF), Chen (ARL) *HRI – H1/H9 (trust)*	AFOSR and Community	Participated in AFOSR Workshop on Trust in Autonomous Systems, Feb 2012
Badler (UPenn) *HRI – H4*	GRASP Lab (UPenn)	Using directional gestures from a smart phone to control a robot (simulation and real) experiments.
Barber, Lackey (UCF) *HRI – H5*	ARL-HRED – Irwin Hudson	ARL sponsored research for evaluation of Soldier-Robot communication interfaces for gesture recognition in NLOS domains
Reinerman, Lackey (UCF) *HRI – H5*	ARL-HRED – Irwin Hudson	ARL sponsored research for measurement of Soldier affective-state for improved robot situation awareness
Reinerman, Lackey (UCF) *HRI – H5*	ARL-HRED – Irwin Hudson	ARL sponsored research for researching methods for training HRI in virtual environments
Likhachev (CMU) *DMUM*	Willow Garage/Sachin Chitta	Search-based planning for single-, dual-arm mobile manipulation

Table 1-3: Selected interactions of the RCTA Alliance with external organizations and researchers.

Relationship to Military Needs

In order to focus on the development of capabilities for ground autonomous systems over the next decade, it is helpful to consider specific ground robotics capability needs that have been identified by Army and Joint Service Future Operating Capability documents, Army S&T Master Plans, COCOM Priorities, Defense Planning Guidance, Prioritized Capability Lists, Warfighter Capability Gaps, Unmanned Systems Integrated Roadmaps, and feedback from current conflicts. These resources use a variety of methods to categorize requirements. For example, COCOM requirements are grouped into JCAs including Battlespace Awareness, Command and Control, Force Application, Logistics, and Protection, while Army Documents such as FM 3-0 Operations consider Movement and Maneuver, Intelligence, Fires, Sustainment, Command and Control, and Protection. TRADOC Pamphlet 525-66 provides perhaps the most comprehensive set of requirements, detailing the mapping between Joint Functional Concepts and Force Operating Capabilities. One way to frame our discussion is in the context of four Joint Functional Concepts most relevant to unmanned systems: Battlespace Awareness, Force Application, Protection, and Focused Logistics. Other examples of military needs come from the operational experiences of commanders in the field.

If we consider the missions of EOD and Route Clearance, ISR, logistics, and combat support, the following fundamental capabilities are important:

- For all mission types – Receive, understand, and acknowledge instructions/orders. This underlying capability is required in order to initiate any mission.
- For route clearance/EOD – Move to a verbally (semantically) described location, find the described object(s) of interest, inspect and/or manipulate the OOI, optionally transport the OOI. This capability is a leap-ahead from the current state of the art where, for example, EOD robots are tele-operated or, at best, perform simple waypoint navigation to go downrange to the vicinity of a suspected IED. In our vision, the robot recognizes and proceeds to the suspected object, inspects it, and even manipulates with permission and possibly some supervision.
- For persistent stare – Move to an appropriate location, survey an area of interest (AOI), recognize salient activity, construct a human-understandable report about the activity, and answer follow-up questions about the activity.
- For convoy/logistics – Identify and locate needed supplies, load the supplies, move to where the supplies are needed, and unload the supplies as needed and/or ordered.
- For robotic wingman – Move in tactically correct manner with unit, provide overwatch as needed, and provide useful response to hostile actions.

The above capabilities overlap a great deal in the technical challenges they impose, and they also differ in important ways. To better understand these differences and similarities, it is useful to decompose each mission capability into a sequence of mission elements. Figure 1-13 summarizes this decomposition for our four mission types.

Matrix of Autonomous Capabilities by Mission				
TRADOC Pamphlet 525-66	**Protection**	**Focused Logistics**	**Battlespace Awareness**	**Force Application**
JONS (Lynch)	**EOD/Route Clearance**	**Convoy**	**Persistent Stare**	**Robotic Wingman**
Receive, understand and acknowledge orders ("Think, Talk")	Receive orders, e.g., verbal or text, translate into internal WM representation, request clarification if needed, acknowledge/restate orders, answer questions if posed			
Report to unit, configure for mission, form up ("Move, Talk")	Load/request needed material for mission, e.g., tools, marsupial robots	Provision vehicle for mission, configure sensor payload, obtain relevant map data	Provision vehicle for mission, obtain relevant map data	Provision vehicle for mission, configure weapon payload, obtain relevant map data
Proceed to next specified location -- phase line,… ("Move, Think")	Move downrange to suspected IED site(s). Area has been cleared for EOD mission	Move in traffic to combat outpost location, either as part of manned unit or wholly unmanned	Move to suitable location for surveillance, taking account of both vantage point and concealment	Move to next objective in tactically correct manner, providing overwatch as needed
Maintain SA and comms -- react to unforeseen events ("Look, Talk")	Look for additional IEDs, react to IED events and possible related small arms attacks	Adaptively replan to disruptions to planned route, respond to IED or other events	Recognize time-critical ISR events and report/react accordingly	Perform 360 SA, recognize threats to self, leader and unit, report salient information
Perform specific mission tasks ("Work, Think, Talk")	Inspect possible IED objects. Manipulate and move objects as necessary	Select, pick up and load/unload needed materials	Observe activity and report salient events	Provide 360 SA, provide and accept targeting information, engage with permission
Conclude mission and return to unit ("Move, Think")	Return to EOD or Route Clearance unit, unload tools as needed, receive/await new orders	Report to commander at logistics mission destination, receive/await new orders	Return to ISR unit commander/HQ, provide more data as requested, receive/await new orders	Return from mission with manned leader, receive/await new orders

Figure 1-13: Matrix of autonomous capability challenges by mission type.

Our research portfolio is driven by capabilities over near-, mid-, and far-term time horizons. Thus, we will assess progress in capability over time. As an example, we consider the integrated capability elements of the EOD mission listed in the first column of Figure 1-14. This mission information was collected based on the interaction of RCTA technical managers and researchers with several EOD Soldiers in late 2010 and early 2011.

- Mission Preparation and Initiation
- Transit Downrange
- Mission Situational Awareness
- Perform EOD Task
- Return to Designated Location

For each of these mission elements, we have worked with the Naval Explosive Ordnance Disposal Technology Division (NAVEODTECHDIV) to identify specifically how RCTA can

impact EOD unmanned operational capabilities (UOCs) in the 3-, 5-, and 7-year timelines. Figure 1-14 summarizes UOCs we envision by EOD mission phase and lists specific foundational elements of autonomy that apply to each element. Here, we summarize the relationship of each of the EOD mission elements both to the RCTA vision outlined in the beginning of this section and to the foundational capabilities of autonomy we are developing to achieve our vision.

Mission preparation. Currently, EOD robots are simply devices that must be operated step-by-step to achieve a mission. They must be configured by operators, transported to where they are needed, and then tele-operated to go downrange and perform the mission. Our vision is to endow the robot with the ability to receive orders or instructions, configure itself for the specified mission, and position itself to begin executing the mission. To achieve this vision, RCTA will give the robot a measure of adaptive tactical reasoning using the cognitive-to-metric world model so that it understands what an EOD mission is, can receive and acknowledge orders, and knows how to prepare itself for the mission.

Transit downrange. Instead of the current practice of tele-operating the robot(s) downrange, our vision is for autonomous movement to the appropriate site downrange. For example, the robot may be instructed to go to a map reference or to a semantically described point, e.g., "the rubble pile behind the blue track." It then navigates through the environment, which may include rugged terrain and heavy urban clutter. It can also transition from outdoors to indoors if needed, possibly using new mobility mechanisms to negotiate stairs or other obstacles. This tactically adaptive movement requires the understanding of terrain, objects, and activity in the environment. We achieve this through parallel efforts in semantic perception and adaptive mobility behavior instantiated in our new world model.

Mission situational awareness. During current EOD missions, SA is achieved only through human "eyes on." This may be direct human vision or vision through robot sensors. The latter requires Soldiers to be "heads down," which is not conducive to overall SA. Thus, we strive for our vision in which the robot builds its own understanding of the environment – both in the immediate vicinity of the object of interest (OOI) and in the broader environment where other IEDs or other forms of threat may exist. The goal is shared SA among both human and robotic team members, combining both the "look" and "talk" capabilities in our vision. The RCTA will achieve this "common ground" through integrated research in semantic perception and meta-cognition.

Perform EOD task. This mission element is more specific to the EOD mission than other elements. It is currently performed through tedious tele-operation that requires very experienced operators. There is now some degree of automation through memorizing some manipulation steps. However, we aim to develop the basis for a much higher level of autonomy where the robot can, for example, search through a pile of rocks and other debris with little or no supervision. To do so, it must understand the objects in its environment as well as its own ability to interact with that environment. The capability will come from our research in adaptive behaviors as they relate to manipulation tasks. Furthermore, our research looks at joint

manipulation and mobility issues to enable a robot, for example, to brace itself to move a heavier object than it could otherwise move. Also, instead of simple imitation learning of simple manipulation steps, we are exploring deep learning – in combination with semantic activity understanding and a hierarchical world model – to enable the learning of more complex tasks.

Return to unit. The final mission element is largely a combination of the previous elements. It requires the robot to exercise judgment about when to return, whether to seek guidance, what it encounters on the way back, and how it should communicate with its unit upon return.

EOD Mission Element (Foundational Capabilities)	**Near-term ~3 years**	**Mid-term ~5 years**	**Far-term ~7 years**
Mission Preparation and Initiation (1, 3a, 5a)	Receive and acknowledge simple mission orders, enabling more "launch and forget" operation.	Receive and acknowledge more complex orders; assemble inside host vehicle.	Deploy (assemble and exit host vehicle) and respond adaptively to changing tactical situation, e.g., hostile action.
Transit Downrange (1, 2a, 2b, 3c, 7b)	Navigate to a single or series of grid coordinates in rugged environments (moderate clutter, positive/ negative obstacles, day/night, GPS/GPS-denied, indoor/outdoor).	Navigate to described location in rugged terrain with heavy clutter (positive/ negative obstacles, day/night, GPS/GPS-denied, indoor/outdoor).	Navigate to described location in rugged terrain with mobility challenges and heavy clutter (positive/ negative/ water obstacles, day/night, GPS/GPS-denied, indoor/outdoor).
Mission Situational Awareness (1, 2a, 2b, 5b, 7a)	Generate 3-dimensional virtual environment annotating user selected locations/items of interest, allowing for 3rd person platform operation and scene analysis (linear distance measurement).	Generate high-res 3-dimensional virtual environment, updated in real-time annotating user selected locations/items of interest, allowing for 3rd person platform operation and forensic scene analysis (distance, angles, volumes).	Generate high-res 3-dimensional virtual environment, updated in real-time, annotating user selected and automatically identified locations/items of interest, allowing for 3rd person operation and forensic analysis.

EOD Mission Element (Foundational Capabilities)	Near-term ~3 years	Mid-term ~5 years	Far-term ~7 years
Perform EOD Task (1, 2a, 2b, 3b, 4, 6, 7a)	Provide user with haptic feedback and intuitive control in a high latency environment. Simultaneous operation of multiple robots by a single operator.	Automate dual arm task based manipulation operations. Automate passive cooperative behaviors with multiple robots. Robots do not simultaneously physically interact with objects.	Learn simple manipulation tasks from Soldier training. Automate active cooperative behaviors with multiple robots. Robots do simultaneously physically interact with objects, e.g., pick-up and carry a payload.
Return to Designated Location (1, 2a, 2b, 3a, 3c, 4, 7a, 7b)	Return efficiently to user-defined location using multi-sensor data. Upon loss of communications, return efficiently to user-defined rally location or until communications are re-established	Return efficiently to user-defined location using multi-sensor data. Upon loss of communications, return efficiently to user-defined rally location or until communications are re-established Demonstrate increased efficiency from near-term years.	Rejoin unit on-the-move using own/mixed initiative; enter host vehicle autonomously.

Figure 1-14: EOD Mission Integrated Capability progress over time.

In order to enhance EOD mission operational capability as indicated in Figure 1-14, we must address the technical barriers to those capabilities. Therefore, we have identified those barriers, examined them in light of our planned research, and verified that our research is targeted to surmount those barriers. While this discussion uses the EOD mission as one example, similar analysis shows that the research is also targeting the other relevant mission areas.

Conclusion

The RCTA program has established a vision for a new level of autonomy in ground robotics. That vision includes five key capabilities: (1) adaptive tactical reasoning, (2) focused situational awareness, (3) efficient and proactive interaction with humans, (4) safe, secure and adaptive

movement, and (5) interaction with the physical world. We anthropomorphically describe these capabilities as “Think,” “Look,” “Talk,” “Move,” and “Work.”

Currently, there are major technical barriers which lie in the way of this vision. At the heart of the problem is the need for a world model, which can be instantiated on robots and which represents the range of entities, spatiotemporal scales, and abstractions that must be reasoned about. There is also a shortcoming in a robot’s ability to sense the environment and understand it at a semantic level; this ability is needed in order to populate the world model with new and relevant information. Robots have only a rudimentary capability to plan behavior, and their planning is focused almost entirely on navigation and is brittle even in that limited domain. Furthermore, cognitive concepts such as mission, enemy, troops, terrain, time, and civilians (METT-TC), which are at the heart of a Soldier’s planning process, are entirely missing in current robots. Finally, the robot behaviors that have been achieved thus far have been programmed for specific applications; any correction or new application requires re-programming. Therefore, without a robust learning capability, our vision of adaptable robots cannot be achieved.

To overcome these barriers, we have identified five major fundamental research thrusts which will lay the foundation for a new level of autonomy. These foundational thrusts are:

- Cognitive-to-metric world model
- Semantic perception
- Adaptive behavior generation
- Machine learning
- Meta-cognition

These thrusts provide the basis for *individual* robotic autonomy. We also build upon that foundation to enable *collaborative* autonomy among a set of robotic and human teammates. We have identified a small set of enhancements that will make autonomy much more effective at both the individual and collaborative levels. Our current plan invests over 60% of our resources in the foundational thrusts listed above; it invests smaller, approximately equal, efforts in collaboration and in selected autonomy enhancements.

We have built a fundamental research portfolio to implement our technical thrusts. The portfolio is structured in four technical areas: Intelligence, Perception, Human-Robot Interaction, and Dexterous Manipulation and Unique Mobility. The following four sections of this Annual Program Plan detail the tasks that comprise our research portfolio in each of those technical areas. The sections include, at the task and subtask levels, a description of the current state of the art, how the currently proposed research advances the state of the art, and the metrics that will be used to assess progress. Each section also contains specific research plans and objectives for 2012.

In order to harness this fundamental research to achieve our vision, we must integrate research outcomes from across the technical areas into cross-cutting technical capabilities. The final section of this document describes our approach to integrating research outcomes and assessing the resulting integrated capabilities. These capabilities, by design, align very closely with the capabilities in our vision for robotic autonomy. Thus, our integrated research assessments will directly measure our progress toward achieving the RCTA vision.

2. INTELLIGENCE OVERVIEW

Objective

The objective of the Intelligence research is to provide the framework, world model, and algorithms for a broad range of robot capabilities. These capabilities include receiving human commands (1)[1], reporting on status, and self-diagnosing problems (2, 11). The research provides context and focus of attention to perception to enable semantic understanding of the world (3). It provides planning capabilities to support navigation through complex and cluttered environments, both urban and natural and with movers present (4, 5). The architecture assembles perception, planning, and execution monitoring into tactical behaviors, such as watching a building for people (6) and responding accordingly (7). It provides support for multiple behaviors running concurrently, allowing for maintenance of situational awareness and responding to emergencies (8). The research supports approaches for improving robot performance, either automatically or via a teacher (9). Finally, the research provides support for a broad range of platforms, including snakes and legged and wheeled machines, and provides for manipulation as well as navigation (10).

State of the Art and Technical Barriers

Robot Intelligence has focused on producing competent navigation capabilities, both on- and off-road, using real mobile robots operating in real conditions. Likewise, cognitive science and artificial intelligence systems have focused on developing human-like capabilities but for simulated or otherwise simple environments. The field of Intelligence requires a framework and methodologies for engaging robots to perform navigation, manipulation, and grasping operations in support of team missions in real environments under real conditions. Intelligence is challenged to deliver this capability in the face of large uncertainty: in the world, in perception, and in the planning/reasoning process, over a larger number of parameters than has been previously addressed. Intelligence cannot be fully engineered – there is neither a quantifiable set of requirements nor the development time and effort to produce it. Instead, intelligent robots must learn efficiently from human experts and from direct experience, at both the component level and at the system level. Finally, intelligent robots must interact with their human counterparts in new ways, including explaining its actions and plans and diagnosing its own problems.

Moving Beyond the State of the Art

We have updated the architecture to address the Intelligence barriers and have crafted a set of tasks and subtasks that continue to flesh-out the needed capabilities. The proposed architecture consisted of three levels: cognitive, relational, and physical without specifics for what modules populate these levels and the functional role and interfaces they assume. The updated architecture revises the level descriptions, introduces a template computational node, defines some of the functional blocks, and details the learning mechanisms.

[1] Items refer to numbered elements of the RCTA Capstone Integration and Assessment Vision.

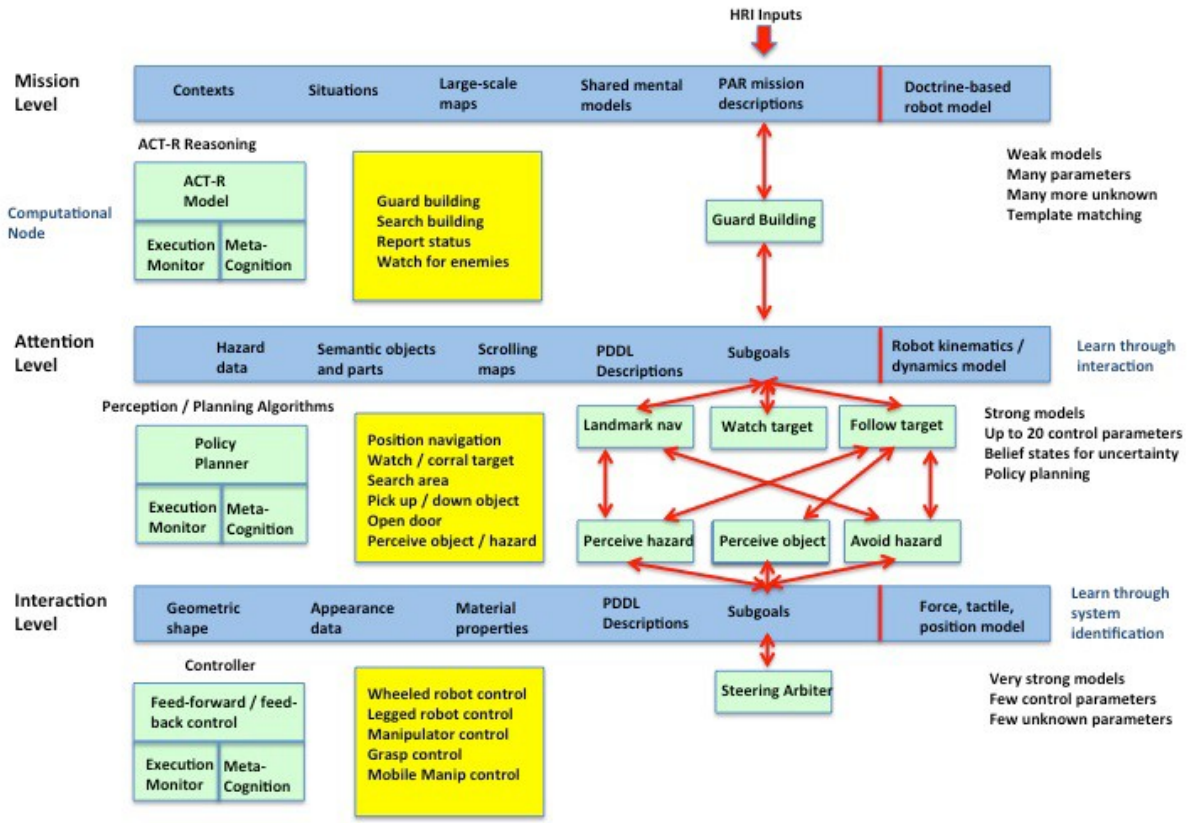

Figure 2-1: Updated RCTA Intelligence Architecture.

The architecture is tightly coupled to the world model. By world model, we mean a model for the world itself, the robot(s), and the interaction between the robot and world. The architecture consists of a hierarchy of *tasks* (Figure 2-1), ranging from those at the *mission* level at the top, to the *attention* level in the middle, and the *interaction* level at the bottom. Each task is a computational node that encapsulates a particular functionality. The node interleaves perception, planning, and execution monitoring to transform the world model from one *belief* state into another belief state. A belief state is a set of actual states, each assigned with a probability. If probabilities are not available, the belief state can consist of just an actual state or an *assumed* state. A task can be invoked either to plan for a particular functionality or to perform it. The difference is that as a planner, a task estimates the effects of perceptual actions and robot motions rather than performing them.

A task includes the following components:

- *Action name:* such as "Navigate-by-Position".
- *Objective function:* cost function to minimize, such as distance, time, visibility, or some combination.

- *Pre-conditions:* conditions that must be true before a task can run; specified in the Planning Domain Definition Language (PDDL); must be true for all actual states in a belief state.
- *Post-conditions:* conditions that will be true after a task runs; specified in PDDL; will be true for all actual states in a belief state.
- *Resources:* robots, sensors, etc. required by the task; requested exclusively or non-exclusively.
- *Estimated cost:* expected cost of performing the remainder of the task, for planning, monitoring, and learning purposes. The value is zero if the task has succeeded and infinity if it has failed.
- *Goal probability:* probability of meeting post-conditions, for planning purposes.

In order to perform a particular action, a task may chain together subtasks at the same or a lower level in the hierarchy by matching their pre- and post-conditions. Essentially, the post-conditions are subgoals that realize the primary goal through a sequence of steps. The execution monitor oversees the subtasking. Each execution monitor includes an ACT-R-based meta-cognition component for diagnostic monitoring and confabulation purposes. Failures at lower levels are propagated up in the usual way for replanning purposes. A subtle yet important point is that because the tasks operate in belief space, the replanning process can consider estimated results of perceptual and robot motion actions *before* performing them. These estimates can be calculated hierarchically by calling subtasks in planning mode.

Tasks can span subtasks that run sequentially or concurrently. If subtasks running concurrently share resources (e.g., imaging sensor or steering control), an arbiter de-conflicts requests and commands. For example, a task that is steering a robot to follow a target runs concurrent with a task that steers to avoid hazards. The arbiter permits factoring of the two capabilities and eliminates the need for a single task that reasons about all steering considerations.

At each level, the world model stores all of the data for matching the task pre- and post-conditions. The world model also stores resource models for the robot, the current task/subtask execution trace (for monitoring and inspection), and the history of this trace (for offline learning). We describe the three task levels below.

Mission Level

The tasks at the mission level implement specific military doctrine. The tasks determine how to sequence atomic actions at the next level down (e.g., navigation, manipulation, and perception) in order to achieve the mission. For example, the task implementing “cover the back door” would instruct the robot to move around the building, detect the back door, and watch for a suspicious person leaving. The challenges include identifying the “back” of the building, finding a “door,” and watching for a “suspicious person” as well as handling the many contingencies that are possible, including a robot failure, an emergency situation, and a new command. Although many of the difficult problems are perceptual, the mission task is able to supply context to assist perception. For example, the back of the building is likely to resemble the side. The mission level can supply context for the robot actions as well, including specifying a stealthy traverse if a shooter has been detected.

Mission-level tasks mimic human functionality in compliance with doctrine. The basic actions are pre-determined, but the challenge is to figure-out when and how to apply these doctrine templates to particular cases, tweaking actions based on situational awareness and context. The world model includes state data in PDDL that represents contexts and situations as well as tasks and subtasks that are planned and executing. Note that uncertainty at the mission level is very high. Not only are there parameters with missing values (e.g., color of back door), but there may be important parameters that are completely undefined and unrepresented (e.g., threatening person in vicinity). ACT-R is used to implement the planning component of a mission task as well as the meta-cognition model in the execution monitor. Much of the learning focuses on adjusting pre-conditions, identifying relevant parameters, and tweaking actions to better stretch and fit the templates to more mission scenarios.

To implement the Capstone Integration and Assessment Vision, we envision mission-level tasks such as the following:

- Guard-building-access
- Search-building
- Report-status
- Wait-for-command
- Watch-for-enemies

Attention Level

The mission-level tasks call attention-level tasks to navigate from place to place, grasp and manipulate objects, and perceive semantic objects and hazards. Rather than encoding doctrine, the reasoning at this level is primarily geometric. For example, a task reasons about how to move a manipulator to avoid obstructions and get into position to grasp an object. The world model includes grids of hazard data (e.g., interpreted relative to a particular robot model); semantic objects such as a door, building, or wall; scrolling maps; and planned/executing tasks at this level and their subgoals.

The robot models are strong (e.g., robot kinematics and dynamics); hence, the planners can forward simulate possible actions and select one with the desired outcome. This planning process includes sensing as well as motion operations, and the reasoning covers the different possible sensing outcomes. Thus, many of the planners produce *policies* (e.g., AO*, PCPP, MDP, POMDP) rather than strictly deterministic plans. For example, a policy for corralling a person precomputes actions for all possible human motions to ensure that the person does not escape.

Even though the robot models are strong, they may contain a relatively large number of degrees of freedom. For example, a wheeled vehicle with two manipulator arms may comprise 15 to 20 degrees of freedom. Thus, the planners include computationally efficient algorithms (e.g., re-using partial solutions, random sampling, projecting problem into smaller space) to plan in real-time.

The world model data is likely to be less strong and rife with uncertainty. Is this patch of terrain a mobility hazard? Is this object light enough for the robot to pick it up? Where is the door on the

back of the building? What color is it? For metric uncertainty, such as a range of possible parameter values, planners can be devised that are robust to the uncertainty. Using the belief space representation, these planners can be chained together to deliver a guaranteed result, regardless of the uncertainty (e.g., funnels). For more severe model deficiencies, the planner may need to interact with the world (e.g., perceive it with cameras, push on it with manipulator) to build a better world model. Fortunately, the robot can use proprioceptive sensing of the interaction as a ground truth signal for building this model. For example, a robot may learn that small, round, black objects can be rolled out of the way.

To support the Capstone Vision, the attention-level tasks include:

- Navigate-by-position
- Navigate-by-landmark
- Avoid-hazards
- Follow-target
- Corral-target
- Watch-target
- Search-area
- Pick-up-object
- Put-down-object
- Push/slide/drag-object
- Open-door
- Climb-stairs
- Perceive-semantic-object
- Perceive-mobility-hazards

Interaction Level

The attention-level tasks call interaction-level tasks to affect motion and interact with the world via contact. Tasks at this level are essentially *controllers*. They typically cycle at 10 Hz or faster. The world model includes robot kinematics and dynamics and metric data, such as geometric shape, appearance, and material property for objects in the world. The model includes robot/world interactions, such a forceful contact, tire-soil effects, etc. The uncertainty is typically confined to a small number of parameters such that standard parameter identification techniques can be used online (e.g., solving for tire slip by comparing commanded trajectory to actual trajectory from IMU). Note that even at this level, we use ACT-R-based execution monitoring to diagnose failures and explain actions.

The controllers are:

- Mobile robot controller: achieve turn radius, direction, and speed.
- Manipulator controller: achieve joint positions, end-effector configuration, and force at end effector.
- Mobile manipulator: same as manipulator but add pose of mobile base.
- Grasp controller: achieve finger positions, force control at fingers.

Learning
As explained, the architecture provides support for online, self-supervised learning. Proprioceptive and interactive sensing is used as a ground truth signal to automatically learn model parameters. Planners use these models at the attention level. Similarly, proprioceptive sensing can be used to automatically train perception modules to interpret imaging sensor data, such as that from cameras or LADARs.

The architecture also supports offline imitation learning. For imitation learning, a human expert controls the robot to perform a mission as he/she would. During this training session, the robot runs its autonomy system and generates its own perceptual interpretations and motion plans, which are likely to differ from those of the teacher. The robot then adjusts its internal parameters until its decisions match those of the teacher. Given enough training sessions, the robot's performance can rival that of the teacher's – achieving a level of competence that exceeds what is practical via human engineering alone.

The architecture supports imitation learning via *inverse optimal control*. The robot's decision space can be represented by a graph, where each decision point is a node and each possible decision is an arc. The robot selects a sequence of decisions that accomplishes the given mission. The idea behind inverse optimal control is to assign a cost to each arc in the graph and to adjust those costs so that the *lowest cost path* through the decision graph matches the teacher's decisions. In the architecture, each task/subtask is a node and its planned actions are arcs. The learning algorithm tells each task how to adjust its estimated cost (e.g., up or down) to match the teacher's decisions. In turn, each task can adjust its own internal parameters so that it better predicts its estimated cost based on input parameters in the world model. This unified, mission-level approach to learning permits *all* tasks in the architecture to adjust parameters – that includes perceptual components of tasks in addition to planning components – and achieves deep learning as originally anticipated.

Relationship to Other Technical Areas
The Intelligence area will draw heavily from the other three areas to populate the architecture. HRI will investigate the means by which humans communicate with robots, not only the form (audio, visual, tactile, etc.) but the content based on trust and common understandings. HRI will be instrumental in determining the types of data in the mission level of the world model, including information that is necessary to define a "situation," "context," or "environment." Of particular importance is the shared mental model, which enables humans and robots to work together effectively. The Intelligence area is investigating computational mental models, and HRI is developing human mental models; the two areas will bring together the best of both.

The Perception area will flesh-out much of the attention level. In addition to pushing hazard detection for mobility for increasingly more complex environments, perception is focusing on semantic labelings, such as person, vehicle, building, door, and semantic maps. This level of perception supports much more than obstacle avoidance; it enables landmark-based navigation, picking-up/putting-down objects, and watching and/or following people. This perception also supports a richer, semantic understanding of terrain that is useful for people as well as vehicles. Perception is not limited to the attention level, however. Rare activity detection, semantic maps,

and behavior prediction are important for situational awareness and understanding of contexts and environments at the mission level. Pressure sensors in skin build a richer representation for understanding objects at the interaction level.

The DMUM area provides the robot models that are used at the interaction level. These include models for snakes, legged robots, manipulator arms, grasping mechanisms, and various combinations. DMUM implements controllers at this level as well, including gait selection for snakes and legged robots and force control strategies for grasping mechanisms. DMUM explores computationally efficient dynamics models that can be used in feedback control schemes. Finally, DMUM explores planning strategies for high-dimensional problems like mobile manipulation and grasping strategies for simple hands.

Summary of Proposed Research

The proposed work is summarized below in Table 2-1.

Task	Primary Research Outcomes	Operational Applications
I1: Framework for Intelligence	Design framework for modules that understand, perceive, and interact Develop shared mental models for robot/human teaming Design of atomic actions that serve as basic building blocks for navigation and manipulation Meta-cognitive ability for architecture to introspect, diagnose, and remedy based on experiences	More robust tactical behaviors due to cognitive-level ability to leverage context and recover from failures Missions that navigate and manipulate objects Meta-level understanding of robot intelligence for purposes of explaining performance and detecting failure
I2: Data Mapping for Inference and Focus	Updated implementation of world map Incorporation of dynamic objects, point clouds, and topological links Incorporation of speech annotations, semantic links, gesture/body motion representations, and episodic memory Algorithms for focus of attention Flexible representations for tightly-coupled perception and action that support learning over a range of tasks Algorithms for recognizing doors and handles and for reaching and manipulating actions	Data representations for connecting the three levels of the intelligence architecture Software for managing data at these levels Capabilities for opening doors, picking up objects, and other highly interactive tasks that require perception and action

Task	Primary Research Outcomes	Operational Applications
I3: Combining Cognitive and Probabilistic Reasoning	Algorithms for planning with multiple hypotheses in perception Algorithms for managing and reducing uncertainty	Understanding how to focus perception and interact with the environment to manage and reduce uncertainty for more robust performance
I4: Generating Adaptive Tactical Behaviors	Planning algorithms for interacting with the world, such as by pushing objects Planning algorithms that re-use old solutions to improve speed over time Algorithms for automatically calibrating robot models	Automatic adaptation of localization, planning, and control algorithms in response to environmental interactions Tactical behaviors that can manipulate as well as navigate
I5: Learning through Experience	Approaches for scaling up learning algorithms to robotics size problems Learning algorithms for improving high-dimensional planning problems Learning algorithms for understanding natural language commands	Understanding how to improve robot performance through imitation learning and experience Deriving maximum benefit from available data
I6: Life-long Learning	Algorithms for learning how to sort through clutter Algorithms for learning how to interact with objects—whether there is an object model or not Algorithms for learning how to automatically detect salient data in streams	Building models for the interaction of robots with the world and using those models in navigation and manipulation tasks
I9: Distributed Intelligence for Human/Robot Teams	Algorithms that generate planning policies suitable for maximizing information gain Approaches for dividing tasks between humans and robots for execution	Collaboration between humans and robots based on capability matching and information gain

Table 2-1: Intelligence Task Outcome and Operational Implications table.

Path to Integrated Research Assessments

In the coming year, the Intelligence tasks will support IRA3 (Intelligent Navigation) and #4 (Physical Interaction). For IRA3, tasks I1 (Framework for Intelligence) and I2 (Data Mapping for Inference and Focus) will support the high-level reasoning (*thinking* thrust) for configuring the perceptual and motion planning behaviors and monitoring progress (*talking* thrust). Task I4 (Generating Adaptive Tactical Behaviors) will produce the low-level behavior (*moving* thrust) that draws from perception to identify and track semantic objects (*looking* thrust).

For IRA4, Task I6 (Life-long Learning) develops models that will be used to reason about moveable objects (*thinking* thrust) using techniques developed in I5 (Learning through Experience). Task I3 (Combining Cognitive and Probabilistic Reasoning) will reason about the uncertainty inherent in mobile manipulation problems (*moving* thrust).

For the remaining IRAs leading to the Capstone Assessment (IRA7), the Alliance will add an Autonomous ISR capability (IRA5) and Tactical Team Movement (IRA6). Task I1 will develop the basic capability to recognize and reason about adversarial and teammate interactions. Task I2 will support cognitive queries into a semantic database, representations for situational awareness, and distributed implementations to support a team. Task I3 will develop the algorithms necessary to reason about uncertainty, including uncertainty inherent in human actions. Task I4 will assemble the capabilities into tactical behaviors that affect sensing and movement. Task I5 will develop the learning algorithms to enable robots to be trained to recognize adversarial situations and to perform a particular role on a team, and Task I6 will support the data collection and management to perform this training. Task I9 will determine the roles for robot teammates in response to actions taken by their human counterparts.

Intelligence Tasks

I1-2012 – Framework for Intelligence

Level: 6.1 Basic Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Tony Stentz	CMU	I1-1
Additional Investigators	Christian Lebiere	CMU	I1-1, I1-2
	Martial Hebert	CMU	I1-1
	Florian Jentsch	UCF	I1-1
	Troy Kelley	ARL	I1-3
	Eric Avery	ARL	I1-3

Objective:
The objective of this task is to investigate the research issues essential for assembling a complete and intelligent robot architecture – one that leverages the best from cognitive science approaches and traditional algorithmic approaches – to support robotic navigation and manipulation operations in a wide range of challenging environments. The objective is to build robots that work well with humans, both as teachers and as teammates, and to directly benefit from the experience of executing missions.

Background:
How do we empower robots to go beyond simple navigation to sophisticated human-like performance? AI researchers have used cognitive reasoning systems to encode domain expertise into sophisticated tactical behaviors. Unfortunately, these systems have almost exclusively been tested in simplified worlds, such as robot simulators with unrealistically reliable perception, navigation, and manipulation primitives, or on actual robots operating in well-structured, robot-friendly environments. Conversely, roboticists have built machines that operate under a wide variety of environments, conditions, and situations, but the performance is primarily limited to navigation, simple manipulation, and data collection. The problem is the enormous gulf that separates the symbolic data employed in sophisticated reasoning systems (e.g., clear building, open door, find hiding places) and the complex perceptual understanding and fine motion/manipulation required to implement them for any realistic instance of the problem. The problem has been addressed in the intelligent architecture literature. Top-down, knowledge-rich architectures start with the symbolic representations and work toward the physical robot. This strategy works well for known, simple environments but tends to break in complex, uncertain worlds. In contrast, bottom-up, algorithmic architectures are simple in form, naturally interact with the environment, and adapt to uncertainties but do not offer performance guarantees.

Research Description:
During FY11, we identified a nominal architecture concept and performed experiments to prove-out previously unexplored aspects. This work led to the creation of the updated architecture described in the introduction to the Intelligence research plan. In the coming year, we will put meat on the bones by specifying the functional blocks, defining the interfaces, integrating with the world model, and steering Intelligence research in the direction of a unified system. We will also shepherd revisions of the architecture which will result organically from exploration of the full scope of Intelligence experimentation and assessment and collaborative interactions within the Alliance.

Metrics:
The measures of the architectural framework's performance include 1) the complexity of the mission that a robot can execute and 2) its robustness to a broad range of environments and conditions. For the former, the complexity increases as we move beyond navigation and incorporate manipulation and fine-grained interaction with human teammates. For the latter, the range includes a variety of semantically rich indoor and outdoor environments.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
We anticipate another Intelligence Workshop to update the Intelligence architecture and plan for integration.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Update the architectural framework, and distribute it to all Alliance members for feedback.
Q2	Identify how each subtask across all four technical areas fits into the architecture.
Q3	Determine an integration path, including necessary interfaces and data support.
Q4	Identify the missing components (and science) to guide development of the next APP.

Related Research:
This task works with all of the other Intelligence tasks to determine functionality and interfaces to the world model (I2), the three primary levels of the architecture (I3 and I4), the learning components (I5 and I6), connections to humans (I1 and I2), and connections to other teammates (I9). This task also works closely with HRI, in particular Tasks H1, H2, H4, and H5.

Integration and Assessment Activity:
Similarly, this task affects all of the Integrated Research Assessments since the architecture forms the basis for integration.

Subtask 1: Connecting Cognitive Architectures to Real Robots (CMU and UCF)

Objective and Benefits:
The objective of this subtask is to investigate the research issues essential for assembling a complete and intelligent robot architecture – one that combines cognitive science approaches with traditional algorithmic approaches and which supports both on- and offline learning. Furthermore, the architecture includes shared mental models for working with a human counterpart. The benefit is a flexible architecture that is able to perform sophisticated tasks with humans in a real environment, without suffering from brittleness.

This subtask develops the overall architecture required to implement the Capstone Assessment by providing the framework for integrating a cognitive architecture with perception, planning, and world map components. Additionally, it develops shared mental models for teaming with humans. This subtask will run multiple years and will be used to continually update the Intelligence architecture from lessons learned.

State of the Art:
Cognitive architectures have been used to implement sophisticated behaviors in simulation, but these simulators do not capture the difficulties of a real robot interacting with the real world. In particular, the failure modes are not properly represented. Conversely, traditional robot architectures have been used on real robots but only for simple behaviors. The end result is robot brittleness when sophisticated missions are implemented on real robots.

Key Barriers and Approach:
The general principle behind our approach is that cognitive architectures are appropriate for human action that is difficult to model but for which there are examples that can be generalized and adapted for different instances of the problem. Traditional robotics algorithms are appropriate for actions with strong models so that simulators can accurately generate the different possibilities and evaluate the outcomes. For example, the “cover the back door” mission requires the robot to perform a sequence of actions that is difficult to derive from first principles but can be constructed from human examples. ACT-R is appropriate for the configuration of those actions; however, algorithmic approaches are better suited for the perceptual and motion primitives themselves, especially the parts that reason about field-of-view and robot kinematics/dynamics.

In FY11, we investigated how to use a cognitive architecture (ACT-R) to assist conventional computer vision algorithms in the perceptual process. In particular, we examined how to provide context to the computer vision system and to compensate for its failures. We also implemented a first version of a shared mental model (i.e., navigation cost map) that was updated by the robot. The robot observed the human’s actions and used Bayesian plan recognition techniques to infer what information was missing from the model. In FY12, we propose to address the *control* side of the architecture as well as the perceptual side (i.e., flow down in addition to flow up).

In FY12, we will design an architecture that embodies the following approaches:

- We will use ACT-R to handle the "cover the back door" mission. The implementation will be based on a cognitive model for how humans perform the task. The ACT-R model will generate perceptual and motion primitives for the robot to execute.
- We will use traditional robotics algorithms to handle the perception and motion primitives. ACT-R will provide both context and configuration parameters for both directives. ACT-R will also provide meta-cognitive monitoring in order to determine when and how the perception and motion primitives are failing.
- The architecture will produce a log of both mission instructions and candidate perceptions and actions for purposes of both offline imitation learning and online self-supervised learning.
- We will incorporate a shared mental model for cooperation with a human teammate. Previously, we inferred missing parts of the model based on robot observations of human behavior. We will produce the complementary behavior of inferring a teammate's actions based on the shared mental model. These inferences will be used to reduce required communication between teammates as well as the chances of misunderstandings.

The architecture will have strong ties to the world model, semantic perception, planning algorithms, tactical behaviors, life-long learning approaches, and HRI shared mental models.

Metrics for Evaluation:
We will evaluate the architectural design by devising different use cases and evaluating robustness to a broad range of conditions and potential failures. The robustness will be compared to competing architectures.

For the shared mental models, we will measure the effectiveness of the human-robot collaboration, including the mission tempo, ability to respond to unexpected situations, and misunderstandings. We will conduct our experiments with and without the shared mental model and measure the performance differences.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Design atomic actions that serve as basic building blocks for navigation and manipulation and serve as the primary interface between cognitive modules and robotics modules in the architecture.
Q2	Develop ACT-R model for "cover the back door" to sequence, configure, and monitor the atomic actions.
Q3	Extend shared mental model to support predictions of human teammate's behavior.
Q4	Conduct experiments with ACT-R model and shared mental model in simulation.

Subtask 2: Architectural Implications for Evaluating and Explaining Performance (CMU)

Objective and Benefits:
The objective of this task is to develop the ability of a cognitive architecture to introspect on its own operations to detect error conditions and provide robust remediation. The benefits include an increased level of autonomy, lessened need for human operators, and increased trust in robotic systems.

This research supports a broad range of basic architectural capabilities, especially Capabilities 4 ("Move autonomously through representative clutter") and 5 ("Moving through mobility challenges") with its emphasis on robust autonomy.

This effort is expected to continue for a number of years as we gradually develop increasingly complex cognitive capabilities and transfer them to the broader framework.

State of the Art:
Typical approaches to intelligent behavior are fundamentally static. Systems are programmed or trained to exhibit a particular set of behaviors. They are incapable of going beyond those initial behaviors because they lack the ability to introspect on their own past and current operations, which is the key to determining whether the current behavior will be successful.

Key Barriers and Approach:
Our technical approach relies on a hybrid cognitive architecture that can learn autonomously from its own operations, especially relying on its range of procedural and declarative memories and the combination of symbolic and subsymbolic inference mechanisms. A key insight of our approach is that the correct inference(s) to make and action(s) to take in a given situation are determined by our experience in similar situations, and most specifically by the discrepancy between our expectations and the actual outcomes. The key to successful introspection is the combination of rich memories, including episodic and semantic as well as procedural memories, at both the symbolic and subsymbolic level that enable robust statistical inferences. We will use as a guide the human ability to quickly detect and even anticipate situations where the current behavior will not be successful and efficiently revise plans to circumvent the problems.

We will define a new episodic memory module to remedy these problems. The function of that module will effectively be that of providing episodic memory traces that can provide knowledge of past operations to guide meta-cognitive introspection. This approach will be useful for all classes of problems related to various parts of the architecture, including visual (perception) modules, declarative memory (knowledge) modules, and procedural (control) modules. We initially focused on the third condition, the detection and remediation of procedural loops and other pathologies, as the most urgent since it concerns the inner loop of the cognitive architecture. We will later extend the approach to the perception and knowledge modules. Episodic memory will consist of a record of past contexts (roughly, the set of buffer contents in the ACT-R architecture), actions selected, and outcome. Introspection will compare the current situation to past instances and base expectations on the consensus of that past experience. Discrepancies from expectations will drive the detection and communication of problems.

This subtask involves collaboration with Troy Kelley at ARL and with Bob Dean at GDRS, especially as it concerns the definition of the structure and protocols of the world model. More broadly, we also anticipate interactions with other Intelligence tasks, especially those focused on behavior generation and mobility tasks including planning.

Metrics for Evaluation:
Progress is measured in terms of the degree of autonomy achieved by the resulting integrated system, especially metrics such as the breadth of situations successfully handled (e.g., range of obstacles or conditions overcome), robustness to disruptions (e.g., noise in sensors or mobility commands), and efficiency of behavior (e.g., time to overcome situations in which the robot is stuck).

We have developed a framework for introspection in the ACT-R cognitive architecture and started its implementation. Additional capabilities including episodic memory are also at the prototype stage.

We will exercise the system both in simulation and on robotic platforms. Capabilities will be initially tested on simple cognitive models, such as by introducing bugs in existing models then in increasingly realistic simulations. We will also support the system on the SS-RICS architecture at ARL and test it there on robotic platforms.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Development of episodic memory underlying basic introspective ability.
Q2	Application to diagnosis and adaptive communication of problems.
Q3	Application to remediation of problems.
Q4	Revision and broadening of introspective ability and improvements to communication and remediation performance.

Subtask 3: Robotics CTA Architectural Development and Integration (ARL)

Background:
As part of a 2004 Director's Research Initiative, the Human Research and Engineering Directorate (HRED) of ARL developed the Symbolic and Sub-symbolic Robotics Intelligence Control System (SS-RICS) in order to increase the autonomous capabilities of the Army's robotics assets. The development of SS-RICS has continued through 2012 and has resulted in a goal oriented cognitive architecture capable of robotic control. As a result of the development of SS-RICS, many lessons were learned related to the integration and use of a cognitive architecture in combination with traditional robotics algorithms. SS-RICS uses a symbolic level production system identical to ACT-R and has combined distributed statistical algorithms, running in parallel, in order to enhance the perceptual capabilities of autonomous systems. This structure (symbolic and sub-symbolic integration) is identical to the overall intelligence framework

proposed by the current Robotics CTA. As part of this subtask, lessons learned from the development of SS-RICS will be transferred to the RCTA architecture.

Objective and Benefits:
The objective is to transfer previous lessons learned from the development of SS-RICS to the RCTA architecture. The benefits are several, but primarily that previous mistakes uncovered during the development of SS-RICS will not be repeated and transferred to the RCTA architecture. Additionally, essential algorithms within the cognitive architecture (working memory, spreading activation) will be emphasized for speeded development within the RCTA architecture, and the interaction of these algorithms with the sub-symbolic components will be tested. The proper functioning and interaction of these key algorithms within the cognitive architecture is a primary objective for this subtask.

State of the Art:
Using a hybrid of a cognitive architecture in conjunction with sub-symbolic robotics algorithms for robotics control is a new concept and has only been tried in a few limited situations – but with some success. ACT-R currently has a few sub-symbolic algorithms which run in parallel for computing memory decay and learning; however, ACT-R does not currently have a variety of perceptual algorithms which run in parallel and feed the symbolic mechanisms. This limitation is being addressed with the addition of the LEABRA architecture into ACT-R.

The Naval Research Laboratory (NRL) has used ACT-R to control a variety of robots, and there is continued development of their hybrid system. Ron Sun has proposed using a hybrid system for the development of the CLARION architecture, but not for the control of autonomous systems. The cognitive architecture Soar has been used for robotics control and is continuing work in this area.

Key Barriers and Approach:
The interactions amongst a wide variety of algorithms which have stochastic outputs can be difficult to assess and evaluate in a timely manner, especially if they cannot be tested using a "faster than real-time" methodology – as with many simulations. However, most of the major algorithms for the hybrid cognitive/sub-symbolic architecture have already been defined within SS-RICS, and it is their further refinement that will be used as a primary approach for integration in the RCTA.

Metrics for Evaluation:
Algorithms will be proposed and tested within the RCTA architecture. Proper functioning and evaluation of these algorithms will be used as metrics for evaluation.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Transfer of lessons learned from SS-RICS to the RCTA.
Q2	Test the core algorithms for interactions within the entire system.
Q3	Iterative development and testing.
Q4	Continued testing within the RCTA architecture.

I2-2012 – World Model and Applied Intelligence Architecture

Level: 6.2 Applied Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Bob Dean	GDRS	I2-1, I2-2
Additional Investigators	Al Kelly	CMU	I2-3
	Dan Koditschek	UPenn	I2-4
	Dieter Fox	UoW	I2-5
	Christian Lebiere	CMU	I2-6
	Randall O'Reilly	UoC	I2-7
	Troy Kelley	ARL	I2-8

Objective:
In this task, we explore approaches to instantiate the World Model and Software Framework to enable the Intelligence Architecture defined by Task I1 and capabilities defined by the Capstone Vision. We investigate data representation, processing and query methods to define the core of the world model. The architecture is to be adaptive, and therefore, we research methods for the world model to automatically scale between platforms of different sizes and capabilities. To assist with moving beyond state of the art metric representations, we investigate how to learn to populate the semantic level using human interaction. We develop and integrate meta-cognitive techniques that use the introspective abilities provided by the world model to give the robot the ability to diagnose its own problems and keep improving upon its performance with minimal external input.

Background:
What is a world model? It is literally how the robot models the world. Most modern robotics architectures define at least a simple “world model” which buffers sensed data, combines the data over time. It is a repository of relevant information representing the robot’s belief of what the environment is. This belief is then used for making behavioral decisions. The world model is a critical component of an intelligent system, yet it is often architecturally a secondary consideration. For instance, it is common in robotics for algorithms to independently process sensor data into internal world model representations to accomplish a specific task. A common example of this is an occupancy grid used for obstacle avoidance. While the grid simplifies the planning operation, how can another operation like manipulation make use of this 2D representation? In many cases, they cannot. As a result, multiple algorithms perform duplicative and distributed calculations upon the same raw data sources, which in turn increase the load upon computing resources.

The state-of-the-art world model for robotics is, unfortunately, metric representations or simple semantic or topology representations implemented to solve point solutions. In very few cases does the state of the art cross between the representational disciplines or share the deeper information necessary to allow functional components to work collaboratively. Metric is like a warm blanket in that it is known, comfortable, and in many cases, this is because it makes specific point solutions easier to implement. This simplicity of design, however, is also a source of brittleness within robotic systems. A change in the underlying metric representation, such as moving a navigation solution to a smaller platform, could break the system, requiring updates to many levels of the software hierarchy and full regression testing taking weeks.

In a traditional robotic architecture, perception processes the inputs and generates a quantified description of the environment; higher-level processes such as planners take that representation of the world together with a description of the goals to be achieved and output a set of plans to be achieved; then mobility and manipulation effectors attempt to implement those plans. The task of each part of the system is exceedingly difficult in large part because of its isolation: perception has to classify inputs in purely bottom-up fashion without the benefit of top-down context and guidance; planning has to solve a combinatorial problem in high-dimensional space; and motor processes have limited ability to adapt the execution to achieve goals in a flexible way. If any part of the chain fails, then the overall system usually fails, with limited ability to adapt to changing circumstances. What is needed is the Meta-Cognitive ability for each part of the system to learn from and interact with the other in a way that leverages their capabilities.

Research Description:
Within this task, we seek to move beyond point-solution world model and architectural concepts. Central to the Intelligence Architecture is the concept of a Common World Model.

Our world model moves beyond the state of the art by representing the world using metric, semantic, and symbolic information. It joins these layers of information to define *objects* in the world which may be reasoned upon jointly using traditional geometric, symbolic cognitive algorithms and new computational nodes formed by the combination of these disciplines. It must understand how these objects relate to each other. Our world model includes the concept of *self-information* about the robot. By encoding current capability, component status, task execution state, and their histories, we track information which enables the robot to reason and adapt its performance using Meta-Cognition and Machine Learning principles. The world model also includes models of how aspects of the environment behave which enable prediction of future world states. These aspects include traditional kinematic and dynamics models of the robot as well as Shared Mental Models enabling interaction with human teammates and opponents.

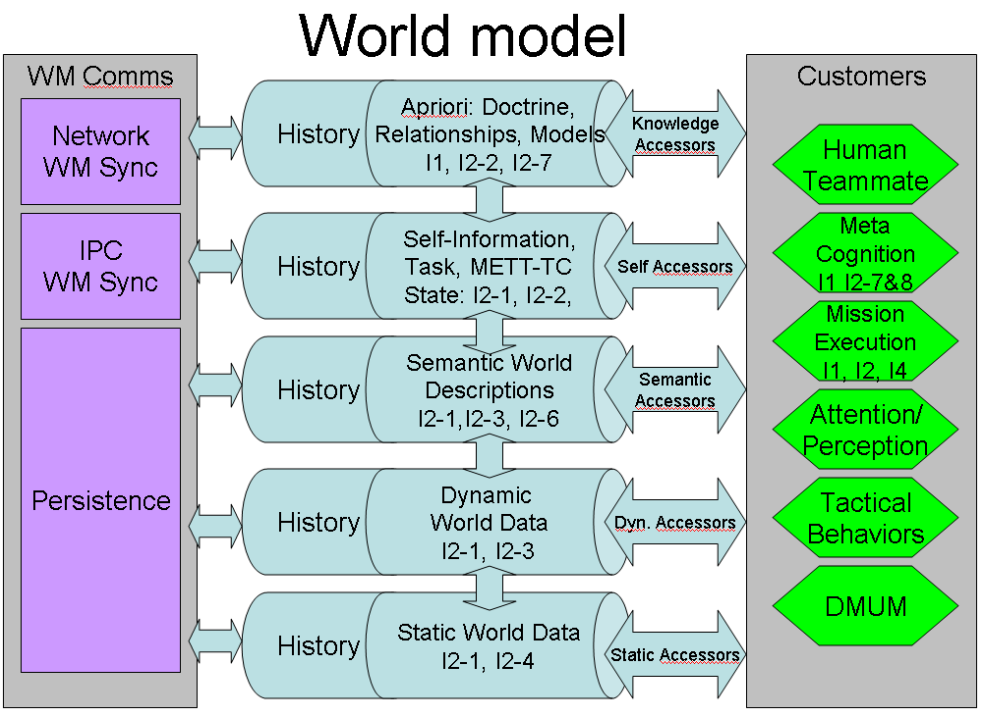

This task defines eight subtasks to research the World Model and Applied Intelligence Architecture. This effort centers around three areas: instantiation of the world model and Architecture, research into the next step of the world model, and cross-task area Architecture concepts.

Subtask	Org.	Provides	Capstone Need
Common World Model Core	GDRS	Core data management of world model state	1, 3, 4, 5
Linking Cognitive Reasoning with the Common World Model	GDRS	Enables cognition to reason on the world model	1, 3, 4
Resolution Independent Raster Map	CMU	Investigates efficient processing of metric information and potential architecture improvements	1, 3, 4, 5, 6
World Model Representations for Unique Mobility	UPenn	Investigates how to link higher-level intelligence with a small platform	4, 5, 10
Interactive Semantic 3D Mapping	UoW	Learning semantic object definitions from human input	1,3, 10
Hybrid Cognitive Architecture for Diagnostics and Confabulation	CMU	Develops principals for applying meta-cognition	3, 5, 10
Focus of Attention, Active Perception, and Learnability	UoC	Investigates using perceptual error signals to learn and improve motor controller	10
Meta-cognitive Development for Cognitive Architectures	ARL	Develops meta-cognitive layer for ACT-R	3, 5, 10

Our approach centers around a blackboard-style implementation. Think of a room of experts who use a blackboard to communicate. A problem is written on the blackboard, and a moderator directs the experts to work on the portion focused on their area of expertise. When a solution is available, the expert places it on the blackboard for others to use. In our Intelligence Architecture, the world model is both the blackboard and the moderator, managing communication between computational nodes (experts) and handling assignment of goals and subgoals.

To instantiate the world model and Architecture, we provide three subtasks. The main element of the world model is the object-oriented data store which embodies not only the raw data management but also a large core of accessor methods which allow the data to be emplaced, read out, or manipulated in scores of different ways depending on the needs of the Common World Model user. Subtask I2-1: Common World Model Core is a multi-year task which provides this basic functionality. During FY11, this task explored interfaces to cognition, low-level planning, and perception. One result of import was the need for a task to focus solely upon the interface between the world model and ACT-R. In Subtask I2-2: Linking Cognitive Reasoning with the World Model, we investigate how to improve the Cognitive/World Model link by providing semantic and symbolic query methods which Cognition can use to understand the world. The remaining aspect of the core world model is interaction with the environment. While this aspect of the architecture is researched in other task areas, such as Task I4: Adaptive Tactical Behaviors, interaction with smaller, computationally limited, and mechanically capable

platforms like RHex requires new methods to map the world model into an appropriate space. Under Subtask I2-4: World Model Representations for Unique Mobility, we investigate how to simplify the world model to a "minimialist" representation, linking smaller, mechanically capable platforms such as RHex to higher-level Intelligence.

There are multiple potential sources of semantic knowledge. Initial semantic understanding of the environment comes from Perception. To reason upon this data, the world model needs to provide methods to infer relationships and interactions between objects and concepts in the environment. Our Cognitive/World Model subtask investigates using semantic relationship databases such as OpenCyc. The Core subtask investigates how to extract topological representations from metric data using techniques developed by ARL researchers. Subtask I2-5: Interactive Semantic 3D Mapping develops approaches to use human input, including speech and gestures, to learn semantic meanings for objects in the world.

Finally, we investigate cross-task architectural concepts related to the world model with three subtasks. Subtask I2-3: Resolution Independent Raster Maps seeks new methods of representation and query of low-level metric information such as point clouds. The promise is to improve efficiency of metric algorithms like Perception and Tactical Behaviors by factoring-out common/redundant processing at those layers. By using resolution independence, we gain the ability for the world model to scale and adapt between platforms.

Subtask I2-6: Hybrid Cognitive Architecture for Diagnostics and Confabulation and Subtask I2-7: Focus of Attention, Active Perception, and Learnability investigate how to apply Meta-Cognition to the overall Architecture in ways that leverage perception, cognition, and action in deeper ways than the traditional stove-piped approach. The advances are two-fold. First, each module has access not only to the final outcome of another module's computation but also to internal information that provides deeper information on the nature of the processing, such as internal representations. That richer interaction results in the ability of cognitive modules to make meta-cognitive judgments on the operation of other modules, such as whether a perceptual module has actually been trained on a particular type of input. Second, it allows for learning signals to propagate throughout the architecture. This is particularly important for end-to-end tasks that integrate perception, cognition, and action in a tight interactive loop. It allows the outcome from external actions to propagate as a reward signal through not only the motor modules that executed the actions but also the cognitive modules that decomposed the task and scheduled the actions, all the way to the perceptual modules that processed the sensory information upon which actions were based.

Metrics:

- The metrics associated with this task are related to the efficiency of the Common World Model with respect to the various clients which include Cognition, Adaptive Behaviors, and Perception. For example, response time to queries form cognition, the rate at which new information is published to cognition, and the time required and correctness of linking between metric and symbolic constructs.

- In some cases, the Common World Model replaces traditional interfaces to raw sensor data and/or Perception outputs. For these, we compare the performance of perception algorithms based on our new representation against performance using more classical techniques.
- For learning links between semantic and metric representations, we will measure progress by the amount of semantic information that can be taught to a robot by a human either guiding the robot through the environment or providing remote annotations.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations: This task, by its nature, is highly collaborative because the Intelligence Architecture is what enables many aspects of Alliance research to come together. First, there is a tight collaboration with Task I1 as we implement the architecture developed there, and by association, with their cross-task collaboration efforts including HRI Tasks H1, H2, H4 and H5. We collaborate with Tasks P4 and P6 to define the data formats output by Perception as well as how to direct Focus of Attention for Perception. We collaborate with Integrated Research Subtask IR5-2 to define representation of moveable objects with pose in the Common World Model. Within ARL, we collaborate with Troy Kelly who has developed the hybrid robotic control architecture SS-RICS, which shares many aspects of our cognitive architecture, and with Dr. MaryAnne Fields on world model efforts.

Milestones:

	Quarterly Milestones and Deliverables
Q1	World Model Core implementation integrating FY11 work. Evaluation of sources of semantic relationships. Define prototype data structures, and examine tradeoffs for dynamic object supporting. Develop and demonstrate robust normal and surface estimation on static point clouds datasets with geometry known a priori. Investigate incorporation of speech annotations (keywords) during mapping to label objects in environment. Development of end-to-end model using symbolic action system. First draft cerebellum model completed and tested on range of perceptually-driven motor tasks, and design and configuration of virtual environment for doorknob task.
Q2	World Model Topology investigation. Adapt point clouds to moving scenes indoors. Speech annotations are parsed based on temporal alignment of speech signal and Kinect mapping data. Introduce "frustration" triggered dynamical replanner [7] for escaping non-convex obstacles in outdoor environments. Integration of neural action model into symbolic architecture. Further evaluation and optimization of cerebellum model, scaling up to harder problems; initial development of integration with ACT-R for outer-loop control in doorknob task.
Q3	Persistent storage of world model information. Investigation of episodic memory within Common World Model. Extend point cloud data structures for use real-time with potential odometry errors. Start adding gestures and human body motion to provide semantic information in the environment. Assess improvement of “frustration” triggered dynamical replanner. Generalization of performance to broader set of doors through perceptual learning. Further development and evaluation of integrated perceptual-motor and ACT-R model on doorknob task; should be able to recognize and manipulate 20 different doors.
Q4	Integration of focus of attention for perception. Perform point cloud segmentation to aid in surface estimation and registration for cluttered and dynamic environments. Combine speech, gesture, and human body motion into semantic labeling system. Generalize task specification for RHex. Scaling up of architecture and optimization of meta-cognitive learning processes. Scaling up of the model to 100 different doors and generalization to 20 novel doors, along with continued performance optimization of all components.

Related Research:

This task works with all of the other Intelligence tasks to determine functionality and interfaces to the world model and architecture defined by I1 to support tactical behaviors (I3 and I4) as well as learning components (I5 and I6). We also work with Perception tasks (P3, P4, P5) to determine the interface for Focus of Perception, representation of shared perception data, and the

bi-directional Perception-World Model link. We assist Integrated Research tasks (IR5) to define representations for dynamic movers, and we interface with DMUM tasks (M6 and M8).

Integration and Assessment Activity:
This task affects all of the Integrated Research Assessments as it implements the architecture, as defined by Task I1, and the architecture forms the basis for integration. Of particular importance for FY12 is IRA3: Intelligent Navigation. In order to meet the needs of IRA3, subtask milestones have been defined, and task-level assessments and integration efforts will be used to define the world model capabilities to support the capability goals of IRA3.

Subtask 1: Common World Model Core (GDRS)
We seek to investigate data representation and processing methods to define the core of the Common World Model, through which we are able to integrate world model concepts developed by other tasks within the Alliance. The core would then be available to enable future research in those areas.

Objective and Benefits:
We seek a world model representation capable of storing and indexing metric, semantic, and cognitive information. The world model should support multiple capabilities for the robot from navigation through manipulation as well as adaptive behaviors supplied by semantic knowledge, learning, and cognitive control. Additionally, it should be resolution and platform independent, allowing it to scale between a variety of platform sizes and computation availability. We believe such a world model will allow us to move beyond the state of the art of task specific world model implementations.

As a multi-year effort, we shall collaboratively integrate the world model principles and techniques developed by 6.1 and 6.2 researchers in Perception, Intelligence, and Human-Robot Interaction into a cohesive world model framework and implementation for the relational and cognitive intelligence levels.

State of the Art:
It is common in robotics for algorithms to independently process sensor data into internal world model representations to accomplish a specific task. A common example of this is an occupancy grid used for obstacle avoidance. While the grid simplifies the planning operation, how can another operation like manipulation make use of this 2D representation? In many cases, they cannot. There is prior work in hierarchical world models linking metric and semantic information. A common theme is using a topology graph to connect separate, smaller metric maps. Usually, the scope of the work is limited to the well-known exploration and navigation problem, with user assisted labeling to provide semantic context.

Key Barriers and Approach:
Development of a Common World Model is a complex and daunting task. During FY11, we began the process of defining how to build the world model. We met with collaborators within

the Alliance to identify requirements for the world model, resulting in a list ranging from low-level geometric queries through higher-level queries of semantic relationships and temporal scale. From this list, we were able to identify the need to support run-time query of different data-types in an automated fashion. In order to meet this need, the Integrated Research code generator developed for RFrame was expanded to include Type Registration, which allows run time introspection and lookup of world model data structures.

We then defined isolated aspects of the Common World Model to identify and evaluate their interfaces. These consisted of interfaces to low-level metric planning, perception history, and feeding cognition and were evaluated independently of each other. A test interface to link ACT-R to a simple world model was implemented for the first Integrated Research Assessment. Lessons learned from that point to the interface being more complex and dynamic than initially expected, resulting in the creation of Subtask I2-2: Linking Cognitive Reasoning with the Common World Model to focus on this link specifically. At the time of this writing, approaches for interfaces to low-level metric planning and perception have been identified, and we are in the process of implementation with a goal of task-level evaluation during the fourth quarter of FY11.

For FY12, we investigate how to bring these separate components together into a single Common World Model and expand upon them to support the next steps of Alliance research. By doing so, the Common World Model itself moves beyond the state of the art into a realm where algorithms may leverage the common world estimation and processing and evolve independently of the underlying data representations.

To further research the potential of the Common World Model, we will additionally investigate in FY12:

- Efficient data representations to link and index metric, semantic, and symbolic information within the world model core. These representations need to support requirements from Perception and Intelligence to support advanced reasoning, such as multi-hypothesis classification, storage of statistical meta-information generated by ACT-R, and statistic outputs from Learning processes.
- Methods to group objects into higher-level semantic concepts like “desk” + “chair” + “walls” = office, and connecting the concepts through a hierarchy to enable topological reasoning. For example, Perception may provide the ability to identify a building and the street in front of it. This goal will allow such inputs to be combined over time to build a topological representation of a city.
- A task may require knowledge of an area the robot has previously seen but is no longer stored in direct memory. We seek representations for persistent storage of collected world model information to allow efficient use of computational resources (such as memory) by applying mission context to the retrieval. A command such as, “Go to Bob's office, and find the stapler” consists of two goals: navigation and search. The first requires semantic/topologic knowledge of the goal location, while the second could depend upon the metric state of the goal location.
- Identification of how to log world model state for later consumption by offline learning algorithms.

Metrics for Evaluation:
The world model is an integration vehicle and instantiation of architecture functions defined by Task I1 and other subtasks. Through it, we are able to integrate disparate research components. As such, its capability is a prerequisite for many integration targets.

Capabilities not covered by the IRA plan will have defined mini-assessments to determine resources metrics and client usefulness. For example, evaluation of the "move to the back door" vignette defined in I1-1 and integration with perception to evaluate "Focus of Attention."

	Subtask 1 Quarterly Milestones and Deliverables
Q1	World Model Core implementation. Integration of FY11 work into shareable software product.
Q2	Topology support.
Q3	Persistent storage of world model information. Ability to select portions of stored information given mission context.
Q4	Integration of focus of attention for perception. Investigation of temporal linking and indexing (time permitting).

Subtask 2: Linking Cognitive Reasoning with the Common World Model (GDRS)

Objective and Benefits:
We seek methods to link Cognitive Reasoning to the Common World Model. This will allow cognition to reason upon the world model and update world model state to be shared with other system components. This link supports the Capstone Assessment Vision by enabling the Intelligence Architecture to use ACT-R to help solve tasks. For example, assisting with task decomposition and understanding by translating high-level commands into executable sequences of Atomic Behaviors (i.e., "move to the back of the building" could be converted by ACT-R into a series of logical steps, such as "find building; locate back of building; navigate to back of building; ..."). In order to generate a list of subgoals, however, ACT-R will require methods to query the world model to reduce its uncertainty as well as to update the world model with results of cognitive planning operations.

This is a multi-year effort, expected to continue to the end of the program.

State of the Art:
In the past, custom perception inputs have been created to feed cognitive reasoning systems which convert environment data into the proper reasoning format, such as CLIPS/R2 User Functions and Facts, ACT-R chunks, or SOAR's I/O links. Systems like SS-RICS (Kelley and Avery, 2011) have been developed specific to the needs of robotic control, while others like SAL (Jilk et al., 2008; Lebiere et al., 2008) are being adapted to the robotic environment. Research by Galindo et al. (2005) has been performed to use "anchor" data structures to link metric and symbolic data structures for indexing and reasoning purposes. In some cases, these systems are stove-piped to a specific environment and model implementations resulting in brittle systems.

Recently, work by Best and Lebiere (2009; 2010) has been performed to investigate using middleware abstractions to enable "plug and play" combinations of simulation environments and cognitive models.

Key Barriers and Approach:
This subtask will advance the state of the art by establishing the bi-directional link between metric, semantic, and symbolic information and reasoning with the RCTA architecture.

During FY12, we propose to investigate how to expand the types of information included in the Common World Model to support further Cognitive capabilities and make this information accessible through a generic cognitive query interface. For example, to determine how to find the back of a building, ACT-R might query the world model for the semantic relationship between "back" and "building," resulting in a chain of connections from "back" to "side" to "front." This knowledge may then result in creation of a subgoal for perception to find the "front" of the building. For the world model to support such a query requires: finding and evaluating sources of semantic relationships (such as OpenCyc for a priori relationships or FY12 research from I2-3 into online learned relationships), expanding the world model Core to include semantic relationship storage, as well as supporting the query through the ACT-R/Common World Model link.

Additionally, we seek methods by which Cognition and Meta-Cognition may update the Common World Model with information it learns from its planning and meta-cognitive processes, including strength of semantic relationships between concepts, changes in concepts based on experience (i.e., paths around a house), likelihood of related concepts, and temporal links, based on experience. This effort should allow Cognition to implement Episodic Memory, allowing application of past experience to current task execution.

During FY11, we successfully integrated ACT-R and perception using a simple world model interface. We know from other research that ACT-R has the ability to respond to various environments; here we apply this ability to the realm of robotics.

Metrics for Evaluation:
These components will be evaluated with efforts in Task I1, using the "move to the back door" vignette described in I1-1, and will assist by implementing the data representations and protocols to allow ACT-R to reason on the world model and enable this task-level integration and evaluation. The metrics associated with this subtask are related to the efficiency of the world model to cognitive link. Specifically, response time to queries form cognition, the rate at which new information is published to cognition, and the time required and correctness of linking between metric and symbolic constructs.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Investigation and evaluation of sources of semantic relationships.
Q2	Integration trials of semantic relationship sources into world model, and trial query using simple ACT-R reasoning modules.
Q3	Investigation of how to support episodic memory within the Common World Model.
Q4	IRA support.

Subtask 3: Resolution Independent Raster Map (CMU)

Objective and Benefits:
We seek to develop an efficient, low-level capacity to process range data that satisfies the requirements of all higher-level processes in perception, world modeling, and intelligence. This module performs low-level processing that can be factored (because it is common) from all higher-level processing.

The research will support the capstone assessment by enabling many of the higher-level cognitive tasks required to a) think, b) look, and c) move. Point cloud features will be a partial basis for detecting motion in the scene; for recognizing the building, the back of the building, the points of egress; and for understanding the activities of any humans. When moving, point cloud features can form the basis for sensing ego-motion in a manner that is not corrupted by wheel slip or inertial sensor drift. Such accurate motion indications will then form the basis for the construction of spatially expansive models that are needed in order to recognize, plan motions, and communicate with humans.

As a research effort that is focused on innovative capacities embedded in the fundamental infrastructure, this effort is likely to continue throughout the program.

State of the Art:
Due to its relatively unique objectives, the previous RCTA program has produced much of the prior art in world modeling for large outdoor and unstructured environments. Our proposed efforts differ from this prior work in several ways. First, we are targeting methods which are more general so as to be more applicable to many related domains, not just specialized problems like detecting concertina wire [6]. While work has been done on fast processing of point clouds [3], we wish to improve general scene understanding through arbitrary queries over space and time and to ameliorate aliasing artifacts from the voxelization process itself. In contrast to the multi-scale interest region work done under the last CTA [2,4], we want our methods to be fully online, suitable for a mobile robot running in real-time. By efficiently encapsulating and abstracting information through segmentation and higher-level geometric primitives, we believe that we can achieve a much richer set of scene recognition capabilities without having to rely solely on training data and PCA descriptors. Some prior art comes from the SLAM community [1,5]. These map representations are usually specialized for SLAM whereas our strategy is to

include these localization capabilities as well as functionality for spatiotemporal queries, object recognition and tracking, and multi-scale scene understanding.

Key Barriers and Approach:
Key barriers include the absence of any science that studies the appropriate low-level geometric vision architecture to support every possible use of perception data in a mobile and autonomous system. We cannot afford to have every algorithm touch the raw perception data, but to avoid it, we must provide an abstraction layer that serves every possible purpose.

Hierarchical techniques will be used to produce feature descriptions across all of scale space. Super resolution techniques are proposed to ameliorate common deficiencies due to aliasing. Canonical features are computed which are common across all applications.

We hope to improve general scene understanding through arbitrary queries over space and time and to ameliorate aliasing artifacts from the voxelization process itself. We want our methods to be fully online, suitable for a mobile robot running in real-time. We hope to achieve a much richer set of scene recognition capabilities and a more generally useful system by computing many more features beyond PCA descriptors.

We believe that we will succeed because the appropriate features to use are already well known. Hardware acceleration will be possible. What has been lacking to date is the systems-oriented viewpoint that the RCTA program provides. The need for efficient low-level processing in perception is now recognized as fundamental. The previous CTA has already invested heavily and with success in these ideas. The proposed work is a continuation along this successful theme.

Metrics for Evaluation:
Progress relative to our goals will be assessed by comparing the performance of perception algorithms based on our new representation against performance using more classical techniques. A basic working instance of the raster world model has been completed. Some perception algorithms (recognitions, surface segmentation) have been tested. Progress relative to our goals will be assessed by comparing the performance of perception algorithms based on our new representation against performance using more classical techniques.

Because the previous RCTA concentrated on data processing for obstacle perception, recognition, and classification of outdoor scenes, the planned next year of activities will focus on low-level processing to support continuous perception aiding to the localization system.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Develop and demonstrate robust normal and surface estimation on static point clouds datasets with geometry known a priori.
Q2	Use results from Q1 to adapt to moving scenes, using a small robot in simple indoor settings.
Q3	Build registration capabilities into the data structure for use in real-time systems with possibly distorted odometry.
Q4	Perform point cloud segmentation to aid in surface estimation and registration for cluttered and dynamic environments.

References:

[1] K. M. Wurm, A. Hornung, M. Bennewitz, C. Stachniss, and W. Burgard, OctoMap: A probabilistic, flexible, and compact 3D map representation for robotic systems, ICRA 2010 Workshop on Best Practice in 3D Perception and Modeling for Mobile Manipulation, 2010.

[2] R. Unnikrishnan and M. Hebert, Multi-scale interest regions from unorganized point clouds, Computer Vision and Pattern Recognition Workshops, 2008. CVPRW '08. IEEE Computer Society Conference on, 2008, pp. 1-8.

[3] J. Lalonde, N. Vandapel, and M. Hebert, Data Structures for Efficient Dynamic Processing in 3-D, The International Journal of Robotics Research, vol. 26, 2007, pp. 777-796.

[4] R. Unnikrishnan, J. Lalonde, N. Vandapel, and M. Hebert, Scale Selection for the Analysis of Point-Sampled Curves, 3D Data Processing Visualization and Transmission, International Symposium on, Los Alamitos, CA, USA: IEEE Computer Society, 2006, pp. 1026-1033.

[5] R. Triebel, P. Pfaff, and W. Burgard, Multi-Level Surface Maps for Outdoor Terrain Mapping and Loop Closing, Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on, Beijing, China: IEEE, 2006, pp. 2276-2282.

[6] N. Vandapel and M. Hebert, Finding Organized Structures in 3-D LADAR Data, Transformational Science and Technology for the Current and Future Force - Proceedings of the 24th US Army Science Conference, Orlando, Florida, USA: 2006, pp. 161-168.

Subtask 4: World Model Representations for Unique Mobility (UPenn)

Develop a mathematically well-formulated dynamical model for Episodic-memory-context-based Localization, Mapping, Focus of Attention, and Active Perception motivated by cognitive neuroscience models that can be integrated into the world model representation of Subtask I1[2] and used as the basis for integrating machines from M6 and M8 tasks into command and control tasks at appropriately flexible levels of the emerging RCTA control hierarchy.

Objective and Benefits:

We seek to develop and apply task/environment representations capable of enabling a mechanically complex platform to operate as autonomously as possible in a complex natural environment. We aim to increase the autonomy of our mobile platforms by relegating to the authority of their mechanical competence as much of the environment interaction as can be achieved while incorporating only as much information about the world as required to achieve a

[2] Our previous 2011 I1-2 subtask has moved to M6-4; this work will help bind our M6 and M8 tasks to the other I1 tasks moving ahead.

task. This research will enable humans to command unique mobility at a much higher level than presently possible, and the algorithms can be used by higher-level planners to assign tasks to DMUM platforms in the capstone assessment. For example, in the "Cover the back door" scenario, the autonomous navigation indoors and outdoors will allow unique mobility platforms to execute the plan without needing constant replanning.

This project represents a continuation of the "minimalist" world model introduced in [1]. The combination of fundamental questions about world models and commitment to provably correct methods imply a multi-year effort, likely exceeding the duration of the RCTA.

State of the Art:
A year ago, we demonstrated the state of the art in stair climbing with a legged robot at the 2011 RMB where a human driver had to aligned the robot and control the gait transitions. The robot was also capable of autonomously climbing hills, but only ones without obstacles. Over the past year, we have demonstrated the beginnings of autonomous legged navigation on both stairs and hills. Specifically, our first refereed publication resulting from this project [1] documents how this is the first time any single platform has achieved autonomous ascent of both indoor and outdoor terrain. Task representation and execution via sequentially composed basins that we introduced to the robotics literature more than a decade ago [6] has become a well accepted, established approach to higher-level robot motor control.

Key Barriers and Approach:
We use machine preflexes [2][3][4] to overcome mechanical barriers and use dynamical task representations [5] to command higher-level goals. This work ties our mobility tasks (M6 and M8) directly into the RCTA assessment scenarios (IRAx) and provides the intelligence tasks (I1) another conduit into those assessments.

Metrics for Evaluation:
We conduct extensive indoor and outdoor experiments to assess our degree of autonomy and failure rate against environmental conditions [1]. We measure the length of traversal through particular environments (characterized by type of slope and density of obstacles) along with failure rates and/or rates of human intervention required. We brought the present working version of this algorithm to the 2011 IRA1 FTIG assessment and are now digesting that experience. For example, our IRA1 trip report shows that overall the robot successfully climbed 100% of the regular hills and 83% of the regular stairs. The hill climbing algorithm had some trouble in the alley with non-convex obstacles, and the stair climbing algorithm had trouble with more challenging stairs and landings. We are still in the process of reviewing the results and are working on ways to improve these results. We expect to learn a fair bit about how much more hardening is required to make it truly useful as well as what extended levels of autonomy in what sorts of environments to aim for as the next step.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Review failure modes discovered in IRA1, and work to refine the algorithms for future tests. Participate in IRA2 with improved behaviors.
Q2	Introduce "frustration" triggered dynamical replanner [7] for escaping non-convex obstacles in outdoor environments.
Q3	Continue implementing "frustration" triggered dynamical replanner, and assess improvement with respect to non-convex obstacles at IRA3.
Q4	Generalize the tasks specification to include extensions such as downstairs, wall following, contour following, and compositions of behaviors. Explore potential for partition-based Bayesian target locator [8] for improved stairwell ascent.

References:

[1] A. Johnson, M. E. Hale, G. C. Haynes, and D. E. Koditschek, "Autonomous Legged Hill and Stairwell Ascent," in *IEEE International Workshop on Safety, Security, & Rescue Robotics, SSRR*, Kyoto, Japan: (to appear), 2011.
[2] D. E. Koditschek, R. J. Full, and M. Buehler, "Mechanical aspects of legged locomotion control," *Arthropod Structure and Development*, vol. 33, no. 3, pp. 251-272, 2004.
[3] I. E. Brown and G. E. Loeb, "A reductionist approach to creating and using neuromusculoskeletal models," in *Biomechanics and Neural Control of Posture and Movement*, NY: Springer, 2000, pp. 148-163.
[4] P. Holmes, R. J. Full, D. E. Koditschek, and J. Guckenheimer, "The dynamics of legged locomotion: Models, analyses, and challenges," *SIAM Review*, vol. 48, no. 2, pp. 207-304, 2006.
[5] D. Koditschek, "Task encoding: toward a scientific paradigm for robot planning and control," *Robotics and autonomous systems*, vol. 9, no. 1-2, pp. 5-39, 1992.
[6] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek, "Sequential Composition of Dynamically Dexterous Robot Behaviors," *The International Journal of Robotics Research*, vol. 18, no. 6, pp. 534-555, 1999.
[7] S. Revzen, B. D. Ilhan, and D. Koditschek, "Dynamical Trajectory Replanning for Uncertain Environments," p. (in preparation).
[8] A. De, W. Moran, S. Howard, and D. Koditschek, "Localization over Partitions," p. (in preparation).

Subtask 5: Interactive Semantic 3D Mapping (UoW)

This effort will develop approaches to incorporate human input for semantic mapping. A person will be able to annotate objects and places in an environment via speech and gesture. This information will be incorporated into a 3D map built with Kinect-style depth cameras.

Objective and Benefits:
Understanding high-level semantic concepts is extremely important for robots to communicate with people in a natural way and to execute high-level tasks. We propose to develop techniques that enable people to **interactively provide robots with semantic information about an environment**. Semantic information relates to objects, places, and spatial relationships in an environment. A person can tell a robot about objects and places using various modalities, such as speech, gesture, and touch displays, thereby augmenting the robot's autonomous recognition

capabilities. Understanding and acquiring semantic information is crucial to execute commands like, "Go to the back of the building, and cover the back door." This work provides key capability elements in the Capstone Assessment Vision, including 1 ("Understand commands") and 3 ("Understand world through perception and WM").

State of the Art:
Researchers start investigating ideas for enabling a person to label objects by explicitly pointing at them. Here, we will incorporate far more complex body motions to indicate objects and places along with speech annotations.

Key Barriers and Approach:
The key idea is to augment a robot's object and place recognition capabilities with an ability to understand human annotations provided interactively either by a person in the environment or by a remote operator. We will develop learning techniques that combine human gesture and speech annotations as the robot moves through the environment and builds a 3D map.

This effort will succeed because we can build on extensive experience in Kinect-based 3D mapping and object and place recognition. Kinect will allow us to extract rich human body motion for indicating desired context.

The effort ties strongly into our research activities in Task I5 on a learning framework for parsing speech and gestures for general human-robot interaction. Here, we focus on applying that framework to the specific task of map annotation which will allow us to develop more constraint solutions.

Metrics for Evaluation:
We will measure progress by the amount of semantic information that can be taught to a robot by a human either guiding the robot through the environment or providing remote annotations. A specific task could be: Given a set of objects and places in a building, how many of these can a person successfully "teach" the robot in a certain amount of time? Success can be measured by the robot's ability to later find those objects and places when they are only referred to by natural language, not xy coordinates.

Progress can be measured by achieving milestone capabilities with regard to the types of annotations, objects, and places that the robot can understand.

	Subtask 5 Quarterly Milestones and Deliverables
Q1	Investigate incorporation of speech annotations (keywords) during mapping to label objects in environment.
Q2	Speech annotations are parsed based on temporal alignment of speech signal and Kinect mapping data.
Q3	Start adding gestures and human body motion to provide semantic information in the environment.
Q4	Combine speech, gesture, and human body motion into semantic labeling system.

Subtask 6: Hybrid Cognitive Architecture for Diagnostics and Confabulation (CMU)

Objective and Benefits:
The theoretical objective is to develop principles for the deep integration of perception, cognition, and action that maximizes the strengths of each component in the context of their interaction in real-world interactive tasks. The practical objective is to apply the resulting architecture to achieve open-ended learning in perceiving and acting upon the world in completely unsupervised autonomy.

This research generally prototypes and enables the integration of a wide variety of perceptual, cognitive, and motor components in the overall program architecture. More specifically, it directly supports Capabilities 3 ("Understand world through perception and WM"), 5 ("Moving through mobility challenges"), and 10 ("Enter building by opening doors and avoiding or moving obstacles") with its combinations of robust autonomous perception and manipulation.

This effort is expected to continue for a number of years as we gradually develop increasingly complex capabilities in a framework combining cognitive architectural with neural perception and motion/action planning and to transfer it to the broader framework.

State of the Art:
Cognition, perception, and motion have generally been seen as independent, operating separately on distinct principles. When integrated into an overall architecture, they exist as separate opaque modules only interacting at a coarse grain scale that prevents deep leverage between their internal operations.

Key Barriers and Approach:
Our framework combines a general-purpose cognitive architecture with neurally plausible perception and motion/action components to enable a deep integration of cognitive, perceptual, and motor functions. Particularly essential is the availability of meta-cognitive introspection signals in the workings of perceptual and motor modules as well as the integration of cross-architectural reward signals that back-propagate from the results of actions in the world through cognition and perception to enable open-ended learning and autonomy.

We will build on the integration of perceptual and cognitive modules performed in this year's task. A symbolic action module will be initially used to separate the development of the overall architecture from that of the neural action module. The high-dimensional problem of action planning is then decomposed into higher-level cognitive organization and low-level movement, simplifying the complexity of the problem and allowing the leveraging of cognitive methods such as learning of subproblem solutions. After integration between symbolic system and action module, feedback from the motor execution can be back-propagated through the cognitive decisions back to the perceptual classifications that led to the actions. For instance, the architecture can then learn that the reason an action failed is that it applied an existing technique to a new type of door handle. After the perceptual system learns to make the distinction, cognitive processes can explore the space of actions that can apply to it, and the motor module can also fine-tune its procedures as it gains experience.

By using an architectural blueprint as the only system to achieve robust autonomy and general intelligence, biologically inspired approaches have a guide that constrains architectural development toward solutions that favor generality and open-endedness over narrow task-specific optimization.

This effort is closely related to subtask I2-7 (Randy O'Reilly, University of Colorado). Our effort focuses on the development of meta-cognitive capabilities to supervise the learning of perceptual representations and motor routines, while the UoC group focuses on the development of neural algorithms for perception and action. Our effort will also guide subtask I2-2 (Robert Dean, GDRS) in providing inputs into the interface between the world model and cognitive, perceptual, and motor components of the overall architecture. More broadly, we also anticipate interactions with other Intelligence tasks such as those focused on Behavior Generation and Learning as well as with efforts in Perception and Manipulation and Mobility.

Metrics for Evaluation:
Progress is measured in terms of the degree of autonomy achieved by the resulting integrated system. Metrics include breadth of situations successfully handled (e.g., diversity of objects recognized), robustness to disruptions (e.g., perceptual variations or randomness in motor operations), and efficiency of behavior (e.g., time to learn successful behavior, deviation from optimal behavior).

We have developed a framework for the integration of the ACT-R cognitive architecture and the Emergent neural simulation. We have integrated ACT-R with a neural visual module and used meta-cognitive signals to differentiate between trained versus novel stimuli as well as to achieve unsupervised learning and labeling of new categories of objects. We are in the process of integrating cognition with a neurally-inspired motor system.

We are exercising the system in a simulation environment, gradually introducing complexity as well as new challenges and capabilities. Intermediate progress will be measured by the growing range and complexity of situations successfully handled.

	Subtask 6 Quarterly Milestones and Deliverables
Q1	Development of end-to-end model using symbolic action system.
Q2	Integration of neural action model into symbolic architecture.
Q3	Generalization of performance to broader set of doors through perceptual learning.
Q4	Scaling-up of architecture and optimization of meta-cognitive learning processes.

Subtask 7: Focus of Attention, Active Perception, and Learnability (UoC Boulder)

Objective and Benefits:
Develop an integrated perceptual and motor system that uses perceptual error signals to learn to improve motor control abilities, without requiring elaborate inverse models of the motor plant and environment. There are many sources of evidence that the human system is similarly perceptually-driven, but leveraging rich high-dimensional perceptual signals for motor control has proven difficult for traditional computer-science-based approaches. Thus, we are leveraging biologically-based computational models with powerful error-driven learning mechanisms that can extract useful motor control signals from high-dimensional perceptual data. These models are based on the biology of the cerebellum, parietal cortex, and motor cortex.

This research attacks a very general problem that is central to any robotic system that acts in the world. We are particularly focused on capability 10 in the Capstone Assessment Vision (enter the building by opening doors, avoid or move obstacles). Other missions that require this form of physical interaction with the world include inspecting possible IEDs or picking-up supply items for logistics missions.

State of the Art:
By focusing on the reduction of an arbitrary error signal, our perceptual-motor system is both computationally efficient and highly adaptive. There are no forward models, inverse-dynamics problems, or combinatorial explosions in trajectory planning. It can leverage high-dimensional perceptual inputs in ways that existing approaches have not been able to do.

Key Barriers and Approach:
We are using a high-capability visual perception system based on our Leabra architecture which can recognize 100 different object categories within the same network and generalize to novel exemplars at over 90% accuracy (O'Reilly et al., under review). This perceptual system then drives motor error-driven learning via a biological model of the cerebellum that is consistent with a wide range of biological data on learning and processing in the cerebellum. This cerebellar model solves the difficult problem of error-driven learning in motor space, where correct target training values are unavailable, by leveraging motor opponency and synergies, which are subject to reorganization under the influence of perceptual error signals. Oscillatory firing dynamics and phase-resetting synchrony provide wide dynamic range within a high-dimensional learning space. At a higher level, this model then integrates with the ACT-R system

as described in subtask I2-6 where meta-cognitive signals are used to provide useful training signals; we anticipate leveraging this same interaction for motor as well as perceptual learning.

We find that problems in motor control can become intractable when a large state-space representation is used to find trajectories which arrive at a goal state. We avoid this problem by instead selecting the error signal which is pertinent to our goal and then using a general perceptual-motor system which acts to reduce that error. Initial results suggest that this approach is highly robust and effective.

Our efforts are directly related to those of subtask I2-6, involving the use of meta-cognitive signals to drive learning in the perceptual and motor models that we are developing in I2-7. More broadly, we also anticipate interactions with other Intelligence tasks such as those focused on Behavior Generation and Learning as well as with efforts in Perception and Manipulation and Mobility.

Metrics for Evaluation:
Progress in perceptually-driven motor control is currently measured at a low level in terms of reduction of cumulative error in a simulation of a manual pursuit task. Progress in the doorknob grasping and manipulation task will be measured in generalization accuracy to novel doors.

We have established prototype models of the cerebellar motor control system inspired by the biological literature and programmed a simulation platform for testing them. We have demonstrated initial success of the basic learning mechanisms.

Progress in the reaching and grasping tasks can be objectively measured by geometrical error measures and the ability to hold new objects, respectively. We will scale-up the complexity of these tasks to approach the difficulty levels associated with the capstone tasks.

	Subtask 7 Quarterly Milestones and Deliverables
Q1	First draft cerebellum model completed and tested on range of perceptually-driven motor tasks. Design and configuration of virtual environment for doorknob task.
Q2	Further evaluation and optimization of cerebellum model, scaling-up to harder problems. Initial development of integration with ACT-R for outer-loop control in doorknob task.
Q3	Further development and evaluation of integrated perceptual-motor and ACT-R model on doorknob task. Should be able to recognize and manipulate 20 different doors.
Q4	Scaling-up of the model to 100 different doors and generalization to 20 novel doors, along with continued performance optimization of all components.

Subtask 8: Meta-cognitive Development for Cognitive Architectures (ARL)

Background:
As part of the integration of a symbolic cognitive architecture with metric level data within a robotics system, a meta-cognitive level is needed to supervise the entire system and check for

erroneous behavior. The meta-cognitive level is important to a cognitive system because it monitors the execution of the production system within the architecture. Additionally, one function of meta-cognition is to use episodic memory as a store of previous experiences in order to recognize patterns in cognition (i.e., boredom) which might be problem areas. These problem behaviors are flagged by the meta-cognitive system and are not stored in episodic memory and not used in future episodic learning episodes. This has implications for the entire system, especially the learning components. Other aspects of meta-cognition involve checks of cognitive processes in order to improve cognition. Subtasks in this area include the development of a meta-cognitive level of understanding, development of episodic memory, development of episodic learning, and the characterization of memorable stimuli.

Objective and Benefits:
The objective of this subtask is to develop a meta-cognitive system that will monitor the execution of the production system within ACT-R and the RCTA architecture. This allows for deeper understanding of the overall cognitive processes by the cognitive architecture. For example, people typically have a meta-cognitive understanding of how well they know something or understand a specific subject matter. This knowledge is then frequently used for decision making and strategy selection. Meta-cognition provides a basis to allow this behavior to happen and is an important component to the cognitive architecture.

State of the Art:
Using a hybrid of a cognitive architecture with multiple robotics algorithms is a new concept and has only been tried in a few limited situations – and frequently without a meta-cognitive layer. ACT-R has sub-symbolic algorithms which run in parallel for computing memory decay and learning; however, ACT-R does not currently have a meta-cognitive layer. Ron Sun has proposed using a hybrid system for the development of the CLARION architecture, but again, this system does not use meta-cognitive processes.

Key Barriers and Approach:
Meta-cognition can be thought of as multiple algorithms which monitor cognition; therefore, the development of *all* these algorithms can be seen as quite labor intensive. We have selected a few key meta-cognitive algorithms to develop which will make the overall cognitive system more robust and responsive. For example, one common problem with autonomous systems is the infinite loop problem, where an autonomous system repeatedly executes a set of behaviors. This problem is one of the main issues that the newly added meta-cognitive system will attempt to solve. Additionally, recognition of patterns of behaviors requires efficient use of episodic memory in order to characterize tasks and chunk them into suitable levels of granularity. While these problems are not complex, they can be time and computationally intensive. We will leverage some of the previous work from the Soar architecture to address some of these problems.

Metrics for Evaluation:
A newly developed meta-cognitive layer will be added to ACT-R. The newly added layer will function appropriately with other layers of the architecture.

	Subtask 8 Quarterly Milestones and Deliverables
Q1	Development of meta-cognitive layer within the ACT-R system.
Q2	Test the meta-cognitive layer with the entire system.
Q3	Iterative development.
Q4	Continued testing within the RCTA architecture.

I3-2012 – Combining Cognitive and Probabilistic Reasoning

Level: 6.1 Basic Research

	Name	*Organization*	*Subtask*
Principal Investigator	Maxim Likhachev	CMU	I3-1
Additional Investigators	Tony Stentz	CMU	I3-2
	Juan Pablo Gonzalez	GDRS	I3-3

Objective:
To build robust tactical behaviors, we combine knowledge-intensive approaches from cognitive architectures with algorithmic approaches from traditional robotics. To bring the two disciplines together, we quantify, manage, and reduce inherent model uncertainty that arises in the interface between them. The combined result advances both fields and enables us to procedurally encode a set of complex adaptive behaviors, generalize them to a wider range of applications, and even learn entirely new behaviors. In particular, in FY12, we study the approaches to deciding when and how to disambiguate "important" uncertainties and how to operate in the presence of uncertainty in human behavior that cannot be disambiguated.

Background:
In theory, our intelligent architecture and world model provide the tools necessary for tactical behavior generation in the real world. The cognitive level provides the knowledge-centric mechanisms for encoding the appropriate procedures. The physical level provides the navigation, sensing, manipulation, and grasping operations to engage the world. Unfortunately, the problem heretofore, and hence the fertile ground for innovation and advancement, is a missing middle level that connects the symbols to the physical world, where knowledge-based approaches meet algorithmic approaches. For example, we can decompose the goal of "clear room X" into other cognitive-level goals such as "enter room X;" "if X contains closet door Y, open Y;" "if X contains window, push curtains aside;" and so on. Similarly, we can implement robot goals such as navigating through the room, grasping and lifting a rigid object, and pointing a sensor and taking an image. The difficulty arises in connecting the two types of goals. The robot may fail to perceive the door, or it may perceive a door where there is none. It may try to grasp the curtains, but they slip through its mechanical fingers. The problem is that the knowledge (e.g., model) used for reasoning and perceptual understanding is insufficient to handle all environments, situations, and tasks. The models may have missing or wrong parameters. The models may be too specific to be applicable or incorrect entirely. We, therefore, need algorithmic approaches that can handle these uncertainties by efficient reasoning over them, taking actions to disambiguate the uncertainties that are "important," and constructing models from scratch whenever existing models are wrong or do not exist.

Research Description:
We hypothesize that the underlying reason robots are unreliable and brittle when attempting to perform complex tasks is incomplete and inadequate knowledge used for understanding and reasoning. At the cognitive level, the ontology may be missing specific objects. A production rule critical for performing a tactical behavior may not be resident in the database, or its conditions might be improperly specified. An episodic memory may not be accessible from the current cues. At the physical level, a control model may be missing parameters crucial for differentiating between problem instances. Our approach builds on the best available techniques in cognitive reasoning and robotics to identify and incorporate missing knowledge to build robust tactical behaviors.

Quantifying and Coping with Uncertainty. Due to missing knowledge, robot performance is not perfect. The unincorporated, hidden data create unexpected outcomes, and robot performance appears to have a random component. We advocate modeling the understanding and reasoning process as stochastic. The robot measures its distribution parameters by observing its own performance over many trials. By computing confidence intervals on the data, we can provide statistical guarantees on task performance. By quantifying this uncertainty, we can properly manage it. We add this stochastic component to the sequential composition framework we previously developed for robotics. Our approach is, thus, the probabilistic generalization of sequential composition which achieves a globally convergent policy by discretely transitioning between continuous closed-loop policy domains. We can reason about alternative courses of action in a decision theoretic sense by considering the utility of each possible action and the probability of each possible outcome. For example, the robot can decide between two strategies for holding a person at bay in a room (e.g., blocking the escape route or ordering the person to remain still) by comparing the difficulty of each operation, the probability of success, and the cost of failure.

Finding Missing Knowledge to Reduce Uncertainty. The missing knowledge may be available (e.g., in the world model) but is not properly linked in to the reasoning or understanding procedure. In our approach, we seek contextual features in the world model that are readily available and are strongly correlated to task performance (e.g., a “jittery” person is less likely to try to flee). These features tighten uncertainty bounds by improving prediction and, thus, increasing task reliability. Because the space of possible features is unlimited, we use domain knowledge to select promising ones. In essence, good features are those that have potential for capturing much of the missing knowledge needed to discriminate between different possible outcomes. Redundant features rather than a minimal set are used to enhance robustness and decrease noise.

Taking Action to Reduce Uncertainty. During mission execution, the robot can reduce uncertainty by considering and evaluating variations on the operation (e.g., position robot at different guarding locations) and by selecting actions to disambiguate (e.g., corral a person into a corner to reduce escape options). For the former, we use a reinforcement learning framework to explore alternative ways to perform a given task, ranging from small perturbations to existing methods to radically different actions. For the latter, we advocate information-based planning

algorithms to represent, manage, and reduce uncertainty. We reason about the perception process in the planning process, modeling perception capabilities and data uncertainty to construct sound plans with performance guarantees. In the general case, this is a partially observable Markov decision process (POMDP) and is computationally intractable for realistic problems; however, using clear preferences for a sensing outcome (e.g., no adversary present), complicated problems can be solved efficiently through a series of low-order searches. Our approach taps this several-orders-of-magnitude speed-up to handle problems rife with uncertainty (e.g., predicting possible actions of pedestrians and vehicles as an autonomous robot encounters a potentially hostile crowd in a crowded marketplace).

Constructing Models from Scratch. In extreme cases, the robot is missing parts of the knowledge needed for performing a given behavior, or the behavior itself is inadequate (e.g., improperly generalized) or even non-existent. In this case, it is not a matter of properly linking symbolic actions to physical actions; the robot must generate the symbolic actions themselves. The robot must build the ontology, create the symbols, and assemble the procedures. This is a new and promising direction in cognitive research. As advocated in Learning and Adaption, a human can teach a robot to perform a particular tactical behavior by demonstration or instruction. By putting the robot through the paces at the physical level, the robot clusters actions into distinct categories and assigns a symbol at the cognitive level. By chunking groups of actions, the robot assembles a hierarchy of actions that constitutes a new tactical behavior. Coordination between actions is accomplished by detecting and generalizing sequential dependencies between concurrent actions.

To push the envelope further, the robot builds a tactical behavior without direct instruction; instead, it operates only from high-level goals. At the physical level, it explores combinations of algorithmic operations and evaluates them in simulation (for efficiency). We tie this algorithmic simulation to similar functionality at the cognitive level. Cognitive architectures are starting to incorporate mental imagery or simulation, such as the imaginal module in ACT-R for performing spatial operations like mental rotations and simulation module in Soar connected to the visio-spatial module for evaluating the result of future courses of action. These combinations of operations are clustered, mapped onto symbols, and chunked into a hierarchy that is evaluated against the high-level goal. Once the basic structure is in place, the new tactical behavior is tuned and adjusted based on real-world data.

Generalizing Models. Over many trials, we track the application of our procedural knowledge to specific problem instances. We use the outcomes from imperfect matches to generalize our models either by expanding the possible instantiations of symbols or by deleting symbols that empirically have no bearing on the outcome. For example, our model may pre-suppose that the presence of a weapon is required for a crowd to be hostile, only to observe there is no correlation. The weapon requirement is deleted from the model, thus generalizing it to other problem instances.

Metrics:

We test our progress via theoretical analysis of the performance guarantees we can provide as well as experimentation both in simulation and on a physical platform. Thus, we will consider

the domain of navigating through a crowded and potentially hostile environment for the experimental evaluation. In such environments, robust mission execution requires a robot to focus its attention on the factors that can affect its performance. In particular, the robot needs to track and understand the intentions of people who may potentially interfere with its course. We will test how our approaches to handling, reducing, and reasoning over uncertainty allow for a more robust behavior that focuses perception efforts and produces actions in such a way as to provide a higher level of robustness.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Max Likhachev has regular collaboration meetings with Juan Pablo Gonzalez (GDRS). We also hope to have collaborations with ARL; in particular, Max Likhachev plans to send a student to Stuart Young's group for a summer internship. He also hopes to host ARL researchers in his lab for short visits in order to assist in the integration of SBPL and new algorithms being developed under this task onto the ARL platforms.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Generalize/extend the approaches to include dynamic probabilistic models of human behavior and scenarios that don't allow full disambiguation of human intentions.
Q2	Finalize the approaches.
Q3	Perform experimental evaluation of the approaches in simulation.
Q4	Perform experimental evaluation of the approaches on a physical platform. Start the development of a plan for transitioning the approaches onto a platform that will be used in the Capstone Assessment.

Related Research:
The research in I3 takes inputs from the following research tasks: methods for achieving situational awareness (H2) and methods for perceiving complex static and dynamic scenes (P3 and P5). At the same time, the research in I3 will provide output to the following tasks: robust decision-making for mobile manipulation under imprecise or imperfect world models (M1), methods for improving situation awareness relevant to the successful mission execution (H2), methods for improving perception for missions in static and dynamic environments (P4 and P6), and building robust adaptive tactical behaviors (I4).

Integration and Assessment Activity:
As mentioned in the previous section, this task will integrate with several perception tasks (P3 and P5) that will provide the perception of complex static and dynamic scenes, particularly the hypotheses of the intentions of people. The results of the I3 task will be integrated into the methods for generative adaptive behaviors that are being developed in the applied research Task

I4. The result of I3-1: Planning under Multiple Hypotheses in Perception may potentially be integrated into some of the DMUM tasks, such as M3.

Subtask 1: Planning under Multiple Hypotheses in Perception (CMU)

Objective and Benefits:
Much of the brittleness in robots is due to the lack of tight integration between planning and perception. Typically, modern robotic systems perform planning under the most likely hypotheses from perception and do not reason about other hypotheses. This results in failures that arise whenever the perception is wrong. In FY 2011, we have been developing algorithmic support for scalable real-time planning under *multiple* hypotheses in perception. In addition to just planning around this uncertainty, our approach was structured to generate plans that explicitly tell the robot how and when to take actions that disambiguate perception hypotheses (for example, by asking "relevant" people to disambiguate their intended trajectories, by focusing perception on them, or by honking at them while navigating through dynamic environments).

In FY 2012, we will generalize our support for tighter integration between planning and perception further. More specifically, we will provide support for the situations where the uncertainty in perception is truly massive, and it may be impossible to fully disambiguate the hypotheses. This will allow us to handle situations with noisy perception and to relax the assumption that "focus" operation successfully disambiguates the uncertainty. Overall, this work is towards achieving robustness in the robotic systems.

In terms of the Capstone Assessment, the machinery we develop under this subtask will enable us to address the problem of robust navigation among people and vehicles. In these scenarios, we need to the robot to deal with inherent uncertainty in tracking and predicting the intentions of moving people, vehicles, and animals. Additionally, the same approaches can be applied to mobility in areas with potentially moveable objects, where there is uncertainty about which objects can be moved out of the way and how. The approach should also be general enough to be useful in performing mobile manipulation tasks under uncertainty in the environment.

In terms of the duration, it will continue for two more years. In FY 2013, we would work with GDRS in transitioning this technology onto a platform that will be used for the Capstone Assessment.

State of the Art:
Typical approach to autonomous robotic systems involves planning under the most likely hypotheses from perception. It does not explicitly reason about other hypotheses due to the computational expense of doing it. We overcome this barrier by capitalizing on the recently developed principle of clear preferences found in many problems.

Key Barriers and Approach:
Our approach is based on clear preference-based planning that we have developed previously. The idea is that for many situations with multiple hypotheses, we can say upfront what hypothesis is "better" for us. Mathematically, it means what hypothesis, if it held true, would result in a provably lower expected minimal cost of an optimal policy. If these preferences hold, then we can decompose the problem into a series of low-dimensional graph searches and solve much larger planning under uncertainty problems in real-time. We build on this clear preference-based planning and, in particular, extend it to handle noisy perception and situations with massive uncertainty.

Our approach would move beyond the state of the art because it would be able to scale to much larger planning under uncertainty problems than currently possible while providing rigorous guarantees on the quality of the solution. We have already applied clear preference-based planning to a number of problems such as path clearance (Figure 2-2(a)), finding a landing site (Figure 2-2(b)), and disambiguating human intentions (Figure 2-2(c)). The last work was done in FY11 under RCTA. In all of these cases, we are solving problems that have billions of states in belief state-spaces. These are many orders of magnitude larger than what can be solved with state-of-the-art techniques for decision-theoretic planning under uncertainty. This makes us believe that we will be able to succeed under this subtask.

This subtask ties closely with efforts in the perception area on detection and tracking of humans in complex scenes as well as understanding of complex scenes. We plan on integrating Jianbo Shi's efforts on tracking humans with our planning for robust navigation in dynamic environments.

(a) Path clearance *(b) Autonomous landing site selection* *(c) Disambiguating human intentions in navigation*

Figure 2-2: Using preference-based planning for various problems.

Metrics for Evaluation:
We assess our progress in several ways. First, we analyze the algorithms that we develop to understand what guarantees they can provide and under what conditions. Second, we use our approaches to build planning modules and run them in simulations as well as on a physical robot like the segbot which is in CMU's lab. The segbot is used to perform navigation among people. We measure metrics including planning times, success rates, and scalability to the amount of uncertainty. We have been able to run limited experiments on the segbot thus far while taking actions to disambiguate the intentions of moving people. Eventually, we will transition the

developed approach into the SBPL library and will assist in integrating it onto a platform used for the Capstone Assessment.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Develop algorithmic generalization to tight integration of perception and planning to handle noisy perception and cases where uncertainty cannot be fully disambiguated.
Q2	Provide experimental evaluation of the approach in simulation.
Q3	Provide experimental evaluation of the approach on a physical platform.
Q4	Provide full analysis of the properties of the algorithm. Start developing a plan for transitioning the approach onto a platform that will be used for the Capstone Assessment.

Subtask 2: Policies for Managing and Reducing Uncertainty (CMU)

Objective and Benefits:
The objective of this subtask is to devise policies for robots to perform tasks in the presence of environmental uncertainty. The subtask focuses on uncertainty that arises from humans in the environment, specifically unpredictable motion. The benefit of the approach is a robot that can negotiate a crowded street without collision, track a person as he flees, or corral a person until help arrives.

This subtask addresses Capability 7 (respond appropriately if/when people leave) in the Capstone Assessment Vision document and is a multi-year effort that will continue for one additional year.

State of the Art:
Robots have a difficult time interacting with humans since human behavior is often unpredictable. For example, a robot moving amongst pedestrians is challenged to avoid collision since the humans may move erratically. The robot cannot simply avoid all locations that could possibly contain a human at a future location since the result could be that the vehicle cannot go anywhere. Current approaches assume either straight-line motion or a probability distribution over expected motion. These approaches steer the robot clear of the human's expected location out to the sensor horizon. Although better than avoiding all possible locations, the approach is too conservative and can cause the robot to slow down or stop unnecessarily. The result is inefficiency in the behavior of the robot and even possible incompleteness. Instead, the robot must reason directly about the possible interactions with the humans, including multiple sense, plan, and act cycles out to a safe horizon, in order to ensure a safe outcome.

Key Barriers and Approach:
The uncertainty in the human's behavior taken in combination with possible robot actions creates a *belief space*. This space is high dimensional and computationally expensive to search. Rather than simply plan sequence of actions, the robot must plan both sensing and moving actions and consider all possible sensing results in order to achieve a desired outcome. This *policy* generation is akin to playing chess where the robot examines what it could do, looks at the expected human

response, plans its response, and so on, out to the sensor horizon. The robot *reduces* uncertainty by planning to sense, and it *manages* uncertainty by accounting for all possible sensing outcomes.

In this subtask, we devise policies for a robot to interact with a human to result in a desired outcome, either from the standpoint of maximum likelihood or to avoid worst-case scenarios. To combat the high dimensionality of the belief space, we order the search so that 1) we expand high probability nodes first, and 2) we re-use results to avoid repeatedly searching the same subspace. In FY11, we developed a local planner that used a dynamic programming implementation of AO* combined with a look-up table and probabilistic anytime node expansion to drive through a crowd of pedestrians without collision in real-time. The approach assumed a *static* probabilistic model for human motion. In some scenarios, however, the model itself may be a function of the individual, the individual's behavior, distance between robot and human, or some other *dynamic* property. In FY12, we will accommodate dynamic probabilistic models in the AO* planner, including efficient calculation of the look-up table and revisions to the node expansion algorithm.

We will conduct experiments with two scenarios. First, a robot is tasked to track a person and not lose sight of him. The robot must deal with occlusions and move to reacquire line-of-sight as needed. The robot follows unobtrusively but becomes more aggressive if the human switches behavior and begins to flee. Second, a robot is tasked to corral a human in a given area. The robot stands guard. If the human attempts to flee, the robot intercedes and blocks the escape route. If the human attempts to escape through a different route, the robot reacts accordingly.

This subtask is tied to human detection and tracking (Tasks P5 and P6) and human behavior amongst robots (Tasks H8 and H9). This subtask complements the other subtasks in this technical area by specifically addressing the combinatorics of the belief space search rather than leveraging preferred outcomes.

Metrics for Evaluation:
We will conduct planning experiments using simulated human behavior data that predicts the human's reaction to the robot in a probabilistic way. We will compare dynamic programming AO* against traditional approaches, such as straight-line motion prediction and one-shot probability distributions, measuring 1) the percentage of time the human is tracked and 2) the frequency at which the human escapes.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Extend AO* planner to include dynamic probabilistic models.
Q2	Devise look-up table to accommodate dynamic probabilistic models.
Q3	Conduct experiments with a robot performing human tracking in simulation.
Q4	Conduct experiments with a robot performing human corralling in simulation.

Subtask 3: Robust Mobility under Uncertainty of the Semantic Labeling of the Scene (GDRS)

The semantic labeling of scenes provided by Task P3 enables higher-level reasoning than previously possible. Semantically labeled scenes can also provide additional information about the environment, including whether an object may be movable or whether it is important for mobility for reasons beyond geometric reasoning. However, much of these capabilities depend on having perfect or almost-perfect labels which are not achievable in the foreseeable future. This subtask seeks to address the uncertainty arising from imperfect semantic labeling in order to achieve robust mobility.

State of the Art:
Most current approaches to mobility perform a purely geometric analysis of a scene. This has allowed significant progress in the past decade for both on- and off-road navigation. Further progress in mobility, especially in human environments, requires higher-level reasoning that is only possible through semantic labeling. Existing approaches to use semantic labeling for mobility are mostly limited to prior data and use maximum likelihood estimates to calculate traversal costs.

Key Barriers and Approach:
Moving beyond maximum likelihood estimates requires modeling mobility as POMDP. POMDPs can handle uncertainty in the environment but are intractable for all but very small problems. We will explore specialized approaches such as Planning with Clear Preferences that can handle uncertainty in a tractable manner if preferred outcomes can be identified. We will also identify subproblems that focus uncertainty only on certain labels, like movable objects and doors.

This research is closely tied to I3-1 which seeks to combine probabilistic and cognitive reasoning as well as to to I4-1 which seeks to develop adaptive tactical behaviors.

Opportunities for Collaboration:
This subtask requires input mostly from Tasks I2 and P3. It will transition to I4 for implementation as a tactical behavior.

Metrics for Evaluation:
Initially, the most important metric will be the ability to complete a mission with a controlled set of parameters. We will combine simulation and field experiments in order to obtain statistically significant measurements on the completion rates for different scenarios. Other metrics that may be useful in understanding the performance of the system are time to complete the task, metrics on the length, speed, and accelerations on the trajectory, as well as human evaluation of the execution of the task. We will compare the approaches developed here with state-of-the-art implementations in order to assess the improvement provided by this approach.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Survey state of the art in mobility with semantic maps. Identify relevant subproblem on which to focus.
Q2	Implement state-of-the-art baseline.
Q3	Survey candidate approaches.
Q4	Initial implementation of robust mobility under uncertainty.

I4-2012 – Generating Adaptive Tactical Behaviors

Level: 6.2 Applied Research

	Name	*Organization*	*Subtask*
Principal Investigator	Juan Pablo Gonzalez	GDRS	I4-1
Additional Investigators	Steve Sablak	GDRS	I4-1
	Maxim Likhachev	CMU	I4-2
	Alberto Lacaze	Robotic Research	I4-3
	Nenad Uzunovic	Robotic Research	I4-3
	Alonzo Kelly	CMU	I4-4

Objective:
To build robust tactical behaviors, we combine knowledge-intensive approaches from cognitive architectures with algorithmic approaches from traditional robotics. To bring the two disciplines together, we quantify, manage, and reduce inherent model uncertainty that arises in the interface between them. The combined result advances both fields and enables us to procedurally encode a set of complex adaptive behaviors, generalize them to a wider range of applications, and even learn entirely new behaviors. While this is a shared objective with Task I3, Task I4 aims to identify those approaches that are more mature and that can be integrated more broadly with other areas. I4 will also focus its metrics and evaluation on integrated assessments rather than on the core concepts.

Background:
In theory, our intelligent architecture and world model provide the tools necessary for tactical behavior generation in the real world. The cognitive level provides the knowledge-centric mechanisms for encoding the appropriate procedures. The physical level provides the navigation, sensing, manipulation, and grasping operations to engage the world. Unfortunately, the problem heretofore, and hence the fertile ground for innovation and advancement, is a missing middle level that connects the symbols to the physical world, where knowledge-based approaches meet algorithmic approaches. For example, we can decompose the goal of “clear room X” into other cognitive-level goals such as “enter room X;” “if X contains closet door Y, open Y;” “if X contains window, push curtains aside;” and so on. Similarly, we can implement robot goals such as navigating through the room, grasping and lifting a rigid object, and pointing a sensor and taking an image. The difficulty arises in connecting the two types of goals. The robot may fail to perceive the door, or it may perceive a door where there is none. It may try to grasp the curtains, but they slip through its mechanical fingers. The problem is that the knowledge (e.g., model) used for reasoning and perceptual understanding is insufficient to handle all environments,

situations, and tasks. The models may have missing or wrong parameters. The models may be too specific to be applicable or incorrect entirely.

Research Description:
For defined tactical behaviors, such as those involving room clearing, urban ISR, self-defense, and logistics, we encode as much domain-specific procedural knowledge as possible into the architecture at the cognitive level. For example, in ACT-R, this constitutes devising production rules to implement the behavior. This procedural knowledge covers contingencies as well, such as detecting and responding to an adversary in a room-clearing operation. We devise the algorithms necessary at the physical level to support robot navigation, sensing, manipulation, and grasping. The algorithms reason about high-dimensional coupling of continuous-valued parameters to achieve given spatial and temporal states. The two levels meet in the middle with symbolic expressions like, "grasp doorknob X, and open" which map onto specific spatial configurations for a mobile manipulator

We devise capabilities that quantify, reduce, and cope with uncertainty in this mapping. Over the course of many trials, we quantify distribution parameters in mapping the symbolic actions to the continuous actions for the purpose of predicting success and selecting between competing actions. We automatically scour the contextual information in the world model to identify knowledge that is correlated with the perceived randomness. The knowledge is added to the model and is used to hone the predicted outcomes and reduce uncertainty. We search for actions, compliant with doctrine, designed to find missing knowledge or disambiguate between key problem instances.

Our tactical behaviors go well beyond fine-tuning models and reasoning about the result. We develop capabilities that leverage training data and real-world experience to construct the tactical behaviors from scratch, populate partial or incomplete tactical behaviors, and validate or invalidate tactical behaviors. For example, in ACT-R, we cluster and chunk production rules to create higher-level behaviors, modify the conditions for firing a given rule, and add/delete rules to match training data. Finally, we generalize the resultant tactical behaviors to as many problems as possible by scrubbing them of irrelevant parameters and widening the match of applicability. The end result is tactical behaviors that marry complex behavior to rich, real-world environments by truly bridging the gap between cognitive and physical reasoning.

The proposed subtasks will investigate possible approaches to create higher-level behaviors required for complex missions, including a logistics resupply and urban passage. They will also address from the bottom up some of the brittleness in the existing implementations by using more principled approaches and will build on this by using statistical learning techniques to generate adaptive tactical behaviors from expert knowledge.

Metrics:
Initially, the most important metric for generating adaptive tactical behaviors will be completion of the task that generated the tactical behavior. We will combine simulation and field

experiments in order to obtain statistically significant measurements on the completion rates for different scenarios. Completion rates will also be decomposed into overall behavior completion and completion of the individual elements that compose the behavior in order to identify which aspects of the behavior limits completion rates. We will also develop other metrics that are relevant for the integrated scenarios. Some of these metrics will be clearly quantitative, such as time to completion, but others will include human factors elements, such as Soldiers' evaluations of the performance of the robot as a teammate.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Initial implementation of basic task executive in ACT-R.
Q2	Integration of task executive with navigation behavior. Develop generalization of approach to planning with learning to tasks that require dealing with changing environments. Develop efficient constrained dynamics models of terrain following.
Q3	Test on a real vehicle in static environments for taking an active role in the environment. Develop models of concurrent terrain following and wheel slip.
Q4	Build a general search-based planning with learning software and integrate it into SBPL with the aim of making it available for integration on hardware platforms (e.g., TALON and/or BigDog) that will be used in the Capstone Assessment.

Related Research:
This task requires input from Tasks I2 and I3. It also interacts with Task I6 and the cognitive aspects of I2. Although not directly, this task would benefit from frequent interaction with the Perception, DMUM, and HRI tasks as they provide information that is essential to the successful generation of tactical behaviors. We will meet regularly with other members of the Consortium working on Perception, Intelligence, DMUM, and HRI in order to exchange ideas and understand the needs of the different areas.

Integration and Assessment Activity:
We will use RIVET for initial integration and assessment of this task, dividing the 3-year capstone vision into small elements that can be tested in simulation. As the task progresses, we will build-up capabilities and start to complement the RIVET testing with field testing in which we can use actual inputs from Perception and HRI.

Subtask 1: Generating Adaptive Tactical Behaviors (GDRS)

Objective and Benefits:
In order to go beyond fine-tuning models and reasoning about the result, we need to leverage training data and real-world experience to construct the tactical behaviors from scratch, populate

partial or incomplete tactical behaviors, and validate or invalidate tactical behaviors. In this subtask, we will investigate possible approaches to create higher-level behaviors. This includes identifying relevant tactical behaviors, creating proof of concept implementations, and generalizing the resultant tactical behaviors to as many problems as possible by scrubbing them of irrelevant parameters and widening the match of applicability. The end result is tactical behaviors that marry complex behavior to rich, real-world environments by truly bridging the gap between cognitive and physical reasoning.

Adaptive tactical behaviors are present throughout the Capstone Assessment as individual behaviors that execute individual tasks and at a higher level as a task execution level in charge of coordinating and triggering different behaviors.

This effort will continue over multiple years, continuously identifying relevant behaviors and adding more complex reasoning to coordinate and trigger them.

State of the Art:
Most implementations of tactical behaviors are brittle and require careful tuning in order to perform adequately. Furthermore, tactical behaviors cannot adapt to changing scenarios or tactical requirements and need to be tuned again for each new situation [1,2].

Key Barriers and Approach:
We will explore different approaches to generating high-level tactical behaviors. Initially, we will explore using ACT-R to bring together subsymbolic approaches and learn better ways to respond to changes in the environment and to failures. This should provide greater adaptability and robustness than hard-coded state machines and static task networks.

We will also develop more robust subsymbolic approaches. Reducing the brittleness of current approaches requires having not only higher levels of cognitive reasoning but also more robust solutions from the ground up. This task proposes to identify critical elements in subsymbolic behaviors that must be addressed in order to successfully support higher-level reasoning.

By leveraging the best aspects of ACT-R with robust subsymbolic behaviors, we expect to combine the strengths of both approaches and get robust and adaptive tactical behaviors.

We will initially implement a behavior executive combined with a subsymbolic navigation behavior. The behavior executive will monitor the progress and context of the navigation task and will produce different outcomes depending on whether the task is progressing as expected. We will also explore different subsymbolic behaviors, such as opening doors or planning to push obstacles out of the way, depending on the information available to ACT-R to evaluate the environment.

This is a multi-year effort that continuously looks at the most relevant tactical behaviors, evaluates their shortcomings, and develops solutions based on a combination of cognitive reasoning and robust subsymbolic behaviors.

This research will enable the "move" capabilities from the Capstone Assessment Vision as well as provide the task execution framework needed to coordinate and execute most tactical behaviors.

This subtask is closely tied to Subtasks I3-3 and I4-2 which seek to improve the robustness of subsymbolic tactical behaviors and make them adaptive at the subsymbolic level. It also works closely with Task I1 in bringing together cognitive approaches and connecting them to real robots.

Opportunities for Collaboration:
We will work closely with Tasks I1 and I2 to identify the architectural requirements to support adaptive tactical behaviors, and to incorporate relevant cognitive approaches. We will also work closely with Tasks I5 and I6 in order to incorporate statistical learning.

Metrics for Evaluation:
Initially, the most important metric for generating adaptive tactical behaviors will be completion of the task that generated the tactical behavior. We will combine simulation and field experiments in order to obtain statistically significant measurements on the completion rates for different scenarios. Once the specific behaviors are identified, we develop other metrics that are relevant for the integrated scenarios. Some of these metrics will be clearly quantitative, such as time to completion, but others will include human factors elements, such as Soldiers' evaluations of the performance of the robot as a teammate.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Initial implementation of basic task executive in ACT-R. Initial implementation of planning with movable obstacles.
Q2	Integration of task executive with navigation behavior. Preparation for IRA3: Intelligent Navigation.
Q3	Evaluation of different progress metrics. Complete IRA 3: Intelligent Navigation.
Q4	Analyze IRA3 results. Evaluation of additional behaviors at the subsymbolic level.

References:
[1] A. Kelly, A. Stentz, O. Amidi, M. Bode, D. Bradley, A. Diaz-Calderon, M. Happold, H. Herman, R. Mandelbaum, T. Pilarski and others, "Toward reliable off road autonomous vehicles operating in challenging environments". *The International Journal of Robotics Research, SAGE Publications,* 2006, *25*, 44.
[2] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer and others, "Autonomous driving in urban environments: Boss and the urban challenge". *The DARPA Urban Challenge, Springer,* 2009, 1-59.

Subtask 2: Life-long Improvement of the Robustness of Tactical Behaviors (CMU)

Objective and Benefits:
Robustness in executing complex behaviors hinges upon identifying and reasoning over salient dimensions of the problem. Real-world environments, however, cannot be predicted perfectly beforehand and, therefore, require the robot to automatically identify the dimensions of the problem that are important to the robust execution of its mission. In this subtask, we will develop principal approaches to identifying these dimensions via incorporating the identification of these dimensions into the planning process itself and/or learning the "right" approaches to planning through previous task-oriented instances of executions. We will show the feasibility of our approaches in the context of several domains including dual-arm high-DOF manipulation of objects in cluttered spaces.

The results of this subtask will be used to build a robust planner for high-dimensional planning problems like planning for mobile manipulation. As a result, these efforts address such Capstone Assessment capabilities as the ability to manipulate objects using one and two arms and the ability to perform mobile manipulation.

We see the work on Life-long Improvement of the Robustness of Tactical Behaviors continuing for approximately three more years and integrating its results with our efforts on the mobile manipulation subtask in DMUM. In there, we study how to develop efficient planners for mobile manipulation *without* learning how to improve their performance. The end result of integration will be to show how a mobile manipulation platform will learn to perform fast and effectively the tasks that it needs to perform mundanely, including picking-up objects, opening doors, etc.

State of the Art:
Unlike previous approaches to learning solutions to planning problems, our approach does not depend on the "right" identification of features and maintains all the guarantees one expects from a planner. These guarantees will not depend on having a representative training set.

Key Barriers and Approach:
Our current work focuses on learning from planning in order to make future planning requests faster. Specifically, we cache some of the information the planner generates every time it is asked to find a path. This information is used to construct a graph which represents the high-level connectivity of the free space. This high-level graph is then used during future queries to accelerate the planner by providing likely routes to the goal. The result is that the planner attempts to reuse parts of the high-level graph in future planning requests in order to eliminate large parts of the search effort. The framework is very general and should be applicable to many domains in order to expedite planning. Furthermore, we do it in a way that preserves guarantees on completeness and bounds on suboptimality. Unlike previous approaches to learning solutions to planning problems, our approach does not depend on the "right" identification of features and maintains all the guarantees one expects from a planner.

We have already implemented the initial prototype of the planner and have obtained very good preliminary results on PR2 (Figure 2-3). In the next year, we will need to make it more robust

and more general in the type of planning tasks it can handle. We will also perform a more careful study of its performance. As mentioned above, these efforts link very well to the efforts in DMUM since we are working on the approach that improves the performance of planners for highly articulated robotic systems such as mobile manipulators.

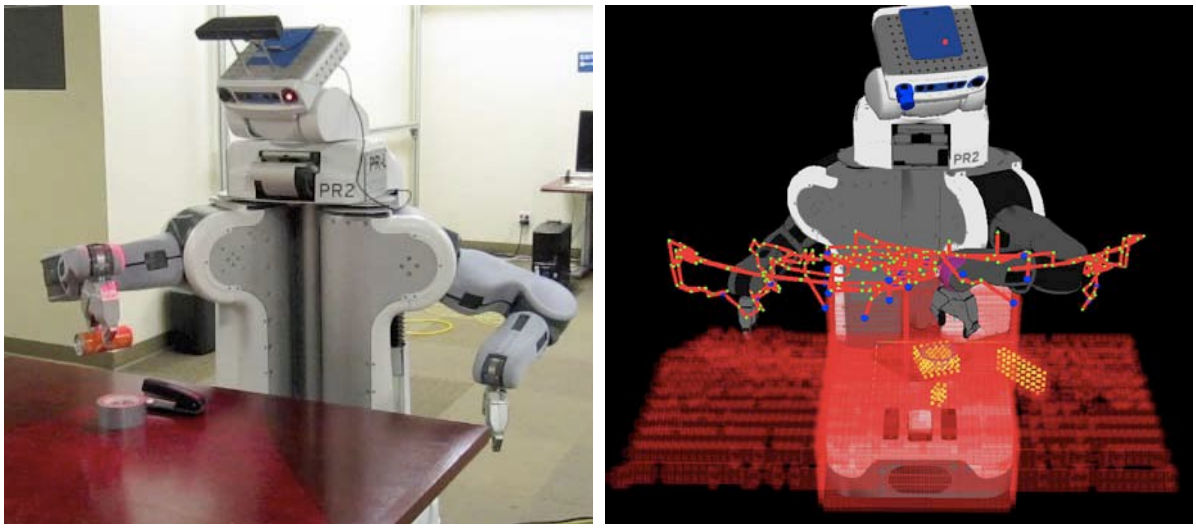

Figure 2-3: Learning to Plan on PR2 (our ongoing efforts in I4-2).

Metrics for Evaluation:
To measure our progress, we apply our learning approach to planning for manipulation on PR2. We, thereby, test the effectiveness of learning in simulation (in ROS) as well as on the physical robot by collecting statistics such as planning times and success rate. In addition, we analyze the algorithm to understand what guarantees it can provide. Our goal is to have fast and effective planning for mundane tasks involving dual-arm mobile manipulation. Eventually, we will release this capability as part of SBPL (Search-based Planning Library) within ROS. We will also integrate it with our efforts under Subtask M1-1 in DMUM and will assist GDRS and other Consortium members in integrating it onto the mobile manipulation platforms like TALON and/or BigDog that will be used for the Capstone Assessment.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Analyze the theoretical guarantees that our current approach to planning with learning provides. Perform preliminary experimental study on PR2.
Q2	Develop generalization of our approach to planning with learning to tasks that require dealing with changing environments.
Q3	Perform extensive experimental evaluation of the approach in simulation and on PR2.
Q4	Build a general search-based planning with learning software, and integrate it into SBPL with the aim of making it available for integration on hardware platforms (e.g., TALON and/or BigDog) that will be used in the Capstone Assessment.

Subtask 3: Planning for Taking Active Role in the Environment (Robotic Research)

Objective and Benefits:
In this subtask, we aim to take robotic systems to a new level by enabling them to reason about taking an active role in the environment while navigating it. Current state-of-the-art systems assume that the environment is static or dynamic, but they do not reason about changes that robot's actions might cause. A simple example would be moving an obstacle out of the way in order to navigate through a cluttered room. Another, more complicated, example would be to reason if and which moving entities will yield or move out of the way so that a robot can safely navigate through a busy street or environment of similar complexity.

This subtask will aid the Capstone Assessment in the following areas:

- Adaptive behavior generation
- Move through mobility challenges

Preliminary results for interacting with a static world are expected by the end of 2012, while preliminary results for interacting with a dynamic world are expected by the end of 2014.

State of the Art:
Current state-of-the-art systems assume that the environment is static or dynamic, but they do not reason about changes that robot's actions might cause in the environment. Also, when robots modify the environment by manipulating objects, it is because they were explicitly instructed to do so. We would like to enable robots to choose whether or not to interact with something depending on need.

Key Barriers and Approach:
These problems are highly dimensional in their nature. The first attempt to solve them will be by using task decomposition methods in which some of the dimensions are grouped into several planners with smaller dimension spaces. Though full and efficient task decomposition is hard to achieve, we believe that there are a few places in which dimensions can really be separated, at least for foveating the attention. For example, if robot reasons in (x, y, time) domain and we have a robot planner and a moving obstacle, we have a six-dimensional system. If we decompose the problem to planning for the robot and for the mover in (x, y, time) as two separate problems, we can determine if their paths interfere with each other. If they don't interfere, the problem is solved. If they do, we can use previous results from the task decomposition to narrow the area of focus for which we will have to solve the six-dimensional problem. Even though we did not eliminate the need to solve the six-dimensional problem space, we are reducing the size of the space by decomposing it at first.

Metrics for Evaluation:
Progress of this subtask will be measured by comparing the latest robot capability with the initial mobility assessment test held at FTIG in August 2011. The assessment test had a couple of different test areas designed specifically to address these issues. For example, we thoroughly tested navigation through the partially open door and busy streets. The advancements in this

subtask will be able to solve complex navigation challenges that the previous systems were not able to solve.

The development process for the task will be divided into several time segments. Algorithms will be tested for feasibility in simulation before they are tested in the real world. Simpler problems will be solved first, and evaluation will be made if the solutions apply to more complex problems.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Work with world model and Integration groups on representations and actions.
Q2	Preliminary simulation results for static environments.
Q3	Tests on a real vehicle in static environments.
Q4	Pathway for complex dynamic environments.

Subtask 4: Uncertainty Adaptive Behavior (CMU)
We seek to develop efficient, online methods that endow robots with the capacity to actively reduce the error and the uncertainty in the models that form the interface between cognition and lower levels of autonomy such as perception, control, and estimation.

The research will support the capstone assessment by enabling many of the higher-level cognitive tasks required to a) think and b) move. Results to date can already endow any wheeled or tracked mobile robot with more accurate models of its own mobility as well as models of the uncertainty in its motion predictions. Such enhanced models improve the accuracy and robustness of control, estimation, and planning.

As a research effort that is focused on innovative capacities embedded in the fundamental infrastructure, this effort is likely to continue throughout the program.

State of the Art:
While there is plenty of work in robotics that addresses the use of uncertainty estimates in order to make better decisions, there is almost none that addresses the question of how to tune predictive models online. There is related work on how to calibrate estimation systems. Estimation textbooks like [4] discuss manual tuning techniques, and these could in principle be automated. The robot state estimation literature in [2] and [1] also addresses the problem of how best to model uncertainty in an offline fashion, and the dynamics identification literature [3] has addressed dynamic model identification for decades. The proposed work differs from all of these prior efforts in two ways. First, it is intended to calibrate predictive models for which there are no measurements as there are in estimation problems. Whereas the noises in estimation systems are primarily sensors nonidealities, our “noises” are wheel slip and other physical disturbance phenomena. Second, the goal is to do this online.

Key Barriers and Approach:
Key challenges include how to measure prediction residuals, how to calibrate online, how to ensure that both fast adaptation and long-term memory can coexist at the same time, and how to parameterize the models to incorporate sufficient degrees of freedom while preventing over-fitting.

We seek to endow robots with the capacity to observe the performance of their own models and to continue to refine those models based on experience. The models could be used for motion prediction, dynamic stability, motion planning, or pose estimation.

To date, we have pioneered a new approach to system calibration called integrated equation error. This formulation has turned-out to enable the calibration of mobile robot motion prediction; whether it also enables other forms of calibration is not yet known.

Our goals are not so ambitious that progress will be difficult. The community is only now beginning to seriously address the auto-calibration problem, and it is still not addressed in its entirety. At this point, we have enough positive results to know that we are pursuing a productive line of inquiry.

The general problem of automatic system calibration is virtually untouched in the robotics literature, even though few would disagree that it is a problem of fundamental importance. Robots need to detect that their models are wrong and then fix those models as a most basic aspect of meta-level reasoning. The literature contains numerous results that calibrate camera intrinsic parameters, sensor poses, etc., and though we have new results in calibration of wheel slip and IMU pose, there are many more things that should be calibrated.

Metrics for Evaluation:
Progress relative to our goals will be assessed by comparing the performance of algorithms based on our new representation against performance using more classical techniques. In many cases for this task, an entirely new capability will be generated for which there is no precedent.

Online stochastic calibration of wheeled mobile robots has been developed and demonstrated on data gathered from multiple platforms.

The work is very fundamental so the mapping to a specific capstone capability is not so direct and clean-cut. The work impacts robustness across the board rather than a particular function in a capstone demo. If the capstone demo takes place on slippery terrain with a tracked vehicle, our impact will become very clear.

The plan for the coming year is to address the question of how our stochastic and systematic models can actually be used in practice. The first quarter concentrates on what value can be extracted from stochastic mobility models, and the remaining quarters cover models of wheel slip and terrain following applied to adaptive control.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Complete stochastic control/estimation work.
Q2	Develop efficient constrained dynamics models of terrain following.
Q3	Develop models of concurrent terrain following and wheel slip.
Q4	Adaptive control demo (e.g., one of door opening, self-righting, entrapment escape, obstacle negotiation, vehicle or manipulator inertial properties).

References:

[1] P. Abbeel, A. Coates, M. Montemerlo, A. Y. Ng, and S. Thrun. Discriminative training of Kalman filters, 2005. In Proceedings of Robotics: Science and Systems.

[2] Kelly A., “Fast and Easy Systematic and Stochastic Odometry Calibration”, In Proceedings of the International Conference on Robots and Systems, 2004.

[3] Minyoung Kim; Pavlovic, V.; , "Discriminative Learning of Dynamical Systems for Motion Tracking," Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on , vol., no., pp.1-8, 17-22 June 2007

[4] Peter S. Maybeck, Peter S.. Stochastic Models, Estimation, and Control, volume 1. Academic Press, 1982.

I5-2012 – Learning through Experience

Level: 6.1 Basic Research

	Name	*Organization*	*Subtask*
Principal Investigator	Drew Bagnell	CMU	I5-1
Additional Investigators	Daniel Lee	UPenn	I5-2
	Dieter Fox	UoW	I5-3
	Jonathan Fink	ARL	I5-4
	Nicholas Fung	ARL	I5-4
	Douglas Summers-Stay	ARL	I5-5

Objective:
This task is devoted to enabling machine learning technologies to be embedded throughout the robot system. We investigate a framework that leverages the best of engineering insight by structuring system components (e.g., perception, reasoning, and world model) and providing features that leverage human expertise while also enabling a deep, systemic learning approach.

Learning through Experience, I5, is devoted to basic machine learning technology that serves other aspects of the program. In particular, Subtask 1: Anytime Prediction and Co-regularization is devoted to enabling learning to scale to robotics problems. The focus here is on trade-offs of computational effort and prediction quality. Additionally, the co-regularization part of this task is designed to improve performance with small amounts of labeled data by leveraging unlabeled data.

Subtask 2: Learning Low-dimensional Manipulation Planning and Control is devoted to learning approaches that can improve high-dimensional planning through experience. Subtask 3: Learning to Understand Human Instructions and Annotations is devoted to understanding language and natural language commands through machine learning technology. Subtask 4: Anytime Learning for Robust Navigation and Control seeks to apply the techniques developed by I5-1 to the problem of robust navigation in militarily-relevant environments. Lastly, Subtask 5: Integrate Visual Texture Recognition will apply learning to allow recognition of traversable terrain from visual textures.

Background:
Machine learning technologies were initially developed for relatively small datasets and for simple classification and regression tasks. In recent years, efforts have been made to make them applicable to, e.g., the web where learning is a core technology.

Task I5 is designed to develop the basic technology to make learning more applicable in robotics. Key differences in this domain include the need to interact with people, the need to compute decisions and plans, and finally, the need to balance computational resources and predictive accuracy while making decisions in real-time.

Prior work is discussed below at the subtask level.

Research Description:
The subtasks address very different concerns and with very different approaches. We defer detailed descriptions to the subtasks.

Metrics:
Metrics vary dramatically by subtask, as they each serve different areas within the RCTA. Subtask 1 focuses on improving prediction time and quality – and most importantly, the trade-off between these – throughout the RCTA program. Subtask 2 focuses on learning planners that are more effective in practice. Subtask 3 is focused on the correct semantic interpretation of people's utterances. Subtask 4 will compare the rate at which the solution increases in optimality versus time spent in computation. Subtask 5 will use navigation success rate and speed at reaching the goal.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
The initial work of subtask 1 will be presented at the conference on AI and Statistics 2012. Additionally, Bagnell's team has submitted an abstract for the SPIE workshop in Baltimore 2012. Some of the basic technology may be discussed as part of a workshop at the International Conference on Machine Learning 2012 as well.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Develop extensions of SpeedBoost anytime prediction appropriate for structured prediction in RCTA. This will be a first step in enabling anytime prediction to be used by perception. We are already in the process of making our existing code available to RCTA projects and other researchers, and we plan to make these extensions and further work available as well. Low-dimensional learning algorithm ported to embedded software. Extend existing CCG (categorical combinatory grammar) parser learning approach to handle input beyond natural language, including gesture and human gaze. Develop texture classification algorithm.
Q2	Enable larger classes of learning algorithms to be used within the SpeedBoost framework. Continue work on anytime prediction to support semantic perception applications. Experimental arm manipulation testbed. Investigate CCG learning in the context of the new data and semantic mapping tasks. Incorporate texture classification algorithm.
Q3	Investigate other algorithms in addition to SpeedBoost for anytime prediction. Implementation of arm planning on testbed. Start developing and implementing efficient learning technique for jointly mapping speech, gesture, and gaze to formal semantic map representation. Perform online training with proprioceptive labels.
Q4	Continue development of SpeedBoost to non-perceptual problems including decision-making and cognitive learning. Investigate use of approach for non-computational resource limitations (e.g., human interaction and labeling) which are expensive. Experimental validation of planning algorithms. Develop approach to interactive object grounding, and expand semantic representation to incorporate novel objects on the fly. Test robot's success in improving navigation.

Related Research:
Task I5 serves principally to provide enabling technology for other aspects of the RCTA. Subtask 1 is developing a basic machine-learning infrastructure that will enable broader and simpler use of learning; it will begin (and continue throughout the program) working with semantic perception which relies heavily on learning technology and where the computational demands are severe. Subtask 2 interacts with the manipulation and planning part of intelligence. Subtask 3 is concerned with HRI efforts.

Integration and Assessment Activity:
An output of this year's effort will be a general SpeedBoost machine learning library which is applicable to generic machine learning tasks, including classification and regression. It will also be special-cased to handle structured (multi-bit) prediction problems. The other subtasks will provide basic technology for HRI and sophisticated motion planning necessary during assessments.

We believe that success on the assessment and capstones will depend on our ability to make high performance predictions in time.

Subtask 1: Anytime Prediction and Co-regularization (CMU)

Objective and Benefits:
Provide a suite of machine learning tools that trade-off computation and learning accuracy. The capstone demo relies on learning algorithms throughout many parts of the system, including notably perception, where the computational demands on learning are extreme. Algorithms must be able to perform *anytime prediction* – providing classifications, regressions, and other more sophisticated prediction tasks (i.e., policies in reinforcement and imitation learning) as they are needed and improving these as time admits. This problem has been almost completely ignored in the learning literature (the most notable exception being the "cascade" popular in the vision literature), and we anticipate its significance, and difficulty will require a multi-year effort. Additionally, we are continuing development of "Deep Inference Machines" that replace graphical models in aiding structured prediction in the RCTA effort. This is in the critical path for our scene understanding efforts in perception.

Technical Approach to Overcome Barriers and Advance SoA:
Our current approach builds on the classical boosting approach to machine learning. It uses a sequential strategy to learn the predictors and provides guarantees on the speed/accuracy trade-off. This effort builds on very recent research in submodular optimization that has enabled a significant new ability.

To improve Deep Inference Machines, we are developing a novel "co-regularization" learning technique that enables efficient use of a small amount of labeled data by augmenting with labeled data and enforcing the constraint that two modality inference machines (e.g., vision, LADAR) should agree on the labelings wherever they overlap.

Our work on I5 ties directly into work on perceptual understanding in both Tasks P3 and P5. It is important to note that anytime prediction is a new capability, only barely addressed in the literature. We are spending this year's effort focused on batch/offline learning which is the largest demand of the consumers of learning technology. (With the exception of the novelty prediction technique of Sofman, Bagnell, and Stentz (previous RCTA), we don't believe any techniques exist that have an anytime character and are online learners.) In future years' efforts, we will look at anytime predictors that can be trained online. It is conceivable that the learning itself can have an "anytime" character, extracting as much from data as is available in the time allotted.

Metrics and Progress Assessment:
Our algorithms will be judged on their performance (log-loss, classification accuracy, etc.) with respect to the time taken to achieve that accuracy. This provides a curve describing any learning algorithms performance. Most classical algorithms only achieve a single point on that curve, taking a long time for good performance or a short time for poor performance. Our efforts are

designed to provide a smooth trade-off and ensure our performance is better for each time period. We have developed the basic theory (submitted for publication) and the initial applications to standard benchmark learning problems in 2011. We are currently planning to integrate into scene understanding efforts over the next year to enable real-time performance on these tasks.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	We will complete efforts to extend to structured prediction problems, including those arising in vision and LADAR interpretation. In this effort, we focus on extending our techniques to predict multiple bits simultaneously. We are already in the process of making our existing code available to RCTA projects and other researchers, and we plan to make these extensions and further work available as well.
Q2	As part of making the software available to others, we plan to improve the learning algorithms further to handle much larger numbers of features and examples, along with exploring approximations that could be used to make the learning more tractable. Additionally, explore more complicated base learners above and beyond the current decision stumps, such as SVM or tree-based learners. We will have developed the co-regularization approach to improving prediction with small amounts of data by this point in the year.
Q3	Explore other approaches to the anytime prediction problem and compare with our existing SpeedBoost solution, further fleshing-out the anytime prediction domain. We believe our current approach for achieving near optimal prediction is the most natural approach, but we also plan to examine alternative ways of quantifying the performance of anytime prediction algorithms and improving our understanding of what types of algorithms work best under specific assumptions or problem conditions.
Q4	Work on additional RCTA application domains for anytime prediction, including other perception problems such as joint segmentation and image labeling problems and problems outside of perception like decision making, inverse optimal control learning, and other prediction problems in real-time domains. Here, we will collaborate closely with other aspects of Intelligence to determine what needs our basic toolbox is unable to meet. Investigate RCTA domains with prediction time constraints, other than computational resources, such as budgets due to data collection time or varying costs of certain tests or actions like exploration.

Subtask 2: Learning Low-dimensional Manipulation Planning and Control (UPenn)

Objective and Benefits:
The manipulation of objects in the presence of clutter requires fine planning of highly articulated actuators. Advances in such algorithms would advance Capstone scenarios like IED manipulation and removal. Unfortunately, current methods for planning the actions of a ground vehicle with a complex arm are inadequate and give suboptimal results. This subtask would use learning to discover a low-dimensional basis for optimal arm planning and manipulation. The proposed effort would require several years for algorithm development and integration into a modified UPenn robot with an arm attachment in order to show successful integration.

State of the Art:
Current state-of-the-art approaches to arm planning involve sampling the high-dimensional configuration space or using heuristic functions to speed the search. Unfortunately, effectively sampling a high-dimensional state space is intractable, and good heuristics generally involve a great deal of human engineered knowledge.

Key Barriers and Approach:
Our proposed approach overcomes the curse of dimensionality by learning an appropriate low-dimensional basis where trajectory planning becomes efficient and tractable. We have already shown that this approach results in optimal trajectories in simulations and on a 7-DOF arm. The learned basis will be useful for visualization for human-robot interfaces as well. The algorithm will be implemented as a module on the ground robots for validation and for the capstone assessment.

Metrics for Evaluation:
Optimality in the planned trajectories will be measured relative to state-of-the-art RRT and other planning methods. Metrics used will include computation time, motion actuation time, as well as arm motion efficiency and collision avoidance. At this point, all of our results have been done using physical simulation, but eventually, the metrics will be measured on ground robots with actuated arms.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Low-dimensional learning algorithm ported to embedded software.
Q2	Experimental arm manipulation testbed.
Q3	Implementation of arm planning on testbed.
Q4	Experimental validation with algorithms.

Subtask 3: Learning to Understand Human Instructions and Annotations (UoW)

Objective and Benefits:
We will develop a framework for **learning to map, or parse, high-level information and commands provided by a human into a formal representation** suitable for a robot control system and semantic world model. Instead of manually building ad hoc solutions that will not scale to the complexity of human body language and speech, our framework will be able to learn to parse human input from examples, and it will map the input into a theoretically well-founded, logical representation.

Ability to parse naturally provided human input into an internal representation is crucial to support missions described at a high level. This work provides key capability elements in the Capstone Assessment Vision, including 1 (“Understand commands”) and 3 (“Understand world through perception and WM”).

State of the Art:
Existing approaches to grounding natural language for robotics typically assume that a language parser is already provided or are not able to parse into a complex, formal representation.

Key Barriers and Approach:
We will build on recent advances in learning for semantic natural language processing. Our approach will learn CCG (categorical combinatory grammar) parsers that map human gesture (beyond pointing) and speech to a formal, first-order logical representation (lambda calculus) that is capable of reasoning about robot control actions and semantic world models. Thus, the system will be able to learn to parse human commands and annotations.

This effort will succeed because our current work gives strong evidence from the domain of learning to parse natural language directions that it is feasible to learn parsers from language to robot control systems. Here, we will extend that work to also incorporate raw speech and gestures and map into complex control systems and semantic maps.

Metrics for Evaluation:
We will measure progress by the ability of the system to correctly parse speech and gesture input into the corresponding formal representation. For instance, can the robot correctly understand and execute commands such as, "Take a picture of every person exiting that building" or information such as, "Bring the small box over here" spoken along with a pointing gesture? Progress can be evaluated based on accuracy on a specific set of test phrases and on the ability to parse more and more complex inputs.

Deliverables will be training datasets consisting of human speech and gesture input along with the corresponding formal semantics. Code parsing unseen human input and expanding semantic representation.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Extend existing CCG (categorial combinatory grammar) parser learning approach to handle input beyond natural language, including gesture and human gaze. Collect first training dataset using off-the-shelf speech recognition system and enhanced approach to human gesture recognition from Kinect data. Coordinate with MIT grounding effort to investigate the possibility of connecting parser learning with MIT perception grounding.
Q2	Develop formal semantic map representation suitable for CCG parsing and annotation. Investigate CCG learning in the context of the new data and semantic mapping tasks.
Q3	Start developing and implementing efficient learning technique for jointly mapping speech, gesture, and gaze to formal semantic map representation.
Q4	Investigate key open research issues related to interactive formal semantic map annotation. Develop approach to interactive object grounding, and expand semantic representation to incorporate novel objects on the fly.

Subtask 4: Anytime Learning for Robust Navigation and Control (ARL)

Objective and Benefits:
The objective of this subtask is to investigate the application of anytime learning techniques that are being developed in Subtask I5-1 to problems of robust navigation and control. We seek to leverage learning-based methods with anytime capabilities to find solutions to problems of robust ground navigation in militarily-relevant environments, i.e., highly constrained indoor/outdoor spaces with a variety of surface materials that lead to uncertain dynamic models.

State of the Art:
Previous research has attempted to address the issue of terrain and dynamic model identification for ground platforms of interest with some success when applied to adaptive control methods. Indeed, this work continues in Task I4. Methods for motion planning and control that consider uncertainty have traditionally been intractable, though recent work has begun to address this problem with suitable approximations. Instead of considering approximation to reduce computational load, our approach relies on learning based on past experience to more efficiently search the complex space of control inputs and uncertain trajectories.

Key Barriers and Approach:
The dimension and complexity of robust navigation and control problems grows quickly as we consider systems that must operate effectively in a wide array of mission scenarios. Consequently, approaches such as model predictive control must explore a large space of control inputs and output trajectories to find even a feasible solution. The pursuit of an optimal or near-optimal solution quickly becomes intractable when computational constraints and operational tempo are considered. Our approach will leverage learning based on past experience to guide the search of possible trajectories. The anytime nature of this approach will allow us to find progressively better solutions as operational tempo allows.

Metrics for Evaluation:
Metrics for the evaluation of this work include the rate at which the solution increases in optimality (where optimality is defined according to robust navigation and control metrics, e.g., control effort, probability of success, etc.) versus the time spent in computation. A more qualitative metric is the ability of this system to operate over a variety of terrain and environment types without “re-tuning” of the control parameters.

Collaboration:
This research task is being conducted in conjunction with Subtask I5-1: Anytime Prediction and Co-regularization, being conducted by Drew Bagnell at Carnegie Mellon University.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Discussion with Drew Bagnell (CMU) to identify possible research collaborations.
Q2	Specification of problem statement and algorithmic approach. This includes a mathematical formulation of the problem to be solved and simple simulations to demonstrate the idea.
Q3	Development of algorithms that can eventually be implemented on an experimental platform. Proof-of-concept in dynamic simulation engine.
Q4	Experimental evaluation of approach in a variety of environments. Compare performance (both computational load and navigation capability) with more traditional adaptive control techniques.

Subtask 5: Integrate Visual Texture Recognition (ARL)

Objective and Benefits:
This subtask will integrate visual texture recognition with existing audio-based proprioceptive sensing to enable intelligent path-planning based upon predicted traversability. By incorporating real-time visual texture recognition algorithms, we plan to allow a robot to learn to associate visual features associated with particular terrains with its own objectively measured success at traversing such terrains. This will make it possible for the robots to perform their own training in the unique environment to which they are exposed and gradually improve their performance at path planning through rough and varied terrain that will have the greatest chance of being successfully traversed in a short time and minimizing battery usage.

State of the Art:
Currently, the CMU team is exploring using acoustic data to recognize when a situation may potentially indicate trouble, such as tire slippage or water hazards. This provides feedback about the traversability of the immediate region surrounding the robot. Long range sensors, such as cameras, can give information about more distant areas, allowing the robot to plan ahead to avoid such obstacles. By using the robot's own experience in attempting to cross the terrains that it sees, it can learn to associate a measure of traversability with a particular visual pattern. This should avoid laborious hand creation of training sets and reduce the brittleness of such training when introduced into new situations. Instead, the robot will be able to continuously adapt to new situations. For example, suppose the normal wheels are removed from a robot and replaced with snow tires. Suddenly, snow becomes a traversable terrain, where before it was judged to be impassable. Normally, this information would need to be provided to the robot along with hundreds of examples of snow images labeled as traversable. However, with the proposed approach, the robot should be able to recognize the new state of affairs automatically and adjust to it.

Key Barriers and Approach:
This subtask consists of the following steps:

1. Develop a robust texture recognition system that is fast and consistent in its ability to recognize similar terrains as similar. This will make use of features such as viewpoint invariant color histograms and data association techniques such as visual filters.

2. Incorporate such data into a spatial map.
3. Make use of the proprioceptive measurements, and associate specific textures with measurements of traversability.
4. Continuously adapt path planning algorithms based on revised estimates of traversability.

Metrics for Evaluation:
The ultimate goal of this research (beyond understanding the continuous learning process) is to increase the ability of the robot to successfully navigate rough, varied terrain. Therefore, the primary metric for evaluation is an improved success rate and speed at reaching a goal through a realistic treacherous course as opposed to a robot lacking such a system. In developing a texture recognition system, we may need to compare with human-labeled segmentations of images to determine whether it is successfully associating similar textures with each other.

	Subtask 5 Quarterly Milestones and Deliverables
Q1	Develop texture classification algorithm.
Q2	Incorporate texture classification into spatial map.
Q3	Perform online training with proprioceptive labels.
Q4	Test robot's success in improving navigation.

I6-2012 – Life-long Learning

Level: 6.2 Applied Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Drew Bagnell	CMU	I6-1, I6-3
Additional Investigators	Daniel Lee	UPenn	I6-2
	Tony Stentz	CMU	I6-3

Objective:
Rather than attempting to design our robots to handle all possible missions, situations, and environments, we train them to mimic their human counterparts and equip them to continue to improve their performance during mission execution. The goal of this task is to extend these learning techniques to efficiently train from a large amount of experience, gathered over the lifetime of the robots. The proposed approaches combine encoding procedures for well-understood phenomena with new techniques for learning physical structures in the robot environment, even in the presence of extreme uncertainty. Thus, we anticipate robot intelligence to continually grow and expand through new and diverse experiences.

Background:
Currently, difficult tasks such as manipulation and mobility in cluttered environments require human operators to intervene and tele-operate individual robots. This task proposes to gather the information from how humans deal with these situations and to incorporate their experience into more autonomous robot algorithms.

The performance improvement in one component of a large system is only weakly correlated with the overall system performance, since most current learning algorithms are only aware of their local effects. The challenge is to fully integrate the various learning algorithms via their coupling through the shared world representation. Some of the subtasks seek to map features between different algorithms so that they can be succinctly incorporated in this shared representation.

Additional challenges exist due to the relatively small size of experiential data available. Existing robotic systems are trained using datasets that, while large by the standards of many machine-learning applications, are tremendously small in relationship to the amount of data available and the diversity of terrains, situations, and activities in which the robots engage. The result is that robots operate with an impoverished dataset of experiences in relation to the potential data available, limiting the depth of their decisions when faced with difficult tasks. Automatic collection of reliable data for the learning algorithms is also critical.

Research Description:
The first step in life-long learning is training from relevant temporal datasets from the tasks at hand, whether for manipulation, mobility, or surveillance. This training requires robot dialog protocols based upon shared mental models that have reasonable memory requirements and that draw upon the models developed by HRI Task H1: Team Shared Mental Models. It also requires integrating information across various robots, drawing on HRI Task H2: Situation Awareness. The training data will be stored by capturing, processing, and querying large datasets of experiences captured from real-world operation. These datasets are shared between platforms and across time and space to achieve maximum performance.

To fully leverage massive datasets of prior knowledge, techniques will be developed to rapidly index from new experiences to old ones and to recall experiences that may be "closest" to a given query. Sensor and map information from complex 3D environments will be fused and segmented according to their utility in the given task. Human knowledge will be leveraged to initially provide partial ground truth information to bootstrap these learning systems. Over time, these learning algorithms will be applied to ever-growing datasets to fully take advantage of the information streams coming from the target robot systems.

Metrics:
With regard to several tasks in the anticipated Capstone assessment, the life-long learning tasks will help to learn the appropriate terrain and environmental object models based upon the robots' previous experience. As the Soldiers input their mission objectives into their control station, the life-long learning module pulls-up the appropriate environmental parameters from its learned database. This program will show the difference in learning with this training data compared to conventional state-of-the-art systems.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
This task will involve collaborations between researchers and students at UPenn and CMU as well as with scientists at ARL.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Data collection from manipulation, mobility, and surveillance tasks.
Q2	Incorporation of data and human labeling of initial life-long learning examples.
Q3	Implementation of learning algorithms for manipulation, mobility, and surveillance.
Q4	Validation with current state-of-the-art algorithms on robot platforms.

Related Research:
This task shares connections with Task I5 as well as with DMUM and HRI tasks.

Integration and Assessment Activity:
Life-long learning models will be transitioned to incorporate the shared world model and architectural framework and will be implemented for the Capstone demonstrations.

Subtask 1: Life-long Learning of Manipulation Programs by Interaction and Demonstration (CMU)

Objective and Benefits:
Our efforts focus on life-long learning of manipulation skills. This is focused on existing failure modes for manipulation programs, which have largely been focused on well-spaced, highly textured objects on flat surfaces (e.g., Willow Garage PR2.) Most such systems fail if objects are slightly too close together, even if not actually occluding each other, making manipulation in real environments impractical.

Without handling the inability of existing manipulation systems to perceive and grasp in clutter and uncertainty, we believe the capstone vision will not be possible. Specific examples including manipulation tasks like unlocking and opening doors, and clearing clutter to find a specific object, e.g., an IED.

Given the criticality of these techniques and the relative immaturity of autonomous manipulation, we anticipate life-long learning techniques will be developed for this task for the length of the entire program.

State of the Art:
Current state of the art in manipulation tasks in clutter and uncertainty is very limited, consisting of mainly human tele-operation. Being able to leverage life-long experience in a learning system would provide many benefits to developing autonomous manipulation systems.

Key Barriers and Approach:
Our technical approach combines a few pieces. Clutter is a major issue for manipulation programs on both the perception and planning/control side. We begin by learning to detect known, important objects in significant clutter. This will be based on a combination of 2D and 3D depth image cues.

Unfortunately, an object – even if well known and modeled in advance – may be difficult to detect due to difficulty extracting it from the background. To address this, we couple the learning-based detection system with manipulation for active detection and localization of objects. We use pre-existing grasping and manipulation primitives to perturb and move objects that are close to one another to detect the separate rigid objects within view.

From this operation, we learn two key things. First, we learn which manipulation actions are effective in manipulating objects in clutter. As current manipulation, e.g., grasping, strategies require isolated objects, it is important to learn which primitives are actually effective through a reinforcement learning-style approach. We anticipate this building on our prior work on optimal control libraries.

Second, we are learning by repeated interaction to automatically segment different rigid objects, to provide an initial guess of potential objects to manipulate in a scene. Interaction provides ground truth, and we are using a new method that collaborating with perception efforts enables researchers to extract rigid bodies from motion.

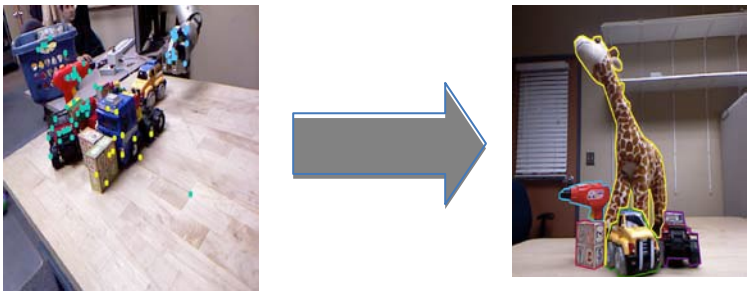

Figure 2-4: Life-long learning of object identity and rigid body segmentation by repeated interaction and perception techniques (2D/3D) to identify distinct rigid bodies.

Finally, we will develop reinforcement learning-based methods to accomplish complex motions that face pose uncertainty (e.g., opening doors and locks) by a new optimization technique to optimize over a library of maneuvers that reduce uncertainty. Our techniques combine both supervised learning and reinforcement learning to improve performance.

Ultimately, our system will provide to the world model segmentation into objects of the perceptual scene (based on both active and passive data), potential objection detections, uncertainty estimates of those detections, and suggested/preferred manipulation actions consistent with this scene. The basic segmentation into rigid objects is a unique combination of perception and learning only possible in robotics as we are using manipulation to determine (e.g., Figure 2-4) what objects are in fact separate and able to manipulated – it is a learning-based approach that is truly an active form of perception. This effort will coordinate with semantic scene understanding in Task P3.

Metrics for Evaluation:
We will measure progress by our ability to segment and our ability to correct grasp and manipulate objects under pose uncertainty. This will be quantified by success rates of each. This effort will leverage hardware and code for motion and perception primitives from the Inter Science and Technology Center Effort and the DARPA ARM-S project.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Develop ground truth signals for use as feedback signals in learning algorithms.
Q2	Develop model features that are useful for characterizing unknown terrain and objects.
Q3	Construct models offline based on robot experimentation.
Q4	Measure robot's ability to recognize and apply the correct model/plan online.

Subtask 2: Feature Learning for Life-long Surveillance (UPenn)

Objective and Benefits:
The long-term covert surveillance of a target, such as the back of the building in the Capstone scenario, will necessitate human operators monitoring the perceptual and world model representations of the surveillance robots. This is an incredible amount of data that can be difficult for untrained operators to manage, leading to decreased perceptual awareness. We plan to use methods to learn salient features present in the surveillance data which will also automatically adapt over time. Automatic calibration of robot sensors as well as alignment of their world representations will be necessary to quickly learn these salient features.

State of the Art:
The flood of perceptual sensor data over time and space from robotic surveillance can be overwhelming. Current state-of-the-art approaches to surveillance do not fully use 3D geometric information and have difficulties distinguishing false alarms with actual event anomalies.

Key Barriers and Approach:
We propose to use semi-supervised learning techniques to learn relevant surveillance features over time and to help register world model representations of different robots. We have started collecting visual and LIDAR datasets from different robots to learn 3D features that can help fuse these representations. Future work will involve adapting these features over longer time scales and implementing them on the robot platforms for surveillance applications.

Metrics for Evaluation:
Event detection rate along with ROC curves will be used to measure the efficacy of the different learned features. Computation time and human interaction time in using these systems will quantify the overall performance of the fielded systems. Short-term progress will be measured using results on collected datasets.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Long-term sensor data collection from multiple sensors.
Q2	Semi-supervised learning algorithm development and testing.
Q3	Application of semi-supervised features learning to merge robot world representations.
Q4	Validation on robot vehicle platforms.

Subtask 3: Self-supervised Learning for Building and Using Robot/Environment Interaction Models (CMU)

Objective and Benefits:
The objective of this subtask is for a robot to automatically learn complex interactions with its environment – without human supervision – in order to benefit directly from experience. The benefits are robots that learn to identify and avoid terrain hazards, manipulate unknown objects such as rubble, and adapt online to changing conditions – much faster than can be achieved through direct programming.

This subtask addresses Capability 5 (i.e., move through mobility challenges on the way due to both natural terrain and urban clutter) and Capability 10 (i.e., enter building by opening doors and avoiding or moving obstacles in doorway or building) in the Capstone Assessment Vision document. We expect this subtask to continue for two more years.

State of the Art:
A robot requires a model in order to plan interactions with its environment. These interactions include safely driving over terrain, pushing an object out of the way, and picking-up and removing an object. In some cases, these interactions are complex and difficult to predict. For example, a robot may high center when driving over a rubble pile, become submerged when driving through water, or become entangled when driving through bramble. When attempting to pick-up a piece of debris or push it out of the way, a robot may be unable to grasp it properly or cause it to slide or roll. The problem is that the robot has poor models for these types of objects and interactions.

The research community has addressed simple navigation interactions, such as identifying and reasoning about body collisions and tip-over hazards. The community has paid far less attention to more complex interactions, including those pertaining to traction, wheel/leg support, and engagement of non-rigid materials like water, bramble, and dense vegetation. Similarly, the literature is replete with research for grasping and manipulating man-made objects with known structure, such as parts for assembly, door knobs, coke cans, etc. The research community has paid little attention to objects with unknown properties, like rocks and debris. Such objects are particularly challenging since each object can be unique, potentially requiring a customized strategy.

Key Barriers and Approach:
Our overarching approach to developing intelligent robots is to engineer what is easy to model and to learn what is difficult to model. The robot can learn via human demonstration or self-supervision. The former is more efficient but requires a teacher. During mission execution, a teacher may not be available, and the robot must learn the appropriate behavior through feedback from its own experiences. This subtask focuses on self-supervised learning.

In order to automatically build and use interaction models, we will develop the following components:

- *Ground truth signal:* in order to learn a model, we need a feedback signal that correctly labels an interaction. For example, we can use proprioceptive data to detect a high-centering condition during navigation (e.g., vehicle is upright, wheels are spinning, GPS reports no motion). Similarly, we can detect that a manipulator is properly pushing an object by observing the motion with cameras or LADARs.
- *Model features and construction:* we don't know the exact physics of an interaction, but we can estimate contributing factors. Robot high centering is related to terrain bumpiness. Submersion is related to splashing noises. Sliding is related to flat objects. Rolling is related to round objects. We can calculate salient terrain/object features from sensor data and learn how to use them to predict robot interactions.
- *Model recognition:* once we have a set of models, when the robot encounters new terrain or a new object, it must decide which model applies. The robot could learn the model from scratch by interacting with it, but that poses a risk if the terrain/object is a hazard, and it generally wastes time and effort. Instead, we can learn to recognize which model applies visually using non-contact sensors such as cameras and LADARs.
- *Model-based planning:* accurate models can be used to plan interactions. A robot can drive around a "bumpy area" to avoid high centering, slow-down when encountering water to avoid submersion, as well as decide whether to pick-up, roll, push, or drag objects in order to remove them quickly.

This subtask addresses both navigation and manipulation interactions to leverage commonality in the models. In FY11, we focused on the interpretation of acoustic data for use as a *model feature* for navigation purposes. We trained the robot to recognize experiences that indicate trouble or the potential for trouble, including the sound of tires rolling through water, the scraping of the robot body against a rock, and the crunching of bramble as the tires compressed it. In FY12, we will fuse acoustic data with vibration and other proprioceptive data (e.g., IMU and GPS) for even greater awareness. We will develop an online learning capability so that once the robot interacts with a potentially hazardous terrain type (e.g., water puddles or mud) and builds a model, it will learn to *recognize* the model based on the appearance of the terrain (e.g., in camera and LADAR data). Thus, the robot can *plan* to 1) slow-down to a crawl to safely explore it or 2) avoid it entirely.

For manipulation of unknown objects, we will select salient *model features*, such as object size, shape, appearance, and sound (when touched) that are useful for characterizing robot interactions with the objects. We will extract these features from a Kinect sensor and microphones that are positioned to observe the robot interaction. Offline, we will use a manipulator with a grasping

end effector to fit a model to the features and learn which *plan* (e.g., pick-up, roll, push, lift, drag, etc.) is appropriate for a given model. The manipulator will learn by experimenting on unknown objects in dozens of piles. Online, the manipulator will *recognize* which plan to use based on the similarity of new objects to objects stored in its database. The recognition process may include simple tests, such as tapping the object to see if it sounds like it is metallic, in order to the choose the manipulation/grasping strategy. The goal is to learn to move objects as rapidly as possible without otherwise disturbing the pile.

This subtask will leverage learning techniques from Task I5 and interactive segmentation techniques (for ground truth) from another subtask in I6. This subtask complements the manipulation work in Task M3 and the navigation work in Task M7 by addressing robot interactions that are model poor (e.g., handling debris and natural objects and driving through water). This subtask builds on the perception work in Task P4 by expanding the set of proprioceptive sensors that are used to classify terrains.

Metrics for Evaluation:
For navigation, we will select a broad range of potentially dangerous situations, such as tires rolling through water or mud, metal scraping against rock, and brush/branches cracking, and measure the ability of the system to properly classify them. We will compare the case of a proprioceptive system consisting of only an IMU to one that adds acoustic and other sensors. We will compare false positives and false negatives for the two configurations.

For manipulation, the metrics for this subtask are the speed at which the manipulator and grasper can remove objects. We will compare the case where the manipulator applies the rules in raw form to the case where the manipulator has learned how to adapt the rules from its own experimentation.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Develop ground truth signals for use as feedback signals in learning algorithms.
Q2	Develop model features that are useful for characterizing unknown terrain and objects.
Q3	Construct models offline based on robot experimentation.
Q4	Measure robot's ability to recognize and apply the correct model/plan online.

I9-2012 – Distributed Intelligence for Human/Robot Teams

Level: 6.1 Basic Research

	Name	*Organization*	*Subtask*
Principal Investigator	Vijay Kumar	UPenn	I9-1
Additional Investigator	Ethan Stump	ARL	I9-2

Objective:
To make robots true team collaborators, we must model and quantify their capabilities, assign them tasks that play to their strengths, and design planners and feedback laws that enable them to cooperate with the appropriate level of communication to execute the task. Further, we must develop models based on HRI research that only require simple and infrequent communication during mission execution to simplify combat dynamics. Our approaches must allow for dynamic tasking, execution in uncertain environments, and adaptation to changes in the environment. Further, the approach must apply to heterogeneous platforms of different sizes, resources, and capabilities and employ computationally efficient algorithms, training one robot from another, and pooling scarce resources across the robot team.

Background:
Researchers have shown simple collaboration among heterogeneous teams of humans and robots using clean task decomposition, and they have applied the results to slow-paced tasks in relatively benign environments. To make collaboration truly effective, we need to address more sophisticated missions and missions that occur under conditions that are hard to predict. At the same time, we need to keep the teaming simple so that Soldiers understand the robots, trust them, and are willing to rely on them.

We, therefore, introduce a framework that combines probabilistic and deterministic reasoning techniques to enable task allocation, planning under uncertainty, and feedback policies that adapt to changes in the environment and actions of other robots or humans. In this approach, robots and humans interact with each other at multiple levels. At the lowest level, each reacts to its team members via implicit communication (for example, presence of neighbors in the vicinity or the need to perform cooperative tasks). At the highest level, the team members interact at the cognitive, human-understandable level. The robots are able to ask for help from other robots as well as human teammates. The end goal is to achieve versatility and multi-functionality but via simplicity in operation combined with high-end robustness in execution.

Our approach includes consideration for real-time performance on robots of all scales, including small robots with fewer computers and sensors, less-accurate positioning systems, and less-

capable mobility. At present, robot intelligence is scaled manually by reconfiguring software and hardware, or worse, by complete redesigning at significant development cost. This task addresses the need to scale-down robot resources without starting from scratch or substantially reducing performance.

Research Description:
Consider a scenario in which a team of humans and robots enters an unknown, complex environment to clear the area of potential targets by detecting, identifying, localizing, and tracking or neutralizing them. A robot must coordinate with other humans to explore the environment to build a world model, a task that involves task allocation, the control of positions (and orientations) to cover a three-dimensional space, and dispersion to appropriate positions to ensure persistent surveillance while tracking possibly moving targets. The development of frameworks and approaches to solve problems engendered by this scenario lends itself to a broad class of robotics applications.

In this task, we will focus on essential components of this problem area: (a) the development of robot control policies based on uncertain estimates of the environment that allows it to explore and cover the environment in a manner that best complements the actions of its human partners; (b) the reasoning about the uncertainty in the environment; (c) the fair allocation of the task across heterogeneous units which includes humans providing advice to human partners; and (d) concurrently performing tasks that involve tracking potential targets.

We derive representations from Intelligence Tasks I1 and I2 in which the continuous world is presented by simplices or grids allowing reasoning in a discrete setting. Uncertainty about each simplex or cell is represented by a probability density function [1-3,6,10]. Entropy-based metrics allow us to model what is unknown about the environment, while at a finer level, robots can reason about their planned actions by computing the mutual information related to their actions [1]. Since different agents have different capabilities, the mutual information gained by the actions of a robot and a human can be different. Thus, the goal of exploring the environment translates into different actions for a robot depending on its capabilities, its locations, and the locations and actions of its neighbors.

Our proposed representation for the group is derived from generalized Voronoi tessellations of the likely non-convex environment [11]. The allocation of tasks is automatically obtained by partitions derived from the Voronoi tesselations [7-8,11]. Robots adopt feedback control policies in each Voronoi cell to coordinate with neighbors. Their actions are natural functions of the nearest neighbors [10]. The uncertainty in the environment, modeled by entropy, and the definition of the cooperative tasks, modeled by appropriately derived cost functions, force the robots to coordinate their positions and cooperate to execute their tasks.

In previous work, we developed algorithms that enabled homogeneous robots to explore an unknown but static environment to generate a map. Further, we addressed the planning for cooperative tasks when robots must maintain constraints relative. In such problems, planning can become difficult because of the explosion in the dimensionality of the joint state space. We

developed a *plan consensus* methodology that builds on our preliminary research in this area [9]. Plan consensus determines how to decompose a coupled high-dimensional state-space planning problem to a series of lower-dimensional searches for plans converging to a locally optimal and *sufficing* solution (i.e., satisfying constraints), characteristic of consensus algorithms. The robots execute searches independently, with communication between the robots limited to exchanging solutions. Thus, at a slightly higher communication cost, the path consensus methods were able to provide rigorous guarantees on performance, an important factor in achieving robustness.

Based on the feedback, the work in 2012 will focus on the adaptation of individual robots to Soldiers and address the human-robot coordination essential to the seamless integration of robots in small unit operations. Specifically, we focus on two problem areas related to cooperative coverage, exploration and surveillance in the capstone vision by a robot and one or more humans. A robot will be able to search in or cover complementary areas or assist in mapping urban indoor or outdoor environments.

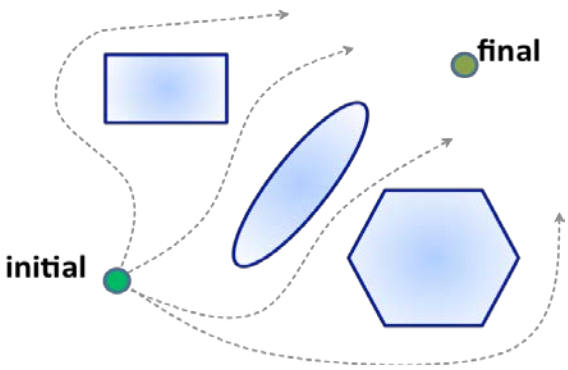

The first problem area concerns *reasoning about topological information about the environment* and planning robot paths or actions that are complementary and synergistic with the human. In the figure on the left, four agents may disperse to take distinct paths to the goals. They are mathematically distinct in the sense that the paths are not homotopically equivalent. We propose to embed in the robot the required intelligence that allows us to *reason about homotopy classes associated with indoor and outdoor maps*, relying on partial or complete topological information. Note that this information can be fairly coarse. For example, the information about the three obstacles in the cartoon can be encoded by three representative points (each homotopically equivalent to one obstacle). If three Soldiers take three distinct paths, the robot should know that it needs to take the fourth path.

Metrics:
Our algorithms will be tested via simulation and experimentation. For example, we will determine the ability to cover complex environments with and without our algorithms where the baseline algorithm will be a D* path planner. We will run mixed team simulations and measure mission success rate as a function of the complexity of the environment, complexity of the task, the need for human intervention, and the information exchanged between the human(s) and the robot. We evaluate the planning and control policies by considering increasing more complex, interconnected subtasks and evaluate the planning/execution time, data exchange, optimality, and completeness of the distributed solution compared to a centralized, fully informed solution.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
We plan collaborations with Dr. Stump on the coverage and mutual information-based exploration policies as well as with Dr. Sadler on estimated the biases and localization based on range sensors. The task includes collaborative work with Dr. Likhachev and Dr. Stentz on search-based planning techniques. UPenn student Ben Charrow will collaborate with both Dr. Sadler and Dr. Stump and visit ARL.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Formulation of the homotopy invariant for human-robot collaboration and the mutual information-based metric.
Q2	Methods for localization of the human Soldiers with time-of-flight radios.
Q3	Developing algorithms for planning in complementary homotopy paths, and maximizing mutual information.
Q4	Testing of algorithms via simulation.

Related Research:
While there is extensive literature on path planning based on metric information (e.g., from GPS), there is very little work that considers plans derived from coarse topological maps. In our own previous work [12], we have developed methods for computing homotopy class invariants based only on topological information about the environment. We will build on this work in our task. There is no other work in this area.

While there are a number of papers discussing active search policies, including those that consider mutual information [7-11], this is the first attempt to develop control policies that allow a robot to collaborate with one or more human partners.

This work complements ongoing work at ARL on estimation theory and localization (Sadler) in addition to planning to enable localization (Fink) and will be conducted in collaboration with both.

Integration and Assessment Activity:
This work can be easily transitioned to RCTA platforms. We will work with GDRS and ARL using the UPenn MAGIC 2010 platform as well as with ARL (Stump and Fink) using the Packbot platforms to realize the algorithms proposed in this work for IRA 2.

References:

[1] C. Stachniss, G. Grisetti, and W. Burgard, "Information gain-based exploration using Rao-Blackwellized particle filters," in Proc. of Robot.: Sci. and Syst., Cambridge, MA, June 2005, pp. 65–72.

[2] W. Burgard, M. Moors, C. Stachniss, and F. Schneider, "Coordinated multi-robot exploration," IEEE Trans. Robot., vol. 21, no. 3, pp. 376–378, June 2005.

[3] C. Stachniss, "Exploration and mapping with mobile robots," Ph.D. dissertation, University of Freiburg, Freiburg, Germany, Apr. 2006.

[4] A. Howard, L. E. Parker, and G. S. Sukhatme, "Experiments with a large heterogeneous mobile robot team: Exploration, mapping, deployment and detection," Intl. J. Robot. Research, vol. 25, no. 5-6, pp. 431–447, May 2006.

[5] M. A. Hsieh, L. Chaimowicz, A. Cowley, B. Grocholsky, J. F. Keller, V. Kumar, C. J. Taylor, Y. Endo, R. C. Arkin, B. Jung, D. F. Wolf, G. Sukhatme, and D. C. MacKenzie. "Adaptive teams of Autonomous Aerial and Ground Robots for Situational Awareness," Journal of Field Robotics, Vol. 24, No. 11-12, 2007: 991-1014.

[6] D. Fox, J. Ko, K. Konolige, B. Limketkai, D. Schultz, and B. Stewart, "Distributed multirobot exploration and mapping," Proc. of the IEEE, vol. 94, no. 7, pp. 1325–1339, July 2006.

[7] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," IEEE Trans. Robot. Autom., vol. 20, no. 2, pp. 243–255, Apr. 2004.

[8] M. Schwager, J. McLurkin, and D. Rus, "Distributed coverage control with sensory feedback for networked robots," in Proc. of Robot.: Sci. and Syst., Philadelphia, PA, Aug. 2006.

[9] S. Bhattacharya, M. Likhachev and V. Kumar, "Multi-agent Path Planning with Multiple Tasks and Distance Constraints," Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Anchorage, AK, May 3-7, 2010.

[10] S. Bhattacharya, N. Michael and V. Kumar, "Distributed Coverage and Exploration in Unknown Non-Convex Environments," 10th International Symposium on Distributed Autonomous Robots, Lausanne, Switzerland, Nov 1-3, 2010.

[11] L. C. A. Pimenta, V. Kumar, R. C. Mesquita, and G. A. S. Pereira, "Sensing and coverage for a network of heterogeneous robots," in Proc. of the IEEE Conf. on Decision and Control, Cancun, Mexico, Dec. 2008, pp. 3947–3952.

[12] S. Bhattacharya, V. Kumar and M. Likhachev, "Search-based Path Planning with Homotopy Class Constraints," 3rd International Symposium on Combinatorial Search (SoCS-2010), July 8-10, Stone Mountain Resort, Atlanta, Georgia.

Subtask 1: Collaboration in Human-Robot Teams (UPenn)

Objective and Benefits:

We will focus on the adaptation of individual robots to Soldiers and address the human-robot coordination essential to the seamless integration of robots in small unit operations. Specifically, we focus on two problem areas related to cooperative coverage, exploration and surveillance in the capstone vision by a robot and one or more humans. A robot will be able to search in or cover complementary areas or assist in mapping urban indoor or outdoor environments.

The first problem area concerns *reasoning about topological information about the environment* and planning robot paths or actions that are complementary and synergistic with the human. Specifically, we may want to have robot paths that are in homotopy classes that are different from those pursued by human Soldiers so we can cover the environment for mapping or for searching for intruders. We can also ask the robot to plan a path that belongs to the same homotopy class as the Soldier to allow for line-of-sight communication and for the robot to watch the Soldier's back.

The second problem area relies on metric information and guides the search based on measures derived from information theory. We assume that the robot acquires information from the Soldiers. This can translate into something. The robot moves to *maximize its mutual information*. In other words, it moves to locate its sensors to gather maximal information incorporating the information acquired by its team members.

State of the Art:

In our own previous work [12], we have developed methods for computing homotopy class invariants based only on topological information about the environment. We will build on this work in our task. There is no other work in this area.

While there are a number of papers discussing active search policies, including those that consider mutual information [7-11], this is the first attempt to develop control policies that allow a robot to collaborate with one or more human partners.

Key Barriers and Approach:

As discussed earlier, we will address two problem areas related to cooperative coverage, exploration and surveillance in the capstone vision by a robot and one or more humans. The first problem area concerns *reasoning about topological information about the environment* and planning robot paths or actions that are complementary and synergistic with the human. In the above figure, four agents may disperse to take distinct paths to the goals. They are mathematically distinct in the sense that the paths are not homotopically equivalent. We propose to embed in the robot the required intelligence that allows us to *reason about homotopy classes associated with indoor and outdoor maps*, relying on partial or complete topological information. Note that this information can be fairly coarse. For example, the information about the three obstacles in the cartoon can be encoded by three representative points (each homotopically equivalent to one obstacle). If three Soldiers take three distinct paths, the robot should know to take the fourth path.

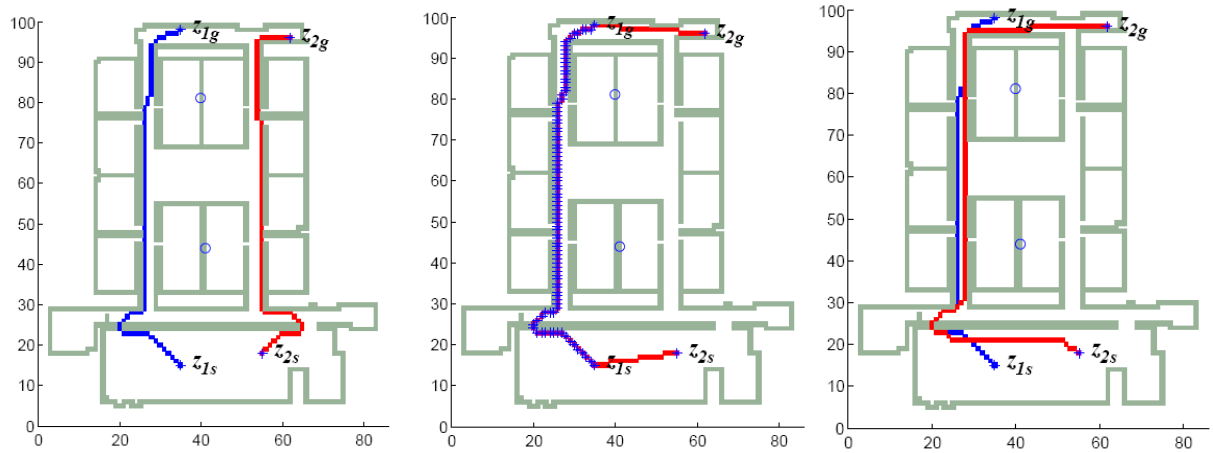

In the above figure, the robot (shown in red) plans its path to its destination independent of the human as shown in the left panel. These paths end-up being in different homotopy classes which is good if we want to cover the environment for mapping or for searching for intruders. However, we may want the robot to keep the human in line-of-sight to watch his back. In such a case, the robot may pursue the path shown in the middle panel where the robot joins the human before he starts moving. In the panel on the right, the robot is asked to plan a path that belongs to the same homotopy class as the human's path but otherwise independent. In all three cases, the

paths were automatically planned by augmenting search-based planning algorithms with topological information encoded by an appropriate 1-form that allows us to define a *homotopy invariant*.

$$H(\tau) = \int_{\tau} \omega \qquad (1)$$

The key point is that coarse topological information can be used to obtain these paths, even if precise metric information is not available.

The second problem area relies on metric information and guides the search based on measures derived from information theory. As in our previous work, the robot uses entropy to measure what it knows and what the team knows. Here, we must assume that the robot acquires information from the Soldiers. This can translate into something as simple as the human Soldiers wearing a belt with an IMU and a small range finder that allows the robot to integrate information acquired by the human Soldier. In the absence of this, a small time-of-flight radio with an IMU can suffice. The robot moves to *maximize its mutual information.* In other words, it moves to locate its sensors to gather maximal information incorporating the information acquired by its team members. In this setting, the robot is agnostic to the number of team members or the type (human, ground robot, aerial robot) of team member.

Metrics for Evaluation:
Our algorithms will be tested via simulation and experimentation. For example, we will determine the ability to cover complex environments with and without our algorithms where the baseline algorithm will be a D* path planner. We will run mixed team simulations and measure mission success rate as a function of the complexity of the environment, complexity of the task, the need for human intervention, and the information exchanged between the human(s) and the robot. We evaluate the planning and control policies by considering increasing more complex, interconnected subtasks and evaluate the planning/execution time, data exchange, optimality, and completeness of the distributed solution compared to a centralized, fully informed solution.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Formulation of the homotopy invariant for human-robot collaboration and the mutual information-based metric.
Q2	Methods for localization of the human Soldiers with time-of-flight radios.
Q3	Developing algorithms for planning in complementary homotopy paths, and maximizing mutual information.
Q4	Testing of algorithms via simulation.

Subtask 2: Cooperative Task Allocation for Surveillance (ARL)

Objective and Benefits:
The objective of this research is to investigate solutions to the general heterogeneous mission planning problem with spatio-temporal, cooperative, and platform constraints, applied to the problem of area surveillance.

Mission planning for distributed intelligence comprises the ability to take a set of tasks and

distribute and order them amongst a team of heterogeneous agents. These tasks can be constrained in both time and space and require the use of multiple agents cooperating to perform them, and the agents themselves may be constrained in their speed, operational time, and capabilities to execute certain tasks. Missions could be defined over a finite horizon or could be indefinite repetitions of a set of time-constrained tasks.

The following specific target scenarios all require robots to periodically visit sites within certain time windows and with certain constraints on simultaneous visits:

(1) Persistent surveillance of a region using aerial robots in coordination with mobile charging stations that must be deployed in support.
(2) Persistent surveillance of a region using aerial robots with mobile computation resources that must be moved in proximity of surveyors to offload the sensor processing.
(3) Persistent surveillance of a region with air and ground robots where certain pairs of sites must be visited simultaneously so that multiple views of a single scene can be obtained.

The intention is to provide a unifying mathematical framework for performing mission planning of heterogeneous assets with cooperative and spatio-temporal constraints. Any multi-robot system implementation will eventually reach the point where such a mission planning capability is required.

State of the Art:
There has been research into cooperative task allocation for robot systems [2], but it generally frames the problem as trying to optimally allocate the next task that each robot performs subject to statements of task utility. Many current implementations rely on a distributed auction mechanism for generating approximately optimal assignments.

Key Barriers and Approach:
For task sets with hard time constraints as we propose, we claim that it is important to plan not only the next task but the next several tasks that each robot will perform in order to make sure that robots assignments are not short-sighted – driving all robots far enough away from time-critical tasks that no robot will be close enough to respond when the task comes due.

We consider the basic version of our problem to be a Vehicle Routing Problem with Time Windows (VRPTW), a variant of the classic Vehicle Routing Problem (VRP) that seeks to design routes for multiple vehicles to efficiently visit all nodes in a graph. With the time window variant, the problem adds the constraints that each node can only be cleared within a certain time window, and we must specify traversal times for all edges.

We will address cooperative tasking by adding new classes of equality constraints that require certain nodes to be visited by two or more vehicles and then further specialize by defining classes of vehicles that are capable of fulfilling certain equality constraints. All of these modifications will naturally lead into adaptations of the basic relaxation framework introduced above and should lead to solvable problems.

To address repetitive task allocation, we modify the basic problem to adding repeat visits of the same site as new sites with different time windows, letting us enforce hard visit frequency constraints. Although the problem size could quickly grow in this case, we control this by only solving the problem optimally within a set time horizon and then executing this locally optimal path for a small amount of time, shifting the horizon, and resolving. This amounts to a receding-horizon control framework and potentially sacrifices global optimality in order to ensure that we can locally optimize a series of more tractable problems. However, the continuous resolving means that we can easily scale the problem to add or subtract agents in real-time, helping to ensure robustness.

The VRPTW is a strongly NP-Hard problem so for worst-case problem instances, it will be virtually impossible to find a solution. The general cooperative spatio-temporal task allocation problem outlined may not tend to produce problem instances that are easily solvable, though preliminary studies suggest it may [1]. Because the solver will probably not run in real-time, the solution is to reach for some of the many heuristic methods that have been developed in recent years.

Metrics for Evaluation:

For comparison and to begin to study the effects of next-best-task allocation versus long-term planning, we will implement the leading auction-based algorithm to evaluate how the final cost and simulated execution compare to our technique.

For extended or ongoing tasks, we have made use of a technique to break the long-term planning task into a series of optimizations over smaller horizons and propose to analyze the suboptimality of such horizon-based planning as a function of horizon length. Together, these analyses will give insight into the tradeoff between look-ahead planning and computational cost.

An initial implementation of an exact solver for VRPTW instances generated in a simulation framework has been documented in [1]. Heuristics and modifications will be based on this codebase.

Collaborations:

This research will be conducted collaboratively with Subtask I9-1: Collaboration in Human-Robot Teams, being conducted by Dr. Vijay Kumar at the University of Pennsylvania. Portions of this work relate to a proposed topic in the MAST CTA, in collaboration with Nathan Michael at the University of Pennsylvania. Prof. Michael proposes to create a hardware system implementation and demonstration of surveillance tasks using a team of UAVs over an extended operational time (~2 hours).

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Define a formal problem structure, and then map the three scenarios into this structure; create software code that can reduce the scenarios to standard form during operation. Research various heuristic techniques for solving large-scale VRPTW instances.
Q2	Expand existing VRPTW solver to operate on standard form. Expand codebase to implement candidate heuristics.
Q3	Develop realistic air/ground simulator with realistic power expenditure models. Implement leading auction-based technique.
Q4	Perform comparison studies of heuristic/exact VRPTW solutions against auction-based technique. Study effects of horizon-based planning on global optimality as function of horizon length.

References:

1. E. Stump, N. Michael, "Multi-robot Persistent Surveillance Planning as a Vehicle Routing Problem," IEEE Conf. on Automation Sci. and Eng., 2011.
2. B. Gerkey, M. Matarić, "A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems," *International Journal of Robotics Research*, 23 (9), pp 939-954, 2004.
3. P. Toth, D. Vigo, *The Vehicle Routing Problem*, SIAM, 2001.
4. B. Golden, S. Raghavan, E. Wasil, *The Vehicle Routing Problem: Latest Advances and New Challenges*, Springer, 2008.

3. PERCEPTION OVERVIEW

Objective

The role of perception is to generate descriptions of the robot's environments from sensor data. While there has been considerable progress toward describing environments for the purpose of mobility, e.g., characterizing terrain drivability, deriving higher-level descriptions of the environment remains a challenge. Such descriptions may include subtle cues and references that distinguish different behaviors and intents, recognition of specific classes of objects and features that are directly relevant to tactical behaviors, labeling of object, features, and terrain classes. The plan addresses the development of these capabilities through a combination of key tools that will be used as building blocks for the perception systems and transformative research ideas for longer-term advent of innovative perception approaches.

In developing this research plan, our objective is to address all of the requirements from the critical elements of the capstone vision, and to address them in one or several research tasks as illustrated in Figure 3-1 and Table 3-1. In converting the Soldiers' directives (go to the back of the building) to an executable internal plan (Element 3 of the capstone vision document), an internal model including semantic labeling of the environment with major features (e.g., building facades) (Tasks P3 and P4) is needed and an encoding of that knowledge in the form a semantic map (P3-5) is needed to communicate with the Soldiers. In moving through a complex environment that includes people and vehicles (element 4), key objects need to be identified. In order to make informed decisions, mere estimation of object locations and motions is not sufficient and identity information is needed, including reliable tracking in cluttered scenes (P6). Further motion through clutter (element 5) requires detailed terrain characterization from both vision and non-vision sensing (P4-1, P4-2). Surveillance functions used for detecting people leaving a building or other important events (elements 6-8) require deep understanding and forecasting of the activities of detected agents, i.e., humans acting in the environment of the robot, in the environments and their interaction with urban features (P5). Further modes of perception are required in close-range, indoor environments in which physical interaction, including pushing and touching parts of the environment, can be used together with vision sensing to acquire a more complete understanding of the environment (element 10, P2). Finally, operations in all of the elements of the capstone involve sensing with the type of surfaces and materials common in urban environments, and integration on small platforms (P1, P2).

Figure 3-1: The perception plan provide a complete coverage of the key challenges involved in the capstone vision through a combination of fundamental research tasks (P1,P3,P5) and applied research tasks (P2,P4,P6) in which the most mature approaches are readied for transition to assessments. The distribution of the perception subtasks to operational challenges is shown in the context of the two key aspects of the capstone vision: Understanding and sharing a description of the environment (top) and understanding and anticipating people's activities in a dynamic environment.

State of the Art and Technical Barriers

The state of the art focuses on generating representations of the robot's environment that are tuned primarily to mobility functions. This includes terrain characterization, detecting mobility hazards, 3D reconstruction, and SLAM. This level of perception provides the representations needed for mobility planning. We go beyond the state of the art by extending the environment representation with a level of semantic understanding by recognizing objects and features in static environments, and behaviors in dynamic environments.

In static environments, the state of the art in recognizing object categories and in scene understanding has made considerable progress in the past decade but performance remains limited in the context of the type of applications envisioned in the capstone. This is true both in terms of recognition performance and in terms of computational performance. While we import many of these state of the art techniques in the RCTA work, we address these performance challenges in three ways. First, we develop new techniques, e.g., for scene understanding, that are *more computationally efficient* than current approaches based on global optimization.

Second, we take advantage of the *contextual constraints* afforded by the mission and the environment. For example, we use prior knowledge about the environment (e.g., urban environment) to constrain tightly the perception process to yield a level of performance that cannot be achieved with general, unconstrained perception algorithms. Third, we use *all of the sources of sensor data* available (video, stereo, LADAR, and even data from physical interaction).

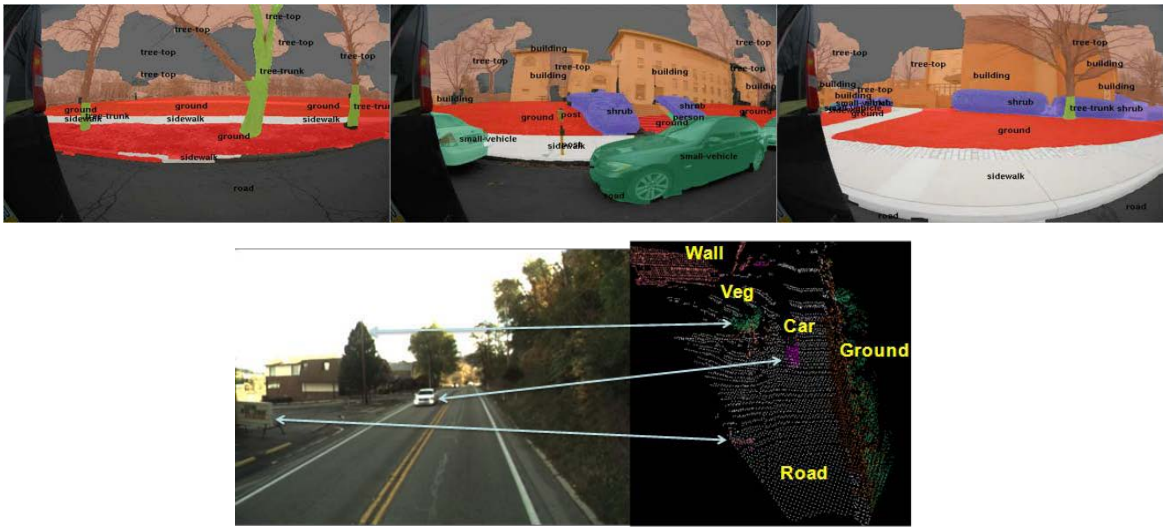

Figure 3-2: Example classification from images (top) and 3D data from LADAR (bottom) using 30 different classes.

In dynamic environments, we also combine contextual constraints and a comprehensive multi-modal approach to tackle the current limitations of the state of the art techniques in detection and tracking. In addition, just like in static environments, our objective is to extend the world representation beyond low-level descriptions of locations and motions and toward higher-level understanding of behaviors and intentions. Here too, we can build on extensive state of the art prior work, especially in the surveillance area. These techniques are limited in their predictive power, in their reliance on restrictive assumptions (e.g., static cameras, or short enough range to discern high resolution cues such as facial expressions). We address these barriers explicitly by developing new algorithms for longer-horizon prediction of motion and behaviors (Figure 3-3), and for making the algorithms suitable for on-the-move, longer range operation. To achieve the latter, our proposed development of algorithms using 3D data (from stereo or LADAR) is critical to beat the limitations inherent to video data, and the extensive use of contextual cues from the static environment (e.g., detecting key building features and their relations to behaviors in the capstone vision) is a key step toward achieving the desired level of performance.

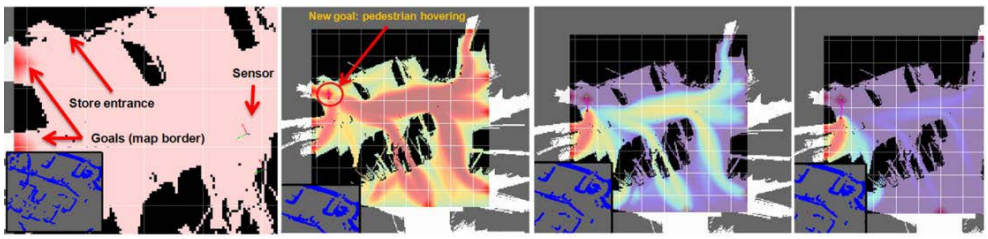

Figure 3-3: (a) Bird's eye view of a typical environment; (b) Prediction generated on the fly at three different times while tracking a person and updating environment model. The techniques developed in the CTA attempt to combine complex prediction models, detailed environment description, dynamically discovered goals and egress points based on observed behaviors, and high-resolution cues.

An important part of the work is to recognize and understand interactions between agents. Here too, while extensive state of the art can be leveraged, the RCTA context requires major advances due to the central role played by the robot (i.e., understanding behavior in the context of interactions with the robot) and the importance of understanding the relations between robot, team members, and non-team members. This aspect is now addressed in a new FY12 subtask (P5-2). These high-level, semantic interpretation elements rely entirely on good quality detection and tracking of people as their inputs. The extensive state of the art in this area is limited in the level of clutter that can be handled and by the fact that we need to track while on the move. On the other hand, in the CTA context, we can tackle these limitations by using rich context cues (e.g., ground surface location, type of objects and terrain) and 3D input from stereo and LADAR (Figure 3-4). Accordingly the plan includes a full complement of tools for robust detection and tracking in cluttered environments and for online execution on-the-move.

Figure 3-4: Example of people through cluttered, crowded environments. This capability is critical to generate the necessary input to semantic interpretation functions.

Finally, this vision of semantic perception can be achieved only if sensing systems are available to generate data of sufficiently good quality. We identified three limitations of the state of the art which define the technical challenges addressed in the RCTA plan. First, the environments of interests (e.g., urban environments in the capstone vision) include difficult materials such as metallic, texture-less, or semi-transparent surfaces, which are not handled well with current solutions. Second, while high-performance sensing solutions exist (e.g., scanning LADAR or high-res stereo), they are not adapted to the small platforms envisioned in this program and significant efforts need to be invested if semantic perception tools are to be exercised with these target platforms. Third, new forms of sensor data are needed for situations that are not covered

by state of the art solutions for mobility, such as operating in difficult environments (e.g., moving through rubble piles) or carrying out manipulation tasks.

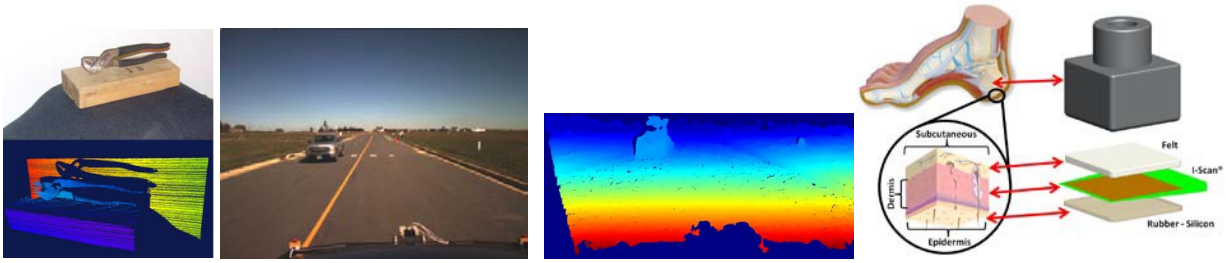

(a) (b) (c)

Figure 3-5: Examples of sensing approaches pursued in the RCTA in support of semantic perception and DMUM on small platforms in difficult environments: (a) precise 3D sensing for manipulation; (b) high resolution stereo; (c) pressure-sensitive robot skin.

Relationship to Other Technical Areas

The perception components interact with the rest of the system through the world model developed in Task I2. Specifically, in the simplest form of the interface, the perception algorithms generate locations of objects and labeling of 3D locations that are used for populating the world model. In more elaborate interactions, the perception algorithms provide distributions of possible labels and locations so that the system can reason about multiple hypotheses, and the world model provides top-down directives, such as guidance from contextual cues, for constraining the perception process. These interactions are developed in close collaboration with Task I1 on cognitive architectures.

The perception tasks leverage heavily the fundamental learning developed in Task I5, including imitation learning, anytime classification, and deep inference. Personnel are shared between I5 and the perception task to implement this relationship.

Several of the perception tasks, in particular in the area of dynamic environment interpretation are directly connected to the HRI vision of shared mental model and deep understanding of interaction between robot and humans. This is especially true of the strengthened emphasis in FY12 on perception tools for classifying interactions between robot and human agents from sensor data.

In addition to the core work on semantic perception, some of our novel sensing approaches (e.g., short-range precise range sensing and tactile sensing) contribute directly to the implementation of the DMUM capabilities and behavior generation.

Summary of Proposed Research

Semantic perception: Understanding the static environment

Our approach to static scene understanding includes building class models, incorporating context, and utilizing large datasets for memory-based learning techniques. The subtasks address different aspects of the scene understanding by using techniques most appropriate to each aspect. This approach breaks down the problem into manageable pieces, each specialized to a different

aspect directly relevant to the CTA scenarios. Specifically, we address five different levels of recognition tasks: 1) We develop techniques for locating *specific 3D objects* based on matching and explicit correspondences between data and 3D model for use in manipulation in which it is necessary to recognize and locate objects to grasp. Beyond manipulation, these techniques are needed for augmenting the environment description with location of known objects. 2) We use the *mission context* to constrain the recognition process. In particular, we focus on *urban features* that are most relevant to operation in urban environments. 3) For dealing with a more general level of description involving detecting all instances of *object categories*, e.g., people or cars, we base our approach on current techniques in which a classifier is trained on a large pool of data. We start refining this approach further by considering *"subordinate" categories*, which involves recognizing fine-grained categories, e.g., distinguish between man, woman, child in addition to categorizing objects as "people". The latter will be particularly important in providing cues to understand friend/foe relations in dynamic environments. 4) We combine vision data (image and 3D) with non-vision input from proprioceptive sensing in online adaptive learning techniques for *fine-grained characterization of terrain*. 5) We integrate image, 3D data (stereo and LADAR) and video in scene understanding approaches for general *scene labeling* in terms of broad classes (e.g., road, building wall) by using *efficient learning and inference* techniques.

We complement these different levels of recognition by investigating representations of the environment that are amenable to communicating descriptions to humans in the form of *semantic maps*. This includes strategy for building informative maps without partial localization and incomplete identification of objects, and techniques for building compact representations of the input visual data tuned to indoor environment.

The algorithms and perception modules developed in this area are directly relevant to the capstone vision elements, including: locating objects for manipulation; identifying building facades when exploring; providing cues for the location of egress points and generating description to Soldier; detecting people while moving in populated areas, including providing cues as to status through subordinate classification; detailed characterization of terrain and adaptation to local environment while en route to destination points; and labeling key classes in the environment while exploring towards destination and performing surveillance. The semantic mapping capabilities will be critical in two ways. First they provide the level of description needed for communication with users (e.g., element 3 "Understand world through perception and WM"). Second, they provide basic capability for representing indoor environments from visual and range data (e.g., element 10 of the capstone scenario). All of these capabilities will be evaluated using standard classification metrics such as precision-recall and ROC curves, and classification accuracy measurements such as AUC and F1 on progressively challenging experiments based on the capstone vision.

Semantic perception: Understanding the dynamic environment

Our approach to activity detection and understanding jointly models dynamic (motions, behaviors) and static (objects, surfaces) aspects of the robot's environment. We investigate robust techniques to identify current behaviors of dynamic objects, predict their future behaviors, and derive detailed descriptions of behavioral cues. Specifically, we explore fundamental algorithms for generating prediction of likely sequences of motions and actions from partially

observed sensor data; modeling and recognizing interactions between agents and robot; and maintaining awareness of relative location of agents with respect to robot, distinguishing between team members and non-team members.

Our technical approach to prediction is based on inverse optimal control developed in earlier work. In FY11 we showed that these techniques could be used in predicting distributions of paths at a rate fast enough for use on the move. In FY12 we extend this approach to the more difficult case of mixing paths with sequences of discrete actions. In addition to motion and coarse-grained actions, we continue our investigation of high-resolution cues, such as facial cues, for understanding behavior and intent. For modeling and recognizing interactions between agents and robot, we introduce in the FY12 plan a new approach based on hierarchical modeling of cause-and-effect relationships composing activities, which will automatically learn semantic structure of the activity in terms of spatio-temporal relationships among its sub-events, representing influences of the activities to the robot itself. We address team localization and tracking issues by using a multi-hypothesis data association technique that is able to accept disparate data types and track asynchronously. The data association part will address the problem of tracking through occlusions and crowded scenes. The multi-modal part of the work emphasizes the use of many different sources of sensor data in addition to vision.

These semantic interpretation functions are supported by a full complement of tools for robust detection and tracking in cluttered environments and for online execution on-the-move. As in the part of the plan on static environment understanding, we attempt to provide a complete coverage of different approaches relevant to different situations, sensors, and platforms. This includes robust detection and tracking from stereo, LADAR, and video, as well as using contextual cues (e.g., scene layout and urban features) and collaborative tracking across multiple sensors and platforms.

Sensing

The above semantic perception capabilities rely on the availability of good quality sensor data. While in benign environments, standard cameras and scanning LADAR are adequate, more complex environments bring in more challenges that are not sufficiently addressed in the state of the art. Given their importance to the success of the proposed semantic perception agenda, we focus on two classes of challenges. First, we investigate the problem of handling surfaces of “difficult” materials such as textureless surfaces (e.g., walls), which are problematic for stereo, and thin structures, metallic and semi-transparent objects, which are problematic for both passive stereo and LADAR. It is particularly important to address these limitations now because these types of surfaces and materials are commonly encountered in the scenarios called in the capstone vision. If these issues are not properly addressed, we run the risk that the semantic perception modules will be starved for high-quality data at critical points in those scenarios. Second, we investigate the migration of the sensing approaches to small platforms and to support additional manipulation capabilities.

Task	Primary Research Outcomes	Operational Implications
P1: Exploiting Novel Sensor Phenomenology	New fundamental approaches to sensing addressing limitations to SoA in difficult environments • Dealing with difficult materials (low-texture, semi-transparent) • Non-visual sensing (e.g., pressure) • Innovative techniques for LADAR sensing	Addresses the limitations of existing sensing approaches for the type of environments envisioned in capstone (e.g., texture-less walls, metallic/glass surface in urban environments) for semantic perception, and mobility/manipulation in difficult environments (e.g., stairs, rubble)
P2: Compact, High Performance Sensors	High performance sensing modules for small platforms in difficult environments • LADAR and stereo for small platforms • Short-range precise sensing for manipulation	Addresses the need for high-performance sensing to move semantic perception capabilities to small platforms
P3: Static Scene Understanding	Fundamental algorithms for recognizing objects and features in sensor data: • Specific 3D objects • Generic object categories and subordinate categories • Context driven features (e.g., urban environments) • Class labeling Tools for semantic mapping	Identification of a larger selection of features for a more complete representation of the environment for execution of more complex missions Communicating environment understanding with human
P4: Perception for Missions in Complex Environments	Perception modules for: Integrating perception and action for mobility through movable obstacles Adaptive fine-grained terrain characterization Online semantic scene labeling from 3D sensing	Robust operation in a variety of environments (unknown terrains, vegetation, etc.) Dynamic updating of environment representation for robust on-the-move operation Automatic adaptation to environment

Task	Primary Research Outcomes	Operational Implications
P5: Dynamic Scene Understanding	Fundamental algorithms for behavior recognition and prediction, and interaction with team members in dynamic environments, including • Long-term prediction of motions and actions from partial observations • Recognizing interactions between agents and with robot • Tracking location of team members from sensor signals	Enhanced situation awareness in dynamic environment beyond location and motion of agents Better model of interactions between agents and environments Better shared mental model through understanding of interaction with robot Toward detailed understanding of activities of team versus non-team members
P6: Perception for Missions in Dynamic Environments	Perception modules for robust detection and tracking in complex environments from different sensor modalities Online, real-time predictive models	Ability to track objects in complex, cluttered environments Ability to anticipate motion while on the move Robustness through multiple modalities (stereo, monocular, LADAR)

Table 3-1: Perception Task Outcome and Operational Implications table.

Path to Integrated Research Assessments

The research products generated by the Perception portion of the RCTA contribute to the IRAs and to the Capstone assessment in all the "Look" functions. We plan to roll-out these perception capabilities on a bi-annual basis, according to the assessment schedule laid out in Figure 1-8. We summarize below how the research planned in our six tasks maps to each of the IRA milestones. All of the capabilities required by the IRA, including the Capstone, are addressed in some form in one of the Perception tasks.

In IRA1, stereo sensing and people detection/tracking capabilities were integrated with reasoning and world model elements. This milestone enabled the initial assessment of our approach to cognitive architecture, in particular around the ACT-R model. Beyond the technology demonstration, this assessment motivated close collaboration between researchers in the Intelligence and Perception areas, paving the way for continued inter-task collaboration in the RCTA.

Moving forward, in IRA2 we will turn our attention to a unique feature of the RCTA program, which is the emphasis on *interacting* with the environment in addition to moving through it. Toward that IRA, we investigate all levels of the "Look" functions needed for dexterous manipulation. At the lowest level, we develop advanced sensing capabilities for precise short-range sensing under the difficult constraints of outdoor illumination, difficult materials, and compact sensor and computation package for small platforms. At an intermediate level, we investigate approaches for representing the geometry of objects and terrain for manipulation, e.g., stereo sensing for robotic trenching. At the highest level, we incorporate knowledge about

specific objects through object recognition capabilities. While a considerable amount of prior work exists in the area of object recognition, in the context of the RCTA, the emphasis is again on harsh observation conditions, difficult materials, and the need for precise localization of the object. This theme of "Look" capability for physical interaction is further explored in IRA4 in which, in addition to manipulation, we consider mobility issues: mobility through difficult terrain and for unique mobility mechanisms. This includes terrain sensing through non-vision means (e.g., pressure sensors, skin, IMUs) as well as learning techniques for terrain adaptation. Current research is already moving toward this IRA through, for example, initial investigation of touch perception for legged robots.

On the semantic perception side, IRA3 greatly expands the initial "Look" capabilities demonstrated in IRA1 in two ways (broadly speaking). First, it incorporates semantic labels for the first time, e.g., automatic detection of building location, doors, or sidewalks. A preliminary version of this capability was shown at the 2011 RMB. Research in this area is central to the perception plan (P3/P4). Second, IRA3 incorporates, also for the first time, semantic elements of dynamic aspects of the environment, e.g., initial cues about people's behavior from their motion (with also a preliminary demonstration at the 2012 RMB), which is a large part of the current perception investment (P5/P6). In addition, this is supported by advances in robust motion tracking from multiple sources like stereo and LADAR. This IRA3 milestone also requires tighter integration across the RCTA areas, specifically for integrating the perception components within the design of the world model (this collaboration is already on the way, based on the work done for IRA1, as documented in the collaboration Table 1-1). Building on this milestone, the plan is then to incorporate more elaborate semantic elements into the system. In short, extension to IRA5 and IRA6 includes the following research products (based on research for the most part already initiated in our 6.1 tasks P3 and P5): more detailed semantic description of the environment (e.g., including object categories, more semantic classes, features such as building elements that can be used to reason about egress points); forecasting of activities based on imitation learning techniques developed in the Intelligence area; and interaction between static and dynamic understanding of scenes. Combined, these capabilities support the "Think" capabilities of advanced reasoning about and prediction of friendly and adversary activities, and they support the construction of a situation model that can be communicated to and shared with Soldiers. As indicated above, the current perception portfolio supports the longer-range goals, e.g., a preliminary view of activity prediction (motion patterns from LADAR tracking) with simple interactions with environment was demonstrated.

Perception Tasks

P1-2012 – Exploiting Novel Sensor Phenomenology

Level: 6.1 Basic Research

	Name	*Organization*	*Subtask*
Principal Investigator	Michael Powers	GDRS	P1-2
Additional Investigators	Larry Matthies	JPL	P1-1
	Roland Brockers	JPL	P1-1
	Emmanuel Collins	FSU	P1-3

Objective:
This task will improve robotic sensing capability by exploiting both physical and mathematical (signal/image processing) phenomena. Capability refers to the perception ability or modalities of a sensor package and also to its practical characteristics, in terms of the methods used and the quality of perception data produced. Unambiguous and high-fidelity perception information is essential to achieve semantic understanding, situational awareness, and accurate mobility. Three subtasks address these needs in the context of the capstone assessment: Subtasks 1 and 2 seek substantial advances in active and passive sensing techniques, respectively, and task 3 explores terrain contact sensing with a particular emphasis on legged mobility.

Background:
This section summarizes the state of the art in the general research area addressed by this task. More focused discussions of state of the art, prior work, and barriers are included in the description of each subtask.

Current sensing approaches deal poorly with many conditions including textureless objects, semi-transparent objects, hidden/obscured objects, and darkness. They lack the robust, reliable, real-time performance needed to support complex military and commercial applications. Although current programs address some of these shortfalls, we must further understand the basic phenomenology of sensing with the goal of discovering low-cost, novel ways to provide rich data to perception algorithms. In addition, we need to mature promising approaches to address limitations in resolution, range, and robustness as well as the size, weight, and power considerations that limit implementation on smaller ground and air vehicles; this aspect is covered in the companion applied research task, P2.

Research Description:
This section describes the key objectives of the task and the corresponding technical approaches and their mapping to subtasks. A detailed discussion of state of the art and metrics associated

with each subtask is provided in a separate subsection for each subtask at the end of the task description.

3D Perception in Challenging Environments.
We investigate issues associated with sensing in difficult environments and on challenging materials (specular, low-texture, thin objects) by exploring the phenomenology of different sensing modalities and designing new approaches. In this plan, we separate the investigation in two parts: active sensing which includes the design of innovative LADAR solutions and passive sensing which addresses the use of cameras.

In active 3D perception, we explore the use of compressed sensing to implement adaptive cross range resolution and frame rate. These techniques can dramatically reduce the power consumption, laser emission, and hardware cost of LADAR imagers. Concepts to be investigated in this work stem from recent work in the mathematics of compressed sampling. One of the most notable applications of these techniques to a practical problem was demonstrated by the "Rice single pixel camera" with results published by Duarte et al. in 2008. With this precedent, we will analyze the trade-offs and impacts this has for active sensors such as LADAR, with particular emphasis on signal-to-noise ratio (SNR), minimum resolvable features, laser power requirements, and compression artifacts.

In passive 3D perception, our FY11 activity exploited recent new insights to address two fundamental, outstanding problems in stereo vision: pixel locking and bias in visual odometry. Pixel locking is a phenomenon in which stereo disparity estimates are biased toward integer values. This phenomenon has been known at least since the late 1970s in the work of Donald Gennery [1] and has been the subject of substantial research in the past 20 years; however, methods have not yet been found that reduce the bias and are very computationally efficient. Our new insight was that bias estimates potentially can be computed efficiently from the reference (e.g., left camera) image before performing stereo cross-correlation, then applied to correct the depth map after cross-correlation. Visual odometry has also been observed to be biased since the work of Hans Moravec in the 1970s; however, to our knowledge, there had been no work to address this prior to 2011. This bias manifests itself in biased estimates of the total distance travelled; where feature points are predominantly on the ground plane, it also produces a biased estimate of vertical motion, making it appear that the vehicle is drifting up or down. In FY11, we are developing methods to reduce this bias. Our proposed FY12 research returns to the theme of 3D perception in challenging environments to improve passive 3D perception of low-texture surfaces and thin structures, which are ubiquitous in urban environments and figure prominently in the vision for the capstone assessment. Whereas the bulk of research in computer vision addresses such problems with algorithms that require global optimization algorithms that are very expensive computationally, we explore algorithms with low computational complexity that are feasible on small, SWaP-constrained platforms. Our intent is to extend this work in FY13 through FY14 to address specular and semi-transparent surfaces which are also ubiquitous in urban environments.

Contact Sensing for Terrain Classification and Dexterous Manipulation.
Contact terrain sensing, to be extended and improved in Subtask 3, supplies perceptive information that is not reliably obtained by passive or active optical inspection. In terms of the use of robot skin on a robot foot, the results of this research give to the world model knowledge or verification of the underlying terrain. This classification may be aided by receiving from the world model contextual information that may aid in the classification (e.g., it is below freezing, so beware of icy surfaces; or there has been substantial rain within the last few hours, so beware of mud; or the robot is indoors, so do not try to classify the terrain as an outdoor surface).

In terms of the use of robot skin on a manipulator end effector, the results of this research give to the world model knowledge or verification of the surface type (e.g., plastic, steel, cloth, etc.) of objects in the environment and the exact nature of the object (e.g., a book, a cup, a small ball, etc.). This classification may be aided by receiving from the world model approximate locations of objects as perceived by vision.

Metrics:
Metrics for evaluation are described quantitatively at the subtask level.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
In 2012 we will explore collaborative efforts with ARL SEDD on LADAR tasks, with the possibility of rotating GDRS staff to ARL Adelphi and vice versa. This will join with active collaboration with the core electronics task in P2.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Develop initial multi-resolution stereo algorithms. Perform initial system and component studies for CS LADAR. Conduct contact terrain sensing experiments on various terrain types.
Q2	Explore edge-preserving stereo methods; reduce computational load of algorithms. Design prototype CS LADAR hardware. Demonstrate contact terrain sensing in real-time.
Q3	Incorporate edge-preserving approach into multi-resolution approach. Measure performance of CS LADAR. Investigate end effector usage of PreSRS system.
Q4	Verify new stereo algorithms on ground-truthed test data. Collect CS LADAR image data, and evaluate results. Conduct object recognition experiments with the PreSRS system.

Related Research:
Tasks in P1 are closely linked to Terrain and Object Classification, Identification, and Reasoning (P3 and P4) in Perception and numerous tasks outside this technical area that require accurate

world models and semantic understanding. The active and passive sensing research explored in Subtasks 1 and 2 is fundamental to the capstone assessment and underpins system performance as a whole. While work in previous years on the RCTA and work produced externally has demonstrated successful results, it is important to recognize that perception information with increased detail and accuracy remains vital to mission performance in practical environments. It is reasonable to expect that adversary groups in actual mission environments will take care to conceal their activities, offering little in the way of ideal robotic perception scenarios. Moreover, common objects not specifically intended to conceal targets pose serious perception problems (specular surfaces, low texture surfaces, semi-transparency). Addressing these challenges is more than a matter of incremental improvement and requires better understanding and incorporation of sensor phenomenology.

Activities in contact terrain sensing are directly related to Task M3: *Sensor-based Dexterous Manipulation*, which considers robust grasping, Subtask I10-4: *Understanding and Manipulating Irregular Objects*, Task M7: *Learning Terrain Interactions*, which requires the robot to sense the underlying terrain, and Task M8: *Dynamic Multi-modal and Reconfigurable Mechanisms*, where the robot gaits and climbing grips should vary with the terrain. It also bears some relationship to the following tasks, which consider terrain classification: Task P4: *Perception for Missions in Complex Environments*, and Subtask I6-4: *Heightened Proprioceptive Awareness and Safe Exploration*. Notably, Subtask 3 is specifically oriented towards legged robots as opposed to wheeled or tracked robots.

P1 tasks are closely associated with applied research tasks in P2, to evolve promising developments through the technology readiness framework (TRL).

Integration and Assessment Activity:
Technologies developed under this basic research task (Subtasks 1 and 2) will be transitioned to the corresponding applied research task, P2, pending successful evaluation of their utility. Testing scenarios will be constructed with recognition of the capstone scenario.

Work conducted under Subtask 3 is expected to be transitioned to Task M7: *Learning Terrain Interactions* and Task M8: *Dynamic Multi-modal and Reconfigurable Mechanisms* in future years.

Subtask 1: Passive 3D Perception in Challenging Environments (JPL)

Objective and Benefits:
The objective of this subtask is to enable high-quality 3D perception on scenes with low-texture surfaces, thin objects, and specular and transparent surfaces. Urban areas are full of low-texture surfaces (e.g., roads and walls), thin structures (e.g., sign posts and vegetation), specular surfaces (e.g., vehicles, metal objects, some floors and walls), and semi-transparent surfaces (e.g., dirty windows, which are also specular). We expect a byproduct of this work will be improved ability to recognize objects and identify the surface materials in urban environments, which will

contribute to semantic perception and improve navigation and manipulation in urban environments. This subtask conducts basic (6.1) research on modeling of phenomenology and improved algorithms for passive ranging in complex scenes, and it complements applied (6.2) research in subtask P2-5 on scaling high performance 3D perception to small platforms.

In FY11, this subtask has addressed outstanding problems with subpixel disparity biases (pixel locking) in stereo vision and with motion estimation biases in visual odometry. In FY12, we will address improving stereo vision-based ranging on low-texture surfaces and thin structures. We intend this task to continue through FY14, progressing to address specular and semi-transparent surfaces after FY12.

Examples of how this work is relevant to the Capstone Assessment Vision include: 1) improved 3D perception of driving surfaces on low-texture roads and sidewalks and on specular floors as well as improved 3D perception of low-texture walls; 2) improved ability to use 3D range data from stereo vision in recognizing objects such as thin tree trunks, sign posts, and chair legs; and 3) the new ability to understand and do 3D perception through windows, despite dirt and specular reflections on the windows.

State of the Art:
Most real-time stereo vision systems operate at only one level of resolution: the highest that can be afforded with the processor available. While performing stereo matching at a high resolution usually translates to better 3D perception of fine details, the fastest class of stereo vision algorithms (based on local optimization) perform poorly on low-texture surfaces and, as a result, only generate sparse stereo results in such image areas (Figure 3-6). Algorithms that embody surface smoothness constraints (e.g., regularization models) are able to fill in range estimates for low-texture areas but at significantly higher computational cost. Multi-resolution approaches have also been used to overcome the issue of poor stereo matching on low-texture surfaces by using low resolution correlation results in an image pyramid to improve results at a higher resolution by limiting disparity search intervals [2] or averaging correlation scores over image resolutions [3]. However, to our knowledge, there has been no work that leads to a practical, real-time stereo system that has demonstrably improved 3D perception on low-texture surfaces while not getting worse results elsewhere, such as at range discontinuities, especially around thin structures.

Additionally, most recent real-time implementations use GPU hardware to achieve high frame-rate results [2, 3], which is not an option for low-power, real-time performance.

Key Barriers and Approach:
In FY12, we will focus initially on low-texture surfaces. Such surfaces create two types of problems for real-time stereo vision: 1) lack of range data, and 2) even where range data is obtained, there tends to be a lot of range data "flicker," where different parts of the surface get range data in successive frames. Flicker can be overcome by integrating data over time, but it would be better to get more complete range data in every frame with reduced flicker, especially if this can be done with little additional computational cost. It is well known that stereo algorithms based on semi-global or global optimization produce better depth maps in these

situations, but at high computational cost that is fundamentally at odds with the need to provide compact, low-power implementations for onboard use. Empirically, we have observed that denser range data for low-texture surfaces is often produced at lower levels of image resolution (Figure 3-6).

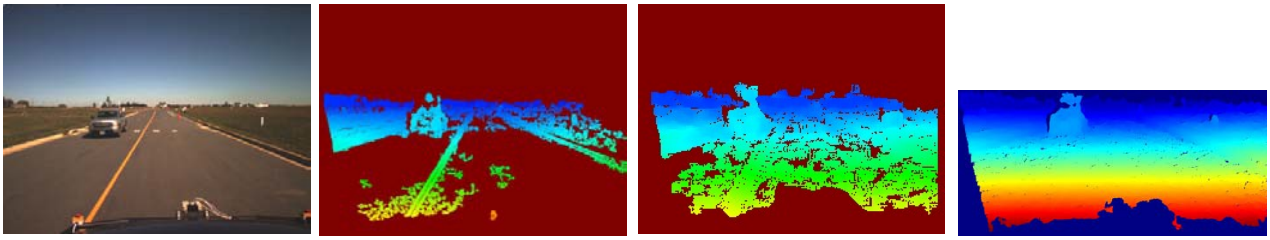

Figure 3-6: From left: left image of stereo pair at 1024x768 pixels; depth map produced at full resolution with standard real-time stereo algorithm; depth map produced at 256x192 resolution with the same algorithm; depth map produced at full resolution with a global algorithm (cp. [3]).

We believe that fundamentally the best approach to improving performance on low-texture scenes is to exploit multi-resolution matching in order to get the kind of benefit seen in Figure 3-6 with the computational efficiency possible with multi-resolution methods. While there is a very long history of research on multi-resolution stereo algorithms, algorithms that are both effective at combining results across resolution and extremely efficient in embedded processors have not emerged. Our research will address this gap.

Getting good stereo results on thin structures requires processing images at the highest possible resolution with the smallest possible correlation window. In the limit, this leads to algorithms that similarly use functions that integrate over only one pixel but make-up for the reduced effective window size with regularization functions that again require expensive optimization algorithms. Nevertheless, some very impressive results have been achieved with novel local optimization algorithms, such as by using image gradient information in what amounts to fast, adaptive window-size algorithms [5] (Figure 3-7). Once we have satisfactorily addressed low-texture surfaces, we will study how to incorporate fast methods appropriate to thin structures into multi-resolution processing schemes, so as to simultaneously address low-texture surfaces and thin structure with the greatest computational efficiency. Key to this overall approach is the recognition that, while improvements over the state of the art are required, results for each frame do not have to be 100% perfect; we seek an appropriate knee in the curve of data density versus computational cost.

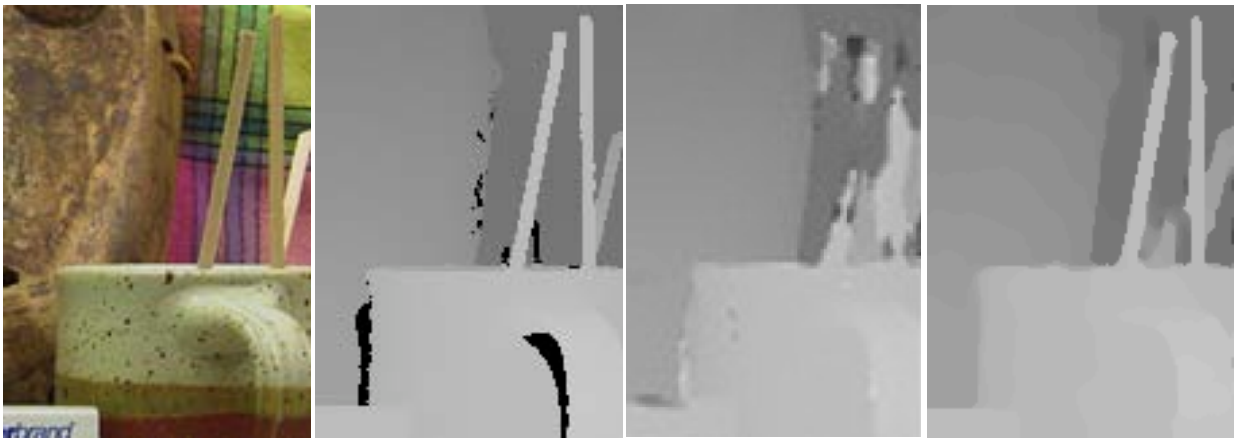

Figure 3-7: Stereo on thin structure: From left: Detail of original left view of stereo pair; ground truth disparity map (lighter grey is closer to observer); real-time GPU-based multi-resolution approach [2]; edge-preserving approach [5].

Results from this subtask will be transitioned to P2-5 for integration in low-power, onboard vision systems, Integrated Research Assessments, and the Capstone Assessment in subsequent years. We expect that improved range data from this task will benefit all areas of Perception.

Metrics for Evaluation:
Metrics for this problem are quantities like data density (percent of the image with correct stereo disparity at pixel resolution) versus computational cost, data density versus measures of image complexity (e.g., image contrast in the simplest case), or variance of the range data versus computational cost and/or image complexity. The perennial challenge in this area is how to get good ground truth range data in complex, outdoor scenes. For important classes of low-texture surfaces, this can be done adequately by using semi-global or global optimization-based stereo algorithms, as seen in Figure 3-6, optionally coupled with other methods to estimate the location of the ground plane to verify the results of the (semi-)global stereo algorithm. With care, good LADARs can also be used to provide ground truth range data for stereo systems, especially for low-texture driving surfaces. We previously refurbished our Riegl LADAR for exactly this purpose. Obtaining ground truth for thin structures outdoors is more difficult because LADAR or other range sensors may also have trouble accurately perceiving the structure. Here, the redundancy provided by large numbers of cameras in a multi-baseline stereo configuration may be the most practical approach to obtain quasi-ground truth. We will acquire datasets with a combination of these methods, and we will measure progress with the metrics described above.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Develop initial multi-resolution stereo algorithm to improve stereo on low-textured surface.
Q2	Explore edge-preserving stereo methods for stereo on thin structures. Develop a fast method targeted for low-power platforms.
Q3	Incorporate edge-preserving approach into multi-resolution approach.
Q4	Verify new algorithm on ground-truthed test data for low-textured scenes.

References:

[1] D.B. Gennery, *Modelling the environment of an exploring vehicle by means of stereo vision*, PhD thesis, Stanford University, 1980.
[2] Y. Zhao, G. Taubin, Real-time stereo on GPGPU using multi-resolution adaptive windows. *Image and Vision Computing* (29), 2011, pp.420-432.
[3] R. Yang, M. Pollefeys, Multi-Resolution Real-Time Stereo on Commodity Graphics Hardware. *Proc. of CVPR*, 2003, pp. 211-213.
[4] R. Brockers, Cooperative Stereo Matching with Color-Based Adaptive Local Support. *Proc. of CAIP*, 2009, pp. 1019-1027.
[5] C. Cigla, A.A. Alatan. Efficient edge-preserving stereo matching. *ICCV Workshop on Live Dense Reconstruction with Moving Cameras*, 2011.

Subtask 2: Compressed Sensing LADAR (GDRS)

Objective and Benefits:

The goal of this subtask is to apply compressive sensing (CS) techniques to small, lightweight, and affordable LADAR imagers. The concept will use a spatial light modulator, single detector avalanche photodiode (APD), and the techniques of compressive sensing. We will adapt and apply established sparse sample reconstruction techniques with a signal/image processor, an intrinsic component of the system. This bears the potential to combine the temporal properties of analog APDs with the spatial selectivity of an array, while achieving a new degree of efficiency, affordability, and dynamic adaptivity.

The scope of this subtask is to study the trade-offs of this approach, especially in quantitatively evaluating the SNR characteristics in the context of active sensing. This has not been addressed in published reports. A prototype will be constructed to demonstrate a proof-of-concept.

This subtask addresses fundamental perception and overall system feasibility in terms of power consumption, size, and weight. Throughout the entire Capstone Assessment Vision, detailed perception information is required to realize semantic understanding and interactivity. This subtask addresses the practicality of systems with respect to real physical limitations (i.e., size, weight, and power).

State of the Art:

Contemporary LADAR sensors use beam scanning or frame and range at once (flash) techniques to cover a desired field-of-view. Scanned LADAR is often based on mechanical beam scanning, the speed of which is a limiting factor on frame rate. Mechanical scanning also requires power consuming motors. Flash LADAR, on the other hand, does not necessarily require any form of mechanical scanning and uses a 2D array of high bandwidth, high sensitivity photodetectors and read-out electronics. InGaAs APD arrays are used in most flash LADAR systems, although these are highly specialized to the niche application and are consequently expensive. APD breakdown voltage uniformity, complex high voltage supply and signal routing issues, and manufacturing yield are imposing barriers to improving quality and reducing cost for small, inexpensive unmanned vehicles. Reliability problems are frequently observed with APD arrays.

Key Barriers and Approach:
To move beyond the current state of the art, an approach based on spatial light modulation (SLM) and single detectors will combine the advantages of scanned and flash type systems. This approach leverages consumer electronics market investment in spatial light modulation as well as telecom investments in single element APDs and receiver electronics. Recent developments in CS techniques will be applied to minimize the number of laser pulses required to sample a full frame image. Literature reports, such as the Rice single pixel camera, suggest that high cross range resolution may be obtained with ¼ of the pulses required for raster or pixel-basis scanning. The cross range resolution of such a system may be dynamically altered to meet current mission profiles and activities. Since it is less difficult to obtain high repetition and high average power from diode lasers than high pulse energy, this technique will allow multiple pulses to be averaged resulting in higher signal to noise ratios.

Metrics for Evaluation:
The progress and value of the work will be measured by comparing the performance of the prototype system and analytical expectations to baseline units, including microLADAR and flash LADARs. Metrics will include maximum detectable range of a target, range accuracy and resolution, laser power per pixel, and effective cross range resolution. We will develop a metric to evaluate the impact of compression artifacts.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Evaluate techniques for spatial light modulation. Simulate sensor link budget and spatial modulation in MATLAB (TRL 1 – TRL 2).
Q2	Select and obtain SLM module. Design prototype system.
Q3	Construct and measure performance of laboratory prototype.
Q4	Collect experimental scene data (TRL 3 – TRL 4). Evaluate performance relative to flash and scanning LADAR.

Subtask 3: Improved Contact Sensors for Terrain Classification and Grasping (FSU)
This research develops robot skins that can be applied to the feet or end effectors of a robot. The feet skin is to be used for terrain identification, while the end effector skin is to be used for object recognition to aid in grasping.

Objective and Benefits:
The goal of our research is to investigate methods of endowing robotic extremities (feet and end effectors) with skin that has a sense of touch that mimics the large number of pressure measurements found in the skin of biological systems, including humans. This skin will allow robot feet to sense changes in the terrain which will trigger changes in the robot gait, and it will allow robot end effectors to identify terrain surfaces and perform more general object recognition via touch, thus aiding in grasping. This refined sense of touch complements vision. It allows robot feet to feel terrain that may have superficial coverings and identify terrain when vision is focused elsewhere. Similarly, it allows a robot to feel things that are hidden to vision (e.g., in a

box) and allows robust grasping when vision is being used for other tasks (e.g., think of a detective holding a gun on a suspected hostage while reaching for evidence on a nearby table).

This subtask is useful in capability element 5 of the RCTA Capstone Integration and Assessment Vision which requires a robot to navigate mobility challenges due to natural terrain and urban clutter. For example, if the robot is a legged robot and encounters loose rocks that are covered with a thin layer of dry grass, the robot will be able to feel the underlying terrain and change its gait to be appropriate for loose rocks. This subtask is also useful in capability element 10 of the RCTA Capstone Integration and Assessment Vision which requires robust grasping to move obstacles and grip doorknobs of various types. This research will allow a robot end effector to obtain better grasps of objects before lifting them.

This research is a multi-year effort that is expected to extend to the 5-year duration of the RCTA.

State of the Art:
This research studies contact sensors on a robot foot or end effector to sense terrain surfaces; to our knowledge, no one else is directly addressing this problem. However, there is related research, primarily in Japanese and European universities, on developing robot skins. To develop skins that can actually sense terrain surfaces requires that the skin has many, literally hundreds, of closely spaced pressure sensors. The largest number of sensors we have seen on one skin (roughly the size of the end segment of a human finger) is 9 (a 3 by 3 matrix). This is not adequate for sensing terrain surfaces. Hence, the key is to develop a skin that has hundreds or even thousands of sensors on an area the size of a hand or foot. This is what is needed for terrain surface classification.

The objectives of this task are: 1) to use smart materials systems to incorporate microarrays of pressure sensors into the “feet” of legged robots and in robot grippers as well as micro-electromechanical systems to design embedded temperature and moisture sensors for terrain classification; 2) to use these pressure measurements to directly classify terrain for a variety of legged vehicles and to classify objects for manipulators; and 3) to integrate these measurements with vision sensors for robust classification involving both vision and proprioceptive sensing.

Key Barriers and Approach:
We mimic the pressure nerves in biological skins with high-density piezoresistive-based pressure sensor arrays and develop mechanisms to classify terrains based on the robot gait. To our knowledge, we are the only group incorporating high-density piezoresistive sensor arrays to design robot skins and the only ones actively using robot skins on the foot of a robot for terrain classification. Referring to robot manipulators, we also expect to take a lead in object identification by touch since we have a robot skin with the necessary pressure measurement density.

The results obtained so far and reported in part below give strong evidence that this approach will succeed. The main barrier in the design of our Pressure Sensitive Robot Skin (PreSRS, pronounced “pressures”) is to adapt the device for legged robots without damaging the many wires of the device or the device’s circuitry. To date, we have overcome these obstacles by only

incorporating the PreSRS system on legs that "swing" such that the wire will not twist. However, to incorporate PreSRS on a legged robot like RHex which has small, circulating legs, a method to both support the device on the leg and keep the many wires (over a hundred) from twisting needs to be devised. We plan to incorporate slip rings to combat the entanglement of wires that will result in a circulating leg.

Metrics for Evaluation:
For robot feet, we measure our progress by the accuracy of the terrain classification as increasing realism is added to the experiments. We began with static tests (a foot drops and remains stationary), then conducted experiments with the robot skin attached to a human foot, and will next conduct experiments with the robot skin attached to the robot foot associated with a "one legged hopper." Classification accuracy will also be used when using the robot skin with an end effector to identify surfaces. However, when used for object recognition via feel, metrics will be developed that are more related to grasping.

With PreSRS, terrain classification was achieved with accuracies of over 95% under both static conditions using a compression rig and dynamic conditions in which PreSRS was placed on the bottom of a human foot. With the success of the static and dynamic terrain classification results, we plan to apply PreSRS to the FAMU-FSU Hopper to investigate the possibilities of terrain classification on a legged robot. A new multi-degree-of-freedom foot needs designing to allow for the attachment of PreSRS. With this new Hopper foot and PreSRS, terrains such as carpet, sand, grass, and rocks will be tested and classified on the legged robot. Our final goal with the terrain classification on the FAMU-FSU Hopper is to incorporate real-time data acquisition software add-ons to the underlying pressure array, allowing real-time terrain classification on the robot. With real-time terrain classification abilities, we will be able to have real-time demonstration of the classification of multiple terrains.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Conduct terrain classification experiments using PreSRS (Pressure Sensitive Robot Skin) with the FAMU-FSU Hopper on six different terrains: concrete, carpet, grass, gravel, loose sand, and hard-packed sand.
Q2	Perform online terrain classification with the FAMU-FSU Hopper using PreSRS and software allowing for real-time data acquisition from the underlying pressure array.
Q3	Provide a demonstration of real-time terrain classification on the FAMU-FSU Hopper. Start initial experimentation of the end effector skin consisting of many pressure arrays located on the fingers and palms. Ways of adapting the system to fit on a manipulator end effector will be investigated.
Q4	Conduct object recognition experiments using the PReSRS system adapted to a manipulator end effector. These experiments will be primitive, using no manipulator path planning; instead the manipulator will be remotely controlled.

P2-2012 – Compact, High Performance Sensors

Level: 6.2 Applied Research

	Name	*Organization*	*Subtask*
Principal Investigator	Larry Matthies	JPL	P2-4
Additional Investigators	Roland Brockers	JPL	P2-4
	Mike Powers	GDRS	P2-2
	Tod Oblak	GDRS	P2-1
	Christoph Mertz	CMU	P2-3
	Martial Hebert	CMU	P2-3
	Barry Stann	ARL	P2-5
	Mark Giza	ARL	P2-5
	William Lawler	ARL	P2-5

Objective:
In this task, we explore approaches to increase sensor performance, reduce SWaP-C, and evaluate sensors on multiple classes of unmanned systems. The overall objectives of the task are in two general areas:

Advanced Spectral Sensors. Develop sensing solutions based on spectral classification, including prototyping of an active spectral LADAR and corresponding classification algorithms.

Scaling Perception to Small Platforms. Develop high performance range sensors with small form factors suitable for use on small-size platforms. Our overall objective is to provide solutions for small mobility platforms as well as for dexterous manipulation.

Background:
Current sensing approaches deal poorly with many conditions including dust, smoke, fog, rain, and darkness. They lack the robust, reliable, real-time performance needed to support complex military and commercial applications. Our basic research task (P1) focuses on understanding the basic phenomenology of sensing. This applied research task is designed to mature promising approaches; to address limitations in resolution, range, and robustness as well as the size, weight, and power considerations that limit implementation on smaller ground and air vehicles; and to develop sensing solutions better adapted to new areas of dexterous manipulation and unique mobility.

Research Description:

LADAR Core Electronics. We (GDRS, Subtask 1) are currently working on highly software-definable LADAR core electronics. This module is designed to suit the purposes of a generic active sensing system and be adaptable to existing and future concepts. The functional components include a GHz bandwidth analog front end feeding a 3+ GSa/s high resolution digitizer, high-speed synchronization lines, FPGA for signal processing, and FPGA for system control and communications. The flexibility of this product will greatly reduce the cost of proof-of-concept demonstration and early prototyping of future active sensing systems. Proposed work in FY12 will finish and deliver the firmware, layout, and complete design data package of the receiver/system controller board. The hardware and software results of this work will be directly and immediately applicable to the ARL MEMS LADAR effort, which may be integrated on a SUGV and participate in the capstone assessment. Additionally, this will support the CS LADAR task (new for 2012) and Spectral LADAR task.

Advanced Spectral Sensors. This task (GDRS, Subtask 2) will extend the development of Spectral LADAR. This technique allows object recognition using spectral signatures in addition to 3D spatial information. Active sensing enables dense point clouds to be obtained in cluttered, low-textured, or obscured environments and at long ranges during the day or night. The range-resolved and multiple-return capability produces unambiguous, intrinsically unmixed spectral signatures of range separated objects. In contrast, state-of-the-art sensors measure separately either shape (through 3D measurements) or spectral properties (through multi-spectral imaging). Proposed work in FY12 will replace the current Spectral LADAR single-element APD receiver with a 32 x 32 pixel APD array, along with other modifications that will increase signal throughput. The layout of the system will be re-factored to enable portability for data collection. The hardware changes will reduce the image acquisition time from about 45 minutes to several seconds. Extensive data collection in real environments is the key step towards analyzing the capabilities of Spectral LADAR, and the present imaging speed is insufficient for this purpose.

Short-range Sensing for Dexterous Manipulation and Unique Mobility. As part of our effort on scaling perception to small platforms, we (CMU, Subtask 3) have developed a compact, high-accuracy short-range sensor, based on light stripe triangulation, suitable for manipulation tasks. Proposed work in FY12 will take two directions. First, we will integrate a 3D sensor on robot platforms to support specific capabilities that will be relevant for the capstone assessment, including inspecting a rubble pile and selecting grasping points on objects of interest, finding openings through which the robot can enter a difficult-to-access location and map areas of interest (e.g., the inside of a pipe), and finding doorknobs, handles, etc. that a robot needs to manipulate to open doors. Second, we will study material classification with this sensor by collecting the full light transport matrix and using it to reconstruct the 3D scene with all its optical properties.

High Performance Stereo for Small Platforms. Another key thrust of our effort to scale perception to small platforms (JPL, Subtask 4) involves developing high performance stereo vision capabilities with small SWaP-C. In FY10 through FY11, we addressed this by developing an implementation of stereo vision and visual odometry on the OMAP3530 system-on-a-chip processor for smart phones which, to our knowledge, achieves the highest energy efficiency of

any stereo vision system published to date. Work started in FY11 and proposed to continue in FY12 is designing a texture projection system to augment passive stereo vision to enable 3D perception of featureless surfaces outdoors, in direct sunlight, at short ranges suitable for manipulation. We anticipate that in future years this task will evolve to address how to affordably achieve very high angular resolution with stereo vision systems on small platforms.

Metrics:

- Capabilities and features of the LADAR core electronics design against the requirements of the ARL MEMS LADAR, CS LADAR, and Spectral LADAR. Key evaluation parameters will include the input referred noise, effective number of bits vertical resolution, and board size and power consumption. Accuracy of range measurement from small sensor on typical materials and observation conditions for DMUM.
- Increase in imaging speed, maximum range, and progress toward portability of spectral LADAR.
- For light stripe triangulation: degree of integration; accuracy and number of missed and spurious points in 3D reconstructions; ROC and precision-recall curves in material classification based on estimated optical properties.
- For outdoor projected texture stereo: range data density and accuracy versus range to object; SWaP-C; degree of eye safety.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Subtasks 1 and 2 are executed in collaboration with LADAR researchers at ARL. Subtasks 3 and 4 involve collaborations between CMU and QNA and between JPL and QNA, respectively. Larry Matthies is lecturing as part of a class on “Vision-based Navigation” at Caltech in 2012.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Complete minor changes to LADAR core electronics board schematic to handle higher speed level converters, among other modifications. Develop data/electrical interfaces to APD array ROIC for spectral LADAR. Light stripe sensor used in tele-operation on TALON; data collection on objects (e.g., various doorknobs). Complete design of projected texture stereo system for outdoor, short-range operation.
Q2	Develop and test-bench VHDL design, develop embedded code for LADAR core electronics. Perform apparatus calibration and performance assessment for spectral LADAR. Light stripe sensor mounted on snake robot and PR2; fast detection of suspicious objects within one category (e.g., doorknobs). Complete fabrication/procurement of components for projected texture stereo system.
Q3	Complete full layout and finalized bill of materials for LADAR core electronics. Collect image data with improved prototype indoors, outdoors, and at remote sites for spectral LADAR. Complete iterative probing algorithm for light striper to improve classification. Assemble and characterize full prototype of outdoor projected texture stereo system.
Q4	Deliverable design data package to ARL SEDD for LADAR core electronics. Develop refined classification algorithms based on large datasets for spectral LADAR. Light stripe sensor used in tele-operation on snake robot and PR2; fast detection and iterative probing combined with real-time application within one object category. Complete preliminary design of stereo system for 360°, 3D situational awareness.

Related Research:
The work on sensing for Dexterous Manipulation and Unique Mobility will eventually contribute to all of the DMUM tasks that require sensing.

Work on small sensors, including LADAR, light striping, and stereo contributes to P3 through P6 in providing sensing solutions that enable migration of these algorithms to small platforms. High resolution stereo contributes to the detection of high resolution behavior cues discussed in P5 and P6 as well as non-verbal communication, gesture recognition, and social cues discussed in HRI tasks and improved high-speed navigation for robot convoy and wingman scenarios.

The sensors developed in P2 are directly used for generating the data in P3 and P4, used to recognize objects and features and build terrain models which, in turn, are used in reasoning about the world model and in executing complex behaviors in Intelligence tasks.

Researchers involved in the RCTA sensor Tasks P1 and P2 will continue collaboration with relevant sensor research at ARL.

Integration and Assessment Activity:
Advanced LADAR capabilities are anticipated to funnel into integration and assessment activities leading up to the capstone assessment.

The light stripe range sensor has already been integrated onto a TALON and will be integrated onto a snake robot and a PR2 in FY12. Integrated Research Assessment 2 (IRA-2) will address manipulation and will incorporate contributions from light stripe and stereo-based range sensing systems. IRA-1 already used stereo vision as part of a basic human activity understanding assessment; future extensions of that assessment will employ improved stereo capabilities from this task.

Subtask 1: LADAR Core Electronics (Software-definable LADAR Architecture) (GDRS)

Objective and Benefits:
This subtask will complete the LADAR core electronics task started in the 2011 APP by finishing and delivering the firmware, layout, and complete design data package of the receiver/system controller board.

This task supports LADAR sensor development and implementation, a fundamental and essential perception modality in the context of the capstone assessment. The hardware and software results of this work will be directly and immediately applicable to the ARL MEMS LADAR effort, which may be integrated on a SUGV and participate in the capstone assessment. Additionally, this will support the CS LADAR task (new for 2012) and Spectral LADAR task.

It is expected that the product of this effort will be delivered to the Photonics Integration Branch in the SEDD at ARL Adelphi. In the future, this board may support RADAR and development of other new sensor concepts since it is based on an extensible architecture.

State of the Art:
Typical LADAR receiver and system control electronics utilize counters and/or analog delay elements to compute time-of-flight. Common receiver architectures which are hardware implemented (e.g., with ACAM time-to-digital converters) are difficult, expensive, and time consuming to modify for specialized purposes or to test new concepts. This new software-based architecture is general purpose, flexible, and friendly to experimentation and rapid proof-of-concept demonstrations. It uses the current generation of digitizer and FPGA hardware for maximum capability.

Key Barriers and Approach:
This board is based on the latest ADC technology and is capable of 5 GSa/s and a maximum of 4 channels. A key feature is its high vertical resolution of 10 bits. This subtask, and the “software defined” architecture on which it is based, serves to break down the barriers of next-generation sensor development. This board will jumpstart state-of-the-art sensor efforts now and in the future.

Metrics for Evaluation:
We will assess the progress of this subtask by evaluating the capabilities and features of the design against the requirements of the ARL MEMS LADAR, CS LADAR, and Spectral LADAR. Key evaluation parameters will include the input referred noise, effective number of bits vertical resolution, and board size and power consumption.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Complete minor changes to board schematic to handle higher speed level converters, among other modifications. With ARL collaborators, define an extensible software architecture. Specify a default software configuration.
Q2	Develop and test-bench VHDL design, develop embedded code.
Q3	Complete full layout and finalized bill of materials.
Q4	Test hardware fundamentals, potentially on surrogate equipment. Deliverable design data package to ARL SEDD.

Subtask 2: High-speed Spectral LADAR (GDRS)

Objective and Benefits:
This subtask will replace the current Spectral LADAR single-element APD receiver with a 32 x 32 pixel APD array, along with other modifications that will increase signal throughput. The layout of the system will be re-factored to enable portability for data collection. The hardware changes will reduce the image acquisition time from about 45 minutes to several seconds. Extensive data collection in real environments is the key step towards analyzing the capabilities of Spectral LADAR, and the present imaging speed is insufficient for this purpose.

Successful execution of missions along the lines of the capstone assessment requires high confidence semantic understanding supported by precise and unambiguous perception information. The multi-modal information supplied by Spectral LADAR addresses this directly. Moreover, results to date have shown its success against concealed hazards and targets. It is reasonable to expect that real adversaries in situations like the capstone assessment will take care to conceal their activities with clutter and camouflage.

State of the Art:
The present Spectral LADAR prototype represents state-of-the-art advances in LADAR and active multi-spectral sensing through the physical unification of these modalities, each of which has separately proven its value in robotic experiments. Spectral LADAR offers essential improvements over the data level fusion of 3D and multi-spectral techniques that are especially relevant in military situations and for the detection of concealed targets in diverse environments. Unlike swept-source active spectral imagers that require the scene to remain static over the sweep time, Spectral LADAR uses single-pulse spectroscopy where an optical pulse is constituted by a nearly flat spectral density in the shortwave infrared. The discriminating spectral

features are read-out in parallel, affording high frame rates and allowing for substantial scene motion or egomotion.

Key Barriers and Approach:
Building on this foundation, this subtask enhances the performance of the prototype to longer ranges and higher imaging speeds through the improvement in optical power throughput and a new receiver array. This dramatic improvement enables the collection of significantly more data than is possible with the slow speed of the current prototype. Enhancements will also allow the system to be portable to data collection sites of interest.

The speed (pixel rate) of the present system is limited largely due to the mechanical tuning of the diffraction grating. This will no longer be necessary with an array-based receiver. Also, several factors such as the multi-mode fiber and aberrations in the transmitter lens adversely affect optical throughput and maximum range. Engineering re-work will address these issues. Following these improvements, a large amount of data will be collected and analyzed. The scenes will correspond closely to the situations illustrated in the capstone assessment.

Metrics for Evaluation:
Successful demonstrations of the Spectral LADAR concept were conducted in previous years on the RCTA program. In 2012, a portable version of the system will be tested in environments with the intent to represent the capstone assessment. Such environments will include imaging outdoor scenes of foliage, structures, vehicles, and other objects around the GDRS' campus. Comparison of Spectral LADAR imagery to imagery collected with other sensor equipment planned for the capstone will be the basis of comparison. Progress on the 2012 APP will be measured by increase in imaging speed, maximum range, and progress toward portability.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Develop data/electrical interfaces to APD array ROIC. Modify mechanical mounting and optical elements to adapt to a new photoreceiver.
Q2	Perform apparatus calibration and performance assessment.
Q3	Collect image data with improved prototype indoors. Collect image data with improved prototype outdoors and at remote sites.
Q4	Develop refined classification algorithms based on large datasets. Conduct image capture with scenes similar to the capstone experiment. Produce final report with experimental findings and benefit study.

Subtask 3: Short-range Sensing for Dexterous Manipulation and Unique Mobility (CMU)
The objective of this subtask is to develop a compact, high-accuracy short-range sensor suitable for manipulation tasks. The general sensing modality will be projecting light patterns and observing them with cameras.

Objective and Benefits:
The objective is to develop a high resolution 3D sensor suite with small size, weight, power consumption, and CPU requirements, capable of operating under all illumination conditions (including bright sunlight) that can detect and classify ordinary and difficult objects like those made out of glass, metal, or translucent materials. The sensor will provide perception for DMUM.

The plan is to integrate a 3D sensor on the following three platforms to support specific capabilities that will be relevant for the capstone assessment. For the TALON (Subtask IR2-3) it is inspecting a rubble pile and selecting grasping points on objects of interest. A sensor mounted on the snake robot (Subtask M5-1) will help in finding openings through which the robot can enter a difficult-to-access location and map areas of interest, e.g., the inside of a pipe. The supported capability of the PR2 robot (Subtask M1-3) is opening doors. The purpose of our sensor is to find the doorknobs, handles, etc. that the robot needs to manipulate. For 2012, the goal is to integrate the sensors so that one can collect data for display, remote control, or offline analysis. For the TALON, we will take an additional step towards autonomous operations by communicating the sensor results to the TALON computer. We will use the world model (Task I2) to guide us in the communication approach. The effort will continue until the Capstone Assessment, where the goal is to do these and other manipulations autonomously.

State of the Art:
The current state of the art is the ability to do 3D reconstruction and classification of challenging objects such as those made of glass and metal and ceramic, using structured light in the lab under controlled conditions. Often the objects have to be manipulated like immersing them in a liquid or placing them in front of cooperative backgrounds.

Key Barriers and Approach:
Our approach is to tackle the most general case by collecting the full light transport matrix and using it to reconstruct the 3D scene with all its optical properties. One gets the light transport matrix by projecting one beam at a time (i.e., the projected image is a single white pixel) and collecting the return image for each beam. This reconstruction method is not suitable for real-time operations, but it will provide the insight on model-driven approaches and active probing. A scene is first scanned to see if it consists of cooperative objects. Any suspicious returns are further probed with specific projects that allows for comparison with models. In an iterative process, the understanding of the scene is increased until we have satisfactory reconstruction. Specific objects on which we will concentrate are derived from the capabilities mentioned above. For example, we will collect models of doors, doorknobs, handles, etc. and find suitable projection patterns to detect the grasping points for door opening in real-time
.
The algorithms that need to be developed are on the one hand the methods that can quickly determine if the objects are cooperative and on the other hand the methods that can iteratively probe the suspicious objects. In the process of developing these algorithms, we will collect and label datasets of objects for training and testing.

When we integrate our sensor on one of the robots mentioned above, we will encounter challenges specific to that platform, e.g., on the snake robot one has very little space, its movements are very different from a wheeled robot, and the location of the robot and the orientation of its joints are not accurately known. We will investigate how we need to change our setup and algorithms to accommodate these challenges.

Metrics for Evaluation:

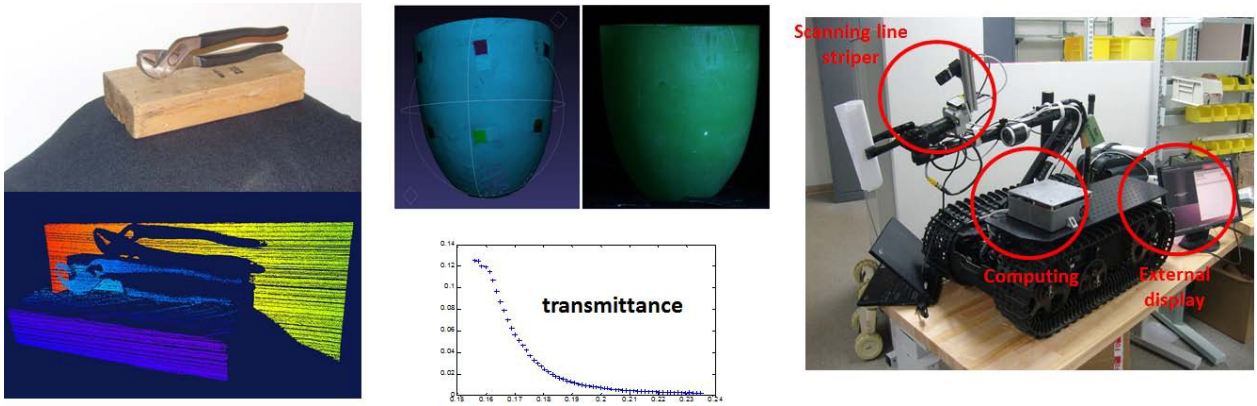

Figure 3-8 Left: 3D point cloud (bottom) of a medium challenging object (top). Center: Point cloud (upper left) of wax bowl (upper right) and the measured transmittance of the material (bottom). Right: Sensor mounted on TALON.

The progress of the sensor integration on a robot can be divided into four stages:

1. Mounting sensor on robot (external power supply and analysis)
2. Increased autonomy:
    a. Power supply and analysis on the robot
    b. Display of sensor results used in tele-operation
3. Sensor used in autonomy

On the TALON, we have accomplished 1 and 2a and have started with 2b. As can be seen in Figure 3-8 right, the sensor and its computing are mounted on the TALON. The resulting data from can be displayed on the TALON control unit.

For 3D reconstruction, the metrics are distance accuracy and missed and spurious points. For classification, the metrics are ROC and precision-recall curves. We have qualitative results from our current work. We are able to do 3D reconstruction of medium challenging objects (metallic or dark, Figure 3-8 left), and we are able to capture the transmittance and reflectance properties of translucent materials (Figure 3-8 center).

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Sensor used in tele-operation on TALON. Collection of object models (e.g., various doorknobs).
Q2	Sensor mounted on snake robot and PR2. Fast detection of suspicious objects within one category (e.g., doorknobs) finished.
Q3	Sensor improvements for snake robot and PR2. Iterative probing algorithm within one category finished.
Q4	Sensor used in tele-operation on snake robot and PR2. World model inspired communication with TALON. Fast detection and iterative probing combined for real-time application within one category.

Subtask 4: High Performance Stereo for Small Platforms (JPL)

Objective and Benefits:
This subtask conducts research on scaling high-performance stereo vision to small platforms. Over the base period of the program, our goals are: 1) to enable stereo vision and visual odometry to run on very compact, low power, embedded processors; 2) to incorporate texture pattern projection to enable stereo-based ranging to featureless surfaces within manipulation range outdoors in direct sunlight; 3) to provide 360°, 3D situational awareness; and 4) to enable 3D perception with much higher angular and range resolution through computationally efficient use of very high resolution cameras and/or zoom optics. We took a major step toward the first objective in FY10 through FY11 by implementing stereo vision and visual odometry on the OMAP3530 system-on-a-chip (SoC) processor for the smart phone market. In FY11, we also initiated a design study for projecting texture onto featureless surfaces outdoors in direct sunlight. Our intent is to enable 3D perception of featureless surfaces in direct sunlight at short ranges for outdoor manipulation, at longer ranges for indoor navigation where featureless surfaces abound, and in the dark indoors, outdoors, and underground. In FY12, we propose to complete a prototype of such an illuminator and to build on our OMAP stereo success by developing a preliminary design of a perception system using four stereo camera pairs looking in the four cardinal directions around a vehicle to provide full 360°, 3D situational awareness. This effort will be done in collaboration with QNA's effort in subtask IR2. We intend this subtask to continue through FY14 by extending the FY10 through FY12 work to build the 360°, 3D situational awareness system and to address computationally efficient methods to obtain 3D perception with very high angular resolution.

The relevance of this subtask to the Capstone Assessment Vision lies in: 1) providing the ability to do 3D perception of featureless surfaces indoors and (at short range) outdoors, which will be important for both navigation and manipulation; 2) providing full 360°, 3D situational awareness; and 3) providing 3D perception with very high angular and range resolution for situational awareness at relatively long range. The texture projector may also become useful for 3D perception and material classification for surfaces with both specular and diffuse reflection. These capabilities will be valuable for opening doors, where both doors and doorknobs may have few features and doorknobs in particular may be specular; this is relevant to doors of buildings and to car doors. Work in FY13 and FY14 on cost-effective integration of very high angular

resolution will support the following functions at much greater standoff distances than otherwise possible: surveillance of doors, windows, and other objects, perception of human gestures and expressions, and obstacle detection for high-speed driving.

State of the Art:
To our knowledge, our OMAP3530-based stereo vision and visual odometry implementation provides the most energy-efficient stereo vision implementation to date, achieving a metric of 58.8 million disparity evaluations per joule (58.8 Mde/J), while simultaneously providing visual odometry [1]. The roadmap for processors in this family includes signification increases in capability within 12 to 18 months. We have concluded that the very high level of integration, small size, and low power consumption make this class of computing system the best strategy to focus on for small robot platforms for the next few years; in particular, we do not anticipate pursuing GPU implementations unless new product data comes out that implies a significant SWaP-C advantage for GPUs in small systems. The OMAP stereo implementation also offers a new, low power approach to 360°, 3D situational awareness, wherein we could bundle a processor with each stereo pair and put four such bundles on a robot to cover 360° with low SWaP-C. Such a perception system would be very valuable in other tasks and in the Capstone Assessment Vision; however, due to anticipated budget constraints in the program, we propose to do only preliminary design of such a system in FY12, to clarify architecture issues, and defer fabrication to FY13. A separately proposed subtask in FY12 (IR-3: Low-power, Low-cost, Real-time Stereo Camera System), if funded, may make progress toward building improved versions of individual OMAP-based stereo systems in the meanwhile.

Subtask P1-1 is addressing algorithms to improve stereo-based 3D perception of low-texture surfaces; however, many surfaces are completely featureless, hence not favorable to algorithmic approaches. Several types of “RGBD depth cameras” have emerged recently that use patterned light projection to enable 3D perception, including a projected texture stereo system developed at Willow Garage [2] and the Microsoft Kinect. Neither of these sensors project enough optical power to work in direct sunlight. Other active sensors for 3D perception that do work in direct sunlight, e.g., LADAR and structured light, have other limitations, such as lower angular or range resolution, slower acquisition, or lack of overlaid color channels. To address funding constraints, we propose that our primary effort in FY12 will address extending projected texture stereo to work at short range in direct sunlight, primarily to support outdoor manipulation tasks. Ultimately, such a system may also address 3D perception for featureless surfaces indoors and enable 3D perception at moderate ranges in the dark. This is related to work JPL has done under the DARPA LS3 program on integrating illuminators into a stereo vision system for night operation; while LS3 does not use texture projection, lessons learned about working with illuminators are still useful here.

Key Barriers and Approach:
The key barrier to projected texture stereo in direct sunlight is to achieve a sufficiently high ratio of projected texture optical power to direct sunlight optical power to make the pattern visible enough in imagery to enable stereo correlation to work, while maintaining eye safety. This leads to questions of what spectral band to use, what degree of spectral filtering is necessary, what illuminator and transmitter/receiver optics to use, and what power level to project. In work to

date under FY11 funding, we have addressed these questions for the visible and near infrared part of the spectrum and reached the conclusion that this is feasible with narrow-band filtering at the camera, for ranges on the order of 1 m. With the remaining quarter of FY11 funding, we expect to extend this study to the SWIR region, where solar illumination levels are much lower and eye safety is much easier, as well as address remaining issues in manufacture of the texture pattern and availability of camera optics that are image-space telecentric, which is needed for narrow-band optical filtering over a wide field-of-view. FY12 will be devoted to fabricating a prototype of the entire sensor and characterizing its performance. Subsequent years will integrate a prototype into a robotic system and address the rest of our agenda for high performance stereo for small platforms

Metrics for Evaluation:
The primary metric for projected texture stereo is the density of range data (i.e., percent of image pixels with valid range estimates) obtained with and without texture projection for sample scenes. Accuracy of the range data is a secondary metric, as are SWaP-C and eye safety. We will characterize data density and accuracy by comparison to ground truth obtained from other sensing modalities; characterization of SWaP-C and eye safety will come naturally from the design and testing of the system.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Complete design of projected texture stereo system for outdoor, short-range operation.
Q2	Complete fabrication/procurement of components.
Q3	Assemble and characterize full prototype sensor.
Q4	Complete preliminary design of stereo system for 360°, 3D situational awareness.

References:
[1] S.B. Goldberg and L. Matthies, Stereo and IMU assisted visual odometry on an OMAP3530 for small robots. *Workshop on Embedded Computer Vision*, held in conjunction with CVPR, 2011.
[2] K. Konolige, Projected texture stereo. *IEEE International Conference on Robotics and Automation*, 2010.

Subtask 5: MEMS-Scanned LADAR for Ground Robot Integration and Data Collection (ARL)

Objective and Benefits:
The proposed subtask objective is to integrate ARL's MEMS-scanned LADAR onto an RCTA-owned ground robot. This work will support data collection and tests of RCTA perception algorithms while providing feedback to ARL on changes to the MEMS-scanned LADAR needed to improve perception performance. Early in this effort, we will supply GDRS with image files collected in the lab with our LADAR breadboard to allow them to perform an initial evaluation of fusing passive and LADAR data for perception.

State of the Art:
State of the art in LADAR imagers for small ground robots relies on products built by MicroSoft (Kinect), Velodyne, and Advanced Scientific Concepts (ASC). These products are flawed in some aspect that makes them likely to be unacceptable for military use; for example, the MicroSoft LADAR will not operate in bright sunlight, the Velodyne LADAR is mechanically scanned and has low vertical resolution, and the ASC LADAR is too expensive.

Key Barriers and Approach:
ARL is researching and developing MEMS-scanned LADAR technologies and architectures that are low-cost, low-power, and compact. Such LADARs are suitable for mounting on small ground robots and can provide real-time and high resolution 3D imagery that is needed for autonomous navigation, driving, mapping, and target detection in highly cluttered environments. Tests with breadboards of the MEMS-scanned LADAR on an Explorer PackBot yielded quality 3D imagery while in motion both indoors and outdoors.

The photo in Figure 3-9 shows the present state of ARL's MEMS-scanned LADAR. Here, the LADAR subassemblies (including the compact laser, the mirror driver board, the new transmitter, the new receiver, the A/D FPGA board, and the Ethernet board) are mounted on a 1-foot by 1-foot optics board. Tests with this integrated system indicate that the imagery is the best quality to date.

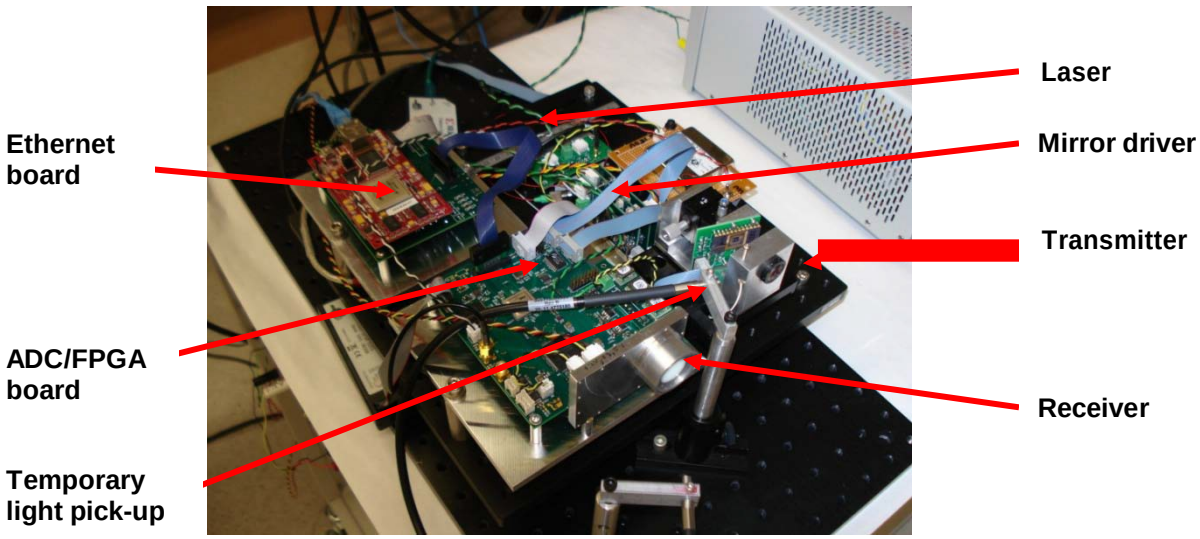

Figure 3-9: MEMS-scanned LADAR breadboard.

At this time, there are no insurmountable problems to mounting the LADAR on a small ground robot for the proposed perception test activities. To make this task easier, we will rebuild the mirror driver board to reduce its size and to allow setting of the scan parameters remotely through the control computer. We will also rebuild the power supply board to reflect reduced power supply requirements and to reduce its size. These activities are underway. To significantly improve the estimate of range, we will provide a means to capture the outgoing laser pulse and use it as a reference in determining target range. Once these changes are proven in the LADAR,

we will build a second set of circuit boards/subassemblies (much of this is done) and package them in a box suitable for mounting on an RCTA-owned ground robot.

Metrics for Evaluation:
We will assess our task progress against the below schedule.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Subtask P2-5 is part of a proposed collaboration between GDRS, Carnegie Mellon, and ARL to integrate and utilize advanced LADAR technology on-board a small robot platform, as part of Task P3: Static Scene Understanding.

	Subtask 5 Quarterly Milestones and Deliverables
Q1	Build mirror driver and power supply boards. Implement means to improve range estimate. Provide GDRS with data from the existing LADAR breadboard.
Q2	Integrate the mirror driver and power supply boards into the LADAR breadboard, and test. Build a second set of boards/subassemblies, and test as a unit.
Q3	Package the second set of LADAR boards/subassemblies in a small box, and integrate on a ground robot.
Q4	Perform LADAR tests, collect data, and start integration with perception algorithms.

P3-2012 – Static Scene Understanding

Level: 6.1 Basic Research

	Name	*Organization*	*Subtask*
Principal Investigator	Martial Hebert	CMU	P3-3
Additional Investigators	Drew Bagnell	CMU	P3-3
	Kostas Daniilidis	UPenn	P3-5
	Jianbo Shi	UPenn	P3-1, P3-2
	Camillo Taylor	UPenn	P3-5
	Pietro Perona	Caltech	P3-4
	Michael Maire	Caltech	P3-4
	Jason Owens	ARL	P3-6
	MaryAnne Fields	ARL	P3-6

Objective:
The overall objective of this task is to design algorithms that identify objects, features, and classes in sensor data, either images or 3D data. Our approach to static scene understanding includes building class models, incorporating context, and utilizing large training datasets. The subtasks address different aspects of the scene understanding by using techniques most appropriate to each aspect. This approach breaks down the problem into manageable pieces, each specialized to a different aspect directly relevant to the CTA scenarios.

Background:
In order to conduct complex reasoning tasks beyond safe driving, robots must have a detailed understanding of the world, including a description of objects, material, and "stuff" in their environment. In addition to "naming" the entities in the scene, perception should also derive qualifiers (e.g., parked car, occluded wall) and relations between scene parts (e.g., car in front of door) from sensor data. This level of scene understanding remains a challenging problem that has focused the attention of the computer vision community in the past decade. In particular, much progress has been made in object recognition, e.g., using trained classifiers from local image features, but these techniques are still brittle. Progress has been made in representing and inferring relations between scene elements and deriving from them scene parsing, i.e., labeling of all the visual input by known classes. These use statistical techniques based on probabilistic graphical models to represent uncertainty. This class of approaches relies on well-established tools based on probabilistic graphical models. While promising, they involve complex optimization problems which are not feasible for online classification. Even the most mature aspect of recognition which is the problem of recognizing and locating specific 3D objects is still limited when applied to data most relevant to real-world CTA-relevant problems, such as objects

with difficult materials like semi-transparent. This task addresses the limitations of the SoA through a complement of subtasks, each targeted to a separate aspect of the problem, complemented by one subtask on semantic mapping.

Research Description:
To achieve this task's objective, we have divided it into a set of subtasks corresponding to different types of semantic information. More precisely, prior work in computer vision has shown that different types of semantic information require very different techniques. The types of information covered in this task, the corresponding subtask, the unique features of each technique justifying the corresponding subtask, and its approach compared to SoA are as follows (approximately from specific to general):

Specific 3D objects: This aspect of the problem addresses scenarios in which the precise location of a specific 3D object is located from sensor data. This type of problem uses techniques based on matching by finding explicit correspondences between data and 3D model. An immediate use of these techniques is in manipulation in which it is necessary to recognize and locate objects to grasp. Beyond manipulation, these techniques are needed for augmenting the environment description with location of known objects. This level of recognition is addressed in Subtask 2. The main idea of our approach is to use shape and viewpoint change as the main cues in detecting and localizing 3D objects. This approach departs from the state of the art in which few approaches address exact segmentation of the object and the pose localization problem simultaneously, and even fewer approaches can deal with optically challenging objects (semi-transparent, semi-reflective) due to lack of sufficient gradient variation and misleading appearance due to refraction and/or reflection of the background.

Urban specific features: The context of the mission in which the perception system should be used to constrain visual interpretation. Here we use the knowledge of the urban context to constrain the process. More precisely, we develop techniques specific to architectural features found in urban environments (e.g., building facades). Those are particularly important as they pertain to the detection of observation or egress points for operation in urban environment, as described in the Capstone Assessment Vision. This level of recognition is addressed in Subtask 1. The key advance proposed in this subtask over SoA is to use a feature representation that is tuned to the problem of recognizing urban features rather than using generic features as is normally the case. If successful, an important product of this approach is the ability to generate a sentence-like description of the environment which is amenable to communication with human operators. This would be an important building block from perception to shared mental model and human-robot interaction.

Object categories: A more general level of description involves detecting all instances of object categories, e.g., people or cars. This level usually involves classification techniques like SVM applied to local features of the input, with the classifiers learned on large pools of data, which are completely different techniques than the ones used at the other levels. This level can be refined by considering "subordinate" categories which involves recognizing fine-grained categories, e.g.,

man, woman, or child. The latter will be particularly important in providing cues to understand friend/foe relations in dynamic environments. This aspect of the work is addressed in Subtask 4.

We complement these different levels of recognition by investigating representations of the environment that are amenable to communicating descriptions to humans in the form of semantic maps. This includes strategies for building informative maps without exact localization or complete identification of objects in the environments and techniques for building compact representations of the input visual data tuned to indoor environment. The techniques proposed in this part of the work depart from the state of the art which relies either on highly discriminative features or on a few, reliably recognized landmark objects.

In the area of category recognition, the proposed approach will provide substantial advances over the SoA through a better approach to segmentation and perceptual organization. In addition, the question of recognizing subordinate categories, while critical to be able to deal with the large variations of appearance in the real world, has not been explored much. Here as well, our approach based on strong prior work on segmentation and perceptual organization will enable significant advances. As a long-term objective, in the case of people detection, our effort in classifying behavior will address an area of the SoA in which very little progress has been made. In particular, our (Caltech) expertise in building datasets will be instrumental in finding good representations from large, realistic, and well-annotated datasets.

Classes: This level addresses the broadest definition of classes to recognize. In particular, it includes “stuff” classes, e.g., wall or road, in addition to object classes. This level, coined “scene parsing” or “segmentation labeling” in the literature, aims to assign a class label, or a distribution of class labels, to every location in the visual data. The class of techniques used at this level is again very different than the ones described above and involves typical inference over graphs of regions. Based on prior work in FY11, we introduce an approach for scene labeling in Subtask 3 which we augment with new ideas for controlling computation through “anytime” techniques and with techniques for using interaction with the environment. Our key tool for advancing the SoA is to use new inference and learning techniques developed in related work and in Task I5. These tools were initially developed and demonstrated in FY11 on image labeling tasks, showing increased performance over the state of the art. We press on the development of this class of techniques as well as import other learning tools in FY12.

Metrics:

The standard metrics for all work on recognition or scene interpretation are based on precision/recall, ROC, or F1 tools, all of which measure the recognition or labeling accuracy of a set of classes or objects. Different types of recognition tasks differ in the type of data used. In that respect, data collection, annotation, and evaluation protocols are part of the research in each subtask. The main datasets that we will use to develop and evaluate algorithms are:

Urban features (P3-1): Collecting a new dataset in Philadelphia Center City containing over four hours of driving in dense urban environment. This will provide an unprecedented source of data for evaluating new and SoA algorithms tuned to urban context.

Category recognition (P3-4): We (Caltech) have developed some of the main datasets currently in use for category recognition. For people detection in particular, we use the standard benchmarks such as PASCAL as well as far more extensive benchmarks such as the Caltech pedestrian dataset (250,000 annotated frames). For subordinate categories, we have planned efforts based on additional annotations of the Caltech pedestrian dataset.

Scene labeling (P3-3): In addition to standard, relatively small benchmarks published in the literature (MSRC, Stanford, etc.), we plan to use the datasets collected in Task P4 in FY11. They include both images and 3D data over long paths (10 miles). Annotating these datasets is a continuing effort.

Semantic mapping (P3-5): We will assess our approach to semantic mapping by varying the following independent variables: number and size of semantic objects in the scene, ambiguity of arrangement of the objects, discriminability of the objects, complexity of space in terms of occlusions, and visual complexity of non-object components of the scene. Dependent variables include the robot trajectory, the 3D locations of objects, and in the future, the performance of the execution of a semantic space commands.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations: We anticipate that the research conducted in these subtasks will be reported in the major vision (CVPR, ECCV) and robotics (ICRA) conferences as appropriate given the deadlines. The PIs give regular invited talks at workshops and conferences in which portions of this work will be represented.

Task P3-5 is part of an ongoing collaboration between UPenn and MaryAnne Fields and Jason Owens at ARL.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Extract large-scale features for architectural details of building recognition in 2D. Implementation of FY11 state of our approach towards capstone experiment. Complete evaluation of co-inference algorithms. Complete evaluation of frame-by-frame algorithms for temporal consistency. Development of GUI to crowdsource image segmentation ground truth. Semantic mapping of indoor space in Fort Indiantown Gap annotating large components (walls, doors, etc.).
Q2	Evaluation of building detection based on large-scale features. Development of new approach (no sliding windows, no bounding boxes) resulting in publication. Algorithms based on (x,y,t) segmentations. Algorithms for semi-supervised learning (e.g., co-regularization). Hand-annotation of 10^3 pedestrians for subordinate category. Object-based localization in large-scale environment (acquisition of a new dataset of a market place).
Q3	Extract large-scale features for architectural details of building recognition in 2D and 3D. Development of accurate localization, and test in FTIG using a gripper. First version of anytime scene labeling. Evaluation of semi-supervised algorithms. Human-assisted tracked trajectories of pedestrians in Moore Video dataset. Semantic mapping of large-scale environment with respect to objects (test on market place).
Q4	Generate sentence-level description of building. Adaptation of new approach for optically challenging objects, and test in FTIG. Evaluation of performance of (x,y,t) approach and comparison to frame-by-frame approaches. Evaluation of anytime algorithm and comparison with baseline. (i) ROC curves comparing pedestrian detection with and without segmentation on Caltech pedestrian dataset and on Pascal dataset. (ii) Taxonomization of behaviors in Moore Video dataset. Combination of large components and objects in localization.

Related Research:
In addition to the linkages internal to P3 as described above, the scope of this task relates to other RCTA tasks as follows:

Applied research component: As documented below, many of the fundamental algorithms developed here will be transferred through the applied research Task P4. In the case of people detection (P3-4), we expect the results to be used in P6 because of their relevance to dynamic environments.

World model and learning: Integration of the fundamental algorithms will be affected through the world model (Task I2). Further details on the transfer to the world model are provided below.

In addition, P3-5 is closely related with I2-4 (Fox, UoW) where the emphasis of semantic mapping is in using new sources of sensor data (RGB-D) and focused on indoor environments. The learning components of P3-3 are closely tied with the research in learning in I5.

Other classification tasks: For complete coverage of the aspects of understanding static environments, P4-1 focuses on fine-grained categorization of terrain. That subtask differs from all the P3 subtasks in the approach (online learning), the representation, the task, and the input (proprioceptive sensors in addition to vision).

The work leverages work conducted outside of CTA in upstream basic research programs. This includes, in particular, basic research on visual recognition through ONR MURI Grant N00014-06-1-0734 and ONR MURI Grant 1015 G NA127.

Integration and Assessment Activity:
The P4 task addresses all the integration issues associated with the modeling of static environments. Most of the algorithms developed in this basic research task (P3) will feed into P4 for integration and further work on the robustness issues associated with complex environments. This mode of transition has already happened in FY11; the prediction algorithms developed in P3-3 formed the basis for the system in P4-2 which demonstrated the scene labeling algorithms on 3D point clouds as well as long image sequences. We expect this level of interaction between the two tasks (P3 and P4) to continue as a mechanism to migrate the capabilities toward the capstone assessments.

Integration of the fundamental algorithms will be affected through the world model (Task I2). Initially, the output of the algorithms from P3 is a set of labels associated to 3D locations in the world model, which we will augment later to including distributions of labels to reflect uncertainty in the interpretation and multiple hypotheses.

Different subtasks operate at different levels of maturity. We expect the approximate ordering of transitions from shorter to longer term: P3-3 (scene labeling, which is already in the process of transfer), P3-2 (specific objects which has been partially demonstrated in the context of manipulation), P3-1 (urban features, which is a more recent approach), and P3-4 (categorization and subordinate category, which addresses challenging and open problem with longer-term horizon). We expect to incorporate elements of semantic mapping (P3-5) in parallel.

The fundamental algorithms developed in this task are directly relevant to the capstone assessment elements, including locating objects for manipulation (P3-2); identifying building facades when exploring, providing cues for the location of egress points and generating description to Soldier (P3-1); detecting people while moving in populated areas, including providing cues as to status through subordinate classification (P3-4); and labeling key classes in the environment while exploring towards destination and performing surveillance (P3-3). The semantic mapping capabilities will be critical in two ways. First, they provide the level of description needed for communication with users (e.g., element 3 “Understand world through

perception and the world model"). Second, they provide basic capability for representing indoor environments from visual and range data (e.g., element 10 of the capstone scenario).

Subtask 1: Scene Context Recognition: Architectural Details of Building with Geometrical Shape Cues (UPenn)

Objective and Benefits:
The primary focus of this subtask is to achieve accurate 3D building geometrical shape estimation from low-level information, including surface orientation and position, to semantic information on building architectural details and architectural configuration. This information is useful for 1) recognizing building and structures of building, 2) reasoning about people behavior in and around building, and 3) communicating descriptions of the building with human operators. For example, a description of "two person are entering through a narrow door into a building with a hexagonal roof and two rows of bay windows" will help a human operator identify the building and the person quickly. The current system relies on small-scale image features. Such local features often ignore shape information that is very salient, semantically informative, and often too low-level for communicating with human operators.

State of the Art:
The current system for building recognition is mostly based on local features such as SIFT or low computation alternative SURF. These feature detection methods have been very effective in recognizing buildings at different distances and orientations and are reasonably reliable under partial occlusion. However, these local features are not semantically meaningful, and it is difficult to use them for communication between robot and human.

Key Barriers and Approach:
The main challenge is that local features are not as highly descriptive for identifying buildings (with many similar buildings). Image regions, obtained from segmentation or co-segmentation across an image pairs, provide strong features for building façade recognition and correspondence. For example, the architectural style of the building, represented by the shape of the roof and the configuration of the windows, doors, and signs, provides a highly informative cue for recognizing buildings from very different angles. Our approach addresses three main areas:

1) Recognize building at the level of architectural details of the building, the shape of the roof, and the configuration of the windows and doors. We will extract this semantic information forming highly informative cues for recognizing buildings from different angles and communicate such details to and from human operators. The main challenge is to include "large-scale" shape features of the image regions in addition to the local features. We will extend our Shape Packing approach that has been used successfully for deformable shape recognition in cluttered scene. We will also combine a method by Barsi and Jacobs (IJCV1996) to achieve better shape-based recognition.

2) Extend the approach above to include stereo 3D depth information. We will enhance the contour shape extract from 2D images with depth information to create 3D contours for architectural detail extraction.
3) Generate sentence-level description of building for communication with human operators. We will explore the vocabulary necessary for such communication. This effort complements well with other efforts from UPenn and CMU on robot localization and scene understanding. The output of our system could be used for inputs to their system, and location information could help our system to generate a shortlist of architectural styles of the buildings nearby.

Metrics for Evaluation:
We have been working in a dataset collected in Philadelphia Center City containing over four hours of driving in dense urban environment. We will generate ground-truth labeling of the architectural details of the buildings. We will measure our performance in terms of detection (precision-recall) and in terms of confusion matrix of the architectural details labeled. We will compare with the current system that uses only local scale image features which does not contain semantic information.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Extract large-scale features for architectural details of building recognition in 2D.
Q2	Evaluate building detection based on large-scale features.
Q3	Extract large-scale features for architectural details of building recognition in 2D and 3D.
Q4	Generate sentence-level description of building.

Subtask 2: 3D Object Recognition for Physical Interaction (UPenn)

Objective and Benefits:
Our goal is to detect and estimate the pose of 3D objects in order to grasp and perform complex manipulation tasks. 3D objects are given as 3D models with optional additional textured views, and they have to be detected in a scene where we have the ability to acquire more views. Detection has to be accompanied by localization with an accuracy depending on the grasping/manipulation task.

This subtask contributes to the world model via semantic labeling and enables behavior generation by supporting physical interaction with accurate object localization. While other tasks in semantic mapping (P3-5) produce a description of large-scale components of the scene, this task deals with:

- Objects that are of interest in terms of a mission like EOD neutralization.
- Tools, ammunition, and supplies (like water bottles) that are involved in logistics for dismounted Soldiers.
- Objects that might be obstacles and have to be removed.

This is a multi-year effort, and its evolution is characterized by increasing the complexity of objects and localization demands for physical interaction.

State of the Art:
State of the art in 3D object recognition can be divided into range-based and/or image-based. We address here approaches that make use of generic 3D objects of models from which we might have very little information about their color or texture and the query is an image. The state of the art may be summarized by taking the set of all possible views of a 3D model and matching it to the gradient histogram representation of the image by sliding a window. This is shown in Figure 3-10 where a 3D model of a drill is sampled, its histogram of oriented gradients is extracted, and matched to a scene in Fort Indiantown Gap.

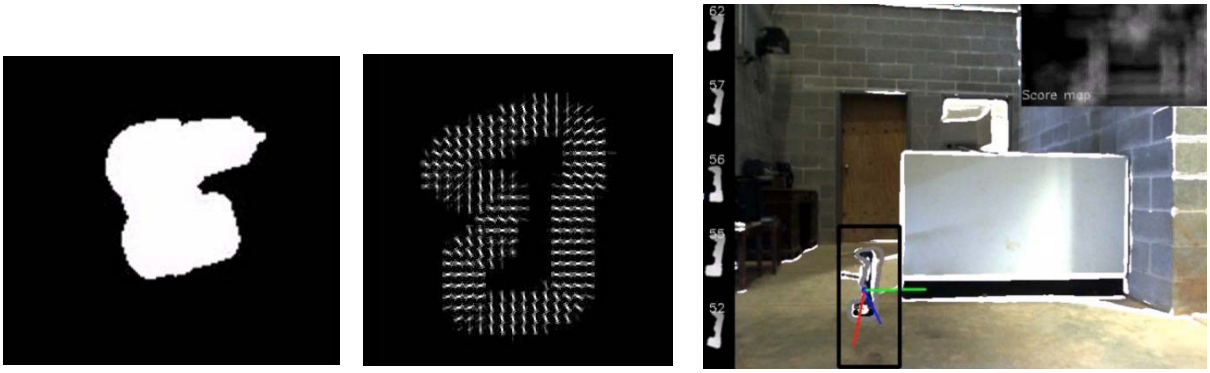

Figure 3-10: One of the views of a 3D model (left), a visualization of its histogram of oriented gradients (center), and detected in a scene in Fort Indiantown Gap (right) (results from FY11).

Approaches differ on the representation of the shape (shape context, histogram of oriented gradients, shape moments) and on the detection inference (bounding box versus hierarchical approaches). Very few works address an exact segmentation of the object and the pose localization problem, and this is the point where our proposed approach makes a difference. Even fewer works can deal with optically challenging objects (semi-transparent, semi-reflective) due to lack of sufficient gradient variation and misleading appearance due to refraction and/or reflection of the background.

Key Barriers and Approach:
Our goal is to detect and localize objects for the purpose of semantic labeling as well as physical interaction. We try to overcome two barriers in the state of the art: 1) Shapes are described through gradients which match many similar contour structures (see Figure 3-11) and result in low precision (many false alarms); and 2) bounding boxes contain still clutter, and they can neither accurately localize the object (such as for grasping) nor allow a 3D pose estimation (rotation and translation).

Figure 3-11: A flashlight is detected as false positive (left) due to resemblance of parallel edges (center) and a proposed segmentation that aligns objects with segments (right).

The main idea of our approach is to use shape and viewpoint change as the main cues in detecting and localizing 3D objects. The problem is formulated as a **representation** and a **selection** challenge. We introduced two new representations: the chordiogram which is a holistic descriptor of a silhouette and the parallax image which is a powerful stereoscopic representation enabling the detection of optically challenging objects (transparent, reflective). Given an over-segmentation of a scene in terms of superpixels and salient superpixels due to motion and stereo discontinuities, we formulate the problem as a selection problem in two spaces: the space of superpixels and the space of silhouettes capturing pose as well as inter-class variation. Resulting selections give simultaneously an exact segmentation and pose required for accurate manipulation tasks.

We still rely on the definition of a 3D object as a set of its views including their internal contour. Given a data-driven hierarchically organized over-segmented image, the proposed approach initially selects promising subsets of spatially contiguous regions determined by a classifier's scores. The scores, based on appearance and holistic scale-invariant shape cues, are computed efficiently by recursively propagating aggregate scores throughout the hierarchy. These regions are then refined by casting the segmentation problem as an integer quadratic program and solved via semi-definite programming relaxation. The obtained solution provides a detection score and a segmentation of an object. While this same semi-definite program is computationally expensive if applied to the set of all regions in the image, the initial pruning stage of the proposed approach yields substantial reductions in the complexity of the problem. This allows for the efficient consideration of multiple high scoring candidates in the hierarchy and a real-time computation. As a subsequent step, we refine the 3D pose estimation with an optimization with respect to translation and rotation starting from the discretized view as initial guess.

We will continue our work on optically challenging objects by using the same detection inference and location optimization as above but by changing the representation from simple views to inverse parallax views. Inverse parallax views can be produced when a supporting surface is detected and are obtained by differencing the back-projections of two images obtained from stereo or a moving platform.

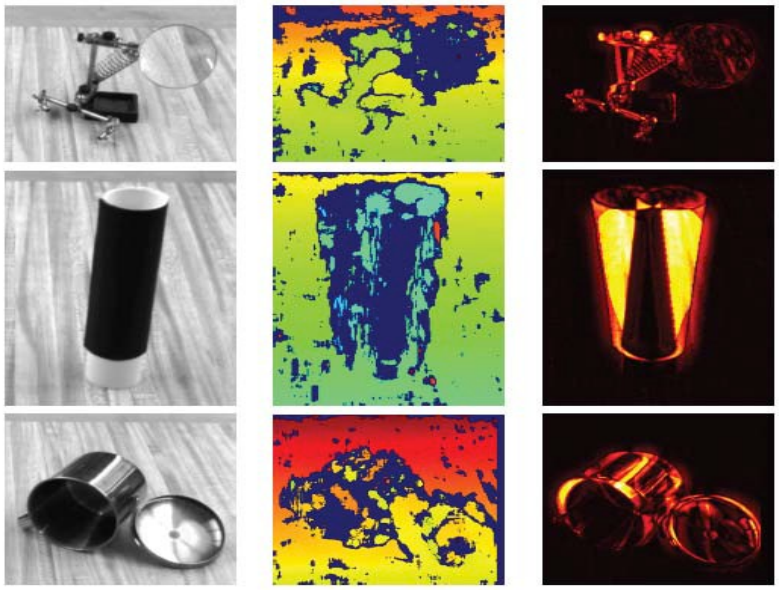

Figure 3-12: Example images of optically challenging objects (left), their disparity images using a standard correlation-based stereo algorithm (center) and their inverse parallax distortion images (right).

Metrics for Evaluation:
We will continue using a set of 3D object classes obtained from the Princeton shape database or by other DoD projects (DARPA ARM-S). As a metric, we plan to use precision-recall of detection and accuracy of pose using a metric in rigid motions.

We assess our approach by varying the following independent variables: scene illumination, number of object classes, number of 3D models per class, number of optional textured views per object, number of online acquired views, and optical complexity of the objects. Resulting plots are given with respect to detection (precision-recall) as well as with respect to accuracy of 3D pose.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Implementation of FY11 state of our approach towards capstone experiment.
Q2	Development of new approach (no sliding windows, no bounding boxes) resulting in publication.
Q3	Development of accurate localization, and test in FTIG using a gripper.
Q4	Adaptation of new approach for optically challenging objects, and test in FTIG.

Subtask 3: Deep Inference for Object Recognition and Scene Interpretation (CMU)

Objective and Benefits:
The objective of this 6.1 task is to develop fundamental algorithms for automatic labeling of sensor data with semantic classes. Integration, online processing, and 3D data are addressed in

the 6.2 Task P4. In FY11, we developed techniques based on inference machines, tested them on image labeling tasks, and migrated them to 3D data (in P4). In FY12, our main objectives are to: 1) generalize the algorithms to incorporate temporal and multi-modal information; 2) develop new segmentation learning tools; and 3) develop strategies for dynamically adapting the performance and speed of the algorithms.

The task contributes to building a semantic labeling to populate the WM. In assessment scenarios, it contributes to any scenario component that requires querying the world model for named elements (e.g., building, cars, etc.), such as element 3 "Understand the world through perception and the world model." Technical goal (1) uses all the sensor data and contextual information available in the IRA system, (2) improves accuracy of classification, and (3) addresses the need for speedy processing and multi-hypothesis representation in on-the-move scene interpretation.

Specific integration issues with the WM and with assessment plans are consolidated in the 6.2 Task P4.

State of the Art:
SoA uses global inference techniques, e.g., Probabilistic Graphical Models (PGMs), which do not scale to large problems and do not address temporal consistency, multi-modal, anytime results. Scene understanding remains a challenging problem that has focused the attention of the computer vision community in the past decade. While promising, existing techniques are faced with two major challenges. First, they involve complex optimization problems to learn the class models from training data and to use them in inferring the labels. This issue limits greatly the amount of data that can be used for training as well as the number of labels. The second problem is that it is difficult for these techniques to handle large sets of unlabelled data, which is necessary in a UGV scenario in which data is continuously received and manual labeling is not possible.

In this subtask, we started addressing the first challenge in FY11 by replacing the complex optimization problem formulated in standard approach by new techniques based on "deep inference" in which the original (intractable) inference problem is broken down in a series of smaller inference problems. Based on this baseline approach, our long-term plan is to address the second problem of unsupervised data.

Key Barriers and Approach:
Our approach, and the corresponding key challenges, to the three FY12 objectives defined above are as follows:

(1) Temporal and multi-modal information: Most of the work so far has focused on processing single images. In fact, the system would get continuous video input, and it is important to use the temporal component of the data. In FY11, we reviewed the current work in scene understanding from time-varying data, and we designed initial algorithms to exploit temporal information. In FY12, we plan to develop temporal consistency algorithms in three directions. First, labeled segmentations can be compared over consecutive frames to enforce smoothness constraints over

time. The basic idea is that the label distribution of a region cannot change drastically from one time to the next. This smoothing operation based on temporal consistency keeps the structure of our current single image interpretation pretty much unchanged. In our second step, however, we propose a substantial redesign of the algorithm by using hierarchical segmentations in 3D (space-time: (x,y,t)) rather than in image space only. This introduces new challenges in being able to perform this computation incrementally and to implement efficient inference algorithms. The benefits can be great, however, since it would provide a principled way to do temporally-consistent labeling in a principled way. Finally, this task focuses on labeling static environments, but the 3D approach can also be used in principle to label actions. We will start investigating this aspect in conjunction with P5 and P6.

In FY11, we started working on multi-modal classification, e.g., combining 3D point clouds and images. We investigated an approach based on co-inference in which hierarchical segmentations are generated for both data sources, e.g., one from a block of 3D point cloud and one from an image, and inference takes place simultaneously across both hierarchies, with communication between them to generate the interpretation that best agrees with both sources of data. One drawback of this approach is that it requires fully labeled data at training time because it optimized the parameters based on comparing the output of the algorithm to ground truth. This can be problematic since the availability of labeled training data is limited. We plan to explore ways to make use of vast amounts of unlabeled data in the training process. The general idea is to enforce the property that the label distributions across levels of the hierarchy must be consistent, even when the ground truth labels are not known. This approach of co-regularization will go a long a way toward enabling semi-supervised learning and using the data continuously recorded by the robot.

(2) *Segmentation and learning tools:* The performance of scene understanding algorithms improves with the quality of the image segmentation. In particular, the ability of separating regions corresponding to at least some objects would greatly enhance performance, and we will investigate ways to learn segmentation to improve overall performance. In particular, while physical interactions with the environment were not used so far in the scene understanding work, current work in I5 shows promise in learning how to segment objects by observing interactions. In this work, observed manipulation of objects (push, move) helps discover which parts of the scene are interconnected and how. We leverage these ideas to support the development of better segmentations. In addition, while our current efforts on scene understanding were focused on mobility and navigation, these techniques will bring us closer to DMUM applications by improving segmentation into physically distinct rigid bodies to enable interaction.

(3) *Dynamic perception computation:* This element addresses the problem that complex scene understanding tasks must be carried out with limited computational resources, requiring algorithms to be forced to report partial results to the WM, e.g., in the form of multiple partial hypotheses. Our approach is to adapt the anytime learning techniques developed in I5 to scene analysis. This part of the work is crucial in trading-off speed and accuracy in a principled manner. Specifically, Drew Bagnell's group's effort has focused on algorithms that trade-off time and classification performance in an anytime manner, meaning that the classification can generate an informative intermediate result at any point during its execution. This is related to

the concept of cascades in which increasingly complex classifiers are applied to the input in a chained manner, i.e., one classifier processes the output of the previous classifier. In contrast, the I5 approach dynamically orders the classification operations and the features used in classification. This is particularly relevant to perception data where classifiers are often applied to >1e6 data points in a single frame, and time must be carefully managed to focus on important areas of an image. Initial results show that this approach to anytime classification is much more effective than the classical cascade approach. For example, Figure 3-13 shows a comparison of the two techniques on a simple object detection task in which a classifier is applied to every location in the image; the anytime computation approach decides on which portion of the image the classification efforts should be focused. In this subtask, we extend this framework to vision task, specifically labeling tasks.

In the context of the scene understanding algorithms, there are two main steps that have to be taken to develop the anytime approach. (1) Develop approach for automatic evaluation of classification difficulty based on the labels on training data and on the distribution of features. This evaluation is used for prioritizing the parts of the input data for processing. (2) Extend the formulation to handle hierarchical representations. More precisely, our baseline scene labeling approach relies critically on a hierarchical segmentation of the image and on propagating label distributions across this hierarchy. Potentially, one can take advantage of this hierarchical structure in many interesting ways in an anytime framework. A key research direction is to determine the algorithm design that best uses this representation.

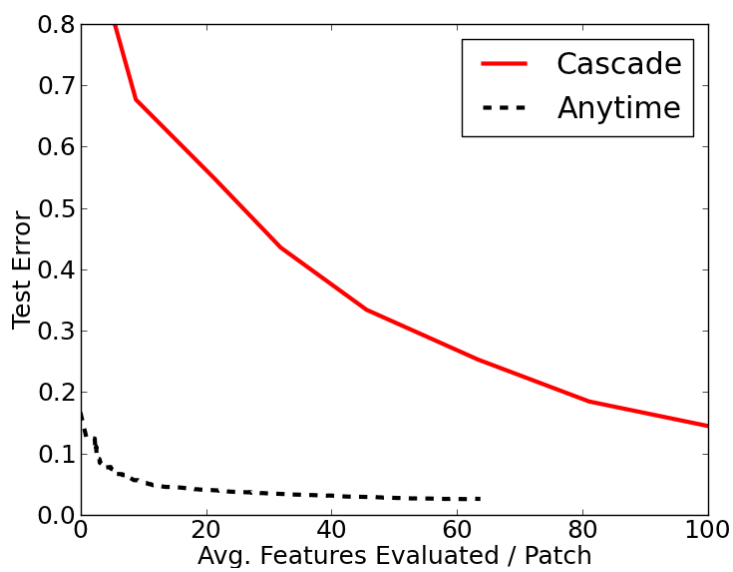

Figure 3-13: Initial comparison of RCTA's anytime computation with a standard cascade approach.

Algorithms for image analysis developed in FY11 were tested on large RCTA datasets and on standard benchmarks with accuracy equal or better than SoA with faster processing. They were ported to multiple 3D sensors (P4). Initial results were obtained (I5 and I6) for the new planned learning tools, in particular for anytime prediction.

Metrics for Evaluation:
We use the standard metrics for classification accuracy, i.e., ROC and P/R curves on labeling accuracy and derived metrics like F1 and AUC, evaluated on a) standard SoA benchmarks and b) large annotated datasets collected in FY11. Figure 3-14 shows an example of output on one of our urban datasets. Our immediate quantitative objective is to show continued performance improvements on standard benchmarks and on data collected for the IRAs. Table 3-2 shows the comparison between three different techniques, including the latest version of our algorithm using the most performant features.

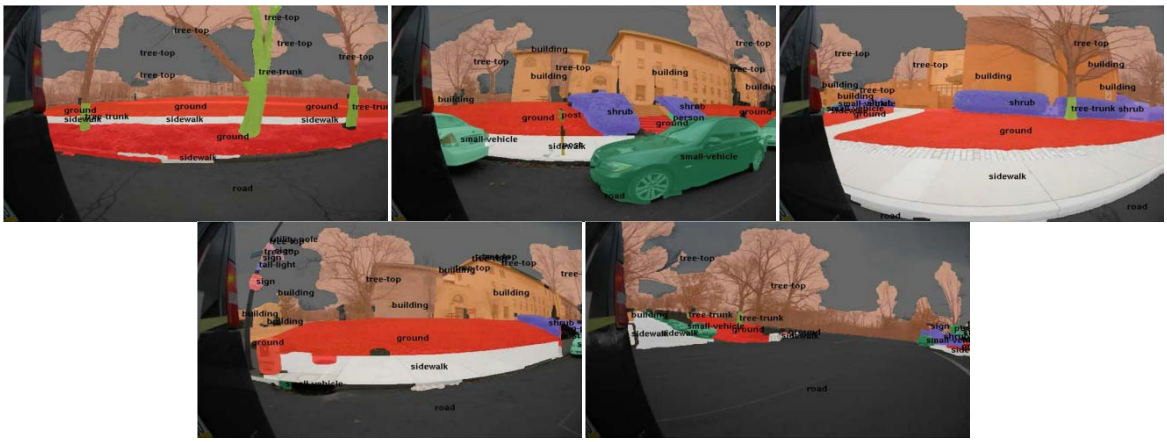

Figure 3-14: Classification output on five frames from the urban dataset.

	Avg. Precision	Avg. Recall	Avg. F1	Total Accuracy
Previous Features	0.727	0.700	0.704	0.801
New Features	0.774	0.751	0.751	0.835

Table 3-2: Example of quantitative evaluation showing improvement in FY11 by using new features for classification.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Complete evaluation of co-inference algorithms. Complete evaluation of frame-by-frame algorithms for temporal consistency.
Q2	Algorithms based on (x,y,t) segmentations. Algorithms for semi-supervised learning (e.g., co-regularization).
Q3	First version of anytime scene labeling. Evaluation of semi-supervised algorithms.
Q4	Evaluation of performance of (x,y,t) approach and comparison to frame-by-frame approaches. Evaluation of anytime algorithm and comparison with baseline.

Subtask 4: Visual Categorization (Caltech)

Objective and Benefits:
The objective of this subtask is to make progress toward making visual categorization and category detection fast, flexible, and accurate. This includes entry-level categorization (e.g., people versus vehicles versus buildings), subordinate categorization (e.g., man versus woman versus child) and behavior (e.g., fighting versus meeting). Our research is organized into three efforts:

1) Detection and segmentation of visual categories from static images (main 2012 aim).
2) Subordinate categorization (aim for 2012 and 2013).
3) Classification of behavior (aim for 2012 through 2014).

The first objective is prioritarian. We will expand upon our recent success in building the most accurate and computationally efficient pedestrian detector. We will maintain pedestrians as our most important category. To support our capability to detect and segment humans and other objects in complex scenes, we will develop methods for perceptual organization, such as figure/ground and occlusion relationships. We postulate that perceptual organization will boost substantially visual category detection performance.

The second objective, subordinate categorization, is a mid-term goal. In the context of pedestrian detection, it consists of discriminating subcategories such as man/woman/child. We wish to develop an approach for this problem and to be able to show significant progress during the next three years.

The third objective is focused on transferring to human behavior the capabilities and the basic understanding we are developing in the context of fly and mouse behavior (under different funding). We foresee stationary camera scenarios where the focus will be inferring individual, individual+object, and social actions and activities.

Our work on single-image detection and segmentation of categories, pedestrians in particular, complements JPL's stereo and motion-based work. We will be exploiting what is called "pictorial cues." While stereoscopic vision offers depth as an extra cue, pictorial cues are sometimes the only source of information. Anyone could escape detection from a stereo-based system by leaning against a building or a tree trunk and remaining still. This is particularly problematic indoors where there is considerable clutter and humans do often sit on chairs, sofas, and beds. Understanding of occlusion relationships from a single image and single-frame person detection from pictorial cues would greatly assist detection in such cases.

Our work on subordinate categorization would allow fine-grained distinctions that are key to quick decision-making; detecting a person is only the first step towards assessing friend/foe/bystander. Our work on behavior will be important in "semantic understanding" of the scene. Determining whether a vehicle is cooperating or is in competition with the humans and other vehicles appears to be a key capability.

State of the Art:
Our understanding of category-level visual recognition has progressed fast over the past ten years. From a state of almost complete ignorance in the late 1990s, we now have a broad-brushstroke understanding of how to approach visual categorization. In particular, we know that object visual categories may be represented by collections of "parts" or "local features" which are glued together by an overall geometric "shape." The details of this representation, as well as techniques to learn representations from examples, are still not settled. On the practical side, performance on visual detection and classification tasks is much improved but not yet sufficient for most practical applications.

We believe that successful segmentation and perceptual organization would enable a quantum jump in the performance of visual recognition methods. We have identified this unexplored question as one of our efforts which will take priority in 2012.

Furthermore, discrimination of similar subcategories is useful in many practical circumstances (e.g., identifying whether a pedestrian is a man, a senior citizen, a woman, or a child is vital to determining the best mode of machine-human interaction). The question has been explored very little, and current approaches to visual recognition are very poor at it. Developing approaches to subordinate categorization is our second effort.

Finally, classifying the behavior of people is a key issue in visual categorization and a topic on which very little progress has been made, mostly for want of good representations and of large, realistic, and well-annotated datasets. This is our third effort which we see as a mid- to long-term objective.

Key Barriers and Approach:
Our strategy to improve category detection is three-pronged:

- Combine segmentation with detection.
- Use multiple visual channels to enrich the information that the algorithms are using.
- Improve our boosting-based learning to take into account spatial correlation.

This approach improves on the state of the art which does not include segmentation and multi-channel features. Our technical approach builds on and combines individual capabilities that themselves are products of years of solid research in the field. Image segmentation and object detection have a history of being tackled as separate vision problems. There is significant opportunity to leverage each to help with the other, while taking advantage of existing algorithms for these individual components. Dr. Pietro Perona is an expert of visual categorization, and Dr. Michael Maire is an expert of segmentation and perceptual organization.

Metrics for Evaluation:
On pedestrian detection and segmentation, use benchmarking procedures described in Dollar et al. PAMI 2011 on the following datasets:

- The Caltech pedestrian benchmark (250,000 annotated frames).
- The Pascal benchmark (~1000 annotated images containing people).

On category detection and segmentation, use benchmarking procedures described in the PASCAL challenge website on the following dataset:

- The Pascal benchmark (~1000 annotated images per category for 20 categories).

On subordinate categorization, use benchmarking procedures described in Griffin et al. 2008 (Caltech 256 dataset) on the following dataset:

- The Caltech pedestrian benchmark. Images to be annotated for subordinate categories.

On behavior, benchmarking procedures are under development. The first dataset is under development, and the second dataset (VIRAT) was published in 2011:

- ~20 hours of video taken from the Moore building at Caltech and involving vehicle and pedestrian traffic. To be annotated.
- The VIRAT dataset (CVPR 2011).

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Development of GUI to crowdsource image segmentation ground truth.
Q2	Hand-annotation of 10^3 pedestrians for subordinate category.
Q3	Human-assisted tracked trajectories of pedestrians in Moore Video dataset.
Q4	(i) ROC curves comparing pedestrian detection with and without segmentation on Caltech pedestrian dataset and on Pascal dataset. (ii) Taxonomization of behaviors in Moore Video dataset.

Subtask 5: Semantic Localization and Mapping (UPenn)

Objective and Benefits:
Efficient communication between humans and robots about space and motions can be achieved only based on semantic descriptions of space even if prior maps (sketches) or robot acquired maps (point-clouds) exist. Low-level maps can facilitate motion planning in terms of obstacle avoidance but are not sufficient in communicating plans, and even more important, they cannot convey high-level situational awareness. The goal of this subtask is to enable communication of navigation commands by exploiting prior maps and by establishing new semantic maps. This is part of the shared cognition of the world model and goes in both directions in establishing the world model as well as localizing a robot or Soldier with respect to this world model. This subtask fits into the effort of understanding the world through perception (element 3 of the Capstone Assessment Vision) and exploiting prior intelligence and learning new semantics and spatial relationships about the world (element 4 of the Capstone Assessment Vision). We work closely with the semantic perception efforts at ARL (Fields and Owens). This is a five-year effort, and its evolution is characterized by increasing the complexity in terms of space (indoor, outdoor, cluttered) and in the size of the ontology involved in the semantics.

State of the Art:
State of the art in semantic mapping and localization relies heavily on visual features and hard decision on detection of objects and other scene entities. Approaches rely heavily on the existence of discriminative objects of space (monitors in office spaces, appliances in kitchens) as well as on dense reconstructions of the environment. Most of the research in object-based localization is tailored to small indoor environments like offices with simple topology and few object categories. Many approaches detect doors or gateways for place recognition as well as for detection of passages in a topological sense.

Key Barriers and Approach:
This subtask comprises two thrusts: an image-based approach relying on objects outdoors and indoors and an RGB-D-based approach characterizing larger components like rooms, walls, doors, and packages in indoor environments.

Thrust 1 (Daniilidis): The main idea of our approach is that semantic localization and mapping should depend less on static understanding of the world by recognizing objects, and instead apply a probabilistic approach on exploiting the robot's motion to acquire new views and establish self-location as well as location of objects without making hard commitment detection in object recognition. In FY 2010 and 2011, we have formulated a novel approach for object-based localization based on prior maps.

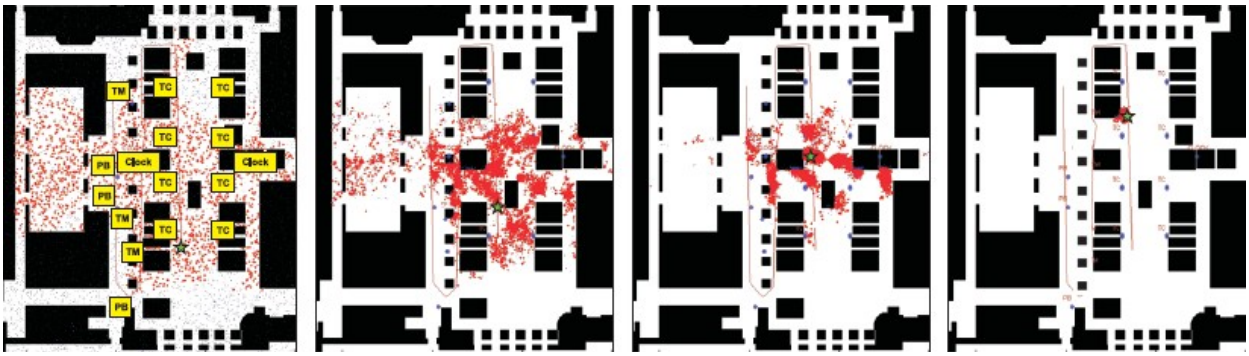

Figure 3-15: Prior map with annotated objects (yellow boxes) and the convergence of Monte-Carlo localization based on the probabilities of detecting an object.

We will continue this approach for larger spaces with more abstract spatial descriptions, and we will start working on establishing semantic maps of unexplored spaces. All semantic mapping approaches rely heavily on performance of object recognition, but ours makes only *soft commitment* on object recognition and exploits probabilistically the acquisition of several views while the robot is moving. Given the still low performance of object recognition approaches, our approach outperforms the state of the art when the final goal is localization and mapping rather than single picture recognition.

Thrust 2 (Taylor): One aspect of our approach has been our application of image segmentation schemes which are used to partition the input RGB image and, hence, suggest groupings to the subsequent surface detection and merging stages. This approach has led to scene parsing

algorithms that are significantly faster than previously proposed techniques. We have also been able to use the results of our plane extraction procedure to suggest high-level floor plans using an efficient optimization-based inference scheme. Further work will expand upon these results so that that the scheme can be used to parse extended scenes and handle a wider range of architectural features.

In previous research efforts, we have developed novel segmentation techniques based on randomized hashing which allow us to subdivide images at video rates on laptop class machines. These algorithms have been used as subsystems in successful scene interpretation algorithms. We are currently developing new techniques that can be used to produce edge-accurate segmentations of depth images. We have been able to produce significant speedups by exploiting the properties of the Delaunay Triangulation and the associated planar graph.

Our thrusts are closely related with I2-4 (Fox, UoW) where the emphasis of semantic mapping is in using advanced sensorial data (RGB-D) and focused on indoor environments.

Metrics for Evaluation:
In thrust 1, we have made significant progress in localization given a map of a large indoor space with a variety of four to ten objects. We will assess our approach by varying the following independent variables: number and size of semantic objects in the scene, ambiguity of arrangement of the objects, discriminability of the objects, complexity of space in terms of occlusions, and visual complexity of non-object components of the scene. Dependent variables include the robot trajectory, the 3D locations of objects, and in the future, the performance of the execution of a semantic space commands. Progress will be measured using the metrics above in the setup of the Capstone Assessment where prior maps about the scene will be used and additional intelligence about position of objects like doors and trashcans will be used for localization. Maps will be automatically enriched with additional objects with the same probabilistic approach.

In thrust 2, we will measure progress by keeping track of two kinds of metrics; the first will relate to the quality of the segmentations produced by the method. One way to measure this is by evaluation on standardized test sets like the Berkeley Segmentation Database. Another approach is by passing the segmentation results from the scheme on to subsequent analysis procedures and evaluating how often the system successfully suggests meaningful groupings which lead to correct interpretations. Another metric which will be monitored is the computational effort required to produce the segmentation. This can be quantified by considering the running time of the procedure and the computational power of the platform (GPU versus CPU versus embedded processor).

	Subtask 5 Quarterly Milestones and Deliverables
Q1	Semantic mapping of indoor space in Fort Indiantown Gap annotating large components (walls, doors, etc.).
Q2	Object-based localization in large-scale environment (acquisition of a new dataset of a market place).
Q3	Semantic mapping of large-scale environment with respect to objects (test on market place).
Q4	Combination of large components and objects in localization.

Subtask 6: Developing Metric/Topological Maps to Promote Common Ground (ARL)

Objective and Benefits:
To effectively use intelligent unmanned systems as components of a military unit requires data structures and algorithms that enable the robots and humans to have a common understanding of the world frequently referred to as *common ground*. In this project, we develop a method to incrementally segment an environment into human understandable regions such as rooms, hallways, atrium, and portals using a RGB-D (Red, Green, Blue – Depth) sensor such as the Microsoft Kinect.

State of the Art:
State of the art in semantic mapping and localization relies heavily on visual features and hard decision on detection of objects and other scene entities. Approaches rely heavily on the existence of discriminative objects of space (monitors in office spaces, appliances in kitchens) as well as on dense reconstructions of the environment. Most of the research in object-based localization is tailored to small indoor environments like offices with simple topology and few object categories. Many approaches detect doors or gateways for place recognition as well as for detection of passages in a topological sense.

Key Barriers and Approach:
We enable this common ground through research in region segmentation and hybrid metric-topological maps. This involves extracting and encoding additional information about the environment in maps with both metric and topological components. Significant features derived from a robot's on-board sensors and interpreted as salient waypoints may then be associated to objects within this map structure to enable navigation without relying completely on metric specifications. Consequently, relationships between entities in space and context awareness through meaningful landmarks become critical elements.

The main idea of our approach is that as the robot navigates through its environment, it needs to run both a short and long timescale mapping process. The short timescale mapping process uses a visual odometry approach to produce a metric SLAM map. Our assumption is that we have no a priori information available about the environment and that global localization sensors are not available. Unlike most SLAM approaches which use image features such as lines and corners, our approach also uses objects segmented from both the image and point cloud produced by the RGB-D sensor to maintain localization. This approach should reduce the ambiguity in the

localization ensuring loop closure. As the research progresses, we intend to consider constellations of objects and recognized objects (utilizing the research in subtask 5) to increase the reliability, and potentially the speed, of the localization process.

The long timescale process will periodically segment and label the SLAM map. There are two approaches that we are exploring: portal identification via Generalized Voronoi Graph (GVG) segmentation and incremental inter-visibility graph partitioning.

Portal identification via GVG segmentation is based on the observation that portals and regions are complementary (or duals): if you know one, then you can derive the other. Where clutter recorded in the occupancy grid, e.g., desks and chairs within a room, may cause the visibility algorithm to over-segment a region, portal classification relies on learning the local features around a segment that indicate portals as defined by human-annotated maps.

Incremental inter-visibility graph partitioning separates the graph into visibility regions. As the robot explores, it updates the cells based on the information it receives from the laser range scanner. At the same time, it uses the range scans to generate the GVG-based topological graph. When the robot discovers new nodes in the graph, it calculates the inter-visibility between the new and previous nodes using the occupancy grid. A single region is expanded until a visibility threshold between all nodes in the current region is compromised. A region may then be defined as the convex hull of the points contained within the region. This method requires one parameter: the threshold which represents the ratio of visible to obscured nodes.

Metrics for Evaluation:
We will assess our approach to semantic mapping by varying the following independent variables: number, size, and type of regions in the scene; the distribution of non-structural objects in the environment; discriminability of the objects; and complexity of space. Dependent variables include the robot trajectory as well as the 3D locations of objects.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Subtask P3-6 is part of an ongoing collaboration between UPenn and MaryAnne Fields and Jason Owens at ARL.

	Subtask 6 Quarterly Milestones and Deliverables
Q1	Develop/Acquire datasets to test segmentation process. Finish implementing Kinect-based SLAM algorithms on K-bot.
Q2	Integrate the segmentation and SLAM processes.
Q3	Initial experiment using the K-bot in office environments at UPenn and ARL.
Q4	Additional field and virtual experiments. Reporting.

P4-2012 – Perception for Missions in Complex Environments

Level: 6.2 Applied Research

	Name	*Organization*	*Subtask*
Principal Investigator	Michael Turmon	JPL	P4-2
Additional Investigators	Max Bajracharya	JPL	P4-2
	Nathan Wedge	QNA	P4-1
	Alonzo Kelly	CMU	P4-3
	Martial Hebert	CMU	P4-3
	Drew Bagnell	CMU	P4-3
	Omead Amidi	CMU	P4-3
	Jianbo Shi	UPenn	P4-4
	Alex Chan	ARL	P4-5
	Shuowen Hu	ARL	P4-5
	Prudhvi Gurram	ARL	P4-5

Objective:
We are developing capabilities in two research thrusts: learning for terrain characterization and real-time semantic mapping of scene components. These capabilities apply to a broad set of missions, and our work is evaluating the capabilities in the context of the platforms and scenarios planned for RCTA. The work is primarily relevant to the "look" and "move" capabilities outlined in the program vision. The perceptual cues we learn also provide key semantic inputs for "think" capabilities. Proprioceptive sensors, trained within a given environment, provide parts of a "work" capability for interacting proprioceptively with the environment.

The semantic associations learned from training data allow scene elements (traversable grass, roads, cars, loitering people) to be identified by the perceptual systems developed in this task and placed into a shared three-dimensional world model for use in translating commands ("follow the road," "watch the people") into actions. The same information supports mobility in complex environments with diverse terrain classes and numerous obstacle types.

Background:
Recent progress in online learning, much of it led by our team, has shown the feasibility of adapting terrain classification models dynamically as the robot explores its environment. This class of approach enables adapting to changing conditions and longer-range perception by learning how to classify long-range sensor data (e.g., from video or SfM) from accurate, shorter-range sensors (e.g., from LADAR or stereo). While we have applied this prior work to estimating terrain traversability, there is now a unique opportunity to build upon these ideas to include scene analysis elements. Based on the current state of the art, the challenges include designing

incremental learning algorithms, selecting subsets of the accumulated sensor data to retain, and reducing the dimensionality of the features from sensor data.

Migrating the fundamental object recognition and scene analysis techniques addressed in Task P3 is challenging because of the computational constraints and the level of classification performance required. Prior work in this area, including our own, suggests two areas that will facilitate this migration. First, temporal integration of the data can be effective in reducing the error rates and in improving classification efficiency. Second, using geometric constraints, e.g., reasoning about visible aspects of the object or enforcing known constraints of the environment, is potentially a powerful way to reduce the search space for context-based recognition algorithms and, therefore, improve error rates.

Research Description:

The subtasks are numbered in approximate order of proximity of the sensor to the environment it measures. Thus, Subtask 1 (QNA, "Interaction with and Displacement of Vegetation for Small UGVs") works directly at the proprioceptive interface between the vehicle and the environment; Subtask 2 (JPL, "Unsupervised Near-to-Far Learning of Terrain Type") extends learned proprioceptive cues into the visual field; Subtask 3 (CMU, "Online Recognition in 3D Point Clouds") integrates visual and stereo/LADAR cues into classified 3D point clouds; and Subtask 4 (UPenn, "Rare Events and Motion-saliency Based People and Scene-context Recognition") identifies rare behaviors and motion in complex visual scenes. All tasks integrate information across time, and are adapted to implementation in an online context.

We proceed to an overview of each subtask. Detailed examination of the state of the art, research strategy, and metrics are given in the individual subtask descriptions.

Interaction with and Displacement of Vegetation for Small UGVs. Last year, this work was situated within P2, but it is a better fit here in P4 due to synergy with Subtask 2. This subtask will pursue a capability to perceive and assess compliant obstacles like tall grass and other vegetation. Activities in the coming year are directed first at assessing trafficability of vegetation using proprioceptive sensors (acceleration, slip, and torque) and visual inputs (visible and infrared cameras, laser rangefinders) on TALON or DR20 robots. Subsequent activities will insert estimated (continuous-valued) trafficability metrics into a cost map to support planning. The combination of well-informed proprioceptive sensing with a continuous-valued cost map will allow a planner to choose to interact with terrain (i.e., push through vegetation) if the cost gain is sufficient.

Unsupervised Near-to-Far Learning of Terrain Type. This subtask uses large corpora of un-labeled vehicle logs containing proprioceptive and visual information (e.g., accelerometer time series and stereo camera imagery) to find terrain classes and insert them into a world model that extends into the far field. It does this by using an unsupervised classifier (a constrained hidden Markov model) to cluster the time series into different proprioceptive classes. The proprioceptively-determined terrain classes are then associated with terrain underfoot. This provides labeled data to train a visual classifier that can label the visible scene according to

deduced terrain type. Work this year is directed at maintaining acceptable end-to-end error rates (< 10%) while accommodating more terrain classes, specifically including vegetation. Future work is directed at ensuring the class decomposition remains stable despite confounding variables like slope and vehicle load. This subtask can use the more refined proprioceptive information provided by Subtask 1 to improve the fidelity of its terrain classes.

Online Recognition in 3D Point Clouds. This subtask will provide semantic labels for 3D point clouds of stereo or LADAR data in an online context. The system operation has proven successful last year in a non-real-time context, and much work this year is directed at identifying and removing bottlenecks to allow fast operation on the robot. First, algorithms will be generalized to allow incremental addition of new data points in a streaming mode. Also, the computation of local image features (like SIFT) will be made more intelligent; for example, by using sampling or quantization strategies. Additional work, interfacing with GDRS, focuses on inter-operation with a world model (export of semantic information to the WM and receiving focus-of-attention instructions from the WM). This subtask interfaces with P3 (classification tools) and Intelligence (machine learning and world-model-directed focus of attention). It is also equipped to integrate terrain-type cues from Subtask 2.

Rare Events and Motion-saliency Based People and Scene-context Recognition. The overall goal of this subtask is to develop maps to recognize people and scene-context attributes in an online, urban context. Fundamental to this subtask is the ability to use static *and* dynamic video information to identify attributes like change and motion. In prior work, detectors for certain activities have been successful, and the proposed work extends this capability to recognizing rare events such as arrival of a truck at a business. Rather than using complex activity models or expert-provided features which would not generalize well, the approach divides the video into segments and associates segments with visual features. The result is a clustering where each cluster contains a set of segments containing a behavior. This subtask is synergistic with Subtask 3: it adds human activity information to the world model, and it can also use the improved semantic scene information provided by Subtask 3.

Metrics:

These subtasks are concerned with detection and classification, so they all are heavily guided by the measures of error probability that are standard in this area: Precision/Recall (P/R) curves, Receiver Operating Characteristic (ROC) curves, and confusion matrices. Additionally, inspection of individual error-causing cases or behaviors is frequently valuable.

These standard statistics provide invaluable information but are only computable with respect to a reference dataset. Obtaining ground truth is a considerable investment in its own right, and subtasks have typically developed reference datasets (JPL, CMU, and UPenn have already gathered these datasets) and have plans to acquire more data through field experiments with high-resolution cameras or extra sensors (from which ground truth can be deduced) or expert annotation. Continual acquisition of ground truth data is necessary to prevent incorrect assumptions on the nature of the sensors or the environment from influencing system design.

Besides these error-probability measures, QNA is also guided by planner-in-loop evaluations, or comparisons of the system with a human tele-operator.

Finally, CMU, and to a lesser extent QNA and JPL, are aiming at real-time operation of some system components, so computation time onboard a robot is a second key metric. A related metric is time-to-completion of a run, used by QNA.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
QNA and JPL (Subtasks 1 and 2, respectively) plan collaboration and interactions by site visits. CMU will use QNA's robot test course for some of the semantic mapping experiments. CMU will also work with ARL to test the classification algorithms on data from the ARL MEMS LADAR.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Characterize perception of vegetation, etc., using proprioceptive sensing hardware. Integrate LADAR and gather field data. Implement streaming computation of point-cloud labels. Construct initial codebook of activity patterns.
Q2	Evaluate classifiers for perception of difficult terrain types. Learn proprioceptive classifier for dense vegetation, and develop co-training approach. Optimize point-cloud feature computation. Construct classifier for detecting rare activities.
Q3	Select a classifier, and define a traversability-cost metric. Estimate end-to-end error rates for vegetation classes. Provide interface to world model for focus-of-attention. Construct classifier for street scene classification.
Q4	Experiment with planners operating on terrain-traversability metric. Use co-training to improve proprioceptive classification. Test point-cloud system integrated with the world model. Improve classifier for street scene classification.

Related Research:
This task provides the online recognition techniques needed to provide perceptual input to the tasks in the other areas, i.e., semantic and geometric information about objects and features. In the DMUM area, image segmentation, object recognition, and 3D model estimation results are used in tasks M2 and M3 to aid in planning grasps. In the mobility area, the online classification methods of P4 are used for selecting the right control model and parameters and to predict vehicle response in task M7. In the Intelligence area, the techniques proposed in P3 and P4 are used in I1 and I2 for reasoning and inference in the world model. They also provide the terrain and environment representations used in I4 for executing complex behaviors. Finally, some P3

tasks (level 6.1) contribute components to P4 as they reach maturity in terms of error probability in realistic settings.

Integration and Assessment Activity:

- Evaluate performance of the proprioceptive system in vegetation using transit-time metrics.
- Evaluate classification performance in unsupervised, near-to-far learning in diverse and multi-class terrains including vegetation.
- Evaluate online classification/recognition algorithm for unstructured 3D point clouds. Develop interface to propagate semantic point cloud representation into a world model. Develop interface to receive focus-of-attention cues from world model.
- Evaluate performance for rare event detection and classification.

Subtask 1: Interaction with and Displacement of Vegetation for Small UGVs (QNA)
This subtask will build a capability to allow small UGVs to interact with vegetation and similar obstacles, to determine their traversability, and thus allow increased mobility and scene understanding. In 2011, this subtask was in P2, but it is more synergistic with P4.

Objective and Benefits:
This subtask effort will pursue a capability to perceive and surmount displaceable "obstacles" like tall grass and other vegetation, leveraging our previous work toward navigation including interaction with movable objects. In the context of the program vision, this effort is situated as a concurrent thrust toward the "look" and "move" capability building blocks, in which we intend to leverage property-based classification to enable a small UGV to intelligently drive through/over cluttered natural terrain that might otherwise be trivially accepted as drivable by a human operator. In particular, this capacity chiefly addresses the RCTA Capstone Assessment Vision element 5 to "move through mobility challenges on the way due to both natural terrain and urban clutter." We anticipate that the fundamental functionality is an achievable goal using the upcoming year to explore and implement the elements of this new capability, while the broader effort will leverage future years to evolve the capability toward fully-interactive navigation with solid, movable objects in addition to displaceable vegetation.

State of the Art:
Relevant technologies to this subtask effort fall into two categories: algorithmic research and fielded system capability. While there is much constructive work in the areas of terrain classification (e.g., [3]) and surface detection (e.g., [5]), implementation of the resulting capabilities on fielded systems are much sparser. Rather, many of the applications of these systems are to overhead/satellite imagery, stationary cameras/sensors, and platforms markedly larger than the typical explosive ordinance disposal robot. Further, many of the classification techniques also provide some manner of confidence in their choice (e.g., in [1]) that is either unused or used in a trivial fashion. One example of this practice is the formulation of the maximum a posteriori decision rule in [2] which concludes that a location is occupied if its occupancy probability is greater than its empty probability.

Key Barriers and Approach:
In this subtask, we seek to explore the capacity of proprioceptive sensing to discern drivability along with its implications on terrain classification and planning as they apply to small UGV platforms. Our primary focus will be on exploiting the ability of this class of sensors to determine the impact of various terrain types on driving cost (e.g., in time). In order to surpass the state of the art, we plan to examine a structure akin to that of [4] which phrases these costs (and plans its paths) directly in the probabilistic space of classification, rather than in the assumed, deterministic space of most likely classification/cost.

In the hardware domain, we have identified value in torque sensing to perceive the interactions that will occur in navigation in challenging terrain based on our previous work in the interactive navigation area. In order to support this concept, we will leverage QNA's IR2-3 subtask to augment a small UGV (TALON or DR20) with any appropriate sensors that will assist in characterizing vehicle-terrain interactions (measuring quantities such as resistance and roll/slip rates). We believe this strategy will facilitate the development of the capability of navigation through vegetation by operating on and directly exposing the uncertainties associated with classification across both perception and planning. This will provide a navigational behavior that explores the unknown and uncertain while sensibly varying its particular strategies.

As there are a myriad of available classification schemes that function via diverse sensor inputs (e.g., visible and infrared light cameras, laser rangefinders, accelerometers, etc.), there are convenient avenues to leverage results from other RCTA efforts such as the previous classification work in JPL's P4 subtask, "Online Learning and Adaptation." In particular, we plan to collaborate in fusing our proposed proprioceptive direction with their upcoming work toward exteroceptive classification of these types of conditions. We have also performed prior classification work based on neural networks at QNA under the DARPA LAGR program [3] that will be beneficial in pursuing our research goals.

Metrics for Evaluation:
Initial demonstrations of functionality for our goal capability can be benchmarked against previous systems that do not account for potential drivability of "obstacles" like tall grass (for a conservative comparison) or against tele-operated systems in which a human operator judges the drivability of such "obstacles" (for an aggressive comparison). As our proposed direction is focused on probabilistic notions, a chief metric of its success is repeatability on equivalent test scenarios. For example, we can tolerate the robot choosing different paths in the same scenario, provided they have similar perceived cost or are motivated by reasonable uncertainty in information. To that end, a primary metric for our experimentation will be both the mean and variance of time taken to complete runs of an individual experimental setting.

As this subtask represents a branch from our current, ongoing work toward interactive navigation, we expect to be able to apply much of what has been accomplished in the development of simulation capability and the measurement of interaction toward this effort. Intermediate progress is measureable via functional metrics of motivated traversal of various types of nontrivial terrain (e.g., tall grass, light brush). In relation to the below quarterly

milestones, we can also validate our overall progress via the completion of the individual components of our effort.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Conduct characterization experiments to capture the ability of currently-fielded proprioceptive sensing hardware to perceive relevant terrain types (e.g., tall grass, sand) and the potential for forward progress through them.
Q2	Explore and evaluate classification methods consistent with established sensing hardware and capabilities, focusing on their ability to discern difficult terrain types.
Q3	Select an applicable classification scheme, and develop a capability alongside it to recognize and contrast robot interaction with displaceable and fixed objects. Include a cost component derived directly from experience with the environment.
Q4	Implement and experiment with planning algorithms that operate on deterministic and/or probabilistic costmaps to test integrated, deliberate interaction with vegetation.

References:

[1] D. M. Bradley, S. M. Thayer, A. Stentz, and P. Rander, "Vegetation detection for mobile robot navigation," Carnegie Mellon Univ. Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-04-12, February 2004.

[2] A. Elfes, "Occupancy grids: a stochastic spatial representation for active robot perception," in *Proc. 6th Conf. on Uncertainty and AI*, Cambridge, MA, July 1990, pp. 136-146.

[3] M. Happold, M. Ollis, and N. Johnson, "Enhancing supervised terrain classification with predictive unsupervised learning," in *Proc. Robotics: Science and Systems II*, Philadelphia, PA, August 2006, pp. 41-48.

[4] L. Murphy and P. Newman, "Planning most likely paths from overhead imagery," in *Proc. 2010 IEEE Int'l Conf. on Robotics and Automation*, Anchorage, AK, May 2010, pp. 3059-3064.

[5] C. Wellington and A. Stentz, "Learning predictions of the load-bearing surface for autonomous rough-terrain navigation in vegetation," in *Proc. 4th Int'l Conf. on Field and Service Robotics*, Lake Yamanaka, Japan, July 2003, pp. 83-92.

Subtask 2: Unsupervised Near-to-Far Learning of Terrain Type (JPL)

This subtask is developing a capability to learn terrain type, such as grass, gravel, or dirt, from proprioceptive signatures and to project this learned class into the far field of the world model by associating it with a visual signature.

Objective and Benefits:

Our objective is to provide a perceptual system that learns, in an unsupervised fashion, the visual appearance and proprioceptive feel of diverse terrains. This perceptual information will enable a robot to make long-range maps and to cross difficult terrains, including vegetation and slopes. There are two points of relevance to the program vision. First, this capability will enable the robot to navigate autonomously, as needed for logistical support of dismounted Soldiers, or for maneuvering to and around the house in the Capstone Vision Assessment. Second, the system categorizes terrain into visual/proprioceptive classes that can be *given semantic tags and put into a world model*. This provides the perception elements needed for commands like "go to the house but avoid the dirt path."

In relation to the Capstone Assessment Vision, this capability addresses element 3, "Understand the world through perception and the world model," and element 5, "Move through mobility challenges on the way due to both natural terrain and urban clutter." The labels produced by this subtask, produced both in the near-field by proprioception and in the far-field by the visual classifier, can be inserted in the world model and used for planning and semantic direction. In relation to the more general RCTA program vision, this work is directed at the "look" and "move" capabilities.

This effort is expected to continue for three years and comprises both fundamental advances in technique as well as engineering improvements that will enlarge the domain of applicability of the system.

State of the Art:
This comparison takes in the joined proprioceptive/visual system because their interaction, especially regarding noisy labels supplied by the proprioceptive classifier, is important in determining end-to-end performance.

Terrain classification systems for autonomous robot navigation have conventionally used rule-based approaches [MB06] or offline, supervised learning [DVH04,MC05,WS03] to segment imagery, laser range finder, or proprioceptive data. For example, [MC05] reports good vision-based classification with color distributions learned from examples taken over varying conditions but points out that a large amount of training data covering all expected environmental conditions is required. Our unsupervised approach circumvents the problem of gathering training data. Other work classifies terrain based on vibration signatures from vehicles on pre-defined terrains such as asphalt, gravel, and grass [WI09,CC08]. The data for supervised learning is normally labeled by a human operator which is tedious and subjective. The unsupervised approach avoids human labeling. Besides the novelty of the unsupervised near-to-far approach, we are not aware of work that is aimed at enhancing the stability of learned object classes by simultaneously using proprioceptive and visual information.

The program of research in this year is distinct from our prior work on visual classification for LAGR [Ba08]. First, LAGR had no provision for use of a LADAR to provide maps in vegetation, where stereo is likely to be unreliable. Extending proprioceptive classification to these "hard cases" is important to real-world navigation. Also, systems issues in the field-test driven LAGR program prevented research on co-training (enhancing segmentation by using proprioceptive features together with visual features). In the proposed work, we intend to demonstrate that significant improvements can be realized in these areas. In general, measures to limit the down-stream effects of the errors produced by the proprioceptive classifier on the visual classifier have not been explored in prior work.

Key Barriers and Approach:
Our approach uses extensive logs of high-rate IMU data, collected over diverse terrain types, to provide proprioceptive information for an unsupervised terrain classifier. The unsupervised classifier uses a hidden Markov model (HMM), determined with advanced optimization

techniques, to exploit temporal coherence of terrain classes as the vehicle moves. The terrain types underfoot are then associated with their visual appearance and projected into a far-field world model for route planning or semantic navigation. Because the approach is based on the *unsupervised* processing of proprioceptive inputs, it can exploit voluminous logs – the more the better – that collectively contain "deep memory" about vehicle-terrain interaction across a range of terrains.

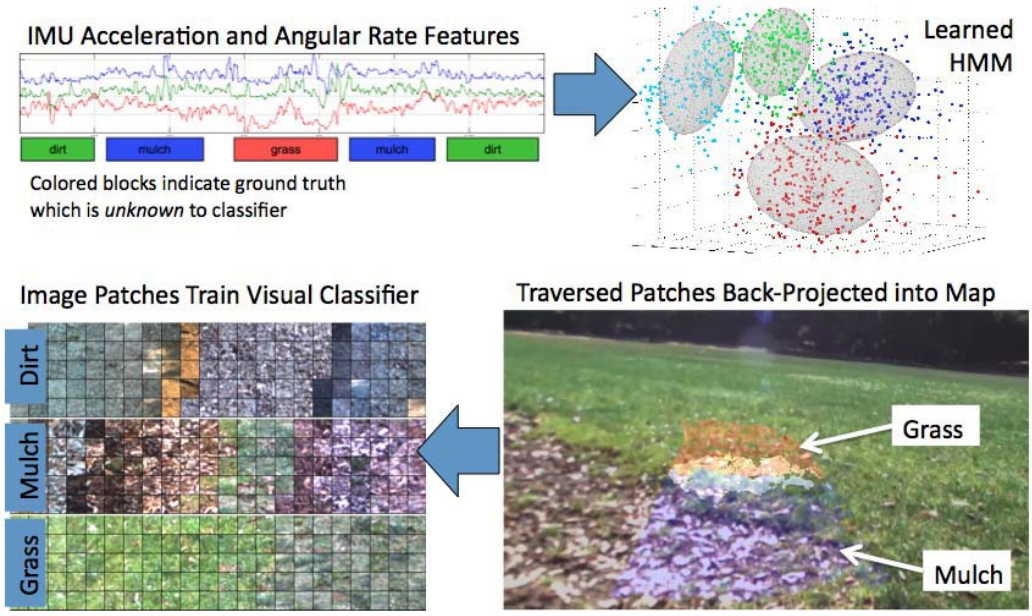

Figure 3-16: Systems operation for unsupervised near-to-far learning. Proprioceptive time series (top left) are clustered (top right) and the classified time intervals back-projected into imagery (bottom right). Catalogs of such image patches form a visual training set (bottom left).

We have shown that this near-to-far learning approach can automatically find the number and type of terrains (e.g., grass, dirt, mulch, and concrete) across a range of velocities. We have also demonstrated a support vector machine (SVM) visual classifier that is trained on the learned proprioceptive classes, runs at 3 to 5 Hz, and achieves 90 to 92% accuracy across these terrains. The main technical challenge we address in future work is to maintain classifier accuracy as the diversity of experimental conditions grows: vehicle load, terrain slope, and number of terrains. In FY12, we will integrate the proprioceptive and visual classifiers into a single system, and future research will advance system capabilities, proceeding along two axes.

Along the first axis, our work will expand the domain of applicability of the system to *more terrain classes*, specifically including vegetation, which is often compliant enough to traverse. In FY12, the existing stereo-vision-based capability will be augmented with a LADAR so that classified vegetation can be segmented and mapped. Additionally, beginning in FY12, we plan to use *co-training* approaches [MM10,Zh08] to jointly use both proprioceptive and visual signatures, which contain complementary information about underlying type, to perform the unsupervised classification.

The second axis of research, beginning in FY13, focuses on fundamental work to *ensure that the class decomposition remains consistent despite confounding variables*, which are known or can be estimated, chiefly vehicle load and slope of the load-bearing surface. We anticipate testing three approaches: 1) learning an empirical load-dependent proprioceptive signature by exploiting the diversity of logs; 2) using feature-normalization approaches to compensate for variable conditions; and 3) estimating Mohr-Coulomb parameters to characterize terrains independent of load. This is a fundamental activity that will carry over into FY14. Additionally, in FY14, we plan other activities to increase performance of the integrated system. Candidates include use of range or segmentation-based superpixels to group pixels into visual patches for classification or random-field-based spatial smoothing on visual classifier outputs.

This work feeds in to the construction of a highly detailed, near-field to far-field world model, with the specific capability of contributing terrain traversability and terrain semantic classifications to the world model. It also provides perceptual information on terrain characteristics in support of the FSU P2 capability to perform energy-minimizing traverses. It provides semantic classification in support of hands-off interfaces to the robot.

Metrics for Evaluation:
For both the proprioceptive and visual classifiers, primary metrics are error rates, ROC (probability of detection versus probability of false alarm) curves, and confusion matrices. The ground truth needed for these metrics is determined from an enhanced sensor package, such as higher-resolution cameras and LADAR. Current state-of-the-art error rates for the visual classifier are around 8%, and error rates for the proprioceptive clustering are around 10%. To measure robustness properties of the clustering, we use stability of the clustering objective function [Be06] as well as information-theoretic statistics like Normalized Information Distance. Besides these cumulative metrics, we also use targeted diagnostics like catalogs of erroneously labeled image-patches, videos of imagery overlaid with terrain-classification masks, and terrain-classification maps from an overhead perspective.

Intermediate milestones for FY12: Maintain 90% accuracy of the integrated system in terrain conditions that include vegetation; and for FY13: Maintain 90% system accuracy across a range of vehicle load or slope.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Integrate LADAR, gather field data in vegetation, and use LADAR to put vegetation into a map.
Q2	Learn proprioceptive classes for dense vegetation and other types. Develop co-training approach.
Q3	Learn visual classes for dense vegetation, estimate system error characteristics.
Q4	Introduce co-training into visual and proprioceptive clustering to lower system error rate to target of < 10% end-to-end.

References:

[Ba08] M. Bajracharya et al., "Learning long-range terrain classification for autonomous navigation." *ICRA*, 2008, pp. 4018–4024.

[Be06] S. Ben-David, U. von Luxburg, and D. Pál, "A Sober Look at Clustering Stability," COLT 2006, 5–19.

[CC08] E. Coyle and E. G. Collins. "A comparison of classifier performance for vibration-based terrain classification," 26th Army Science Conf., 2008.

[DVH04] C. Dima, N. Vandapel, and M. Hebert. "Classifier fusion for outdoor obstacle detection." *ICRA*, 2004.

[H97] J. Huang et al., "Image indexing using color correlograms." *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 1997, pp. 762–768.

[MB06] M.Maimone, J.Biesiadecki, E.Tunstel, Y.Cheng, and C.Leger. "Surface navigation and mobility intelligence on the Mars exploration rovers." *Intelligence for Space Robotics*, 2006.

[MC05] R. Manduchi, A. Castano, A. Talukder, and L.H. Matthies. "Obstacle detection and terrain classification for autonomous off-road navigation." *Autonomous Robots*, 18:81–102, 2005.

[MM10] G. S. Mann and A. McCallum, "Generalized Expectation Criteria for Semi-Supervised Learning with Weakly Labeled Data," *Journal of Machine Learning Research* 11 (2010), 955-984.

[WI09] C. Ward and K. Iagnemma. "Speed independent vibration-based terrain classification for passenger vehicles." *Vehicle System Dynamics*, 47(9): 1095–1113, September 2009.

[WS03] C. Wellington and A. Stentz. "Learning predictions of the load-bearing surface for autonomous rough-terrain navigation in vegetation." *International Conf. on Field and Service Robotics* (FSR), 2003.

[Zh08] X. Zhu, "Semi-Supervised Learning Literature Survey," C.S. TR 1530, University of Wisconsin, Madison, 2005, revised 2008.

Subtask 3: Online Recognition in 3D Point Clouds (CMU)

The objective of this subtask is to generate automatic labeling of sensor data with semantic classes from 3D point clouds (from LADAR or stereo) in online operation.

Objective and Benefits:

This applied task focuses on the long-term objectives:

1) Processing of 3D point clouds fast enough for on-robot operation from sensors, including 3D scanners and passive stereo.
2) Integration of multiple sources of sensor data.
3) Formulation and implementation of interpretation algorithms for online, incremental processing on different platforms.
4) Develop practical implementations of learning tools to adapt the system to new platforms and sensors (longer-term objective not addressed in FY12).

In FY11, we showed labeling on large 3D point clouds/color datasets *offline* and started migration of the algorithms to online processing. In FY12, we aim for 1) full implementation of online classification 3D/color with interface to world model; 2) additional components specific to 3D and UGV scenarios such as visibility and occlusion reasoning; and 3) migration of next-generation algorithms from P3 and Intelligence tasks.

The task contributes to building a semantic labeling to populate the WM. In assessment scenarios, it contributes to any of the scenario components that require querying the world model for named elements (e.g., building, doors), such as element 3, "Understand the world through perception and the world model." We expect the staged migration to WM interfaces and capstone elements following roughly the following progression:

- Small number of labels to WM: The first version uses a small number of labels (e.g., walls, cars) and provides to WM a 3D point cloud with regions annotated by their labels. This stage connects to item 3 of the Capstone document in which initially labeling of street/walls is needed.
- Label distribution to WM: The natural next step is to generate distributions of labels, thus enabling ambiguous interpretation and multiple hypotheses reasoning. Successive versions and experiments will handle increasing numbers of labels for more complex scene descriptions. This stage connects to items 4 through 6 of the Capstone document in which more object classes (and accordingly increased ambiguity and errors in classification) need to be used.
- Directives from WM: Interfaces in the reverse direction involve WM generating focus of attention directives to a) limit the part of the sensor data to be processed and b) limit the number of labels. Both are critical in taming the computation time issues and in increasing accuracy by focusing on the information needed for the current task at hand. This part of the interface implements mission context. This stage connects to items 6 through 10 of the Capstone Assessment Vision in which replanning, focus on surveilled area, etc. constrain the operation of perception.
- Other knowledge from WM: Longer-term, we expect perception to be required to re-train models or to incorporate prior knowledge from external datasets based on directives from WM. This will not be addressed in FY12 beyond initial design with I2. This part of the work will start addressing these of context from external knowledge sources.

State of the Art:
Inferring the labels based only on local 3D features is very difficult for a variety of reasons. For example, the viewpoint from which objects are perceived can widely vary, the sensor irregularly samples points from objects, and there is often local ambiguity in appearance. Progress has been made in fundamental research to address this problem in the general area of 3D scene analysis in many types of environments, including street-level, indoor, and aerial. Existing approaches include statistical graphical models and memory-based learning approaches. In both cases, the techniques do not scale well to large data and online processing. Furthermore, they cannot easily be used to learn how to combine features from multiple modalities from training data. In contrast, we use deep inference, in which the original (intractable) inference problem is broken down in a series of smaller inference problems, and we extend these approaches to 3D point clouds with new features and feature learning techniques better adapted to fusion of 3D point clouds with other modalities. FY11 results show (offline) the result of processing large data with good performance with respect to SoA.

Key Barriers and Approach:
Our approach and the corresponding key challenges to the three FY12 objectives defined above are described below. In addition, we will continue to develop and experiment with versions of the perception system with different sensors, in particular continuing the efforts in using stereo and ARL LADAR.

(1) *Implementation of online classification 3D/color with interface to world model*: Our aim is to be able to use the classification algorithms on the robot on the move, at a fast enough rate to keep-up with sensor data and with the rest of the architecture. This requires progress in the fundamental redesign of the algorithms. The first direction is the design of algorithms for streaming 3D points. Specifically, by the end of FY11, we will have implemented a "moving map" version of the system for classification from 3D point clouds. We will generalize this approach to handle fully incremental computation, i.e., the ability to add new data points with the minimal amount of computation required. Following the FY11 approach, we will complete first the 3D processing part and then add the image information. The latter involves incorporating the image features and segmentation in the online algorithm. A final step (to be considered in FY13) is to incorporate temporal consistency, i.e., propagating labels from one frame to the next in a video/3D sequence. We will investigate this addition based on results of the 6.1 work in P3 and other related projects.

During our FY11 work, we identified the computation of the local features as the main bottleneck in terms of computation, especially for the image features. In fact, even if we reduce the image segmentation time to almost zero by replacing the expensive state of the art region segmentation by a simple pyramid, we found the computation time to be quite substantial. This is due in large part to new developments in feature design which boost the classification accuracy but slow down the processing considerably. Our first direction of investigation is to design strategies for intelligent sampling to reduce the number of atomic features computed. For example, one feature requires a SIFT vector to be computed at every node of a dense grid over the image. In fact, high density may not be required, and we might be able to subsample the grid based on image content. The same applies to other features. Our second direction is to use and improve on the methods used to encode the atomic features, e.g., vector quantization of the features. We propose to leverage the extensive prior work on effective coding for vision to develop more efficient algorithms. Finally, we will investigate integrating fast implementation of the low-level features (e.g., SIFT, HoG) on GPUs or FPGAs. This activity will be conducted in close consultation with our GDRS partners.

A key aspect of online processing is the ability to continuously produce labels to be provided to the world model (or more precisely, distributions of labels representing multiple hypotheses about the understanding of the environment). We will develop the necessary interfaces to the world model to provide this data. A richer set of interactions will take place in the reverse direction: The world model must be able to specify directives to this visual analysis module. The two most important directives are a) the area in which processing should focus and b) the limited number of labels that are of interest for the specific task at hand at this particular time. These are particularly important because they implement the use of mission context in perception and because they limit the computational load to what is required for the current task (e.g., only a

small number of semantic labels are useful at any given time). This activity will be driven by the world model design in I2.

(2) *Additional components specific to 3D and UGV scenarios such as visibility and occlusions reasoning*: Most of the work so far has focused on using the data directly generated by the sensors (image pixels and 3D points), without taking advantage of information that can be inferred from the 3D structure of the environment. The first step is to incorporate visibility information from the 3D measurements. We have conducted initial work in FY11 in incorporating that information in low-level 3D features, which we will extend in FY12. A much richer approach is to infer occlusions from combinations of observed data and other knowledge, e.g., prior models of objects from training data, information for the WM, or prior databases of 3D models. The basic idea would be to regress the occluded part from the observed data and the prior knowledge. This work leverages recent work in occlusion reasoning in object recognition and in use of 3D model repository in scene understanding. We plan to conduct initial work in this area in FY12.

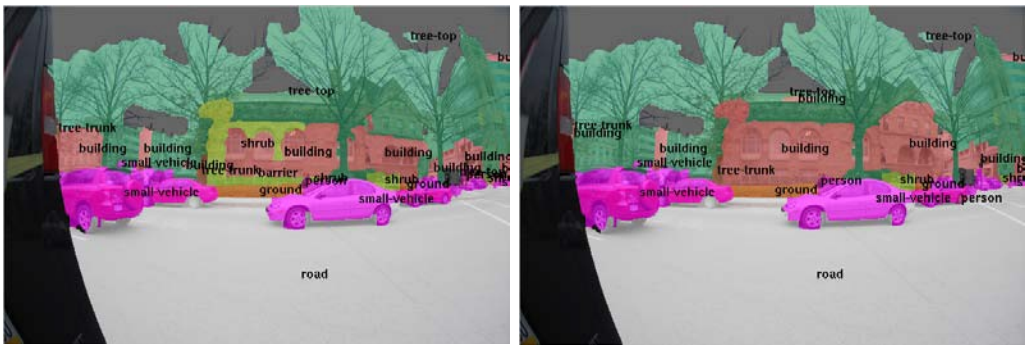

Figure 3-17: Output of 2D –only segmentation and the 2D+3D interpretation.

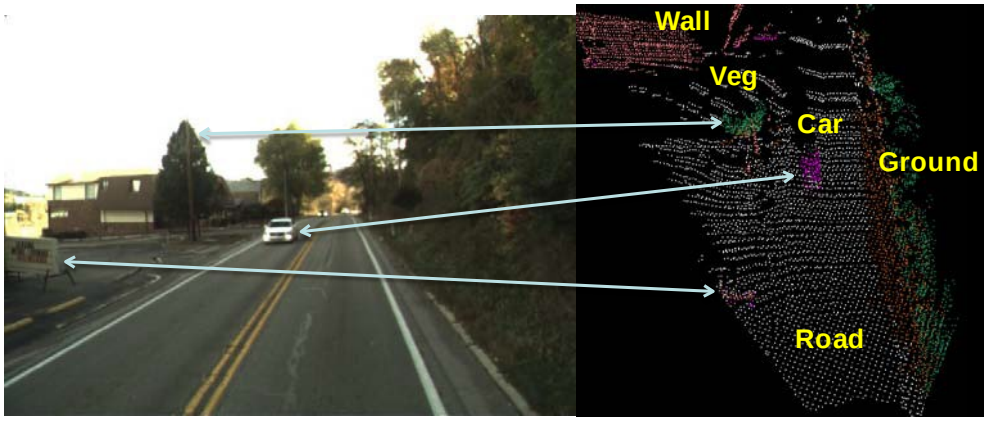

Figure 3-18: Example 3D classification from the Gascola dataset (30 labels, 10 miles, 230 Mpts, 1570 images).

(3) Migration of next-generation algorithms from P3 and Intelligence tasks: Much of the work in this 6.2 task involves the integration of the basic algorithms developed in P3 and the machine learning tools developed in the Intelligence tasks. In particular, so far we have worked on integrating online versions of the 3D point cloud labeling algorithms, but we have not yet

investigated the learning phase of the system. This will become important as we move across different platforms and sensors, requiring a *standard training procedure*. Also, new tools for using *unsupervised data* open-up new challenges in managing large amounts of data. Finally, long-term, we expect directives from the world model, e.g., indication that the system has poor performance on a specific object class, to trigger *dynamic training* of part of the system. We will start investigating these issues in FY12, in consultation with I2 to design the necessary interfaces.

Successful completion of FY12 objectives relies on the results of the FY11 work: the basic classification tools were implemented and tested successfully in FY11. They were compared with SoA. Computation time shows feasibility of online processing. More precisely, we produced a complete implementation of 3D/color labeling using direct fusion. We collected and annotated large (10 miles) 3D/color datasets and evaluated the current implementation. We developed and integrated additional 3D features and segmentation tools. We started the design of the algorithms/data structures for online processing and the integration with other sensors (such as stereo).

This effort uses the fundamental classification tools from P3 machine learning tools developed in the Intelligence tasks; Use of 3D point clouds directly connected to JPL's efforts in P1 and P2 on passive sensing on small platforms and on difficult materials (free of texture, or transparent); Connection with ARL's research in MEMS scanner (initial work with ARL conducted in FY11; will continue in FY12); The output of this task feeds into the world model in I2. The techniques used in this task are complementary to those of the other Perception tasks focusing on narrower semantic classes. For example, UPenn's P4-4 on recognizing objects from silhouettes for manipulation focuses on recognizing specific objects; Caltech's P3 subtask focuses on category and subordinate category recognition; and JPL's P4-2 addresses fine-grained terrain classification.

Metrics for Evaluation:
We use two metrics: 1) classification accuracy, embodied in ROC and P/R curves evaluated on a) standard SoA benchmarks and b) large annotated datasets collected in FY11; and 2) computation time.

In the short term, we will measure intermediate progress as 1) progress in increasing throughput (real-time and online processing) and 2) classification accuracy on data from FTIG assessment site (data collection underway).

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Implementation of incremental computation with 3D and image input. Interface to WM for individual labels.
Q2	Optimization of feature computation. Interface to WM for label distributions.
Q3	Interface to WM for focus of attention directives. Refinement of classification from other data sources (e.g., stereo, ARL scanner).
Q4	Testing of integrated system. Study of occlusion reasoning and initial algorithms. Design of re-training.

Subtask 4: Rare Events and Motion-saliency Based People and Scene-context Recognition (UPenn)

Objective and Benefits:
This task considers the use of information from the dynamic part of the environment as contextual cues to help in detailed understanding of the static part of the environment. Specifically, we focus on detecting "rare" action events in video that may be discriminative for scene attribute recognition and extract detailed attributes of pedestrians, vehicles, and background for improved feature representations. For example, building classification performance may be improved by including rare or unusual activities with detailed attributes: daily 18-wheeler deliveries to a grocery store or construction trucks in the parking lot of a hardware store. The main challenge is detecting such unusual events in days of video. Salient motion can also drive our attention to certain parts of the video, creating "hot spots" for scene context recognition. This is a multi-year effort.

State of the Art:
The overall goal of this subtask is to develop attribute maps for online people and scene context recognition in urban video. The primary research challenges for this subtask are ones of representation and generalization. Most existing work has focused on using static image features (2D and 3D) to infer scene context recognition; for example, recognizing which part of the road is more suitable for driving versus walking. In contrast, our approach utilizes dynamic and static video information to recognize scene context.

Key Barriers and Approach:
The key challenge is that "rare" events are difficult to predict and highly context dependent on scene-activity-people involved. This feature of "rare" event makes it highly informative for scene context recognition. Our approach is to use spectral graph co-clustering to detect "rare" activities in video. No complex activity models and no supervised feature selections are used. We divide the video into equal length segments and classify the extracted features into prototypes, from which a prototype–segment co-occurrence matrix is computed. Then, we pose the unusual event detection as a co-clustering of segments and prototypes, for which an efficient, globally optimal algorithm exists using spectral methods.

The goal of this development was to provide tools to locate rare activities in a long video that may be useful for classification. These small events, such as delivery personnel arrival/departure, loitering, window-shopping, double-parking, bus arrival/departure, and pedestrians walking in unusual locations, are difficult to detect in long videos due to the frequency of events occurring. The output of this pipeline is a set of *K* clusters of unusual segments, where a *segment* is an interval of *N* sequential frames. These segments contain unusual or rare activity patterns that may be useful for building classifiers. This effort complements efforts from CMU on scene understanding using image cues, by addition of human activity information. We will also benefit from better scene understanding to include spatial context for action recognition.

- **Approach**
  - Construct localized descriptors for all motions on "drivable" region
  - Histogram of Oriented Gradients (HoG) descriptor
  - Run algorithm to detect "unusual vehicles" locally to roadways
- **Results**
  - Detected delivery truck, contractor truck, landscape truck, garbage truck, police car, street sweeper
  - Vehicles entering/exiting parking spaces
  - Pedestrians entering/exiting vehicles
  - Vehicles performing U-turns

Metrics for Evaluation:
For this task, we have been working on the Ocean City dataset. Our system is able to detect the rare events of cars double-parked at night with people exiting; cars leaving parking spots with people exiting cars and walking across the street; a bicyclist and pedestrians crossing the street at unusual locations; cars parking and people exiting cars to enter buildings. We have also used human action information to provide a background attribute map based on motion profiles. This attribute map provides the likelihood for "roadway" and "sidewalk." This approach used the same software for detecting unusual events, and instead detected usual or common events.

We will collect an additional dataset of a similar type. We will measure both the "rare" event detection ability and how such "rare" events can improve street scene context recognition.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Unsupervised rare events detection in video to construct initial code book of activity patterns.
Q2	Construct discriminative classifier for detecting rare activity patterns.
Q3	Construct discriminative classifier for street scene classification using dynamic rare features.
Q4	Construct classifier for street scene classification using dynamic and static features.

Subtask 5: Integration of Human Detection Capability for Ground Robots and Data Collection (ARL)

Objective and Benefits:
The proposed subtask objective is to develop and integrate real-time human detection capability for small ground robots. This work will supply accurate and robust human detection, specifically optimized for robotics sensors, to subsequent RCTA perception algorithms, such as activity recognition and pattern of life determination. We will perform a data collection as well using small robots to obtain a set of imagery that we can use to train and optimize the real-time human detection algorithm.

State of the Art:
State of the art in human detection approaches mainly use histogram of oriented gradients (HOG) as features for human detection. The most recent and well-known approaches include Felzenszwalb's part-based detector and Yan's HOG-local binary pattern detector. However, these approaches are computationally expensive and impractical for small ground robots due to size, weight, and power constraints. Furthermore, these algorithms are limited to a range of human sizes (i.e., if human size is <50 pixels or too large within the image, detection performance severely deteriorates).

Key Barriers and Approach:
ARL is researching and developing robust human detection algorithms for integration on small robots. The human detection algorithms will be capable of operating in real-time within the size, weight, and power constraints of small robots. Such algorithms can provide robust human detections for subsequent perception stages, such as activity recognition and pattern of life determination. Figure 3-19 shows some sample human detections using ARL's present human detection algorithm which yields accurate detection with few false alarms.

Figure 3-19: Human detections in sample imagery.

Although there are no insurmountable problems to integrating the human detection algorithm on small ground robots, there are several challenges/barriers. One barrier is real-time human detection approaching the frame rate of sensors mounted on robots (typically 1 to 30 frames per second). We will overcome this barrier by implementing the algorithm in C/C++ for GPU processing and, furthermore, developing a cascade of classifiers approach that can reduce runtime while decreasing false alarms. Another barrier is the variation in human sizes that could appear in imagery acquired by small ground robots (e.g., data collected on a human at a distance versus data for a close-in, large image of that person). We will investigate other features to complement HOG features for robust detection with respect to human size.

Metrics for Evaluation:
We will assess our task progress against the below schedule.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Subtask P4-5 is part of a proposed collaboration between GDRS and Carnegie Mellon and Alex Chan, Shuowen Hu, and Prudhvi Gurram at ARL.

	Subtask 5 Quarterly Milestones and Deliverables
Q1	Develop robust HOG-based approach for human detection.
Q2	Collect set of data using small robots for algorithm training and optimization.
Q3	Implement approach in C/C++ for CPU and GPU for real-time capability.
Q4	Perform tests, and start integration with perception algorithms.

P5-2012 – Dynamic Scene Understanding

Level: 6.1 Basic Research

	Name	*Organization*	*Subtask*
Principal Investigator	Martial Hebert	CMU	P5-1
Additional Investigators	Luis Navarro-Serment	CMU	P5-1
	Drew Bagnell	CMU	P5-1
	Jeff Cohn	CMU/UPitt	P5-1
	Michael Ryoo	JPL	P5-2
	Larry Matthies	JPL	P5-2

Objective:
Our objective is to investigate robust approaches to detecting, tracking, and identifying objects in general configuration, to use the resulting intermediate description to identify behaviors of individuals and groups (Subtask 2), and to predict distributions of likely behaviors based on learned models (Subtask 1). In addition to the motion cues traditionally used in these approaches, we incorporate features that describe the interactions of the people with their environments, with each other, and with the robot. Finally, we design techniques to identify finer-grained behavioral cues such as facial cues, which are critical to refining behavior analysis and predicting behaviors/intent.

Background:
Robots must understand the dynamic nature of the real world in which people and vehicles are acting in their environment. Much of the work to date has focused on understanding the motion of objects for the purpose of safe navigation, e.g., detecting and tracking moving objects in LADAR data or in image sequences. For complex missions, especially in urban environments in which groups of people may interfere with their progress, robots need a more detailed semantic model including predictions of people's intent and future behaviors as well as the ability to recognize detailed behavioral cues. While a vast literature exists for short-term prediction based on classical temporal filtering techniques, longer-term prediction of motions and behaviors remains open, and recognition of behavioral cues has been largely limited to well-structured environments.

Our approach to activity detection and understanding jointly models dynamic (motions, behaviors) and static (objects, surfaces) aspects of the robot's environment. We investigate robust techniques to identify current behaviors of dynamic objects, predict their future behaviors, and derive detailed descriptions of behavioral cues.

We made progress in FY11 in basic algorithms for prediction and their integration with the semantic interpretation algorithms of P3 and P6, and we had conducted initial studies of the feasibility of using facial cues. In addition to extending these techniques to more complete models of activities, we introduce two important directions of research for FY12.

Recognizing that the work thus far focused mostly on the behavior of individual people, we introduce a new task focusing on understanding interactions between people and between people and the robot. Of particular importance to the capstone assessment scenarios is the ability to recognize actions directed at the robot and to understand reactions of people to the robot, in particular identifying threats.

Research Description:
While SoA is concerned primarily with tracking for mobility in dynamic environments, this task goes further by exploring tools for deriving a semantic representation of the environment beyond describing motion. Specifically, this 6.1 task explores the fundamental tools for 1) generating prediction of likely sequences of motions and actions for partially observed sensor data, and 2) modeling and recognizing interactions between agents and robot. This task takes input from the 6.2 tasks in the form of robust tracking in video and LADAR data. The objective is for the fundamental algorithms developed in this task to migrate to capstone assessment and to Task P6 which addresses the 6.2 aspects of understanding dynamic environments. The connections with the other tasks and with the assessments are described at the end of this summary and in the detailed task descriptions.

Our technical approach in (1) is based on inverse optimal control developed in earlier work. In FY11, we showed that this technique could be used in predicting distributions of paths at a rate fast enough for use on the move, and we transitioned to integration under the corresponding 6.2 task. In FY12, we extend this approach to the more difficult case of mixing paths with sequences of discrete actions. This part of the work departs from the state of the art in two ways. First, current techniques for predictions, e.g., based on Kalman filters or Hidden Markov Models, can handle only short-horizon predictions, whereas the basic tools used here can handle forward-looking horizons since they use distributions over paths rather than over transitions. Second, incorporating discrete actions provides a much higher level of interpretation than in existing approaches which are mostly targeted to mobility in dynamic environments. In addition, we investigate the use of other cues which are not normally used in the UGV context.

To address (2), we develop a technique based on hierarchical modeling of cause-and-effect relationships composing activities. The method will automatically learn semantic structure of the activity in terms of spatio-temporal relationships among its sub-events, representing influences of the activities to the robot itself. We plan to formulate the hierarchical recognition problem as a Bayesian probabilistic inference. More specifically, conditional probability of the robot's movement (i.e., effect) given the action of the human (i.e., cause) will be estimated, and their grammar patterns will be learned. This approach is supported by extensive prior works which demonstrated potential use with hierarchical grammar-based representation and recognition of objects and activities. These SoA techniques have not been tested in the context of robot-targeted

activities. More generally, the proposed approach extends the SoA in which there has been a significant amount of previous research on human activity recognition from static third-person point-of-view cameras but not recognition of human activities from the point of view to a moving platform. The aforementioned previous works focused either on the recognition of human movements from a camera far above the actors (e.g., on top of a building) or on the estimation of simple movements of the camera.

Metrics:

We use two classes of quantitative metrics:

- *Sequence comparison*: This class of metrics is used to evaluate the performance of the prediction algorithms. The metrics evaluate the accuracy of the distribution of predicted sequences defined as combinations of paths and sequences of actions. More precisely, they compare the most likely sequence predicted with the actual observed sequence and compute the percentage of route and actions shared, and they evaluate the percentage of the observed sequences matching at least 90% with the predicted sequence.
- *Recognition accuracy*: Reliability (i.e., recognition accuracy) of the activity recognition activities will be measured in terms of precision-recall, computing true positives and false positives of detecting human activities from continuous video streams. To use these metrics, we will establish a new dataset. True positives as well as false positives will be measured, and precision-recall rates will be presented as the final result.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:

We anticipate that the research conducted in these subtasks will be reported in the major vision (CVPR, ECCV) and robotics (ICRA) conferences as appropriate given the deadlines. The PIs give regular invited talks at workshops and conferences in which portions of this work will be represented. Collaborative activity with ARL can be explored.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Design dynamic reward adaption. Increase set of labels in prediction model. Design the activity recognition algorithm, and explore motion features for the detection.
Q2	Evaluate adaptation on controlled scenarios. Incorporate goal information extracted by perception modules. Intermediate results on generic expression detection. Initial dataset collection (10 to 20 videos per activity with different conditions).
Q3	Evaluate performance on scenarios with complex features and goals (egress points, etc.). Completion of algorithms for incorporating discrete actions. Initial implementation of the methodology, and evaluation of the system on a classification task using segmented videos from the established dataset.
Q4	Evaluation of mixed trajectory/action prediction. First experimental evaluation of parameter influence on social interactions. Improve the algorithm to increase its reliability, and test it on a detection task using continuous streams from the dataset.

Related Research:

Several Perception subtasks address aspects of the problem of understanding dynamic environments. The relationships amongst them (and the rationale for selecting this complement of subtasks) are as follows:

- P5-1 and P5-2 address two complementary aspects of the problem. P5-1 addresses the problem of forecasting sequences of actions based on current observations and learned models and on using high-resolution cues such as facial cues. As part of its scope, P5-1 investigates models of interaction between agents and static environment, while P5-2 focuses on interactions between agents and agent/robot.
- The P6 task addresses integration issues, including robustness issues for operation in cluttered environments.
- The prediction techniques of P5-1 are transferred to P6-3 for integration with robust tracking and other components.
- Conversely, P6-1 provides robust tracking tools that generate tracking information used in all of the P5 subtasks.
- P6-4 addresses a similar prediction problem as P5-1 but with a completely different method based on forward planning.

This task ties to RCTA tasks outside of the perception area. Integrating descriptions of the static environment in P5-1 ties directly with all of the semantic labeling tasks in P3 and P4. The high-resolution cues will be used as input to HRI functions. All of the tools for dynamic situation awareness will be used to provide information to the world model, Task I1, and be integrated in the architecture developed in Task I2.

Integration and Assessment Activity:
The P6 task addresses all the integration issues associated with the modeling of dynamic environments. Most of the algorithms developed in this basic research task will feed into P6 for integration and for further work on the robustness issues associated with complex environments. This mode of transition has already happened in FY11. The prediction algorithms developed in P5-1 formed the basis for the system in P6-3 which demonstrated the generation of prediction distributions on the UPenn K-bot, the CMU car, and a Pioneer testbed, using e.g., the Hokuyo scanner. We expect this level of interaction between the two Tasks P5 and P6 to continue as a mechanism to migrate the capabilities toward the capstone assessments.

Subtask 1: Behavior Prediction from Motion and Other Cues (CMU)
The objective of this subtask is to develop algorithms for behavior prediction in dynamic environments. The output of this task is to be used for planning in dynamic environments and for behavior recognition and forecasting (e.g., in the I1 architecture).

Objective and Benefits:
The objective of this subtask is to develop algorithms for behavior prediction that include 1) the ability to dynamically adapt to the environment; 2) detailed semantic description of the environment; and 3) observation uncertainty and discrete actions in addition to continuous motion. In addition, we 4) develop high-resolution cues useful for prediction.

This effort addresses the need for a semantic description of dynamic environments. This research is particularly suited to support capabilities 4, 6, 8, and 9 of the Capstone Assessment Vision.

State of the Art:
In the area of predictive models, a vast literature exists on prediction based on classical temporal filtering techniques. Approaches include Kalman Filter approaches, which model parametrically the motion assuming Gaussian noise and Gaussian uncertainty on state estimates and non-parametric approach like Hidden Markov Models (HMMs) and particle filters. These approaches model the likelihood of transitions between states based from training data. All of these classical approaches have two fundamental limitations. First, because they model only local transitions (as in HMMs) or, essentially, time derivatives (as in Kalman filters), they are unable to propagate prediction forward in time further than a short time horizon. Second, they do not provide convenient ways to account for higher-level knowledge about interactions between motion patterns and environments.

To address the limitations of the standard, state-of-the-art techniques, recent work leveraged concepts developed in the context of reinforcement learning and imitation learning. These techniques form the basis of the approach used in this task. They address these limitations by learning a model of the distribution of paths given observations rather than individual transitions and by learning from training data relationships between path cost and features in the environment. They take into account explicitly preferences in motion models such as destinations and preferred routes. In essence, they model the agent as moving with a purpose as opposed to following a random walk with a local transition model.

Key Barriers and Approach:
Described below is our approach to the four objectives previously mentioned and the corresponding challenges:

(1) *Dynamic adaptation*: Our current approach uses a reward function to compute the prediction of pedestrian trajectories. This function is obtained after training from a number of examples and is capable of capturing the features in the environment that best explain the pedestrian behavior. The reward function is selected for a given environment and remains constant during normal operation. However, if the set of environmental features that were abundantly present in the training examples is sparsely present in the environment, the accuracy of predictions will deteriorate. Likewise, a constant reward function may not be able to take into account additional cues from the environment when they become available (e.g., a semantic cue from a Soldier identifying a landmark not detected by the robot's sensors). To overcome this limitation, we plan to develop algorithms to perform dynamical identification of prediction accuracy and a set of measures of feature richness. We also plan to develop algorithms to incorporate semantic labels into our prediction algorithm. This approach ties with research done in semantic labeling and robust detection and tracking in dynamic environments. For example, this research is particularly relevant to the parts of the capstone scenarios in which new information is provided by the higher level of cognition or by the human team members. In those cases, the configuration of the prediction algorithm needs to be dynamically adapted.

(2) *Detailed description of static environment:* In FY11, we showed how semantic labels from the static environment could be incorporated in the prediction model for a small set of categories (e.g., walls). Figure 3-20 shows one example from the FY11 work. We plan to develop algorithms for a larger set of classes as well as more detailed information such as specific objects (UPenn), door/stairway (JPL), and subordinate categories (Caltech), by extending the feature representation and the learning algorithm. Extensions to the feature representation are necessary to handle the effect of distant label in the scene description. For example, the location of a window might affect behavior of an agent trying to stay out of line-of-sight even though the agent might be arbitrarily far. The learning algorithm must be extended to account for different types of goals and for the different features. This aspect of the work is important in handling situations (as in the capstone scenario) in which egress points and detailed architectural features play a central role in understanding the behaviors of interest to the mission.

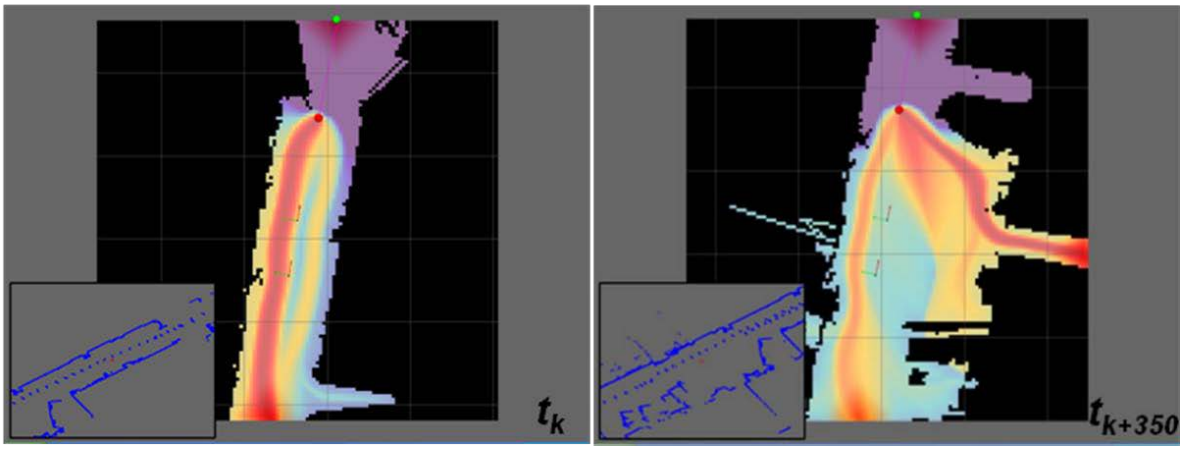

Figure 3-20: Prediction distribution from a partial trajectory observation (red is higher probability) using the walls as additional features.

(3) *Observation uncertainty and discrete actions:* The existing algorithms work directly with motion in the plane. We need to handle discrete actions (e.g., pick up object, look up, etc.). We plan to use recent results in HIOC (Hidden state Inverse Optimal Control), which generalize the current approach to handle uncertainty and representations such as discrete actions. In that approach, we model human activities as a sequential decision-making process (a Markov decision process) and solve the inverse optimal control problem to recover the so-called reward function that guides human activity. Specifically, since we assume a vision-based system, we do not have access to the true state, true action, or the true environmental attributes. Instead, we only have noisy tracking results (state and action of the actor) and noisy vision-based observations of the environment (e.g., distributions of labels from the semantic segmentation). Our hidden variable inverse optimal control (HIOC) algorithm deals with sensor uncertainty by adding all observations as the features and allows the IOC framework to assign the proper weights (confidence) to each of the observations. In related work, we have shown results on surveillance videos and demonstrated how the underlying reward function that drives human activity can be recovered from noisy vision inputs. Our objective is to show how this approach can be used to extend the current prediction methodology to handle observation uncertainty and symbolic actions.

(4) *High-resolution cues:* Much of the information needed to understand behavior and intent is far more detailed than motion only; it includes all the aspects of facial and body expression. We continue to address this area in this task with the objective of combining these cues with the overall prediction approach and combining with the other behavior and interaction cues developed in JPL's new P5 task. This work is conducted by Jeff Cohn (CMU adjunct) at the University of Pittsburgh. We plan to make progress in two areas; the first involves further work on automatic expression recognition. The work on person-dependent expression and tracking shown at the October 2011 PI meeting could be extended to person-independent expression detection. Automatic generic expression detection would be an important capability for robots in social groups. The second is computational modeling of social interaction. In computer-mediated interactions, parameters such as camera placement and signal delay are likely to influence social perception and consequent feelings of rapport and trust. To understand the influence of such

constraints on social interaction, we propose to experimentally investigate the influence of camera orientation and transmission variability on measures of movement synchrony, self-reported rapport, and indices of trustworthiness. A first example experiment is to manipulate the perception of the partner's gaze by varying effective camera orientation without the knowledge of the participants. We hypothesize that movement synchrony (e.g., back-channeling), self-reported rapport, and readiness to trust the other person will be higher when shared visual regard is possible. In a second example, variability of audio-video delay will be varied, again without either interlocutor's awareness that perception of their behavior has been altered. Observations and some theory suggest that variable delays will attenuate movement synchrony, self-reported rapport, and trustworthiness. Our long-term goals are to develop computational models of social interaction with people, agents, and social robots.

Success of this subtask is based on the promising results from FY11. The algorithms have been tested extensively on several testbeds, including on unsupervised data from street scenes using the CMU testbed. Initial implementations of the planned extensions have been tested in the context of video analysis.

Integrating descriptions of the static environment ties directly with all the semantic labeling tasks in P3 and P4. The prediction techniques are transferred to P6 for integration with robust tracking and other components. The high-resolution cues will be used as input to HRI functions.

Metrics for Evaluation:
To evaluate prediction performance, we plan to use two metrics. The first individually compares the most likely sequence predicted with the actual observed sequence and computes the percentage of sequence distance shared. The second is the percentage of the observed sequences matching at least 90% with the predicted sequence. The baseline system uses trajectory prediction in which the sequences are paths in the environment. As we progressively add capabilities in dealing with discrete actions, we augment the sequences with discrete elements. The metric for computational performance is the time it takes to calculate a prediction. We measure short-term progress by the increased complexity of the environment description, measured as the number of labels and the accuracy of predictions as defined above.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Design dynamic reward adaption. Increase set of labels in prediction model.
Q2	Evaluate adaptation on controlled scenarios. Incorporate goal information extracted by perception modules. Intermediate results on generic expression detection.
Q3	Evaluate performance on scenarios with complex features and goals (egress points, etc.). Completion of algorithms for incorporating discrete actions.
Q4	Evaluation of mixed trajectory/action prediction. First experimental evaluation of parameter influence on social interactions.

Subtask 2: Understanding of Human Interactions and Reactions (JPL)
The goal of this subtask is to provide a robot semantic understanding of events and activities occurring in its environment. Based on video perceptions from cameras, the robot will infer ongoing behaviors of nearby persons including intended activities of its team members, hostile reactions of insurgents, and actions of other civilians. Such ability will not only enable more natural interactions between the robot and humans (HRI) by making the robot aware of surrounding events and their context but also allow the robot to perform an intelligent mission by recognizing intention of unknown persons/group from their reactions to the robot (e.g., insurgents versus civilians).

Objective and Benefits:
Our desire is to develop a mathematical framework for the recognition of high-level human activities performed to the robot. More specifically, the objective is to address the following three research challenges toward the construction of a robot semantically aware of complex human behaviors:

1) Recognition of actions targeted to the robot from a robot-mounted camera.
2) Recognition of human reactions triggered by the robot's existence and/or actions.
3) Context-aware recognition of human intentions (e.g., hostile intention to damage the robot) based on their reactions.

Our plan is to focus on each of these challenges individually with three years of research. Thus, the goal of this subtask in FY12 is the recognition of human activities targeted to the robot from first-person point-of-view videos. These include friendly actions (e.g., waving and talking) as well as hostile activities such as punching the robot, throwing a stone to the robot, and kidnapping the robot. Ultimately, after three years of research, this subtask will provide a methodology necessary for the robot to automatically perform human intention-aware missions, such as guarding the back door of a building by observing a potential insurgent's behavior and recognizing his/her intention of escaping.

State of the Art:
There has been a significant amount of previous research on human activity recognition from static third-person point-of-view cameras (e.g., CCTV cameras) [1]. Particularly, in the work done by Ryoo and Aggarwal [2], stochastic activities performed by multiple groups were recognized. In their work, a video-based matching of group activities (e.g., two groups having a fight and a group of thieves stealing an object from owners) was proposed, illustrating potential for an automated understanding of military operations and group behaviors from videos. In addition, there have been recent attempts to recognize ego-action during sport plays from a wearable camera [3].

However, recognition of human activities targeted to a moving platform has not been studied in depth. The aforementioned previous works focused either on the recognition of human movements from a camera far above the actors (e.g., on top of a building) or on the estimation of simple movements of the camera. In addition, the previous works made little consideration on human intentions and reactive behaviors, making them inappropriate for robots actively exploring their environment.

Key Barriers and Approach:
The key contribution of this subtask in FY12 will be the development of methodologies to recognize human activities targeted to the robot, which has not been attempted previously. The system will recognize friendly/neutral activities that may occur commonly in an urban environment, including a human waving, pointing, and calling another person to show the robot. In addition, hostile activities will be recognized.

The essence of our approach will be the hierarchical modeling of cause-and-effect relationships composing activities. The method will automatically learn semantic structure of the activity in terms of spatio-temporal relationships among its sub-events, representing influences of the activities to the robot itself. For example, the activity of a person throwing a stone to the robot will be decomposed into the following three sub-events: a throwing action (i.e., arm swing motion) of the person, the event of a stone flying to the robot, and the action of the robot being hit as a consequence. Conventional motion recognition features (e.g., spatio-temporal local features) for analyzing target movements as well as optical flow-type features for the estimation of the robot movements caused by the target will be considered for the recognition of each of the atomic-level actions.

We plan to formulate the hierarchical recognition problem as a Bayesian probabilistic inference. More specifically, conditional probability of the robot's movement (i.e., effect) given the action of the human (i.e., cause) will be estimated, and their grammar patterns will be learned, constructing activity representations for robots. The preliminary works [2,4] showed a potential with hierarchical grammar-based representation and recognition of objects and activities, but they have not been tested with robot-targeted activities. Our direction is to extend the paradigm of stochastic grammar-based recognition for the robot recognition of activities.

Our previous experience in recognition of human-human interactions using sequential structure representation [5] will benefit our research. Our approach will also benefit from the Subtask P6-1; 3D object segmentation and tracking methodology will provide locations of humans and objects that serve as cues for activity recognition.

Metrics for Evaluation:
Reliability (i.e., recognition accuracy) of the system will be measured in terms of precision-recall, computing true positives and false positives of detecting human activities from continuous video streams. Particularly, in FY12, we will focus on an offline evaluation of the human activity recognition approaches. The idea is to test the validity of the developed methodology by establishing a new dataset. The videos of humans performing activities to the robot will be collected in several different conditions. This will be done by mounting a camera to an imitated model of the robot and making human participants perform activities. The videos of pedestrians will also be collected to measure false positives. At the initial stage, the system's classification accuracy using segmented videos will be measured. Later, the system will be asked to locate all ongoing activities from the continuous video streams. True positives as well as false positives will be measured, and precision-recall rates will be presented as the final result.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Design the activity recognition algorithm, and explore motion features for the detection.
Q2	Initial dataset collection (10 to 20 videos per activity with different conditions).
Q3	Initial implementation of the methodology, and evaluation of the system on a classification task using segmented videos from the established dataset.
Q4	Improve the algorithm to increase its reliability, and test it on a detection task using continuous streams from the dataset.

References:

[1] J. K. Aggarwal and M. S. Ryoo, "Human Activity Analysis: A Review", ACM Computing Surveys (CSUR), 43(3), April 2011.

[2] M. S. Ryoo and J. K. Aggarwal, "Stochastic Representation and Recognition of High-level Group Activities", International Journal of Computer Vision (IJCV), 93(2):183-200, June 2011.

[3] Kris M. Kitani, Takahiro Okabe, Yoichi Sato, Akihiro Sugimoto, "Fast Unsupervised Ego-Action Learning for First-Person Sports Videos", CVPR 2011.

[4] R. Girshick, P. Felzenszwalb, D. McAllester, "Object Detection with Grammar Models", NIPS 2011. [5] M. S. Ryoo, "Human Activity Prediction: Early Recognition of Ongoing Activities from Streaming Videos", ICCV 2011.

P6-2012 – Perception for Missions in Dynamic Environments

Level: *6.2 Applied Research*

	Name	*Organization*	*Subtask*
Principal Investigator	Mike Wellfare	GDRS	P6-2
Additional Investigators	Mike Hoffelder	GDRS	P6-2
	Max Bajracharya	JPL	P6-1
	Martial Hebert	CMU	P6-3
	Luis Navarro-Serment	CMU	P6-3
	Jianbo Shi	UPenn	P6-4
	Kostas Daniilidis	UPenn	P6-4

Objective:
The objective of this task is to robustly detect, track, and predict future locations of targets of interest in highly cluttered and dynamically changing environments. Targets are primarily humans and vehicles, although can include other semantic objects labeled in Tasks P3 and P4. Detection and tracking need to be achieved despite high levels of occlusion from pedestrians, vehicles, and other objects common in urban and natural scenes. We also seek robustness due to sensing degradation, such as from weather.

Background:
The current state of the art in autonomy is insufficient to enable large robots to operate without close supervision around people due to safety concerns. An important reason for this is that perception is currently not sufficiently robust or accurate in detecting, tracking, and predicting trajectories of critical objects (namely vehicles, people, and animals). In addition, robots are not able to effectively assess when their sensors have degraded and so cannot adapt their behavior accordingly; therefore, safety requires operating under worst-case degradation which greatly limits autonomous missions.

Research Description:
This task focuses on the perception needed for a robot to operate safely and intelligently amongst other movers and in an environment degraded by poor weather. It is not attempting to solve all aspects of scene understanding. Tasks P3 and P4 address the problems of static object analysis and understanding what is the “stuff” surrounding the robot. In this task, the focus is on other movers including vehicles, people, and animals. These need to be reliably detected and accurately and robustly tracked, all in real-time. While a good deal of work has been done on this from static sensors or in open environments, what is new here is achieving highly reliable detections and tracking in very cluttered and degraded environments. There is some overlap with object categorization in P3 and P4 as both seek to detect objects, including people and vehicles.

However, both the goals and approaches differ. For instance, here we not only need to detect people, we need to recognize the same person after he/she is temporarily occluded.

Subtask 1 seeks to develop robust, real-time (5Hz) mapping and object segmentation, tracking, and classification in complex dynamic scenes, including outdoor, indoor, urban, and cross-country environments, and leading into higher resolution analysis of human behavior. FY12 focuses on computing high-fidelity appearance and motion cues of objects and extending the mapping range.

Subtask 2 seeks to achieve semantic scene understanding in environments degraded by poor weather, high clutter, and low sensor resolution. We develop parts-based object and vehicle discriminators that achieve robust detection. By assessing the impact of environmental factors on our recognition algorithms, we will develop ways to improve their capabilities and, as necessary, adjust mobility constraints to ensure safe operation.

Subtask 3 aims to achieve long-term prediction in part by integrating algorithms from P5. Immediate goals are 1) integration issues such as computation speed and interface to different sensors and tracking algorithms, and 2) combine different motion models to handle clutter.

Subtask 4 aims to detect and track people in very high detail in cluttered street scenes. This includes achieving dense optical flow on human limbs, precise boundary segmentation, and long-term motion trajectory estimation for each person in the scene.

An important component of this effort is to both contribute to and leverage the world model. This will enable sharing of information between robots separated spatially or temporally. For instance, detections and identifying characteristics of target vehicles will be stored so that the same vehicle can be identified at a later time or location.

Metrics:
Metrics for evaluating the subtasks include precision-recall curves on detection performance. Labeled data are used to evaluate algorithms and include both hand labels and automatic labels generated for simulated data. In addition, achieving real-time operation is important, and so algorithms are evaluated on what frame rate they achieve.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Task P6-3 involves collaboration with ARL's LADAR group (Barry Stann). This collaboration started in FY11 (initial results on tracking from MEMS LADAR) and is expected to expand in FY12. We expect results from this work to be submitted to some of the main Robotics conferences (IROS, ICRA, RSS).

Milestones:

	Quarterly Milestones and Deliverables
Q1	Implementation of combined Cartesian and polar-perspective 3D representation; API to world model for surfaces and objects. Create labeled sensor test dataset. Integration of prediction with the world model. Design of algorithms for model switching for dealing with clutter. Combining dense fine-grained tracking with bounding box people detection.
Q2	Computation of high-fidelity, range scaled appearance cues and investigation of motion cues. Develop weather models for sensors. Create target and clutter models to enable discrimination. Integration of small number of labels (from P4 algorithms). Implementation and testing of model switching algorithms in moderate clutter. Extending short motion trajectory into long trajectory through occlusion.
Q3	Computation of high-fidelity 3D motion cues and object classifier integration, with API to world model to send object categories. Show individual sensor detection performance. Initial version of using tracking input from stereo. Design of top-down directives in tracking/prediction and initial experimentation. Detailed tracking and pose recognition of body including limbs in clutter.
Q4	Complete integration with world model, and support IRA with terrain mapping and object detection. Show multi-sensor detection performance. New version of combined tracker/prediction with complex labels. Collaborative tracking across multiple robots.

Related Research:
This effort uses the semantic environment labels from P3 and P4 (labeling static environment) and the results of CMU's P5 (prediction models). Integration efforts connect with tasks from JPL on stereo and people tracking/detection as well as from GDRS on robust tracking. Integration efforts started in FY11 with the ARL scanner will continue. Long-term, the integrated systems will input to the world model (I2) and to the ACT-R cognitive model (I1).

Integration and Assessment Activity:
This research supports capabilities 4, 6, 7, 8, and 10 as outlined in the Capstone Assessment Vision. We aim to enable the use of multiple platforms and sensors for IRAs along with the techniques of P4 and P5 for labeling and prediction, at speed and robustness required by the IRAs.

We plan to show increased capabilities for tracking and prediction to support the capstone scenario. We propose the following steps to migrate capabilities to the assessments:

- Tracking and prediction from LADAR data (if using the SICK, GDRS, or Hokuyo sensors, this capability is available in FY11).
- Integration of the prediction information with the world model (following design for WM task I2).
- Integration of simple semantic labels (using algorithms from P3 and P5).
- Integration with control directives from world model (currently developed in the architecture task (I1)).
- Tracking and prediction with other inputs (e.g., stereo).
- Integration of detailed complex semantic labels (e.g., detailed identification of features such as doors and stairs in reasoning about behavior in capstone element 6 in which more complex reasoning about people entering/leaving a building is needed).

Subtask 1: 3D Mapping, Segmentation, and Tracking in Complex Dynamic Environments (JPL)

This task is developing robust, real-time (5Hz) mapping and object segmentation, tracking, and classification in complex dynamic scenes, including outdoor, indoor, urban, and cross-country environments, and leading into higher resolution analysis of human behavior. FY12 focuses on computing high-fidelity appearance and motion cues of objects and extending the mapping range.

Objective and Benefits:

The objective of this task is to develop robust, real-time (5Hz) mapping and object segmentation, tracking, and classification in complex dynamic scenes, including outdoor, indoor, urban, and cross-country environments. Over the course of the next three years of the program, this task results in the capability for an autonomous vehicle to segment, track, and compute high-fidelity motion, geometry, and appearance features of all moving objects in the environment, while simultaneously mapping the static aspects of the environment. Object classifiers being developed under the P3 task are integrated to then label all moving objects, and object tracks and labels are sent to the world model to be used by other Perception and Intelligence tasks to recognize actions, behaviors, and intention. A map of the static terrain including features used to compute traversability (such as slopes and steps) is also sent to the world model for use by autonomous navigation algorithms. As part of the Capstone Assessment, this system is used to segment and track people and other moving objects (such as vehicles or animals) and map the environment containing objects that may move (such as doors and chairs), allowing a robot to autonomously navigate dynamic environments (elements 4 and 10) and watch for people, classify them, and respond appropriately (elements 6 and 7).

State of the Art:

Currently, most robotic systems build and process 2.5D or 3D maps assuming a static environment or by processing only a single frame or short time period of data. Because the latter approaches do not temporally accumulate data, they can only process the instantaneous view of the robot with limited statistical accumulation of surface geometry, appearance, or motion properties. This issue can be addressed by explicitly computing pixel-level motion of elements in the scene, but dense optical flow computed from either monocular or stereo imagery remains

computationally expensive (accelerated GPU implementations on a single camera achieving non-regularized results at 15Hz [7] and regularized results at 1Hz [13] on 640x480 imagery and joint stereo-flow with a variety of implementations [2,6,8,9,11,12], with the fastest running at 1 fps on 640x480 imagery or 5 fps on 320x240 imagery). Scene motion can also be segmented in a 3D representation by projecting range data into an occupancy grid and estimating changes in the map [14]. While sufficient for tracking objects in simple scenes, this approach does not scale well to complex, cluttered environments. Motion is a strong cue for detecting pedestrians and other objects in a scene [10], but approaches that are limited to stationary cameras or are single frame detection with tracking of detected objects [5], like some of the current best appearance-based approaches [3,4], still do not achieve acceptable false alarm rates for autonomous navigation (with 0.1 false alarms per frame, translating into a false alarm every few seconds, for a 0.9 probability of detection).

(a) An image taken by a stereo camera pair of a cluttered dynamic urban scene.

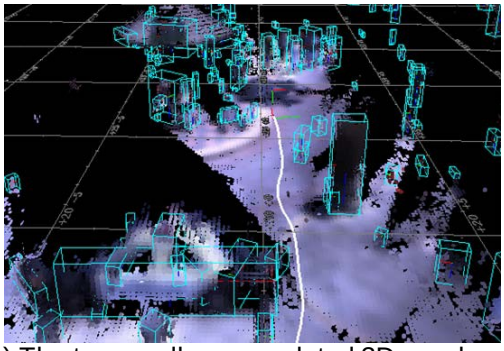

(b) The temporally accumulated 3D voxel map showing the path of the vehicle.

(c) The segmentation of the 3D map back-projected into the image, showing region bounding boxes and a pixel-level color coding of the regions.

(d) The output of the system as used in IRA1 to detect people and classify their behavior.

Figure 3-21: (a) Stereo imagery produces colorized 3D point clouds that are accumulated in (b) a spatio-temporal representation that handles dynamic and static elements in the environment and is used to produce (c) a segmentation of the scene into regions-of-interest; (d) shows the output of the integrated system used in IRA1.

Our approach to dynamic scene mapping, segmentation, and object classification supersedes state-of-the-art approaches by temporally accumulating data in dynamic environments and overcoming the computational costs associated with computing high-fidelity appearance and

motion cues, enabling accurate terrain mapping and learning substantially better classifiers for objects such as pedestrians and cars.

Key Barriers and Approach:
During the first two years of the program, we developed a spatio-temporal representation that addressed the key barrier of efficient temporal accumulation and segmentation of 3D point-cloud data in dynamic scenes. The representation results in a 3D voxel map of the static elements of the scene and a segmentation of all dynamic elements in the scene. In FY12, we focus on addressing the barriers of efficiently computing high-fidelity motion and appearance cues, identifying a unified way of using these features to perform classification of many object categories across large-scale changes, and extending the approach to handle multiple stereo camera pairs covering 360 degrees around the vehicle and longer ranges. We make computing high-fidelity motion and appearance cues and classifying many object categories efficient by using geometry to generate initial regions-of-interest, coarse object motion estimates, and object classifications, each of which are then refined with more expensive algorithms and features, such as histograms of oriented gradients, optical flow, and parts-based classifiers. We extend our representation to handle longer ranges with a 360-degree field-of-view by investigating the use of a hybrid Cartesian and polar-perspective data structure.

Metrics for Evaluation:
Our current system runs at 5Hz on 612x510 resolution 90-degree field-of-view stereo range data, with 10cm voxels in a 30m map (centered on the vehicle) on a single core of a 2GHz Intel Core2 Duo. In addition to being tested on data collections from crowded urban environments, it was fielded as part of IRA1 at the FTIG MOUT site. For IRA1, geometry features are computed for each object and used to perform simple classification of upright pedestrians, which are then fed via the world model to a cognitive reasoning algorithm that recognizes pedestrian behaviors (Figure 3-21 d). A 2D scanning LIDAR was used to simultaneously collect data to ground-truth the pedestrians and evaluate the detection and false positive rate of the stereo-based system. In order to generalize the system evaluation to more complex terrain and object categories, in FY12 we investigate using the RIVET simulation engine to generate stereo imagery and evaluate the accuracy of our underlying representation and features as well as the object segmentation, track length, and classification rates. We focus on urban environments as they are easier to model and render realistically. We augment the simulations with data collections and hand labeled ground-truth in lightly cluttered urban environments. We aim to keep the total system processing rate at 5Hz on a standard embedded computer.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Implementation of combined Cartesian and polar-perspective 3D representation; API to world model for surfaces and objects.
Q2	Computation of high-fidelity, range-scaled appearance cues, and investigation of motion cues.
Q3	Computation of high-fidelity 3D motion cues and object classifier integration, with API to world model to send object categories.
Q4	Complete integration with world model, and support IRA with terrain mapping and object detection.

References:

[1] M. Agrawal, K. Konolige, and L. Iocchi. Real-time detection of independent motion using stereo. IEEE workshop on Motion, 2005.

[2] T. Dang, C. Hoffmann, C. Stiller. Fusing optical flow and stereo disparity for object tracking. In proceedings of the IEEE 5th International Conference on Intelligent Transportation Systems, 2002.

[3] P. Dollr, S. Belongie and P. Perona. The Fastest Pedestrian Detector in the West. Proceedings of the British Machine Vision Conference, 2010.

[4] M. Enzweiler and D. M. Gavrila. Monocular Pedestrian Detection: Survey and Experiments. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol.31, no.12, pp.2179-2195, 2009.

[5] M. Enzweiler, P. Kanter and D. M. Gavrila. Monocular Pedestrian Recognition Using Motion Parallax..IEEE Intelligent Vehicles Symposium, Eindhoven, The Netherlands, 2008.

[6] U. Franke, C. Rabe, H. Badino, S. Gehrig. 6D-vision: Fusion of stereo and motion for robust environment perception. Pattern Recognition, 2005.

[7] J. Marzat, Y. Dumortier, A. Ducrot. Real-time Dense and Accurate Parallel Optical Flow using CUDA. International Workshop on Computer Vision and Its Application to Image Media Processing (2009).

[8] N. Slesareva, A. Bruhn, J. Weickert. Optic flow goes stereo: A variational method for estimating discontinuity-preserving dense disparity maps. Pattern Recognition, 2005.

[9] A. Talukder, L. Matthies. Real-time detection of moving objects from moving vehicles using dense stereo and optical flow. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2004.

[10] P. Viola, M.J. Jones, D. Snow. Detecting pedestrians using patterns of motion and appearance. International Journal of Computer Vision, 2005.

[11] A. Wedel, A. Meisner, C. Rabe, u. Franke, D. Cremers. Detection and segmentation of independently moving objects from dense scene flow. Energy Minimization Methods in Computer Vision and Pattern Recognition, 2009.

[12] A. Wedel, C. Rabe, T. Vaudrey, T. Brox, D. Cremers. Efficient Dense Scene Flow from Sparse or Dense Stereo Data. Proceedings of the 10th European Conference on Computer Vision (ECCV), 2008.

[13] M. Werlberger, W. Trobin, T. Pock, A. Wedel, D. Cremers, H. Bischof. Anisotropic Huber-L 1 Optical Flow. Proceedings of the British Machine Vision Conference, 2009.

[14] D.F. Wolf, G.S. Sukhatme. Towards mapping dynamic environments. In Proceedings of the International Conference on Advanced Robotics (ICAR), 2005.

Subtask 2: Multi-Sensor Perception and Understanding of Complex and Dynamic Environments (GDRS)

This subtask seeks to achieve semantic scene understanding in environments degraded by poor weather, high clutter, and low sensor resolution. We develop parts-based object and vehicle discriminators that achieve robust detection as well as the ability to differentiate between different vehicle types. By assessing the impact of environmental factors on our recognition algorithms, we will develop ways to improve their capabilities and, as necessary, adjust mobility constraints to ensure safe operation.

Objective and Benefits:

This research supports capabilities 3, 4, 6, 7, 8, and 10 as outlined in the Capstone document. The key capabilities we develop will support the Capstone Assessment in a number of ways:

1. Safe mobility of the robot as it navigates to its goal.
    a. Using LADAR and video to jointly detect and discriminate both targets and obstacles.
    b. Detecting hazards including obstacles and those caused by adverse weather.
    c. Assessing these hazards and adjusting mobility constraints to achieve its mission safely.
2. Categorization of vehicle types and recognition of a target vehicle.
    a. Enable commands such as “Follow that vehicle,” and reacquire it after an occlusion.
3. Sharing knowledge between robots using the world model.

State of the Art:

Over the past few years, improvements have been made in algorithms to detect humans and other objects for robotic perception. However, algorithms still need to attain better performance, especially in mitigating false alarms before an autonomous robotic vehicle can reliably complete missions in all real-world conditions. Our method moves beyond the state of the art in two areas: 1) by incorporating sensor-impact modeling of weather, clutter, and other factors into the discrimination, and 2) by fusing and cross-cueing LADAR and video data. This will enable more reliable object detection and safer mobility for real-world conditions.

Key Barriers and Approach:

The essence of our approach is to detect and recognize target objects and obstacles affecting mobility under adverse conditions. While we have shown promising results in discriminating obstacles from traversables, results were demonstrated in good weather conditions. It is known that rain and snow can significantly affect both LADAR and video cues. We intend to investigate the impact of these conditions on our discrimination algorithms. We also intend to investigate how using multi-sensor cues can mitigate perception degradation and improve its effective range.

Our initial results on categorization using a basic grid-based map are very promising. We expect that by combining multiple features and using relative spatial locations, we will be able to significantly increase the number of categories and also enable target identification. This ties in

closely to semantic scene understanding via the world model. We plan to implement this inside the world model to enable robots to automatically share relevant target information.

Categorization is illustrated in Figure 3-22 and Figure 3-23.

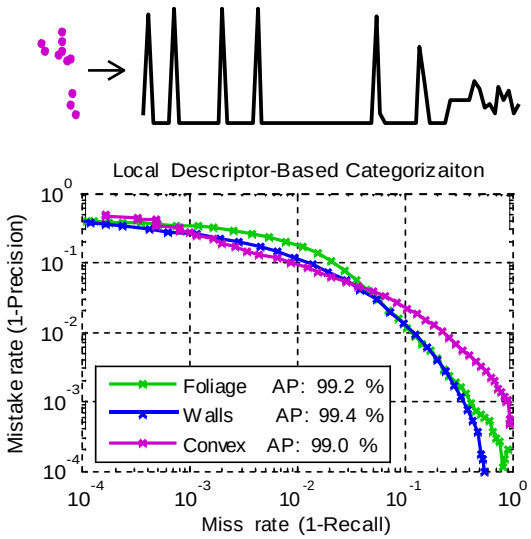

Figure 3-22: Local surface is encoded in a high-dimensional vector.

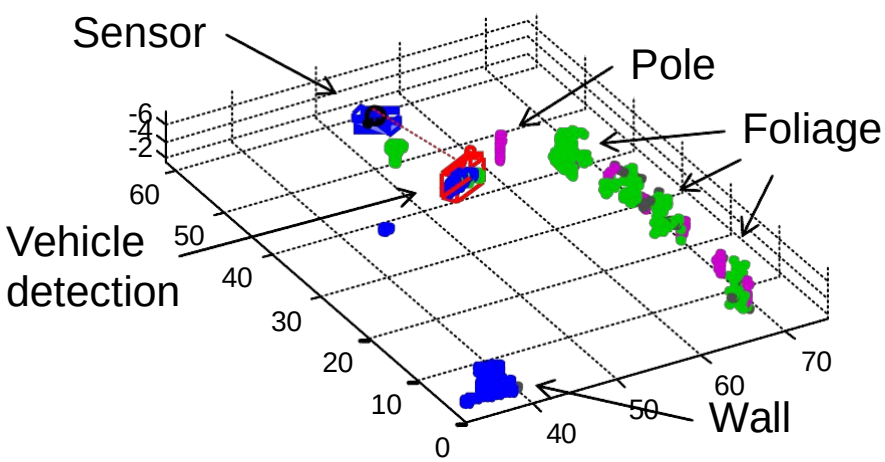

Figure 3-23: These vectors are used to discriminate obstacle types. The average precision (AP) for three categories is very good.

Metrics for Evaluation:

We are measuring our progress through a quantitative Precision-Recall analysis. We have created a ground-truth dataset on which we can assess our progress. Initial quantitative results from 2011 are shown in the above chart.

To measure progress towards supporting the capstone capability, we intend to build a labeled sensor dataset that captures a scenario similar to the capstone. We can then quantitatively assess how each component of our work supports the necessary perception needs of the capstone.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Create labeled sensor test dataset.
Q2	Develop weather models for sensors. Create target and clutter models to enable discrimination.
Q3	Show individual sensor detection performance.
Q4	Show combined sensor detection performance.

Subtask 3: Integration of Tracking and Prediction (CMU)
The purpose of this subtask is to develop algorithms for robust perception in dynamic environments, including integration of the long-term prediction algorithms (from P5). Immediate goals are 1) integration issues such as computation speed and interface to different sensors and tracking algorithms, and 2) combine different motion models to handle clutter.

Objective and Benefits:
This research supports capabilities 4, 6, 7, 8, and 10 as outlined in the Capstone Assessment Vision. We aim to enable the use of multiple platforms and sensors for IRAs along with the techniques of P4 and P5 for labeling and prediction, at speed and robustness required by the IRAs.

We plan to show increased capabilities for tracking and prediction to support the capstone scenario. We propose the following steps to migrate capabilities to the assessments:

- Tracking and prediction from LADAR data (if using the SICK, GDRS, or Hokuyo sensors, this capability is available in FY11).
- Integration of the prediction information with the world model (following design for WM task I2).
- Integration of simple semantic labels (using algorithms from P3 and P5).
- Integration with control directives from world model (currently developed in the architecture task (I1)).
- Tracking and prediction with other inputs (e.g., stereo).
- Integration of detailed complex semantic labels (e.g., detailed identification of features such as doors and stairs in reasoning about behavior in capstone element 6 in which more complex reasoning about people entering/leaving a building is needed).

State of the Art:
The state-of-the-art techniques for the detection and tracking systems that rely on active range are able to track individual targets (people or vehicle) in environments with few occlusions and a limited number of moving entities. In contrast, to generalize these capabilities to real-world UGV scenarios, two key challenges need to be addressed: 1) the ability to deal with many

simultaneous tracks while maintaining accurate data association along each track, and 2) the ability to recover from extended occlusions from the environment or from other moving entities. The ability to use long-term prediction information and to dynamically switch between different motion models using cues from the environment has not been demonstrated in SoA. The concept of adaptive motion models has been explored using particle filters and probabilistic ARMA models but without using cues from the scene.

Key Barriers and Approach:
We have two approaches to the aforementioned objectives:

(1) Our current approach integrates the IOC prediction from P5 with tracking from LADAR. We plan to continue integration with other sensors and other sources of tracking information. Specifically, in FY11, we have completed the integration with the Hokuyo LADAR. We have started the integration with stereo input and the ARL MEMS LADAR. We will continue and expand these efforts in FY12 toward capstone integration.

The current implementation is fast enough for operation on a robot with a limited map size (40m) and few targets. We plan to explore faster ways to implement our algorithms, particularly through the use of GPUs to scale to larger maps and more movers, and to integrate online labeling techniques.

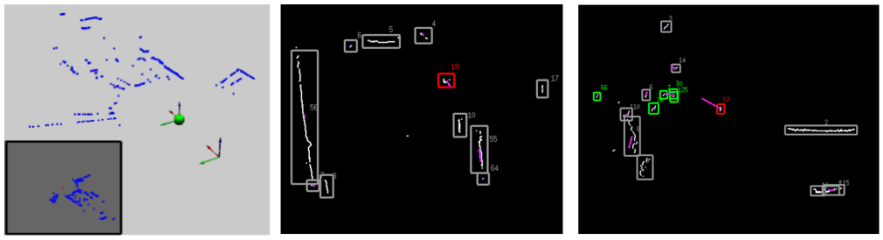

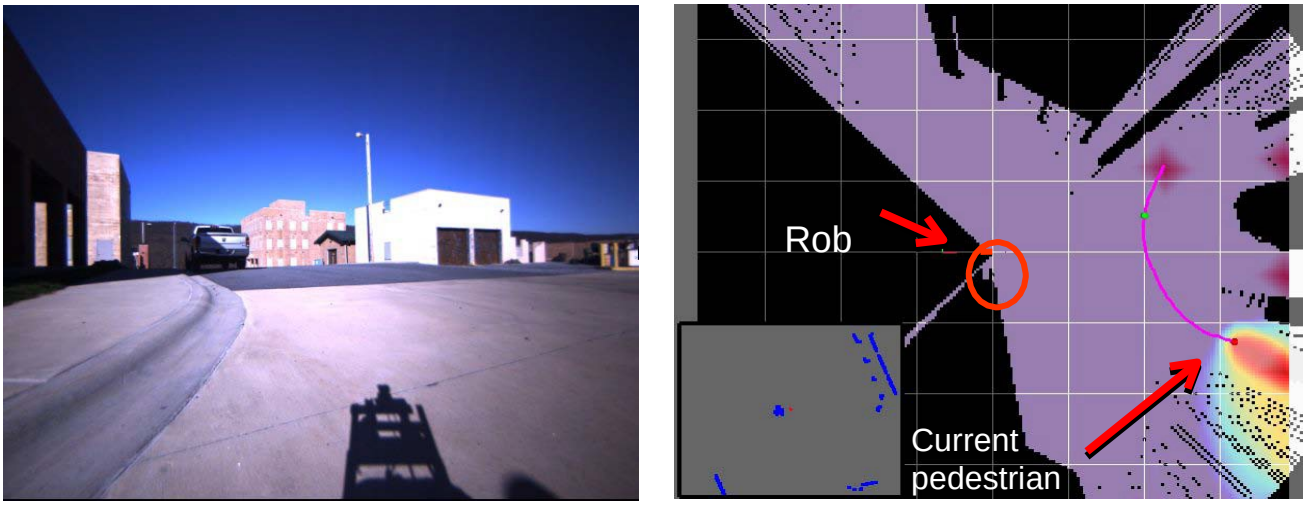

Sample scene at CACTF (FTIG) from K-bot robot

Prediction of pedestrian trajectory

Figure 3-24: Example data and tracking results from Hokuyo (top), prediction of pedestrian trajectories using a Hokuyo mounted on a UPenn K-bot (center), and sample data and tracking using the ARL LADAR (bottom).

Another important aspect of reasoning about tracking and prediction is to interact with the world model. Our first step is to represent the predictions as a distribution of likelihood values that can be queried by the world model. We will extend the interface to handle top-down directives (as investigated in the architecture task I1).

(2) Current tracking approaches use a fixed motion model to estimate position and velocity for tracking in open spaces with moderate clutter. We plan to develop an algorithm that adapts to high clutter environments by combining the motion model with the path predictions. Specifically, in current tracking models, target motion estimates loose accuracy in cluttered environments, since new measurements often fail to be associated with the correct target. To overcome this, we plan to develop an algorithm that is able to adapt to the environment by selecting the most suitable motion model. Specifically, we propose to combine the standard motion models in combination with a target's predicted path (produced by the algorithms explored in Task P5). Since these models generate outputs at very different rates, we will explore ways in which the system can smoothly switch from one motion model to another according to expected target maneuvering. The advantage of this approach is that environmental information, which is taken into account to predict the target's trajectory, is used to model the target's motion more realistically, and therefore, more accurate estimates are produced.

In FY11, we have created a system for tracking objects in real-time tested with several LADAR sensors and trajectory prediction fast enough for operation on robot. These provide solid infrastructure for further development and basis for successful development.

This effort uses the results of CMU's P5 (prediction models) and P3 and P4 (labeling static environment). Integration efforts connect with tasks from JPL on stereo and people tracking/detection tasks as well as from GDRS on robust tracking. Integration efforts started in FY11 with the ARL scanner will continue. Long-term, the integrated systems will input to the world model (I2) and to the ACT-R cognitive model (I1).

Metrics for Evaluation:
We plan to evaluate tracking accuracy by comparing position and velocity estimates with ground truth data and the prediction performance from CMU's P5. We evaluate the action recognition performance by standard P/R and ROC curves. We plan to evaluate the system periodically at the FTIG site on the target platforms and sensors based on the above metrics. Initial testing has been conducted already.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Integration of prediction with the world model. Design of algorithms for model switching for dealing with clutter.
Q2	Integration of small number of labels (from P4 algorithms). Implementation and testing of model switching algorithms in moderate clutter.
Q3	Initial version of using tracking input from stereo. Design of top-down directives in tracking/prediction and initial experimentation.
Q4	New version of combined tracker/prediction with complex labels.

Subtask 4: Detection and Tracking in Densely Cluttered Environments (UPenn)

Objective and Benefits:
This task is primarily focused on detailed tracking and detection of people in cluttered street scenes. By detailed, we mean 1) dense optical flow, 2) precise segmentation boundary, and 3) long motion trajectory; for each person in the scene. Detecting and tracking people in a cluttered environment are capabilities crucial for autonomous vehicles operating together with humans. To enable collaborative tracking across humans and robots, we need to establish metric and topological relations between cameras so that detected targets can be pursued across a network of robots and Soldiers. Overlap in the field-of-view between cameras enables the estimation of relative position and orientation with an associated uncertainty as well as with a topological relation that they share a common field-of-view. The goal is to compute the pose and orientation of each mobile camera as well as to establish a topological description which will enable distributed tracking without necessarily knowing the exact camera location.

This is a five-year effort, and its evolution is characterized by increasing the visibility and clutter complexity of the environment.

State of the Art:
Recent progress of object detectors [JS1] has made tracking-by-detection most popular, achieving state-of-the-art results in pedestrian tracking applications [JS2, JS3]. A hierarchical scheme for detection linking is presented in [JS4] where new features for computing association costs are extracted at each step of tracklet expansion. Authors of [JS5] represent pedestrians as ensembles of body parts for tracking under frequent occlusions. Work of [JS3] estimates scene geometry online through visual odometry to filter car and pedestrian responses through hypothesis selection in explaining 3D points. Authors of [JS6] propagate head detections through KLT trajectories. Correspondence through trajectories anchors detection bounding boxes and produces stable pedestrian tracklets.

We see two downsides in detection-based tracking works:

1) Tracking performance is upper bounded by detector accuracy. Detectors need to be trained (and retrained) to adapt to the appearance of the objects to track for each application setting.
2) Detection-based approaches represent the target with a bounding box, which is not a good representation for people in cluttered scenes. Under partial occlusions or close interaction between the moving targets, the bounding box will leak across multiple targets.

The features aggregated from the bounding box are corrupted from surrounding objects and lead easily to drifting. Bounding box representations have a key weakness in cluttered scenarios as they cannot adapt to the changing image support that the target occupies while interacting with the environment.

Figure 3-25: Comparison with bounding box trackers. We show with solid lines detection responses and with dashed lines interpolated bounding boxes. We see how trajectlet point clusters capture accurately the targets while boxes cannot even indicate if a person is visible or not since often interpolation happens due to deformation or partial only occlusions. Our basic model does not track through full occlusions (see the color switch in purple circles), which can be done in a second data association of the basic detectlets, benefiting from their accurate support.

Figure 3-26: Integrating fine-grained tracking with bounding box tracking. Trajectory clusters adapt flexibly to the changing visible masks of the targets. We track through partial occlusions without complicated reasoning about occlusions/disocclusions of body parts or recovery from wrong level set segmentations. Thanks to the large time context of trajectlet units, we can correctly separate interlocked objects as long as they unlock (have separate motion or disparity or are captured by separate detections) for at least one frame during their appearance in the camera.

Regarding collaborative tracking, state of the art is based on distributed SLAM relying on metric reconstructions of the environment and/or line-of-sight measurements. We avoid any mapping of the environment, and we rely only on topological relations between cameras accompanied by uncertain pairwise metric descriptions.

Key Barriers and Approach:
Current systems treat each of these tasks in separation, and as such, the overall performance is upper bounded by each module in isolation. By getting into the details tracking and detection, at multiple levels of granularities, we show how to improve the performance of the entire system for our goal (Figure 3-25 and Figure 3-26 illustrate the operation of our proposed algorithm with respect to state-of-the-art box tracking algorithms). This is a multi-year effort and will continue at least through this year.

1) Combining dense fine-grained tracking with coarse level bounding box people detection, to break the tracking performance bottleneck upper bounded by detector accuracy. Our approach uses a two-level graph to encode multiple levels of tracking information. We establish one Trajectlet and one Detectlet graph, encoding affinities in each space and associations across. We formulate pedestrian detection, tracking, and segmentation as a grouping problem in the joint space of Detectlets and Tracklets, seeking object specific non-accidental grouping alignment in the two spaces. We introduce a graph "mediate" process to propagate highly confident Detectlets affinities to Trajectlets and vice versa to remove false alarms while linking correct detections across partial occlusions.
2) Extending short motion trajectory into long trajectory going through persistent occlusion due to crowd clutter. Our approach uses multiple-hypothesis path planning to generate distinct homotopy path class to piece together partial short motion trajectory. We plan to include a learning approach to rank long motion trajectories based on prior information of the scene context.
3) Detail motion tracking on a person's body limbs. In image regions with one-dimensional structures (human limbs) or little texture information, point-wise correspondences cannot

be established. To overcome this problem, we seek correspondences on contours rather than single points.

4) Our approach for collaborative tracking tries to establish a graph where connectivity between two mobile cameras is characterized by a partial overlap of the field-of-view. This connectivity can be enhanced by uncertain estimates of the relative position and orientation. Since a metric connectivity cannot be established in a large-scale environment, we aim the detection of robot cliques that are visually connected and the planning of motions that would re-establish the missing links and enable tracking throughout the robot camera network. State of the art is based on distributed SLAM relying on metric reconstructions of the environment and/or line-of-sight measurements. We avoid any mapping of the environment, and we rely only on topological relations between cameras accompanied by uncertain pairwise metric descriptions.

Metrics for Evaluation:

We have been working in a dataset collected in Philadelphia Center City containing over four hours of driving in dense urban environment. We have established ground truth labeling of people detection and segmentation for over 13 sequences. We have measured our current performance and showed improvement over the current state of the art for both people detection (precision-recall) and tracking measure (miss-false match-ID switches). In the collaborative tracking setting, we will measure our approach by testing the "visual" connectivity between robots in a distributed tracking task. The independent variables will be the visibility complexity of the environment and the percent of field-of-view overlap between robots.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Combining dense fine-grained tracking with bounding box people detection.
Q2	Extending short motion trajectory into long trajectory through occlusion.
Q3	Detailed tracking and pose recognition of body including limbs in clutter.
Q4	Tracking across multiple cameras with varying visibility and overlapping field-of-view.

References:

[JS1] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained partbased models. TPAMI, 32, 2010.

[JS2] M. D. Breitenstein, F. Reichlin, B. Leibe, E. Koller-Meier, and L. V. Gool. Robust tracking-by-detection using a detector confidence particle filter. In ICCV, 2009.

[JS3] B. Leibe, N. Cornelis, K. Cornelis, and L. V. Gool. Dynamic 3d scene analysis from a moving vehicle. In CVPR, 2007.

[JS4] C. Huang, B. Wu, and R. Nevatia. Robust object tracking by hierarchical association of detection responses. In ECCV. 2008.

[JS5] B. Wu and R. Nevatia. Detection and tracking of multiple, partially occluded humans by bayesian combination of edgelet based part detectors. IJCV, 2007.

[JS6] B. Benfold and I. Reid. Stable multi-target tracking in real-time surveillance video. In CVPR, 2011.

4. HUMAN-ROBOT INTERACTION OVERVIEW

Objective

In line with the Army vision for future robotic platforms, we are pursuing research to support the development of highly effective *Soldier-Robot Teams* where team members, Soldier or robot, understand each other and their actions at a level sufficient to ensure effective mission performance with minimal Soldier tasking. Key to achieving this is leveraging what we know about expert human teams and the development of technology to represent and effectively process this kind of team information. Included in this is the capability to communicate between Soldiers and robots to a level sufficient to support effective coordination. Human-level cognitive capability is not achievable in the near term, nor is it necessary to produce effective teamwork. The model we are pursuing is more modest yet still represents a crucial step forward in robots if successful: development of cognitive models and interaction modalities that could function at the cognitive level expected from service animals or for human-automation interaction in critical application domains.

A great deal is known about how human teams operate and the kind of information and level of communication required for effective human teamwork. Our research approach to embedding robots within Soldier teams is to model and implement within robots the type of knowledge referred to as "shared mental model." Closely related to this is providing the means for Soldiers and robots to communicate naturally and with minimal specialized training and devices. Among the products of this research will be development, implementation, and evaluation of a limited but mission-relevant communication vocabulary and the creation and validation of internal knowledge representations that allow reasoning about team behavior, other humans in the environment, and effective inter-team communication.

State of the Art and Technical Barriers

There have been significant advances in robot capability to operate autonomously in some important but focused areas. However, tele-operation using a dedicated and specialized operator control unit remains the mode of operation for fielded robots and, in particular, for small- to medium-size vehicles. The DARPA Grand Challenge and Urban Challenges demonstrated that relatively large and well-instrumented vehicles could achieve fairly high degrees of autonomy in a relatively rich world. Prior Army-sponsored research has resulted in competent commercial autonomous robots for perimeter patrol and similar tasks. There has also been some successful research on collaborative robots and human-robot interaction conducted by NRL, ARL, CMU, and MIT. Others have shown the feasibility and great potential for the capabilities that we seek; however, only limited success has been achieved in fielding these capabilities to combat Soldiers due to limitations in robustness and the lack of suitable platforms. Key barriers we will address include: how to achieve adequate levels of integration, how to make representations robustness, how to gain human trust, methods for real-world reliable communication with minimal or no operator control unit, and models of team behavior such that robots can function effectively within human teams.

Moving Beyond the State of the Art

The advances proposed for HRI are coordinated with the advances in the Intelligence and Perception technical areas. In particular, the major HRI challenges are related to development, initial population, and dynamic maintenance of the internal representations of human, team, cultural, and communication models. The primary modes of research we proposed directly mirror that used for human performance and team modeling. We will conduct research to model appropriate levels of behavior, implement experiments to validate the models both in simulated systems and in software, and work closely with Intelligence and Perception to embed the models. A significant part of this is developing internal APIs for integrating HRI derived information within the overall decision architecture. Details of experiments and evaluation methods and information exchange mechanisms are addressed within each subtask description.

Relationship to Other Technical Areas

This HRI research will provide new communication modalities for Soldier-robot interaction, the knowledge bases necessary for robot reasoning about team and task execution, and experimental validation of their accuracy and effectiveness. Among our products will be definition and validation of the external and internal communication protocols required for interaction between Soldiers and robots as well as between Intelligence and Perception processing and HRI. We will initially focus on collaboratively defining information format and planned content within the Intelligence world model and the team and task Shared Mental Models.

HRI expects to provide Intelligence with data for creating and maintaining internal models for team, basic military tactical, and low-level non-combatant cultural behavior. This will enable effective reasoning about robot actions and behaviors in the context of team objectives, the actions of other team members, and in situations where non-team humans may be present. These new elements within the world model are needed to achieve real team-like behavior and, in the longer term, more effective operation within cultural environments. There are two model elements to be considered: relatively static representations for teams and culture, and dynamic model elements dealing with unfolding events during operation. There are important design details for representing and maintaining these models that require continuation of current collaboration between Intelligence and HRI. Among these are architectural implementation considerations, agreement on interaction APIs, internal format requirements, and agreement on division of processing responsibility. HRI focuses first on experimental development of models, then simulation-based validation of the models, followed by integration on robots for both laboratory and field evaluation. At each stage, active collaboration is needed, but particularly as experimental validation is undertaken to ensure that models can be effectively incorporated within the Intelligence architecture.

HRI will provide Intelligence with the knowledge bases for reasoning where team and other human cultural elements are a factor. A significant collaboration effort will jointly define the format and populate these knowledge bases. Similarly, HRI expects to be able to provide Perception with information that will enable what in human research is called “priming.” In this context, priming is the presentation of information useful for labeling of perceptual information and direction to seek particular information useful in disambiguation of situations. Conceptually, this instantiates a robotic form of the Recognition-primed Decision (RPD) Model pioneered by

Klein (Klein, Calderwood, & Clinton-Cirocco, 1986) to understand and train rapid decision making in humans. Obviously, the representational form for such information within Intelligence will be important for the linkage with Perception. In particular, HRI and Perception will collaborate in defining APIs and appropriate levels of information exchange, definition of formats, and division of responsibility for acquisition and processing of external information, such as visual and auditory data.

Summary of Proposed Research

The HRI research program is organized into three general research areas: Team Shared Mental Models, Soldier-robot Team Communications, and Social Dynamics Models. There is very significant collaboration between the thrust as well as active interaction with researchers in Intelligence, Perception, and Mobility and Manipulation. These include a weekly thrust leader meeting, a biweekly ARL HRI teleconference, and cross discipline workshops (currently about two per year). Collaboration between UPenn and BYU within HRI spans several tasks within the program, and most tasks are interdependent to varying degrees. During the first two years of execution of the RCTA, several subtasks were completed, and research results across the program have resulted in better focusing of the research tasks. For our HRI research, we will focus on teams consisting of four to five Soldiers using one to two robots, operating in spaces where there may be non-team humans present. Agreement on the capstone vision for the RCTA has also facilitated consolidation of research tasks and identification of specific capabilities that are required. The ten tasks from the FY 2011 HRI program have been consolidated into six 2012 tasks, two within each thrust as outlined below. Table 4-1 summarizes planned outcomes and their implications for each HRI task.

Thrust 1 – Shared Mental Models

Task H1 6.1: Shared Mental Models for Soldier-Robot (SR) Teaming

This task is designed to enable implementation of algorithms that will provide effective teamwork between Soldiers and robots. This task is expected to provide specific guidance concerning what the Soldier needs to know about robots, what the robot needs to know about Soldiers and the environment, and how Soldiers and robots can best make decisions within the context of military tasks.

Task H2 6.2: Situation Awareness in Human-Robot Teams

Situation Awareness and Shared Mental Models are the key issues in effective human mission performance and teaming. Similarly, these are also expected to be crucial for effective Soldier-robot teams. The objective of this task is to provide the foundation for shared situation awareness (SA_w) in Soldier-robot teams. The approach to accomplishing this is development of measures of Soldier situation awareness as a step toward identifying gaps in individual and team-level SA_w and assessment of requirements for robotic contributions to team situation awareness.

Task H3: Trust in Human-Robot Interaction (2010 and 2011 APP)

Task deleted in this APP. Selected subtasks revised and incorporated in Tasks H1 and H9.

Thrust 2 – Soldier-Robot Team Communications

Task H4 6.2: Soldier-Robot Tactical Communication Protocol

This task consists of two subtasks aimed at different elements of team communication. The first subtask addresses creation of a representation methodology and processing capability for converting Soldier and robot interactions into a canonical internal format suitable for interacting with the Intelligence and Perception architectures. This capability is key to maintaining dynamic shared mental models and is expected to significantly contribute to potential robot effectiveness in teaming. This task will leverage advancements to the Parameterized Action Representation model completed in 2010-2011 to develop a tactical communication protocol for SR teams. The second subtask is more applied. It addresses experiments in using spoken commands to direct the operation of a robotic forklift.

Task H5 6.1: Evaluating Tactical Command and Coordination Vocabulary and Protocols

This task will experimentally evaluate the command and coordination section of the HRI Tactical Vocabulary being developed as part of HRI thrust 1 through a series of empirical experiments. The efficacy of the HRI Tactical Vocabulary content will be validated and team interaction requirements further defined. This will also allow assessment of the HRI Tactical Communication Protocol from task H4 as the interface between Soldiers and the internal representation of information within Intelligence, Perception, and Mobility.

Task H6: Implicit Communication (2010 and 2011 APP)

Task completed so not included as a separate task in this APP.

Thrust 3 – Social Dynamics Modeling, Simulation and Experimentation

Task H7: Integrating Multi-modal HR Communications in Live and Virtual Environments (2010 and 2011 APP)

Selected subtask efforts incorporated into Task H5 in this APP.

Task H8 6.2: Social Dynamics Modeling and Simulation

The objective of this task is to produce theoretical models and develop simulations that capture the social dynamics associated with Soldier-robot, Soldier-civilian, and robot-civilian interactions. Among the products will be identification of the most critical social parameters that permit effective operation of robots in the presence of both team and non-team humans. We expect to be able to implement simulations using these results, most likely in ACT-R, as an HRI experimentation application but ultimately as part of an Intelligence implementation.

Task H9 6.1: Social Dynamics Experimentation

This task consists of an integrated set of experiments that evaluate social dynamics issues within Soldier-robot teamwork and in the context of Soldier-robot, Soldier-civilian, and robot-civilian interactions. Expected results include: identification of data parameters for Soldier-robot behaviors, specification of means for integrating perception and context through behavior identification, and metrics for social processes that can be used in robotic decision architectures.

Task H10: Dynamics of Operating within Social Environments (2010 and 2011 APP)

Selected subtask efforts incorporated into Task H8 in this APP.

Task	Primary Research Outcomes	Operational Implications
H1: Shared Mental Models (SMM) for Soldier–Robot (SR) Teaming	Methodology for SMM measurement of interactions among human robotic team members. Operational guidelines for SR teaming, including robot design recommendations to increase SR trust. Computerized mental models and decision-making algorithms.	Models of individual and shared mental models among SR teams will provide ability for true H-R teaming through system transparency and more human comprehensible robot behavior.
H2: Situation Awareness in Human-Robot Teams	Effective working model of Soldier's situation awareness. Metrics of Soldier situation awareness. Guidelines for robot communication to Soldiers.	Correct selection, use and updating of shared mental models will provide effective Soldier-robot communication and coordination. Higher situation awareness between Soldier and robot will make operations more effective, efficient, and safer.
H4: HRI Tactical Communication Protocol	HRI Tactical Communication Protocol API layer for exchange of information among HRI, Intelligence, and Perception. This tool will be made available for individual experimentation, IRAs, and capstone activities. Empirically evaluated natural language tool.	Translation of raw communication data sources (e.g., accelerometer and tactile devices) into relevant mental and world models will allow robots to more effectively communicate within team. This will significantly reduce Soldier workload and as well as internal processing within the robot. Adding natural language capability will enable a robot to execute verbal commands from Soldiers.
H5: Evaluating Tactical Command and Coordination Vocabulary and Protocols	Algorithms, methods, prototype software and hardware for bi-directional and team command and coordination in SR teams. Requirements and design recommendations for the development of the HRI Tactical Communication Protocol. Verification, test and evaluation of the HRI Tactical Communication Protocol developed under H4 to permit experimentation and implementation within a robot.	H-R tactical communication vocabulary(s) verified through experimentation will permit more natural human-robot interaction in combat situations. Development of standardized interaction protocols modeled on human-human communication will reduce the requirement for specialize communication methods and training for effective H-R interaction.

Task	Primary Research Outcomes	Operational Implications
H8: Social Dynamics Modeling and Simulation	Produce models and simulations that capture social dynamics associated with Soldier-robot, Soldier-civilian, and robot-civilian interactions to a sufficient level to enhance robot operation. Implement these within an ACTR model to evaluate ability to enable more effective robot operation in inhabited spaces.	This basic research will allow experimental verification of the accuracy and effectiveness of modeling that will ultimately permit robots to operate effectively in conjunction with Soldiers and in realistic environments where non-team humans are present.
H9: Social Dynamics Experimentation	Evaluate and validate models and simulations with human subjects. Provide algorithms that can be implemented within physical robots and simulation environments, and evaluation and validation of issues in trust and confidence for both Soldier and non-combatant interactions. The research findings may also be useful to improve Soldier ability to predict human behavior both Soldier and civilian, within a specified cultural situation.	The results of this research will make possible implementation of software that will permit reliable and safe operation of robots teaming with Soldiers in real world environments where non-team humans are present.

Table 4-1: HRI Task Outcome and Operational Implications table.

Human-Robot Interaction Tasks

H1-2012 – Shared Mental Models for Soldier-Robot (SR) Teaming

Level: 6.1 Basic Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Florian Jentsch	UCF	H1-1, H1-2, H1-4, H1-5
Additional Investigators	Randall Shumaker	UCF	H1-1, H1-2, H1-3
	Stephen Fiore	UCF	H1-1, H1-2
	Scott Ososky	UCF	H1-1, H1-2, H1-4
	Elizabeth Phillips	UCF	H1-1, H1-2, H1-4
	Peter Hancock	UCF	H1-3
	Deb Billings	UCF	H1-3
	Kristin Schaefer	UCF	H1-3
	Christian Lebiere	CMU	H1-4
	Unmesh Kurup	CMU	H1-4
	Stephanie Lackey	UCF	H1-5
	Daniel Barber	UCF	H1-5
	Seth Teller	MIT	H1-6
	Nicholas Roy	MIT	H1-6

Objective:
The objective of this task is to support the enabling of teamwork between Soldiers and robots. In particular, this task provides input to: (a) what the Soldier needs to know about the robot, (b) what the robot needs to know about the Soldier and the environment, (c) how the Soldier and robot make decisions within the context of military taskwork, and (d) how the Soldier will communicate to the robot in order to carry-out commands.

Subtask 1 addresses measurement of Soldier-Robot (SR) team shared mental models (SMM). Subtask 2 investigates analogs to the SR team as a source of inspiration for robotic technical

capabilities as well as guidance for cultivating accurate mental models of robots within Soldiers. Subtask 3 examines the impact of trust and confidence on mental model development and sharedness with respect to the SR team dynamic. Subtask 4 investigates cognitive decision making frameworks within robots and the similarities/differences to human decision making processes. Subtask 5 seeks to develop a shared vocabulary for interaction between humans and robots. Subtask 6 is focused on the development of natural language understanding within robots that will allow them to receive commands from Soldiers.

Background:
Mental Models (a.k.a. scripts, schemas) are of foundational importance in the way humans interact with their environment. Specifically, mental models of team members, systems, taskwork, and teamwork shape the ways in which people cooperate, collaborate, and coordinate. Measuring, assessing, and being able to describe mental models is, therefore, of critical importance for the design of systems, procedures, and operations, regardless of whether the focus is on individuals or teams. When teaming is the focus, however, Shared Mental Models (SMM) must be assessed because they provide the foundation of efficient teamwork; for example, in the standardization of tactical communication between team members, the scaffolding of situation awareness (SA) within teams, and the cultivation of trust among team members.

In the military domain that forms the focus of the Robotics CTA, sources of *Soldier* mental models exist within Field Manuals, in Tactics, Techniques, and Procedures (TTPs), and in other extant materials, but we do not yet know what form *robot* mental models will take. Furthermore, as autonomous and intelligent capabilities are built into robotic platforms, the need arises to structure compatibility between human mental models and robot "thinking." While we know that sharedness of mental models is predictive of human-human team performance, the impact of shared mental models on *human-robot* team performance is currently unknown.

Research Description:
The proposed program plan will focus on enabling the technical capabilities in order to execute the mission primitives outlined in the Capstone Integration and Assessment Vision document, through the development of Solder-Robot mental models. Although we will focus our investigations on squad-size teams of Soldiers operating collaboratively with one (or two) robots, in operations such as Cordon and Search, our approach is theoretically grounded in cognitive, behavioral, and computer sciences and will provide foundational knowledge to system design of the next-generation military robots. Our program plan under this task seeks to (a) develop a reliable methodology for measuring SR shared mental models, (b) investigate appropriate SMM in teams thought to be analogous to SR teams, (c) conduct experimentation on the role of trust in SMM development, (d) use cognitive architectures, such as ACT-R, to instantiate SMM in tactical HRI, and (e) develop the means by which a Soldier can communicate a command which is understood and executed by a robot. A supplementary approach to this research will exist in the form of collaboration with colleagues from intelligence, perception, and the integration/world model teams to support the instantiation of an assessment-ready cognitive architecture into a robot.

Metrics:

- Development of SMM measure methodology.
- Development minimum of 10 SMM attributes relevant to tactical HRI based on human-robot team metaphors.
- Guidelines for improving tactical HRI trust through robot/agent transparency.
- Reliability analysis of decision-making tendencies between humans and cognitive architecture.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Submit, present at, and attend:

- SPIE Defense Security and Sensing Symposium (Unmanned Systems Technology XIV Conference DS114), Baltimore, MD (April 2012).
- Human Factors and Ergonomics Society Annual Meeting, Boston, MA (October 2012).

Conduct "Tiger Team" meetings to address instantiation of SMM representation into world model.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Description of initial SMM measurement methodology. SMM feature list derived from HRI surrogates and metaphors. Methodology and deployment plan for experiment. Initial cognitive model of SMM structure with likely decision outcomes based on shared world representation. Initial communication vocabulary set. Collect and analyze a corpus of representative supervisory command utterances for autonomous mobile manipulation within a semi-structured or unstructured environment. Demonstrate embodied collection of crowd-sourced utterances of representative mobile manipulation commands.
Q2	Methodology evaluation and analysis report. Initial taxonomy of SMM attributes. Experimentation. Validation of cognitive model of SMM structure and inference against human performance data. Feedback from SRI Bi-directional Communication evaluation experiment, and refinement of vocabulary set. Extend learning process to unannotated corpora. Demonstrate scalability of learning with considerably less labeling effort.
Q3	Revised SMM measurement methodology. Initial report applying taxonomy to SR teams. Experimentation and initial data analysis. Development of cognitive capability to perform inference of world knowledge based on SMM and observed decision. Refinement of communication vocabulary set. Collect and analyze a corpus of representative narrated demonstrations of manipulation within a semi-structured or unstructured environment. Demonstrate embodied collection of representative narrated mobile manipulation actions.
Q4	Prototype methodology, reports, and publications. Final report/taxonomy with guidelines, publications. Robot design recommendations based on empirical research findings. Technical report on research findings to date. Validation of inference of world knowledge from SMM against human performance data. Feedback from SRI Team Communication evaluation experiment, and refinement of communication vocabulary set. Demonstrate execution of learned mobile manipulation actions from both verbal instruction and narrated demonstration within a semi-structured or unstructured environment.

Related Research:

- Provide guidelines and suggestions for HRI tactical communication protocols based on HRI metaphors (H2-1)

- Provide measurement methodology supporting experimentation with team SMM (I9-3)
- Provide guidelines for instantiating SMM in SR teams based on HRI metaphors
- Provide foundational knowledge on team SMM to HRI (H8: Social Dynamics)
- From task H2-2 and H1-3 (previously H3): incorporation of SA and trust metrics into H1 experimentation
- From perception and intelligence: receive input on tactical capabilities to improve the validity of deployable cognitive decision-making architectures

Integration and Assessment Activity:

- Provide SMM assessment methodology to support Integrated Assessment elements 1 through 3 and 9
- Provide recommendations for improving trust through appropriate robotic action transparency
- Provide guidelines instantiating Soldier-compatible MM representations in world model based on current technical feasibility

Subtask 1: Measure Shared Mental Models in Soldier-Robot (SR) Teams during Tactical HRI (UCF)

Objective and Benefits:
We plan to develop a methodology to reliably measure shared mental models (SMMs) of the interactions among human and robotic team members. The ability to accurately assess individual and team mental models enables the ability to improve their richness and accuracy.

The Soldier's mental model contains knowledge of how a team member operates and achieves tasks. Complimentary to this, the robot's *thinking* must be engineered in a way that is predictable to humans. The mental model also provides a context for *talking* with the robot and a shared understanding of the robot's perceptual and movement capabilities. An accurate mental model allows the Soldier to know what (a) a human team member is/would be doing, (b) what a robotic team member is/would be doing, and (c) how the two models are different.

This research will directly support Capstone Assessment Vision, Item 1. For example, if the Soldier commands the robot to "cover the back door," as a part of a Cordon and Search tactical task (FM 3-24.2, 5-44) an accurate mental model would be able to predict if the robot was "guarding," "blocking," or "covering from concealment," without needing to monitor the robot itself. It is expected that this effort will continue through 2013 (and extensions beyond) as methods and techniques are tested and refined.

State of the Art:
Even within highly structured domains, it has been found that multiple mental model measurement techniques do not agree within the same participant. This extends to the examination of SMM in human-human teams, where no single "best" method exists for

capturing these models. As we presented at the RCTA Joint Area Workshop on Shared Mental Models/Shared Cognition in Orlando in December 2011, unsolved problems remain with respect to (a) the measurement of individual mental models, (b) the aggregation of mental models in teams, (c) the assessment and indexing of sharedness among team member mental models, and (d) the representation of mental models and sharedness at the computer/machine level. In this subtask, we will focus on the first three of these aspects.

Key Barriers and Approach:
Methodology, specific to SR teams, will be developed from existing techniques, research into human-robot team analogs, and newly developed mental model elicitation techniques. Methodology will eventually be tested in experimentation, where the SMM metaphors (described in H1-2) can also be exercised. While even these measures are intended to be applied across domains, we will strive to develop a reliable methodology specifically tailored to tactical HRI. This research can support other HRI tasks (e.g., H2: Situation Awareness, H4: Communication, H8: Military Cultural Features) as well as Intelligence tasks like I9-3. Mental models provide for the activation of situation awareness, establish a shared foundation for communication, influence trust, and enable meaningful teaming.

Metrics for Evaluation:
We anticipate techniques that are developed will be incorporated into experimentation internally as well as in experimentation conducted by other task areas. Ideally, progress will be measured by the development of the methodology itself and subsequently by the validation of that methodology across multiple studies.

In FY 2011, we evaluated multiple mental model assessment techniques with respect to their suitability to the specific domain of HRI in a military context. UCF also hosted a joint workshop on shared cognition in HRI, which included members of the Intelligence and Perception teams, to better understand the state of the art with respect to robotic systems.

The ability for a Soldier to understand and predict the actions of a robot is paramount to teaming; Soldier training methodologies and robot design guidelines can be explored within the *think-work-talk* framework. We will first refine our proposed methodology through internal validation. Prototype measurements, tools, and guidelines for measurements produced by this subtask will be distributed to the SDE for use in joint experimentation.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Description of initial SMM measurement methodology.
Q2	Methodology evaluation and analysis report.
Q3	Description and draft tool set for revised SMM measurement methodology.
Q4	Prototype methodology, reports, and publications.

Subtask 2: Identify and Leverage Appropriate SMM Metaphors for Tactical HRI (UCF)

Objective and Benefits:
The objective of this subtask is to identify and leverage metaphors and analogs to the interactions in Soldier-Robot (SR) teams that can be used to guide the design of systems, interactions, and procedures. Identifying metaphors for emerging SR teams will enable the ability to predict the impact of robotic design features on the Soldier's mental model and may yield operation guidelines to minimize the need for additional or specialized training. This subtask will describe why specific metaphors serve as useful surrogates for SR teams and how these metaphors can be leveraged to evoke useful mental models in humans.

When Soldiers interact with robotic team members to execute tactical mission elements, such as in Cordon and Search, aperture analysis, or target identification/tracking, the Soldiers need to have appropriate mental models of what the robot *can* do, *cannot* do, and *likely will* do. Similarly, robots should be engineered to *think* and maintain a world model congruent to that of their human teammates. One way to represent and to teach these mental models is to use familiar *metaphors* of similarly capable systems. Human-animal teams are one likely inspiration for desirable qualities within SR teams; others are human-automation interaction (e.g., in the automated flight deck of the transport-category aircraft) or even interactions with avatars and social robots. This effort will select the appropriate metaphors and will also help to develop an appropriate set of commands and interactions (*work-talk*) for teaming with robots.

This research will directly support Capstone Assessment Vision, Items 2 and 7. For example, again applying the 'working animal' metaphor, we observe that human-animal teams are able to perform complex, interdependent tasks, without relying on natural language, but instead using an often small, very specific vocabulary including speech, tone, and gesture. Using existing Soldier mental models for communication and action, a similar command set can be adapted to HR teams, with a minimal workload impact to the Soldier (i.e., *"cover, wait and report"* accompanied by appropriate modalities for indicating direction and location). Communication is just one attribute that may be inspired by investigating human-robot team metaphors.

State of the Art:
It is yet unknown exactly how humans and robots will communicate and share information or *what information* should be shared. As we presented at the RCTA Joint Area Workshop on Shared Mental Models/Shared Cognition in Orlando in December 2011, we do know, however, that humans tend to employ metaphors for understanding unfamiliar systems. For example, humans tend to (a) anthropomorphize technology, (b) overestimate the capabilities of robots, (c) be more critical of unreliability in technology than in biological systems, and (d) apply social rules to technology interaction. This subtask seeks to provide guidance on critical aspects of SR teaming through the lens of appropriate metaphors.

Key Barriers and Approach:
The structure of robot thinking and the mechanisms of adaptive behavior generation are still being developed. How human-animal teams complete tasks and *how* human-animal teams are similar to, and different from, human-robot teams should provide a source of focus during these

developmental stages. We will extend our taxonomy of HR metaphors that we began with a limited number of working-animal teams in FY 2011 and will include communication techniques (supporting shared vocabulary efforts in H2-1 and H4) and shared model information suitable for small team SRI. Our approach will also inform robot thinking and working, allowing them to support development and communication of shared mental models.

The interaction techniques and teaming characteristics extracted from HRI metaphors provides input to the design of the robots themselves and the design of interaction between humans and robots. This directly relates to intelligence, perception, and (within HRI) to shared cognition (e.g., trust, situation awareness), communication, and social cognition. These findings may also contribute to considerations for training.

Metrics for Evaluation:
In FY 2011, we conducted a review of human-animal teams which revealed that there are many levels of team interdependence and a varied set of interaction modalities (from tool-like to teammate-like) enabling many different technical abilities. We also identified an initial taxonomy of characteristics to classify human-working animal teams by task, capability, and team-level interaction. We will leverage that work to demonstrate how HRI metaphors and surrogates can inform SR interaction.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	SMM feature list derived from HRI surrogates and metaphors.
Q2	Initial taxonomy of SMM attributes.
Q3	Initial report applying taxonomy to SR teams.
Q4	Final report/taxonomy with guidelines, publications.

Subtask 3: Investigate the Role of Trust and System Transparency in SMM Development during Tactical HRI (UCF)

Objective and Benefits:
Following recommendations from ARL collaborators, mental models and transparency of robot behavior were identified as two critical areas needing empirical evidence as to their roles in trust development. There is an important need to empirically test the association between mental models, system transparency, and trust in order to ascertain the extent of the relationship between these variables, to determine how they affect Soldier behavior in HRI, and to identify specific design features that can improve a Soldier's trust in a robot. Experimentation focusing on transparency and trust may provide means to mitigate possible degradation of trust when there are changes in performance due to failure or changes in environment. This idea of system transparency becomes especially important if a robot is adaptive to its environment; adaptive robots will not respond or act in the exact same way every time. The goal of this subtask is to increase our understanding of these factors in developing and maintaining trust in human-robot

team collaborations through both simulation-based and live experimentation (using platforms such as VBS2, MIX test-bed, RIVET, and UAS simulators) and to link it with the investigations of SMM in SR teams. Human-in-the-loop studies will be planned to examine these critical aspects of HRI and trust. In addition, we will begin preliminary exploration of potential objective trust markers (i.e., physiological markers and behavioral cues). These objective measures may be vital in identifying when degradations in trust occur, therefore enabling mitigation strategies to be implemented so as to avoid negative outcomes associated with loss of trust in a system.

State of the Art:
Cring & Lenfestey (2009) define transparency as "the amount of knowledge an operator has about why the automation [robot] performs as it does" (p. 38). They suggest that designing systems with transparency can help facilitate appropriate trust, but the most effective level or type of transparency has yet to be researched. What degree of system transparency is necessary to facilitate appropriate levels of trust in the robot? Information relating to the system needs to be relayed to the human without the risk of information overload, which can overwhelm the human and have negative effects on a human's ability to perform tasking. Transparency can refer to information relating the process by which automated systems perform (e.g., making decision rules explicit), or it can refer to information about the product or outcome of the system (Adams et al., 2003). Stubbs and colleagues (2007) performed research in which they observed human interactions with a robot in the field. Some of the problems encountered by the human team members related to common ground, or shared mental models, and in particular the "lack of transparency in behavior" was cited as an impediment to successful collaboration. Thus, transparency is essentially a method by which mental models can be "tuned" or corrected, and thereby influences the degree of trust a human has in a robot. The likelihood of a human to trust a robotic partner is directly impacted by the degree to which the human understands the system's/robot's actions and the reasoning or purpose behind such actions (robot behaviors).

Key Barriers and Approach:
Participants will collaborate with a robot to complete a mission in a simulated environment, which will likely consist of one or more of the tasks outlined in the integrated assessment (communication, moving through mobility challenges, etc.). Different aspects of the robot's behavior as well as the system transparency will be manipulated. Performance and trust variables will be measured and analyzed to determine the impacts of different approaches to system transparency. Additionally, mental model prior to interaction and post-interaction (i.e., post-transparency manipulation) will be assessed.

Potential Independent Variables (Manipulations):

- Malfunctions of the robot (scenario with malfunctions versus scenario without malfunctions)
- Level of details of communication/information (low versus high)
    - Humans need to know what it going on, but we don't want to overload them with information so performing their own tasking becomes more difficult

- Types of information robot that the robot conveys (robot-centric [ability/reliability] versus situation/mission-centric); this deals with some of our previously identified antecedents of trust
- Availability of video/images from robot (yes versus no)
- End results (robot can perform task versus robot can't perform task)
- Queries for information versus information given periodically (when errors, status updates, certain time points) versus all the time

Potential Outcome Measures (Dependent Variables):

- Human-Robot Subjective Trust Questionnaire (H3 product)
- Human's mental model of robot
- Human performance/success of collaboration
- Negative Attitudes towards Robots Scale (NARS): developed by Nomura et al. (2004) is a 14-item self-report inventory, each item is being rated on a 5-point Likert scale, with anchors 1 = strongly disagree and 5 = strongly agree. NARS measures the degree of humans' attitudes towards communication robots in daily life.

Metrics for Evaluation:

- Improved understanding of how system transparency influences trust development in HRI and how system transparency can assist in providing a more accurate mental model to the human
- Increased understanding (and empirical support) for how aspects of human-robot communication and mental models influence trust in robots
- Specific robot design recommendations (relating to system transparency) to mitigate issues associated with trust in HRI

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Methodology and deployment plan for experiment.
Q2	Experimentation.
Q3	Experimentation and initial data analysis.
Q4	Robot design recommendations based on empirical research findings. Technical report on research findings to date.

Subtask 4: Use Cognitive Architectures to Instantiate SMM in Tactical HRI (CMU and UCF)

Objective and Benefits:

The objective is to develop computerized mental models and decision-making algorithms and validate them against that of human mental models. As robot intelligence becomes more sophisticated, it will eventually be capable of *thinking* through and making decisions based upon internal and environmental factors. These factors will comprise the robot's "mental model." To

facilitate interaction with human team members, this model will have a degree of sharedness with the human members of the team; therefore, it is important to examine the similarities and differences between decision making in humans and robots.

This research will directly support Capstone Assessment Vision, Items 7, 8 and 9. For example, the robot must make a decision to "follow that guy" (i.e., *work*) if someone leaves a building that is being covered during an operation such as Cordon and Search (Army FM 3-06.20). How does the robot *think* about what to do, and is the decision predictable by the human team members? How does the robot determine which observations are unexpected and important enough to be worth reporting?

This effort is expected to continue for a number of years through an iterative process of software model testing, comparison (to human counterparts), refinement, and validation.

State of the Art:
The development of mental models and cognitive frameworks for robotic platforms is in an early state. Furthermore, human decision making is often unpredictable; humans use heuristics, take shortcuts, and employ satisficing. While we do not desire unpredictability in robot teammates, we will strive to reconcile differences between the two in order to create an overall reliable decision-making team structure.

Key Barriers and Approach:
We implement a general approach to mental models grounded in theoretical frameworks that goes beyond ad hoc task-specific models to provide broad predictive power to the practical use of shared mental models. That approach captures all aspects of shared mental models including individual differences in situation awareness, decision-making knowledge, and effects of practice and expertise.

In collaboration with *intelligence (CMU)*, the HRI members of this subtask will provide inputs to the ACT-R modeling process with respect to the *content* of the mental model upon which the decision making algorithm makes decisions. We will also validate ACT-R modeling efforts against human performance data, initially in a constrained scenario. Grounding the representation and use of mental models in a cognitive architecture provides useful constraints for both capabilities and biases in representing and sharing models between human and robots.

This effort ties into other HRI and intelligence activities as well as the world model and also responds directly to the results of the RCTA Joint Area Workshop on Shared Mental Models/Shared Cognition in Orlando in December 2011. To improve the linkage between Intelligence, Perception, HRI, and Integration within the RCTA, we will form a "tiger team" consisting of members from the aforementioned areas. The team will focus on advancing the instantiation of a cognitive architecture toward an assessment-ready state. Meanwhile, we will continue to provide guidance with respect to what information is most relevant in the world model and how a team can act on, share, and process that information.

Metrics for Evaluation:
We initially intend to train human "experts" on the decision-making techniques used by the software model for a narrowly defined, but operationally relevant, scenario. We will then measure the reliability of human novices and human experts and compare that against the decision-making frequencies of software models.

In FY 2011, we constructed a quantitative SMM for a "follow that guy" task and then developed variations of that scenario. Our colleagues in Intelligence constructed a preliminary ACT-R model to process the data and decide "who should follow." While the model has only been applied to that scenario, it can receive a description of any mental model expressed in the proper decision logic and learn to apply it in a cognitively plausible manner similar to human trainees. In future efforts, we plan to have the ability to rank a set of possible actions and weigh their desirability. Additionally, we will investigate alternative techniques to conduct the decision consideration itself. The development of our efforts will be guided using human subject data obtained in experimentation.

Ultimately, we hope to provide useful guidelines and decision-making abilities to the world model. To that end, our work enables the capstone capabilities by turning perception data (*look*) into actionable output by the robot (*think-work*). Those actions, in turn, would be both expected and understood by the Soldier teammate.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Initial cognitive model of SMM structure with likely decision outcomes based on shared world representation.
Q2	Validation of cognitive model of SMM structure and inference against human performance data.
Q3	Development of cognitive capability to perform inference of world knowledge based on SMM and observed decision.
Q4	Validation of inference of world knowledge from SMM against human performance data.

Subtask 5: Develop Communication Vocabulary for Supporting Shared Mental Models in Tactical HRI (UCF)
This effort contributes to the vocabulary development for shared mental models between humans and robots. This subtask focuses on developing shared vocabulary for interactions between humans and robots. Examples include command and coordination vocabulary (e.g., acknowledging receipt of command, "copy;" ending a specific message, "over;" ending a conversation, "out") and other standard and emerging mission-relevant communication events.

Objective and Benefits:
Based on the findings from 2010 and 2011 work of the communications thrust within HRI, and in specific response to the results of the RCTA Joint Area Workshop on Shared Mental Models/Shared Cognition in Orlando in December 2011, the primary objective of this subtask is

to define a vocabulary set that addresses communication messages passed between Soldier and robot team members. A specific requirement for this subset of the HRI Tactical Vocabulary is the need to clearly delineate bi-direction communication events, messages, and conversation components.

This effort directly supports elements 1, 2, 7, 8, and 9 of the Capstone Assessment Vision document. Specifically, this task enables methods for bi-directional communication between SR team members.

This task is expected to be a one-year effort with empirical experimentation intended to evaluate the utility of the communications vocabulary set in a separate task under HRI.

State of the Art:
This subtask advances the current understanding of effectiveness of complex communication in dynamic mission-relevant situations. An HRI Tactical Vocabulary does not currently exist, and that gap is addressed by this subtask.

Key Barriers and Approach:
A systematic approach to vocabulary development will be applied to enable integration of the communication vocabulary set and other vocabulary subsets as well as to facilitate interactions among SR team members to contribute to shared mental models.

Ties to other research activities include:

- Provide communication vocabulary content for the HRI Tactical Communication Protocol.
- Provides capabilities for supporting Shared Mental Models and Team Situation Awareness.

Metrics for Evaluation:

- Development of a minimum of 20 vocabulary terms

	Subtask 5 Quarterly Milestones and Deliverables
Q1	Initial communication vocabulary set.
Q2	Feedback from SRI Bi-directional Communication evaluation experiment, and refinement of vocabulary set.
Q3	Refinement of communication vocabulary set.
Q4	Feedback from SRI Team Communication evaluation experiment, and refinement of communication vocabulary set.

Subtask 6: Human Implications for Evaluating and Explaining Performance (MIT)

Objective and Benefits:
One key challenge in robotics in general is to create natural interfaces in which robots can interpret and execute ordinary spoken commands. We are trying to develop natural language understanding that allows a mobile manipulator, such as a robotic forklift, to take commands from a human teammate and execute the appropriate action sequence. In FY 2011, we developed a new model of learning new concepts (actions, objects, relationships between objects in the world) called the Generalized Grounding Graph and demonstrated execution of simple tasks. The model learns a mapping from English phrases to perceptual and motor features from heavily annotated datasets. In FY 2012, we propose to address the challenge of learning the same mappings with unannotated data to allow much more rapid concept learning. This is intended to be a multi-year effort. The length of the effort will depend on the complexity of the task. This will address capabilities 1, 2, 7, 9, and 10 of the Capstone Assessment Vision document and capabilities 3 and 8 at a secondary level.

State of the Art:
There are broadly three different ways people have approached the symbol grounding problem in robotics. Starting with Winograd (1970), many have manually created symbol systems that map between language and the external world, manually connecting each term onto a pre-specified action space and set of environmental features. This class of systems takes advantage of the structure of spatial language but usually does not involve learning, has little perceptual feedback, and has a fixed action space. A second approach involves learning the meaning of words in the sensorimotor space (e.g., joint angles and images) of the robot (Sugita, 2005). By treating linguistic terms as a sensory input, these systems must learn directly from complex features extracted by perceptual systems, resulting in a limited set of commands that they can robustly understand.

Key Barriers and Approach:
The essence of our technical approach is probabilistic inference. We have shown previously how to use learning to convert from language onto aspects of the environment. Our current approach depends on learning from datasets that are heavily annotated after the fact by a human supervisor. We propose to reformulate the inference as an optimization of a planning problem, which allows us to then reformulate the learning problem as a reinforcement learning problem. Reinforcement learning has been demonstrated to learn symbol grounding models well with a much sparser signal, eliminating the need for the same level of annotation. This approach has also been shown to work well for learning of natural language models in non-robotic domains. However, these previous approaches learn symbol grounding in purely discrete environments (e.g., artificial game worlds, operation of desktop computers) and have not been extended to the continuous, unstructured domains that robots inhabit. It is strongly linked to the work in inverse optimal control being performed at CMU under Task I7. This research also leverages results from previous research programs in our group.

Note that all approaches, including ours, rely on an intermediate symbolic representation. We have previously used a very simple representation known as “Spatial Description Clauses” that

were only capable of representing simple navigation commands. Under previous RCTA efforts, we extended the representation to the "Generalized Grounding Graph (G^3)," which allowed us to represent complex relationships between entities in the environment. However, we assume the existence of a parser that can translate from the natural language text to this intermediate symbolic representation and are focused on the problem of learning the relationship between the symbols and the low-level sensor data and motor actions.

In contrast, Intelligence Task I5 is focused on learning the relationship between the natural language text and the intermediate representation and assumes the existence of a model that maps the symbols to the low-level sensor data and motor actions. Although some of the underlying mathematical machinery is very similar, there is a natural complementarity to the two approaches; we have shared our data collected under FY 2011 with the group at University of Washington working on I5, and the PIs have already begun discussions about integrating these two technologies.

Metrics for Evaluation:
Our primary metric is the correct execution of natural language commands. Our primary form of evaluation is to collect example commands from untrained people and determine if the robot correctly executes the command. We are able to collect large corpora of test data using Amazon Mechanical Turk, and we are additionally able to use Amazon Mechanical Turk to assess whether the executed robot command matches the text of the instruction. Our current performance is approximately 60% to 70% (depending on the specific corpus), where failures relate predictably to specific limitations of our model.

	Subtask 6 Quarterly Milestones and Deliverables
Q1	Collect and analyze a corpus of representative supervisory command utterances for autonomous mobile manipulation within a semi-structured or unstructured environment. Demonstrate embodied collection of crowd-sourced utterances of representative mobile manipulation commands.
Q2	Extend learning process to unannotated corpora. Demonstrate scalability of learning with considerably less labeling effort.
Q3	Collect and analyze a corpus of representative narrated demonstrations of manipulation within a semi-structured or unstructured environment. Demonstrate embodied collection of representative narrated mobile manipulation actions.
Q4	Demonstrate execution of learned mobile manipulation actions from both verbal instruction and narrated demonstration within a semi-structured or unstructured environment.

References:
Winograd, T. (1970). *Procedures as a representation for data in a computer program for understanding natural language.* Ph.D. Dissertation, Massachusetts Institute of Technology.

Sugita, Y., & Tani, J. (2005). Learning semantic combinatoriality from the interaction between linguistic and behavioral processes. *Adaptive Behavior - Animals, Animats, Software Agents, Robots, Adaptive Systems*, 13:33--52.

H2-2012 – Situation Awareness in Human-Robot Teams

Level: 6.2 Applied Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Florian Jentsch	UCF	H2-1, H2-2, H2-3
Additional Investigators	Randall Shumaker	UCF	H2-1, H2-2, H2-3
	Stephen Fiore	UCF	H2-2, H2-3
	David Schuster	UCF	H2-1, H2-2, H2-3
	Joseph Keebler	UCF	H2-1, H2-2, H2-3
	Stephanie Lackey	UCF	H2-1
	Daniel Barber	UCF	H2-1
	Jessie Chen	ARL	H2-4

Objective:
Situation Awareness and Shared Mental Models are the two cornerstones of effective mission performance and teaming. They mutually feed into and support one another: Without situation awareness, correct mental models cannot be selected, used, and updated; without the correct and appropriately shared mental models, a team cannot build, and share, good situation awareness. The objective of this task is, therefore, to provide a foundation for situation awareness (SA_w) in the human-robot team and, specifically, Soldier-Robot (SR) team by developing methods to: measure the Soldier's situation awareness, identify gaps in individual and team-level SA_w, enable the robot to provision operationally relevant inputs to SA_w, and assess the SR team's situation awareness (i.e., ensuring that the unit has sufficient knowledge to meet objectives).

Background:
SA_w is a mission-critical need at two levels: that of the individual Soldier (who must be aware of mission elements relevant to individual performance) and also for the SR team, as all mission elements must be accounted for and communicated appropriately. In missions such as Cordon and Search (FM 3.06-20), for example, SA_w is needed to establish speed, surprise, isolation, target identification, timeliness, and to minimize collateral damage. In a SR team, it is not

sufficient for robots to gather information; robots must also be able to adapt to the Soldier's changing needs, identify gaps in the Soldier's SA_w, and provide operationally relevant information without explicit direction (see also H1-1 and H1-3) while doing so in a way that is compatible with the Soldier's SA_w and mental models. Measurement of the human team member's SA_w is needed to achieve this objective. Given the robot's developing abilities to *Think*, *Work*, and *Talk*, our approach promotes the robot to the level of team member and investigates optimal strategies for combining individual-level SA inputs to arrive at a team-level model of SA_w.

Research Description:
When robots engage in non-invasive diagnoses of performance and make attributions of human SA_w, they will be able to dynamically support the Soldier's SA_w by providing relevant and necessary information at the right time. We will develop a model of the Soldier's SA_w within the HR team. With this information, robots will know what the Soldier needs to know, allowing them to become contributors to shared SA_w through the integration of robot *Thinking* and human SA_w.

Metrics:

- Development of minimum 20 vocabulary items
- Concurrent validity of SA measures
- Speed and fidelity of situation representation by SA_w metrics
- Robustness of SA_w metrics when environment changes

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Submit, present at, and attend:

- Human Factors and Ergonomics Society Annual Meeting, Boston, MA (October 2012).

Milestones:

	Quarterly Milestones and Deliverables
Q1	Initial communication vocabulary set. Integration of existing, compatible SA_w measures. Integration of explicit communications and combinatory metrics of SA_w. Methodology and deployment plan for experiment.
Q2	Feedback from SRI Bi-directional Communication evaluation experiment, and refinement of vocabulary set. Integrated SA_w measure and empirical study of measure. Summary of existing combinatory metrics of aggregate SA_w. Experimentation.
Q3	Refinement of communication vocabulary set. Dataset from investigation of SR team SA_w measure efficacy. Description of laboratory study comparing aggregate SA_w measurements. Experimentation and initial data analysis.
Q4	Feedback from SRI Team Communication evaluation experiment, and refinement of communication vocabulary set. Report describing empirical results, publications, and dissemination among RCTA collaborators. Dataset from laboratory study comparing aggregate SA_w measurements. Robot design recommendations based on empirical research findings. Technical report on research findings to date.

Related Research:

- Provide communication vocabulary content for the HRI Tactical Communication Protocol
- Provide capabilities for supporting Shared Mental Models and Team Situation Awareness
- From tasks H2-1 and H4: Explicit, especially verbal, communications will provide the 'how' of team member communication that will enable this task to investigate the 'what' of SA. Explicit communications is the medium for teamwork that is critical to building SA.
- From Task H1: Mental models are a critical input to SA. Task H1 will provide mental model measurement techniques.
- *Perception* and *Intelligence* will provide robotic inputs required to (a) build robotic SA and (b) assess Soldier SA
- Provide foundational knowledge on SA in teams to *HRI* (e.g., H3, culture, H4-H6, H9)
- Provide list of cues for robot SA_s to *Perception* and *Intelligence*
- Provide data, tools, and simulations of SR shared SA to *Intelligence* and *ATOs*

Integration and Assessment Activity:

- Provide improved methods for assessment and diagnosis of SA in SR teams to *HRI, Intelligence, Alliance, ATOs,* and *Army*

- Provide guidelines for creating Soldier-compatible robot representations of mission elements to *Alliance* and *Army*

Subtask 1: Applying Communication Vocabulary to Support Situation Awareness (SA) in Tactical HRI (UCF)
This subtask focuses on applying vocabulary for specifying communication and interactions between humans and robots to the development and support of situation awareness (SA). Examples include command and coordination vocabulary for Cordon and Search operations and other mission-relevant communication events.

Objective and Benefits:
Based on the findings from 2010 and 2011 work of the communications thrust within HRI, and in specific response to the results of the RCTA Joint Area Workshop on Shared Mental Models/Shared Cognition in Orlando in December 2011, the primary objective of this subtask is to develop and apply a vocabulary set that addresses communication messages passed between Soldier and robot team members so that complex communication can support SA in dynamic mission environments.

The results from this effort will support Intelligence, Surveillance, and Reconnaissance (ISR) missions, routing to a back door, pursuing an enemy exiting a building, and Cordon and Search. This effort also supports elements 1, 2, 7, 8, and 9 of the Capstone Assessment Vision document through the application of bi-directional communication between SR team members.

This task is expected to be a one-year effort with empirical experimentation intended to evaluate the utility of the communications vocabulary set in a separate task under HRI. However, the deliverables of this subtask are intended to be used by RCTA performers for the life of the program.

State of the Art:
Critical elements of operations are command and coordination communication events. An HRI Tactical Vocabulary does not currently exist but will be developed as part of HRI Task H1-5. This task will enable application of the HRI Tactical Vocabulary to both general classes and specific instances of human-robot team military operations

Key Barriers and Approach:
The approach for this subtask is to survey field manuals such as the Cordon and Search – Multi-Service Tactics, Techniques and Procedures for Cordon and Search Operations (Air Land Sea Application center, 2006) or the Army Field Manuals for Cordon and Search (FM 3.06-20) and Combined Arms Operations in Urban Terrain (FM 3.06-11) for insight into relevant communication requirements and events that will drive the identification and generation of a set of mission elements that require command, control, and coordination communication. From there, a message set for SR bi-directional and team communications will be derived.

Ties to other research activities include:

- Provide communication vocabulary content for the HRI Tactical Communication Protocol
- Provide capabilities for supporting Shared Mental Models and Team Situation Awareness

Metrics for Evaluation:

- Development of a minimum of 20 vocabulary terms

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Initial communication vocabulary set.
Q2	Feedback from SRI Bi-directional Communication evaluation experiment, and refinement of vocabulary set.
Q3	Refinement of communication vocabulary set.
Q4	Feedback from SRI Team Communication evaluation experiment, and refinement of communication vocabulary set.

Subtask 2: Improve SA Measurement for Tactical HRI (UCF)

Objective and Benefits:
The objective of this subtask is to improve situation awareness (SA_w) of the SR team by creating and testing improved metrics of the Soldier's SA_w, identifying gaps in the Soldier's awareness, and supporting the robot's provision of operationally relevant information to improve the Soldier's SA_w.

Situation awareness is a mission-critical need at two levels: that of the individual Soldier (who must be aware of mission elements relevant to individual performance) and also for the SR team, as all mission elements must be accounted for and communicated appropriately. This subtask aims to see robots support Solider SA_w by participating in basic levels of situation assessment (SA_s). It is not sufficient for robots to gather information; robots must also be able to adapt to the Soldier's changing needs, identify gaps in the Soldier's SA_w, and provide operationally relevant information without explicit direction (see also H1-1 and H1-3) and do so in a way that is compatible with the Soldier's SA_w and mental models. Measurement of the Soldier's SA is needed to achieve this objective.

This research will directly support Capstone Assessment Vision, Item 9. Measurement of the Soldier's SA_w through mechanisms such as: assessing the Soldier's field-of-view, extracting meaning from explicit team communications, and assessing Soldier workload. Robot assessments will guide the *Talk* of the robot; the robot will determine the criticality and urgency of the information that needs to be communicated to the Soldier.

This effort is ongoing and will continue through FY 2013. Within FY 2012, this component of our research will provide guidance for measurement of SR team SA and develop metrics that can assess the Soldier's SA.

State of the Art:
Situation Awareness (SA) is a complex phenomenon that requires the integration of multiple individual- and team-level processes (Durso and Gronlund, 1999; Endsley, 1995; Hartel, Smith, and Prince, 1991; Salas, Prince, Baker, and Shrestha, 1995). Although measures of situation awareness exist, they are cumbersome and not suited for live environments. Little experimental work has been done to identify what is needed to develop SA_w in teams with a robotic team member.

Key Barriers and Approach:
Standard measures of SA_w and SA_s are individual, intrusive, and usually not field-ready. The essence of our approach is the analysis and integration of existing measures of SA and their applications to SR teams. We will build on existing methods in the development of a measure of SA applicable to SR team tasks such as the Capstone Assessment Vision document.

We will develop a model of team member SA_w by studying how agents could engage in non-invasive diagnoses of performance and make attributions of Soldier $SA_{w.}$ In doing this, we will establish robots as contributors to shared SA_w by integrating robot *Thinking* and Soldier SA_w.

In FY 2011, we have conducted a literature review of extant SA measurement techniques, which revealed the need for improved, mission-ready measures of SA_w and SA_s. Our FY 2010 and 2011 laboratory research has revealed that individual measurement techniques capture different facets of SA. Leveraging of these measurement techniques collectively is not only a necessary but also an achievable goal.

<u>Ties to other research activities include:</u>

- From Task H1: Mental models are a critical input to SA. Task H1 will provide mental model measurement techniques.
- *Perception* and *Intelligence* will provide robotic inputs required to (a) build robotic SA and (b) assess Soldier SA
- Provide foundational knowledge on SA and role in teams to *HRI* (e.g., trust, culture, communications)
- Provide list of cues for robot SA_s to *Perception and Intelligence*
- Provide data, tools, and simulations of SR shared SA to *Intelligence* and *ATOs*
- Provide improved methods for assessment and diagnosis of SA in SR teams to *HRI, Intelligence, Alliance, ATOs,* and *Army*
- Provide guidelines for creating Soldier-compatible robot representations to *Alliance* and *Army*

Metrics for Evaluation:

- Concurrent validity of SA measures

- Speed and fidelity of situation representation by SA_w metrics
- Robustness of SA metrics when environment changes

Over the past year, we have assessed suitability of individual-level SA_s and SA_w measurement techniques. This analysis lead to a laboratory study which investigated differences across SA measurement techniques.

We will continue laboratory experimentation with an emphasis on transitioning to cross-task simulation platforms. Intermediate progress will be demonstrated by successful collaboration with related research areas using common simulation platforms.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Integration of existing, compatible SA_w measures.
Q2	Integrated SA_w measure and empirical study of measure.
Q3	Dataset from investigation of SR team SAw measure efficacy.
Q4	Report describing empirical results, publications, and dissemination among RCTA collaborators.

Subtask 3: Analyze Team-Level SA in Tactical HRI (UCF)

Objective and Benefits:
The objective of this effort is to determine the conditions under which the metrics describing situation awareness (SA_w) of Soldier and robot can be aggregated into meaningful, process- and performance-predicting metrics of SR team-level SA_w. Situation awareness is a mission-critical need at two levels: that of the individual Soldier (who must be aware of mission elements relevant to individual performance) and also for the SR team, as all mission elements must be accounted for and communicated appropriately. This subtask supports SA_w at the SR team level by integrating the SA_w of the Soldier (developed through Soldier cognition) with the SA_w of the robot (developed through integration of Intelligence and the world model).

According to Army FM 3-21.75, “every Soldier is a sensor” (p. 9-1). This concept appropriately extends to robots – both robots and Soldiers must fight for knowledge in order to gain and maintain SA_w. In bounding overwatch, for example, mission knowledge is distributed across the two units, with each unit having unique and complementary information. We seek to greatly improve the integration of mission knowledge across the Soldier and robot by reducing the amount of explicit acts of communication required from the Soldier.

This research will directly support Capstone Assessment Vision, Item 1. Integration of the robot’s world model, the status of the environment, and knowledge across team members will support the robot’s *Thinking*. Specifically, the robot will understand commands in context. This research will also directly support the robot’s *Work* as described in Capstone Assessment Vision,

Item 8. The robot will need to integrate information from other team members in order to maintain SA in a broader context.

This effort builds upon research across the Alliance and within Task H2. Thus, this work should continue through 2014. Within FY 2012, this component of our research will develop best practices for the combination of individual-level SA measures.

State of the Art:
No standard way exists to conceptualize, aggregate, or scale team-level SA and its measures. The two currently advanced notions of SA in teams are those of *shared SA* (e.g., Endsley et al.) and *distributed SA* (Salmon et al.), which are usually presented as competing models. Further, most measures of SA_w are focused on individual-level processes. Finally, robots have been considered tools, not collaborators/teammates in research on team SA_w. Given the robot's developing abilities to *Think*, *Work*, and *Talk*, our approach promotes the robot to the level of team member and investigates optimal strategies for combining individual-level SA_w inputs to arrive at a team-level model of SA. We will advance the state of the art by developing a model and metrics of team-level SA in SR teams.

Key Barriers and Approach:
Our review of measurement techniques in FY 2011 revealed a much greater theoretical foundation for individual-level SA_w than for team SA_w. Researchers have suggested that SA in a team is more than the sum of its parts. Further, past research in other areas has shown that mere agreement is not sufficient for understanding SA of teams. Instead, our approach combines elements of *shared SA* (e.g., Endsley et al.) and *distributed SA* (Salmon et al.) and sees them as complementary, rather than competing, models.

Ties to other research activities include:

- From tasks H2-1 and H4: Explicit, especially verbal, communications will provide the 'how' of team member communication that will enable this task to investigate the 'what' of SA. Explicit communications is the medium for teamwork that is critical to building SA.
- Provide foundational knowledge on SA in teams to *HRI* (e.g., H3, culture, H4-H6, H9)
- Provide list of cues for robot SA_s to *Perception* and *Intelligence*
- Provide data, tools, and simulations of SR shared SA to *Intelligence* and *ATOs*
- Provide improved methods for assessment and diagnosis of SA in SR teams to *HRI, Intelligence, Alliance, ATOs,* and *Army*
- Provide guidelines for creating Soldier-compatible robot representations to *Alliance* and *Army*

Metrics for Evaluation:

- Scalability of metrics to teams of robots and two or more Soldiers (accuracy, predictive validity)
- Concurrent validity of different SA_w measures
- Speed and fidelity of situation representation by SA_w metrics

We are currently investigating whether indexing of SA_w measures in human teams by agreement, overlap, or correlation (similar relational patterns) is most effective. Initial results suggest that a model of shared situation understanding that combines agreement on key elements of the situation (i.e., *team SA*) with the aggregate of unique assessments by team members (i.e., *distributed SA*) is most promising.

We will continue laboratory experimentation with an emphasis on transitioning to cross-task simulation platforms. Intermediate progress will be demonstrated by successful collaboration with related research areas using common simulation platforms.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Integration of explicit communications and combinatory metrics of $SA_{w.}$
Q2	Summary of existing combinatory metrics of aggregate $SA_{w.}$
Q3	Description of laboratory study comparing aggregate SA_w measurements.
Q4	Dataset from laboratory study comparing aggregate SA_w measurements.

Subtask 4: Investigate the Role of Trust and System Transparency in Shared Mental Model (SMM) Development during Tactical HRI (ARL)

Objective and Benefits:
There is an important need to empirically test the association between mental models, system transparency, and trust in order to ascertain the extent of the relationship between these variables, to determine how they affect Soldier behavior in HRI, and to identify specific design features that can improve a Soldier's trust in a robot. Experimentation focusing on transparency and trust may provide means to mitigate possible degradation of trust when there are changes in performance due to failure or changes in environment. This idea of system transparency becomes especially important if a robot is adaptive to its environment; adaptive robots will not respond or act in the exact same way every time. The goal of this subtask is to increase our understanding of these factors in developing and maintaining trust in human-robot team collaborations through both simulation-based and live experimentation and to link it with the investigations of shared mental models in SR teams. Human-in-the-loop studies will be jointly planned by ARL and UCF personnel to examine these critical aspects of HRI and trust. In addition, the researchers will begin preliminary exploration of potential objective trust markers (i.e., physiological markers and behavioral cues). These objective measures may be vital in identifying when degradations in trust occur, therefore enabling mitigation strategies to be implemented so as to avoid negative outcomes associated with loss of trust in a system.

State of the Art:
Trust in automation has been a key area for research and development at DoD labs. For example, the Air Force Research Lab is in the process of launching a major research program on human

trust in automated systems. According to Stubbs et al. (2007), "lack of transparency in behavior" was cited as an impediment to successful human-robot collaboration. Essentially, transparency is a method by which mental models can be "tuned" or corrected and, thereby, influences the degree of trust a human has in a robot. The likelihood of a human to trust a robotic partner is directly impacted by the degree to which the human understands the system's/robot's actions and the reasoning or purpose behind such actions (robot behaviors).

Key Barriers and Approach:
Participants will collaborate with a robot to complete a mission in a simulated environment, which will likely consist of one or more of the tasks outlined in the integrated assessment (communication, moving through mobility challenges, etc.). Different aspects of the robot's behavior as well as the system transparency will be manipulated. Performance and trust variables will be measured and analyzed to determine the impacts of different approaches to system transparency. Additionally, mental model prior to interaction and post-interaction (i.e., post-transparency manipulation) will be assessed.

Potential Independent Variables (Manipulations):

- Malfunctions of the robot (scenario with malfunctions versus scenario without malfunctions)
- Level of details of communication/information (low versus high)
  - Humans need to know what it going on, but we don't want to overload them with information so performing their own tasking becomes more difficult
- Types of information robot that the robot conveys (robot-centric [ability/reliability] versus situation/mission-centric); this deals with some of our previously identified antecedents of trust
- Availability of video/images from robot (yes versus no)
- End results (robot can perform task versus robot can't perform task)
- Queries for information versus information given periodically (when errors, status updates, certain time points) versus all the time

Potential Outcome Measures (Dependent Variables):

- Human-Robot Subjective Trust Questionnaire (H3 product)
- Human's mental model of robot
- Human performance/success of collaboration
- Negative Attitudes towards Robots Scale (NARS)

Metrics for Evaluation:

- Guidelines for improving tactical HRI trust through robot/agent transparency

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:

- Communicate with RCTA collaborators (H1-3; Dr. Hancock's group) regularly about the project (e.g., experimental design, literature review, data collection and analysis, report generation, etc.).

- Present research findings at the 2013 ACM/IEEE International Conference on Human-Robot Interaction (March 2013).

Related Research:

- Provide guidelines and suggestions for HRI tactical communication protocols based on HRI metaphors (H2-1 and H4)
- Incorporation of SA and trust metrics into H1 experimentation
- Safe Operations for Unmanned Reconnaissance in Complex Environments (SOURCE) ATO

Integration and Assessment Activity:

- Provide recommendations for improving trust through appropriate robotic action transparency

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Methodology and deployment plan for experiment.
Q2	Experimentation.
Q3	Experimentation and initial data analysis.
Q4	Robot design recommendations based on empirical research findings. Technical report on research findings to date.

H4-2012 – Communication Protocol and Language Processing

Level: 6.2 Applied Research

	Name	*Organization*	*Subtask*
Principal Investigator	Norm Badler	UPenn	H4-1
Additional Investigators	Seth Teller	MIT	H4-2
	Nicholas Roy	MIT	H4-2
	Michael Wellfare	GDRS	H4-3
	Steve Sablak	GDRS	H4-3
	Ellen Haas	ARL	H4-4
	Chris Stachowiak	ARL	H4-4

Objective:
The objective of this task is to deliver preprocessed tactical communication data to relevant elements of the Intelligence and Perception architectures in order to calibrate shared mental models. Specifically, this task will leverage the advancements to the Parameterized Action Representation (PAR) model completed in 2010-2011 to develop a tactical communication protocol for SR teams.

Glossary:
Key Terms: “Gesture capture,” “Gesture encoding,” “Gesture discrimination,” “Gesturer authentication,” “Human Motion,” “Dynamic Object,” and “World Model”

Parameterized Action Representation (PAR): The PAR bridges the gap between natural language and simulation by giving an expressible, computer-understandable description of an action. PAR uses a hierarchical action representation model which enables a parser to map the components of the instruction directly for agent control.

Background:
Developing models for bi-directional communication between a heterogeneous team of Soldiers and robots requires data interchange and unification between many inter-disciplinary modules. Raw data perceived by the robot using multiple sensor modalities must be appropriately represented for cognitive processing. Knowledge acquired by individual team members (Soldiers or robots) must be communicated using a universal standard that can bridge the gap between individual mental models, a centralized world model, and cognitive architectures (e.g., ACT-R software elements).

This communication language must satisfy several requirements:

- **Expressive**: It should be able to express the wide vocabulary of raw tactical data communicated between Soldiers and robots. These include arm gestures (recognized using cameras or hand-held devices), text messages, simple voice commands, non-linguistic audio, and other multi-modal sensor data.
- **Compact**: The syntax needs to be compact for efficient data-exchange.
- **Human and Computer Readable**: The data representation should allow both human experts as well as computational tools to create, interpret, analyze, and modify the information for data generation and empirical evaluation.
- **Universal:** The data format must serve as a standard interface for data exchange between individual mental models, world models, and cognitive architectures.

To our knowledge, no paradigm of data representation and information exchange exists which satisfies the above requirements.

Research Description:
This research proposal outlines tasks that apply concepts, results, methods, and technologies developed under previous basic and applied research conducted in the first 18 months of the RCTA HRI communication and social/cultural efforts. By leveraging these research results and products, the proposed task will bridge the data exchange gap between raw tactical communication data and relevant mental model, world model, and ACT-R software elements. Empirical testing will inform iterative development of the protocol and reduce the risk to other RCTA Consortium members interested in testing the protocol or applying the protocol in IRAs or capstone assessment activities.

Metrics:
The evaluation metric for this task is the ability to capture raw communication data from multiple modalities, regardless of interface device, distill pertinent message information (e.g., content and frequency), and make that data available for processing by other RCTA architecture elements such as mental models, world model, ACT-R model, etc.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:

- Collaboration between Human Modeling and Simulation Lab (UPenn) and GRASP Lab (UPenn): Using directional gestures from a smartphone to control a robot in both simulation and real experiments.
- Collaboration between UCF and UPenn: Deliverables from above mentioned collaboration will be sent to UCF to perform human factors experiments to evaluate the efficacy of the smartphone interface.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Develop initial HRI Tactical Communication Protocol (v0.5) based on tactical communication vocabulary from Task H2-1.
Q2	Extend and refine representation based on preliminary experiments completed in Task H1, and incorporate the ability to represent individual as well as team mental models. Perform initial investigation into integration with cognitive architectures that are being developed by the Intelligence thrust.
Q3	Complete HRI Tactical Communication Protocol v1.0.
Q4	Modify HRI Tactical Communication Protocol v1.0 based upon experimental results.

Related Research:
This task leverages the implementation of the HRI Tactical Vocabulary under the Cognition Thrust and the empirical evaluation from the previous task. This task also supports interactions between shared mental models and Intelligence and Perception.

The task deliverables will provide a unifying construct for data interchange between both human and robot mental models, the centralized world model, cognitive architectures (e.g., ACT-R software elements), as well as an RCTA-relevant simulation environment (e.g., RIVET).

Integration and Assessment Activity:
This effort directly supports elements 1, 2, 7, 8, and 9 of the Capstone Assessment Vision document. Results of the task will directly be useful within the "talk" portion of the Integrated Research Assessment. This task will deliver a protocol to deliver preprocessed tactical information data to relevant elements of the Intelligence and Perception architectures and will enable bi-directional communication between SR team members. The two-year effort will involve iterative development and experimentation that will result in multiple IRA events with other RCTA performers from Intelligence, Perception, Mobility, and Integration that will further vet the HRI Tactical Vocabulary within the Capstone Assessment Vision

Subtask 1: HRI Tactical Communication Protocol (UPenn)
The subtask will enable the preprocessing of raw tactical and mission-relevant communication data to inform Perception and Intelligence (e.g., world model, ACT-R). Additionally, this capability will directly support mission-relevant shared mental models. The technical approach advances the current Parameterized Action Representation (PAR) model as follows: (1) enabling representation of different communication modalities including non-linguistic audio and simple language commands, (2) integration and implementation of HRI vocabulary, and (3) provision of interface for delivering preprocessed raw tactical communication data for use by others.

Objective and Benefits:
The primary objective of this subtask is to develop and empirically evaluate the HRI Tactical Communication Protocol. This effort directly supports elements 1, 2, 7, 8, and 9 of the Capstone Assessment Vision document. Specifically, this task enables methods for bi-directional communication between SR team members. The results from this effort will support Intelligence, Surveillance, and Reconnaissance (ISR) missions, routing to a back door, pursuing an enemy exiting a building, Cordon and Search.

This subtask is expected to be a two-year effort with iterative development and testing (experimentation) in the second year that results in multiple IRA events with other RCTA performers from Intelligence, Perception, Mobility, and Integration that will further vet the HRI Tactical Vocabulary within the Capstone Assessment Vision.

State of the Art:
Traditional knowledge representation models can be classified into two categories. Human-authored models (e.g., finite state machines, rule-based systems) are easy to hand create but cannot be easily extended to represent shared mental models between heterogeneous agents. In addition, adding new information into an existing knowledge representation often requires far-reaching changes which cannot be automated. Statistical models (e.g., neural networks, decision trees) use machine learning to emulate the behavior of real-world data. However, the generality of the models greatly depends on the representative nature of the training data. Also, the models are not human-readable and do not serve as an appropriate medium of information exchange between human Solider-robot teams.

There exists no paradigm of data representation that offers a compact, expressive, and universal language for information exchange between human and robot mental models, shared world representations, and cognitive processing architectures.

Key Barriers and Approach:
Parameterized Action Representation (PAR) offers a flexible, extensible approach to define action capabilities for single agents. However, they are affordance centric, where each PAR represents a parameterized action capability for a single agent. As part of this task, we propose the following PAR extensions:

- Extending PAR formulation to enable representation of multi-modal sensor capabilities, in addition to affordance capabilities. Sensor capabilities include arm gestures (recognized using cameras or hand-held devices), text messages, simple voice commands, non-linguistic audio, and other multi-modal sensor data.
- Implementation and integration of HRI vocabulary in PAR representation. This will enable efficient preprocessing and representation of raw tactical communication data pertinent to Solider-robot teams for use by others.

Knowledge representation using situation-centric approach. This will facilitate representation of agent-centric knowledge (individual mental models) as well as event-/situation-centric knowledge (shared team mental models). This task will leverage the advancements to the Parameterized Action Representation model completed in 2010-2011. We currently have a

simulation framework that enables the use of the extended PAR architecture. We will work closely with Task H2 to incorporate the proposed tactical communication vocabulary and conduct experiments to evaluate the vocabulary as well as communication interface.

The subtask deliverables will provide a compact, expressive, standard communication interface to help bridge the gap between human and robot mental models (HRI thrust), shared world representations, simulation environments, and cognitive architectures (Intelligence thrust).

Metrics for Evaluation:

- Current progress: Simulation framework for experimenting with PAR models has been implemented and will soon be made available open source.
- Scripting interface available for developing and designing PAR models. This will soon be extended to a graphical user interface for ease of use.
- Virtual middle-eastern marketplace scenario with villagers with social- and culture-specific parameters ready for running simulation experiments.

We will work closely with the first task to embed the HRI tactical communication vocabulary within PAR and perform human factors experiments to evaluate the validity of the vocabulary as well as the communication protocol. Experiments in the marketplace scenario to exercise the communication capabilities within the Soldier-robot team will be designed, and experiment-specific metrics (e.g., time and effort required to satisfy experiment goals) will be used to measure performance. The simulation environment can also be used to perform offline exploration and optimization of variations in vocabulary and communication protocol.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Develop initial HRI Tactical Communication Protocol (v0.5) based on tactical communication vocabulary from Task H2-1.
Q2	Extend and refine representation based on preliminary experiments completed in Task H5, Subtask 3, and incorporate the ability to represent individual as well as team mental models. Perform initial investigation into integration with cognitive architectures that are being developed by the Intelligence thrust.
Q3	Complete HRI Tactical Communication Protocol v1.0.
Q4	Provide suggestions for integration with shared world representation and cognitive architectures (e.g., ACT-R). Modify HRI Tactical Communication Protocol v1.0 based upon experimental results.

Subtask 2: Dialogue Management for Robust HRI (MIT)

Objective and Benefits:

One key challenge in robotics in general is to create natural interfaces in which robots can interpret and execute ordinary spoken commands. We are trying to develop natural language understanding that allows a mobile manipulator, such as a robotic forklift, to take commands

from a human teammate and execute the appropriate action sequence. In FY 2012, we propose to address the problem of online concept learning by using clarifying dialogue with human teammates and by using demonstration, such as by exercising environmental affordances (e.g., handles, knobs, switches, slots). This is intended to be a multi-year effort. The length of the effort will depend on the complexity of the task. This will address capabilities 1, 2, 7, 9, and 10 of the Capstone Assessment Vision document and capabilities 3 and 8 at a secondary level.

State of the Art:
For robots that interact with people using speech, the main purpose of automated speech recognition (ASR) is in a spoken dialog system (SDS) module that receives speech input and translates it into action commands or messages for other parts of the system. In today's state-of-the-art, large-scale dialog management systems, such as those used in automated call centers, the construction of dialog managers is a major engineering effort, with many person-hours devoted to developing specifications for the spoken dialog system's behavior, developing the actual system, testing extensively, and monitoring actual usage to make improvements. Given the highly rule-based, hand-crafted nature of these dialog management strategies, machine-learning researchers have, for more than a decade, focused on formulating dialog management as a model-based reinforcement-learning problem. Under this framework, the actions of a spoken dialog system can be optimized, typically by maximizing a reward function, in a data-driven manner. These approaches, it is envisioned, could automate the construction of dialog managers, lead to improved dialog behaviors, and ensure system robustness even as ASR performance degrades. Increasing attention has focused on modeling a dialog as a partially observable Markov decision processes (POMDP). The spoken dialog system POMDP (SDS-POMDP) model is perhaps most distinct in that it treats the teammates' intent in the dialog as the underlying, hidden "state" and speech recognition hypotheses as noisy observations of that intent.

Key Barriers and Approach:
The essence of our technical approach is decision-theoretic planning for human-robot dialogue. Our approach depends on a distribution over the intentions of human teammates. If the distribution is uncertain, then robot actions will be expected to perform poorly (i.e., will have a high expected cost) relative to asking clarification questions. We compute the expected cost of actions using a partially observable Markov decision process. While the SDS-POMDP framework has many attractive properties, many open problems need to be addressed. We propose to address two of these. First, the existing state of the art assumes that a single observation is received at each time step and all observations are conditionally independent given the (hidden) state. However, ASR systems provide a distribution over possible hypotheses, and the observations are *not* conditionally independent in that environmental context and history are often valuable clues to the accuracy of the recognized text. We propose to extend the POMDP model to include continuous observations in the form of confidence scores, which allow the robot to determine if the ASR system has returned a reasonable recognition or if clarification questions are required to reduce the uncertainty in the estimate of the teammate intentions.

Secondly, the existing state of the art assumes a small, finite number of possible teammate intentions, in contrast to the probabilistic inference model described in H2-4 that can capture an unbounded number of possible intentions. We propose to integrate the probabilistic inference

model with the decision-theoretic planner for dialogue, allowing the robot to automatically determine where the spoken instructions are ambiguous or uncertain, leading to poor expected performance. This integration will allow us to generate clarification questions when the inferred grounded of the teammate instructions from H2-4 are potentially ambiguous. Additionally, this model will allow the system to acquire new concepts from demonstration by human trainers, including gestures and narration.

This research also leverages results from previous research programs in our group.

Metrics for Evaluation:
Our primary metric is the correct execution of natural language commanded plans, where the meaning of some words is unknown or ambiguous. Our primary form of evaluation is to collect example commands from untrained people and determine if the robot correctly executes the command or generates an interaction question for more information. For example, we will present an ambiguous scene and instructions to Mechanical Turk, along with the robot's proposed response, and query whether the robot's response is reasonable.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Extend data collection to interactive dialogue for a corpus of representative supervisory command utterances for information delivery and autonomous mobility within a semi-structured or unstructured environment.
Q2	Demonstrate confidence scoring model on corpus of supervisory command utterances.
Q3	Extend POMDP model to incorporate confidence scoring. Demonstrate improved performance in simulated dialogue scenario.
Q4	Extend POMDP model to incorporate dialogue with supervisory command understanding (H2-4). Demonstrate improved performance on symbol grounding and instruction following.

Subtask 3: Gesture Commands For UGVs with Identification, Authentication, and Feedback (GDRS)

Key Terms: "Gesture capture," "Gesture encoding," "Gesture discrimination," "Gesturer authentication," "Human Motion," "Dynamic Object," and "World Model"

Objective and Benefits:
Gesture recognition is just one of many robotic functions which requires precision measurement and tracking of moving bodies. As a result, robots will be continually monitoring the humans in their vicinity. We will explore segmentation of human activities intended to contribute to gesture-based communications from other human activities. Our plan does not include direct processing of sensor data but begins with summarized data from a prototype Common World Model data structure. This Common World Model is intended to be universal in nature and not focused on a specific gesture recognition task. Low-cost mechanisms provided as part of the

UGV will supply appropriate feedback to ensure that the communications loop is closed and that gestures are being monitored.

State of the Art:
Critical elements of operations are command and coordination communication events. A gesture-based HRI Tactical Vocabulary does not currently exist but will be developed by the HRI thrust area in related subtasks. These methods are focused on gesture recognition during discrete experiments and do not address periods of communication and non-communication.

Key Barriers and Approach:
We will adopt available gesture recognition software, available either from RCTA HRI and Perception research or from open source libraries using the Kinect sensor. This sensor is used as a surrogate for stereo vision and/or LADAR imaging by the robot. We will evaluate various gesture vocabularies developed by the HRI thrust area (Task H5) for robustness to different types of terrain, lighting, noise levels, and crowd configurations. We will perform studies and develop databases of communicative and non-communicative behavior. The development of an unambiguous "start command" signal, perhaps a gesture or action involving a smartphone button press or a voice command, will be investigated. Low-cost and secure methods for identification of the individual giving the command will be developed, contributing to development of a method to authenticate, i.e., ensure that the person giving the commands is authorized to do so and is trustworthy.

Current techniques focus on recognition of a continuous series of a few gestures and do not address false detection of gestures during a non-communicative period. Currently, the vocabulary is ad hoc and does not address potential issues with ambiguity, periods of other human activity, and authentication of the human's command authority.

Much of the research value of this subtask lies in pulling together existing capabilities in a well thought out and systematic way. The overall technical risk is low, but the value of the effort is high. This is a Human-Robotic Interaction task, which has some association with Perception activities as well as some association with the Intelligence task of following a dismounted Soldier.

Metrics for Evaluation:
The primary metric will be the false communications rate versus time against several baseline human activities. Specifically, examples of Warfighter non-communicative activities will be investigated to determine if special issues arise with this population. Other metrics include the contributions of unambiguous communication and non-communication periods to authentication. This subtask will be closely coordinated with other HRI and Perception tasks that are studying the process of detecting the gestures.

Related Research:
Collaborating tasks include:

1. Intense collaboration with Dr. Tony Stenz's Intelligence team, especially with Mr. Robert Dean in the areas of Common World Model development under I2-1, to which we are

specifically providing a detailed interface with the emerging Dynamic Object capabilities.

2. Collaboration with Dr. Martial Herbert's team in Perception, especially with the detection of human motion models through the use of Kinect and stereo pair sensors in conjunction with Dr. Bradley Stuart and Mr. Robert Landle of GDRS (Integration and Assessment) and Dr. Max Bajracharya of JPL (P6-1).

3. Research collaboration with other participants in Task H4, including subtasks 1 and 2, and if practical, adoption of common gesture and vocal vocabularies and authentication systems. Methods are investigated to associate data from auxiliary devices (such as the smartphones used in related work) with direct perception of humans and their locations. Data from experiments will be provided to MIT, UPenn, UCF, and others to help assess the completeness of the planned Dynamic Object representations as a solid baseline for enabling human motion recognition in the RCTA Framework with the Common World Model.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Develop databases of communicative and non-communicative human activities in coordination with Task H5. Define algorithm architecture in the context of the Common World Model in coordination with Intelligence thrust area.
Q2	Demonstrate ability of Common World Model motion detection and tracking models to contribute to gesture recognition.
Q3	Demonstrate gesture period segmentation and authentication capability. Integrate capability onto an RCTA platform.
Q4	Integrated assessment demonstration.

Subtask 4: Multi-modal Controls and Displays for Soldier-Robot Interaction (ARL)

Objective and Benefits:
This subtask seeks to determine the extent to which multi-modal displays (visual, audio, tactile) and controls (speech and touch) interact in a human-robot interface. Three elements of interest are involved in this research: 1) Which multi-modal display/control combinations provide the shortest response time between display and control action; 2) Which control/display combination provides the shortest inter-command time (time between the onset of a first control action, relative to the onset of a second, sequential, dependent control action); and 3) The effect of map objects of different sizes (large static objects such as roads versus small static objects such as intersections) on user touch command response time and accuracy. In addition, research will explore whether there is a preferred control modality (a specific multi-modal control action is performed first more often) and whether this is a function of display type and/or individual

differences. Answering these questions will determine which multi-modal display and control commands go together most efficiently and effectively.

Answering these research questions will help define multi-modal control and display parameters that will allow a smoother fusion of display and commands, which can allow the human-robotic interface to more easily interpret operator intent, as well as support and reduce system error. Knowledge of these parameters can result in enhancing the performance of Soldier-robot teams in using multi-modal display/control interfaces.

State of the Art:
Researchers like Oviatt, Coulston, and Lunsford (2004) as well as Oviatt, Lunsford, and Coulston (2005) explored multi-modal controls, but few researchers have explored the use of multi-modal controls and displays in a dynamic map display interface.

Key Barriers and Approach:
Multi-modal displays and controls can be used in Soldier-robot team operations. Multi-modal audio and tactile displays, alone or together, can supplement visual displays to signal alarms, elements of interest (i.e., targets or waypoints), and critical battlefield elements. Multi-modal controls such as speech or touch can be useful as inputs on dynamic map displays that show changes in battlefield objects over time. For example, Soldier speech can define which objects are placed on the dynamic map, while touch can be useful in specifying the location of map objects.

Few researchers have explored issues that might affect the use of both multi-modal displays and controls in a human-robot dynamic map display interface. One issue involves the extent to which different multi-modal displays influence the use of different multi-modal controls. Another issue involves the potential role of individual differences in user preference in order of use of touch and speech commands. Answering these questions will help define control and display parameters that will allow a smoother fusion of display and commands, which can allow the robotic interface to more easily interpret operator intent as well as support and reduce system error.

Independent variables include the type of multi-modal display prompts (audio, tactile, audio and tactile) as well as different types of map objects (roads, intersections) that can involve different levels of difficulty in providing touch commands.

Participants will use multi-modal displays and controls to perform a robot control task in a simulated environment. Response time and accuracy will be measured, and individual differences in use of modality to input commands will be explored.

Metrics for Evaluation:
Research metrics include the following dependent variables: 1) User response time (time between the onset of initial multi-modal alarm and the onset of a first multi-modal control action); 2) Inter-command time (time between the onset of a first control action relative to the onset of a second, sequential, dependent control action); 3) The number of accurate touch and

speech command inputs; and 4) The type of the first control action most preferred by each individual for each trial.

Improved understanding of the extent to which multi-modal displays and controls interact in a human-robot interface. Answering these questions will determine which multi-modal display and control commands go together most efficiently and effectively. Data from this research will also help define multi-modal control and display parameters that will allow a smoother fusion of display and commands, which can allow the human-robotic interface to more easily interpret operator intent, as well as support and reduce system error.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Conferences include the 2012 International Conference of Multi-modal Interface. Collaborative activities will include sharing and discussing experimental results with HRED and RCTA researchers, including Dr. Stephanie Lackey and Mr. Daniel Barber.

Related Research:
For this research, we are considering the novel approach of using multi-modal displays and controls, with which the user is prompted by multi-modal displays, to perform commands using more than one modality. This is a new approach that can enhance human performance over non-flexible control interfaces, such as pushbuttons (Oviatt, Darrell, and Flickner, 2004), and that will allow support of variable sets of control tasks. Voice and touch controls (touch screens) are considered because multi-modal interfaces that fuse speech and touch input have been shown to substantially reduce speech recognition errors (Oviatt, 2000). Oviatt found that users prefer to interact multi-modally with speech and touch when given a choice (multi-modal or uni-modal), especially in spatial map tasks. Oviatt (1996) noted that control interfaces using combined spoken and touch input may be particularly effective for interacting with dynamic map systems.

Integration and Assessment Activity:
Research can be integrated into the Robotics CTA Human-Robotic Interaction tasks, specifically tasks H4: Communication Protocol and Language Processing and H5: Evaluating Tactical Command and Coordination Vocabulary and Protocols.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Attend International Conference on Multi-modal Interface, and present paper describing previous multi-modal research.
Q2	Perform FY12 multi-modal experiment.
Q3	Analyze FY12 experimental data.
Q4	Write report of FY12 research results.

H5-2012 – Evaluating Tactical Command and Coordination Vocabulary and Protocols

Level: 6.1 Basic Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Stephanie Lackey	UCF	H5-1, H5-2, H5-3
Additional Investigators	Daniel Barber	UCF	H5-1, H5-2, H5-3
	Lauren Reinerman-Jones	UCF	H5-1, H5-2, H5-3

Objective:
The first objective of this task is to vet the command and coordination section of the HRI Tactical Vocabulary through a series of empirical experiments. The purpose of these experiments is two-fold. First, the experiments will evaluate the efficacy of the HRI Tactical Vocabulary content. Second, the experimental results will contribute to the refinement of the content for team interactions. Initially, command and coordination vocabulary (e.g., acknowledging receipt of command, "copy;" ending a specific message, "over;" ending a conversation, "out"), and as the HRI Tactical Vocabulary develops, other mission-relevant vocabulary terms will be added.

The second objective of this task is to empirically assess the effectiveness of the HRI Tactical Communication Protocol for digitally interfacing between raw communication data sources and other HRI and RCTA entities such as Intelligence, Perception, and Mobility.

Toward the Capstone Vision:
This task integrates into the RCTA Capstone Vision through participation in Integrated Research Assessments (IRAs). Results from this task will be used within the "Talk" portion of the IRA for Intelligent Navigation, Autonomous ISR, and Tactical Team Movement, enabling team members to send commands to robot team members and be informed about team and mission status for shared situation awareness. In 2012, experimental results will provide "accept and understand" movement orders and will provide robot status. Follow-on efforts in 2013 and beyond will enable more complex information sharing from robot to Soldier on an as-needed basis.

Background:
Traditional HRI and the current state of the art in military operations involves a human operator explicitly controlling or supervising an unmanned asset using a Human Computer Interface (HCI) (Barnes & Jentsch, 2010). Within this paradigm, humans do not interact with the unmanned asset as another team member when co-located in a real-world dynamic environment (e.g., combat). Previous efforts by the U.S. Army Research Lab (ARL) show the additional burden this model places on operators and the team. For example, having a remotely located commander can affect the team member's confidence in decision making and reduce the ability to fully understand the commander's true intent (Grey, Redden, Coovert, & Elliot, 2008). Moreover, adding robotic team members further increases the complexities of communication by adding an additional step in the process (Cosenzo, Capstick, Pomranky, Dungrani, & Johnson, 2009). That is instead of being able to communicate directly to the team members, the commanders must communicate through a robotic operator first.

In order to decrease the cognitive load required from the human team members and subsequently reduce the time it takes to issue a command, intuitive multi-modal communications (MMC) must be implemented to replace some of the traditional control features (Pettitt, E.S., & Carsten, 2009). Identification of different methods of communication across visual, auditory, and tactile modalities resulted from an examination of the literature, communications with ARL researchers (e.g., Dr. Ellen Haas), and a scientific workshop on MMC held in December 2010. For example, voice commands have shown reduced operation time in discrete robotic tasks (Redden, Carstens, & Pettitt, 2010). Additionally, auditory cues were shown to increase performance within robotic tasks when used separately and in conjunction with tactile cues (Gunn, Warm, Nelson, Bolia, Schumsky, & Corcoran, 2005; Haas, 2007). Arm and hand gestures are a natural and intuitive method of communication between humans (Wexelblat, 1995) and can be quicker and feel more intuitive than manual controls (Guo & Sharlin, 2008). Tactile communication via skin stimulation is a method under investigation because it is hands-free, stealthy, and shown to increase situational awareness, decrease sensory overload, and reduce response time due to its egocentric directional cueing (White, 2010; Elliott, Coovert, Prewett, Walvord, Saboe, & Johnson, 2009; Hutchins, Cosenzo, McDermott, Feng, Barnes, & Gacy, 2009). Although there has been extensive research within the tactile, visual, and audio modality separately, there is limited work and research performed where they are combined into multi-modal displays for dismounted operations (Haas & van Erp, 2010). Leveraging the state of the art in visual, auditory, and tactile communication research, this task investigates bi-directional modalities using an HRI Tactical Vocabulary for integration into HR teams.

Research Description:
The research conducted under this task is composed of basic research tasks that result in a series of experiments that will inform theoretical constructs as well as deliver recommendations, processes, and prototype capabilities that can be instantiated and tested by other Consortium members individually and at Capstone or Integrated Research Assessments. Additionally, this task will provide documented recommendations for requirements and design, and content recommendations for HRI Tactical Vocabulary and the HRI Tactical Communication Protocol.

Metrics:

- Evaluation of up to 20 vocabulary terms
- Prototype software and hardware capabilities that can be made available to the larger Consortium
- Algorithms for channel(s) selection (robot to human) and classification (human to robot)

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:

- Collaborative task between UCF and UPenn
- Collaborative consultation to identify potential gaps in SoA for improved perception of HRI Communications in future tasks
- Seminar conducted at Human Factors and Ergonomics Society Conference 2012
- Seminar/demo provided at Army Science Conference 2012

Milestones:

	Quarterly Milestones and Deliverables
Q1	Bi-directional SRI Experiment: Design experiment, and collect data.
Q2	Bi-directional SRI Experiment: Analyze data, and develop recommendations document for HRI Tactical Vocabulary development task.
Q3	Small Team SRI Experiment: Design experiment, and collect data.
Q4	Small Team SRI Experiment: Analyze data, and update recommendations document for HRI Tactical Vocabulary development task. Complete Effectiveness of HRI Tactical Communication Protocol in Bi-directional and Small Team SRI experiment.

Related Research:

This task supports the requirements, design, and content development of the HRI Tactical Vocabulary under Task H2. Experimental results will provide requirements and design input for the HRI Tactical Communication Protocol. It is expected, through collaborative consultation with Perception, that gaps in the SoA of static and dynamic scene understanding in relation to HRI will be identified for inclusion in future RCTA task areas.

Integration and Assessment Activity:

This task will supply the second task in the HRI Bi-directional and Team Communication Thrust with data, recommendations for requirements, design, and implementation of the HRI Tactical Vocabulary elements into the HRI Tactical Communication Protocol. The HRI Tactical Vocabulary development task will also benefit from this experimentation since the primary objective is to evaluate the content. Finally, the results from the experiment will inform RCTA

researchers how to use technologies (e.g., accelerometers, compass, and audio) provided by devices such as a smartphone within Capstone and IRA activities.

Subtask 1: Evaluating HRI Tactical Vocabulary for One-on-One SRI (UCF)
This experiment will include one human and one robot collaborating to complete a mission-relevant task. Multiple modes of communication will be evaluated for use in this experiment, including gestures from human to robot and tactile and/or text from robot to human. The use of a smartphone as a communication device is anticipated and will leverage the previous work of UPenn and UCF. This experiment will be conducted in a simulation environment. Specifically, this experiment includes a human using a device (e.g., smartphone) to communicate with a robot (live or virtual) in a simulated environment. Results from this experiment will be delivered to the HRI Tactical Vocabulary content development task in the form of a recommendations document. Communication methods (e.g., algorithms) and prototype technologies used under this experiment will be made available for use by other RCTA researchers and IRA events.

Objective and Benefits:
The primary objective of this subtask is to empirically evaluate the initial command and coordination section of the HRI Tactical Vocabulary. Additionally, this subtask serves to investigate audio, visual, and tactile interface device(s) enabling SR teaming and robust, natural bi-directional communication methods for noisy or non-line-of-sight (NLOS) situations.

This effort directly supports elements 1, 2, 7, 8, and 9 of the Capstone Assessment Vision document. Specifically, this subtask enables methods for bi-directional communication between SR team members. The results from this effort will support Intelligence, Surveillance, and Reconnaissance (ISR) missions, routing to a back door, pursuing an enemy exiting a building, and Cordon and Search.

This subtask is expected to be a two-year effort with progressive experimentation in the second year that results in multiple IRA events with other RCTA performers from Intelligence, Perception, Mobility, and Integration that will further vet the HRI Tactical Vocabulary within the Capstone Assessment Vision.

State of the Art:
Critical elements of operations are command and coordination communication events. An HRI Tactical Vocabulary does not currently exist but will be developed by another HRI task. Thus, the systematic verification of this vocabulary within bi-directional SR team communication will advance the SoA of the vocabulary itself and its implementation within SR team environments. Additionally, this subtask advances the current understanding of effectiveness of complex communication in dynamic mission-relevant situations.

Key Barriers and Approach:
The approach for this subtask is to conduct an independent experiment to assess the HRI Tactical Vocabulary pertaining to command and coordination.

- Develop or leverage communication capabilities using smartphone technologies to support human-to-robot (e.g., gesture) and robot-to-human (e.g., text, audio) communication.
- Develop experimental design implementing one-to-one bi-directional SR communication task using innovative communication modalities and information found from initial HRI Tactical Vocabulary development.
- Conduct experiment using up to 60 human participants using simulated robotic platform.
- Summarize experimental results and recommendations for HRI Tactical Vocabulary and HRI Tactical Communication Protocol.

This thread of research provides systematic verification of the HRI Tactical Vocabulary within bi-directional SR team members as it pertains to the Capstone Assessment Vision. In order to minimize the risk of this complex problem space, an interdisciplinary team of human factors researchers, computer engineers, culture experts, and robot technology specialists has been assembled. This team has the experience necessary to execute well-designed experiments with the best technology available. Further, this task extends results from research performed under the 2011 APP and foundational work completed by researchers at the Army Research Lab (ARL).

Ties to other research activities include:

- Provide communication methods for integration in HRI Tactical Communication Protocol under HRI Subtask H4-1
- Support HRI Tactical Vocabulary developed under HRI Tasks H1 and H2
- Provide capabilities for supporting Shared Mental Models and Team Situation Awareness
- Provide requirements and prototypes for Perception, Intelligence, and Integrated Research Assessments

Metrics for Evaluation:

- Evaluation of up to 20 vocabulary terms
- Prototype software and hardware capabilities that can be made available to the larger Consortium
- Algorithms for channel(s) selection (robot to human) and classification (human to robot)

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Design experiment, and collect data.
Q2	Analyze data, and develop recommendations document for HRI Tactical Vocabulary development task.
Q3	None – task complete.
Q4	None – task complete.

Subtask 2: Evaluating HRI Tactical Vocabulary for Small Team SRI (UCF)
This subtask represents the logical progression following Subtask 1. This experiment will include multiple humans (up to four) and one robot collaborating to complete a mission-relevant task. This subtask will leverage heavily from the previous subtask. Thus, several similarities in format exist, but it will be tailored to the team environment rather than one-on-one bi-directional interaction. Multiple modes of communication will be evaluated for use in this experiment including gestures from human to robot and audio/tactile and/or text from robot to human. The use of a smartphone as a communication device is anticipated and will leverage the previous work of UPenn and UCF. This experiment will be conducted in a live or simulated environment. Results from this experiment will be delivered to the HRI Tactical Vocabulary development task, and communication methods used under this experiment will be made available for use by other RCTA researchers and IRA events.

Objective and Benefits:
The primary objective of this subtask is to empirically evaluate the initial command and coordination HRI Tactical Vocabulary within an SRI team of up to four humans and one robot. Similar to Subtask 1, this effort serves to investigate audio, visual, and tactile interface device(s) enabling SR teaming and robust, natural bi-directional communication methods for noisy or non-line-of-sight (NLOS) situations.

This effort directly supports elements 1, 2, 7, 8, and 9 of the Capstone Assessment Vision document. Specifically, this task enables methods for bi-directional communication between SR team members within ISR and/or Cordon and Search missions.

This subtask is expected to be a two-year effort with progressive experimentation in the second year that results in multiple IRA events with other RCTA performers from Intelligence, Perception, Mobility, and Integration that will further vet the HRI Tactical Vocabulary within the Capstone Assessment Vision.

State of the Art:
An HRI Tactical Vocabulary does not currently exist but will be developed by another HRI task. Thus, this subtask extends the work under the previous subtask to extend the evaluation of the HRI Tactical Vocabulary to SR teams with multiple humans. Advancement of SoA of HRI vocabulary terms and how to implement them within team constructs and dynamic missions will result.

Key Barriers and Approach:
The approach for this subtask is to conduct an independent experiment to assess the command and coordination HRI Tactical Vocabulary within SR teams.

- Develop or leverage communication capabilities using smartphone technologies to support human-to-robot (e.g., gesture) and robot-to-human (e.g., text, audio) communication.

- Develop experimental design implementing multiple-to-one bi-directional SR communication task using innovative communication modalities and information found from initial HRI Tactical Vocabulary development.
- Conduct experiment using up to 120 human participants using live or virtual robotic platform with simulated intelligence.
- Summarize experimental results and recommendations for HRI Tactical Vocabulary and HRI Tactical Communication Protocol.

This thread of research serves as phase two of the systematic verification of the command and coordination HRI Tactical Vocabulary within SR teams as it pertains to the Capstone Assessment vision. Similar to Subtask 1, an interdisciplinary team has been assembled. This team has the experience necessary to execute well-designed experiments with the best technology available. Further, this task extends results from research performed under the 2011 APP, foundational work completed by researchers at the Army Research Lab (ARL) and Subtask 1.

Ties to other research activities include:

- Provide communication methods for integration in HRI Tactical Communication Protocol under HRI Task H4
- Support HRI Tactical Vocabulary developed under HRI Tasks H1 and H2
- Provide capabilities for supporting Shared Mental Models and Team Situation Awareness
- Provide requirements and prototypes for Perception, Intelligence, and Integrated Research Assessments

Metrics for Evaluation:

- Evaluation of up to 20 vocabulary terms tailored to team interaction
- Prototype software and hardware capabilities that can be made available to the larger Consortium
- Algorithms for channel(s) selection (robot to human) and classification (human to robot)

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Consult with HRI Tactical Vocabulary development team to identify initial vocabulary set.
Q2	Consult with Subtask 1 for experimental recommendations.
Q3	Design experiment, and collect data.
Q4	Analyze data, and update recommendations document for HRI Tactical Vocabulary development task.

Subtask 3: Evaluating HRI Tactical Communication Protocol for Bi-directional and Small Team SRI (UCF)

This subtask represents the logical progression following Subtasks 1 and 2 above and Task 2 in this research thrust (HRI Bi-directional and Team SR Communication) and the previous subtask.

This experiment will include a condition with a one-to-one human-to-robot relationship and a condition that includes multiple humans (up to four) and one robot collaborating to complete a mission-relevant task. This subtask will leverage heavily from the previous task experiments. Thus, several similarities in format exist, but it will be tailored to the evaluate hypotheses pertinent to the testing of the protocol. Multiple modes of communication will be evaluated for use in this experiment including gestures from human to robot and tactile and/or text from robot to human. The use of a smartphone as a communication device is anticipated and will leverage the previous work of UPenn. This experiment will be conducted in a live or simulated environment. Results from this experiment will be delivered to the HRI Tactical Vocabulary development task, and communication methods used under this experiment will be made available for use by other RCTA researchers and IRA events.

Objective and Benefits:
The primary objective of this subtask is to empirically evaluate HRI Tactical Communication Protocol within an SRI team of up to four humans and one robot. This effort directly supports elements 1, 2, 7, 8, and 9 of the Capstone Assessment Vision document. Specifically, this task enables methods for bi-directional communication between SR team members within ISR and/or Cordon and Search missions.

This subtask is expected to be a two-year effort with progressive experimentation in the second year that results in detailed refinement of the protocol.

State of the Art:
An HRI Tactical Communication Protocol does not currently exist but will be developed by the previous subtask. Thus, this subtask extends the work under the previous subtask to evaluate the HRI Tactical Communication Protocol. Advancement of SoA of HRI communication protocols and interfaces in dynamic SRI missions will result.

Key Barriers and Approach:
The approach for this subtask is to conduct an independent experiment to assess the HRI Tactical Communication Protocol in one-on-one HRI and team interactions as well.

- Develop experimental design leveraging Subtask 1 designs
- Conduct experiment using up to 120 human participants using a live or simulated robotic platform
- Summarize experimental results and recommendations

Similar to Subtask 1, an interdisciplinary team has been assembled. This team has the experience necessary to execute well-designed experiments with the best technology available. Further, this task extends results from research performed under the 2011 APP, foundational work completed by researchers at the Army Research Lab (ARL) and Subtask 1.

Ties to other research activities include:

- Provide requirement, design, and implantation recommendations to HRI Tactical Communication Protocol under HRI Task H4-1

- Support HRI Tactical Vocabulary developed under HRI Task H1 and H2
- Provide capabilities for supporting Shared Mental Models and Team Situation Awareness
- Provides requirements and prototypes for Perception, Intelligence, and Integrated Research Assessments

Metrics for Evaluation:

- Verification of the HRI Tactical Communication Protocol performance for bi-directional and team interaction

	Subtask 3 Quarterly Milestones and Deliverables
Q1	None – task not started.
Q2	Design empirical evaluation experiment based upon preliminary review of Protocol v0.5.
Q3	Receive HRI Tactical Communication Protocol v1.0.
Q4	Complete Effectiveness of HRI Tactical Communication Protocol in Bi-directional and Small Team SRI experiment.

References:

Air Land See Application Center, (2006). *Cordon and Search.* Multi-Service Tactics, Techniques and Procedures for Cordon and Search Operations FM 3-06.20, MCRP 3-31.4B, NTTP 3-05.8, AFTTP(I) 3-1.62.

Barnes, M., & Jentsch, F. (Eds.). (2010). *Human-Robot Interactions in Future Military Operations.* Ashgate.

Cosenzo, K., Capstick, E., Pomranky, R., Dungrani, S., & Johnson, T. (2009). *Soldier Machine Interface for Vehicle Formations: Interface Design and an Approach Evaluation and Experimentation.* (ARL-TR-4678). *Aberdeen Proving Ground, MD: U.S. Army Research Laboratory* , 1-40.

Elliott, L. R., Coovert, M. D., Prewett, M., Walvord, A. G., Saboe, K., & Johnson, R. (2009). *A review and meta-analysis of vibrotactile and visual information displays* (ARL-TR-4955*).* Aberdeen Proving Ground, MD: U.S. Army Research Laboratory, 1-36.

Grey, A. A., Redden, E. S., Coovert, M. D., & Elliot, L. R. (2008). Empowering followers in virtual teams: Guiding principles from theory and practice. *Computers in Human Behavior, vol 24* , 1884-1906.

Gunn, D. V., Warm, J. S., Nelson, W. T., Bolia, R. S., Schumsky, D. A., & Corcoran, K. J. (2005). Target acquistion with UAVs: Vigilance displays and advanced cueing interfaces. *Human Factors*, 47(3), 488-497.

Guo, C., & Sharlin, E. (2008). Exploring the use of tangible user interfaces for human-robot interaction: A comparative study. Proceedings of *CHI'08*, ACM, 1-10.

Haas, E. C. (2007). Integrating auditory warnings with tactile cues in multimodal displays for challenging environments. *Proceedings of the 13th International Conference on Auditory Displays*, Montreal, Canada, 127-131.

Haas, E. C., & van Erp, J. B. (2010). *Human-Robot Interactions in Future Military Operations.* (M. Barnes, & F. Jentsch, Eds.) Ashgate.

Hutchins, S., Cosenzo, K., McDermott, P., Feng, T., Barnes, M., & Gacy, M., (2009). An investigation of the tactile communications channel for robotic control. *Proceedings of Human Factors and Ergonomics Society Annual Meeting,* 53(4), 182-186.

Pettitt, R., Redden, E.S., & Carsten, C. (2009). *Scalablity of Robotic Controllers: Speech-based Robotic Controller Evaluation* (ARL-TR-4858). *Aberdeen Proving Ground, MD: US Army Research Laboratory* , 1-46.

Redden, E. S., Carstens, C. B., & Pettitt, R. (2010). Intuitive speech-based robotic control (ADA519652). *Aberdeen Proving Ground, MD: U.S. Army Research Laboratory*, 1-50.

Wexelblat, A. (1995). An approach to natural gesture in virtual environments. *ACM Transactions on Computer-Human Interaction*, *2*(3),179-200.

White, T. (2010). *Suitable body locations and vibrotactile cueing types for dismounted soldiers*. (ARL-TR-5186*).* Aberdeen Proving Ground, MD: U.S. Army Research Laboratory, 1-38.

H8-2012 – Social Dynamics Modeling and Simulation

Level: 6.2 Applied Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Michael Goodrich	BYU	H8-1, H8-2
Additional Investigators	Ron Jenkins	GDRS	H8-2
	Stephen Fiore	UCF	H8-3
	Jonathan Streater	UCF	H8-3
	Ladislau Boloni	UCF	H8-4

Objective:
The overarching objective of this thrust is to produce models and simulations that are able to capture the complex social dynamics associated with Soldier-robot, Soldier-civilian, and robot-civilian interactions. In Subtask *H8-1: Design and Analysis of Information-Based Behavior Management in Soldier-Robot Teamwork*, our objectives include the identification of the spatio-temporal information needs and representations suitable for coordinating behaviors between squad-level Soldiers, robots, and other remote Soldiers and who use semi-automated resources to identify behaviors of those around them. In Subtask *H8-2: Soldier-Robot Team Organization*, our objective is to help characterize organizations in existing and planned military HR teams and to define how our findings can fit into and inform evolving U.S. Army operations with realistic robots. In Subtask *H8-3: ACT-R for Modeling Dimensions of Social Intelligence in Human-Robot Teams*, our objective is to develop ACT-R/machine learning based guidelines, algorithms, data structures, and analyses for addressing social dimensions of Soldier-robot, Soldier-civilian, and robot-civilian interactions. In Subtask *H8-4: Simulating Dynamic Environments with Social, Cultural Parameters*, our objectives include the creation of a real-time, interactive simulation of a socially and culturally accurate middle-eastern marketplace scenario that is populated with event-driven personnel who exhibit rich interactions and the use of this simulation environment as a testbed to perform human factors experiments. In Subtask *H8-5: Robot-Soldier Coordination in the Presence of Civilians*, our objective is to develop a model of interaction between a robot with a group of Soldiers and civilians. We will develop different behavior models of the robot so that it is able to recognize Soldier and civilian movements and consider the cultural sensitivities of the civilians.

Background:
This work is focused on understanding and supporting collaborative human-robot interactions in a complex operational environment. To contextually ground our collective efforts, we are using the Cordon and Search task as a notional scenario. The Cordon and Search task is an ideal environment in which to explore our models and simulations because it allows for the

instantiation of the research constructs we are exploring. Further, it is of sufficient complexity such that we can examine modeling and simulation involving proximate Soldier-robot interactions, Soldier-civilian interactions, and civilian-robot interactions. In short, small-unit action is embedded in a larger tactical context with more human stakeholders than are present in just the small-unit. Working with GDRS, we will characterize existing military organizational structures to ground this work in U.S. Army operations. In this context, we consider two key organizational assumptions that drive our modeling and simulation approach and how information is represented and shared among multiple humans: (a) at least one remote Soldier may have input to robot behavior and (b) other Soldiers, especially commanders, will have access to sensor information obtained from the robot.

In addition to examining the dynamics of the military squad structure, we also address the more complex social dynamics which emerge in operational environment. Our research will investigate how to implement decision/inference trees within the ACT-R cognitive architecture and train it via supervised learning to make operationally successful team-level decisions and use this to make conclusions about how to use socio-cognitive models that inform robot design. Further, prior work in simulating virtual agents has focused on simulating the macroscopic flow of homogeneous crowds or dealing with complex fully-articulated virtual humans. There exists no work that can simulate heterogeneous autonomous agents which exhibit complex interactions with the environment, other agents, as well as human avatars. Last, our approach allows us to examine how robots are able to recognize the activity performed by the team. In particular, robot team members must have a model of the activity such that it can infer the role of the individual Soldiers. The complement to this is that such models allow the Soldiers to simply communicate the nature of the activity to the robot and its own role in it. As such, research must help us understand how to recognize more than just teamwork patterns.

Research Description:

This task addresses a set of inter-related research questions associated with modeling and simulating social dynamics associated with human-robot teams. This includes identifying and designing representations that allow remote robot supervisors and information consumers to understand the spatio-temporal situations and tasks of the robots and proximate Soldiers and developing metrics for classifying civilian behaviors to “actively sense” intent. Additionally, this task will provide machine learning based guidelines, algorithms, data structures, and analyses for addressing social dimensions, especially social cues and related context factors, of Soldier-robot, Soldier-civilian, and robot-civilian interactions. Similarly it will create a real-time, interactive simulation of a socially and culturally accurate middle-eastern scenario with event-driven personnel who exhibit rich interactions with the environment, other members of the virtual populace, and human-robot Solider teams in a manner that conforms to social and cultural norms of the region.

This Modeling and Simulation of Social Dynamics task moves beyond the state of the art in the following ways: (1) It helps to identify canonical spatio-temporal needs and representations for joint proximate and remote supervisory control of heterogeneous robot teams; (2) It seeks to manage social variables not yet addressed in the RCTA cognitive models; (3) It pursues tasks

that build a framework for simulating dynamic environments with social and cultural parameters and performing human factors experiments in a simulation environment in order to extract models of interaction of a robot with a group of Soldiers and civilians; and (4) It extends existing work for the robot to participate in the Soldier's teamwork, adapt to the human teamwork, while taking into consideration the civilians and their cultural sensitivities.

Metrics:
Each subtask within H8 describes the metrics it will use to evaluate progress.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
There will be substantial collaboration between UCF, UPenn, BYU, GDRS, and ARL given the interaction required for this work. While no staff rotations are anticipated, there will likely be site visits to foster an understanding of particular research environments. Additionally, these visits will provide an opportunity for guest seminar lectures that foster an increased understanding of each team's research efforts.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Develop simulation of context understanding within a defined vignette using ACT-R.
Q2	Simulation of increased complexity context understanding (robot identifies its role and position with a fixed formation).
Q3	Code, train, and experiment with initial models (simulated robot supports Soldier movement in a defined mission).
Q4	Experiment with additional models and scenarios.

Each subtask within H8 describes the milestones it will use to evaluate progress.

Related Research:
Our subtask H8-1 is related to recent advances in robot autonomy which have caused a dramatic shift in how humans interact with these robots. This shift in interaction has two facets: For remote operations, tele-operation of remote vehicles has been replaced by supervisory control wherein much of navigation is delegated to autonomous algorithms. This allows the human to emphasize mission-level interactions, such as managing payloads [Cummings, 2007]. For proximate interactions, unmanned ground vehicles (at least in laboratory settings) can follow, explore in near proximity to, or act as a "mule" for a human through simple activity or gesture recognition. This opens-up the possibility of an unmanned ground vehicle acting as a collaborator in realistic military situations.

These recent advances are enabling, but most prior work has focused on isolated interactions between one human and one vehicle. These may be appropriate for operation-level military interactions, but tactical-level and strategic-level considerations suggest that there will be multiple *stakeholders* in what the robot senses and what tasks it performs [Woods]. These stakeholders include both proximate and remote Soldiers who may share responsibility for robot operations and tactical sensor interpretation, but it may also include other *information consumers*, such as a commander, who can use sensor information to make tactical decisions affecting multiple Soldiers [Goodrich & Schultz, 2007]. To support effective teaming, all "interactants" will need a common vocabulary for interpreting, for example, spatio-temporal behaviors and objects. As such, this effort directly links with the "Shared Cognition" and "Comms" thrust in the development of this common vocabulary. In the case of H8, these spatio-temporal behaviors include immediate things like tasking a robot so that its sensor footprint samples areas of interest at desired resolutions and regular intervals, and they also include more far-reaching things like interpreting or modifying the behavior of *bystanders* who may exhibit differential behavior depending on whether they are civilians or enemy combatants.

Our efforts in H8-3 build upon FY11 work that is still developing in H9 in elucidating the concept of social intelligence for human-robot teamwork. Thus, H8-3 will be instantiating frameworks, scenarios, and data developed in H9 as well as providing modeling and learning results to H9. Additionally, H9 is related to both I3 and H1-3 in their collaborative attempt to instantiate and learn shared mental models using the ACT-R cognitive architecture. As such, there will be close ties between H8-3 and these tasks. Finally, as the developing framework hopes to be extremely relevant both to HRI communications tasks and to HRI shared mental model/shared awareness tasks, there are direct linkages here as well in terms of conceptual frameworks and experimental and modeling results.

In other elements of H8, the interpretation of the team action performed by the team members will be done using probabilistic models (primarily Dynamic Bayesian Networks), which need to be coordinated with the other approaches used in tasks I3 and P5. The behavior of the civilians is "gain probabilistic" with a distribution which can be handled only with mixture models, as it is likely that the civilians are divided into several sub-populations (friendly, adversarial, etc.). This modeling must take input from and can provide output to tasks I3 and P5. Similarly, the work in H8-4 is related to previous work at UPenn where they developed CAROSA, a software framework for creating and simulating virtual populations. CAROSA used the PAR ontology to define agents, which is currently being extended to support the new requirements in various RCTA thrusts. A new software foundation (ADAPT) is currently under development which will allow larger virtual populations and more detailed personal attributes, gestures, communication signals, and interpersonal behaviors. The work in H8-4 is closely linked to H9-5. In particular, the experiments carried-out as part of the H9-5 task will leverage the simulation framework being developed in H8-4. The simulation framework will be used to experiment and evaluate new mental models which can directly be used by the Intelligence thrust.

Integration and Assessment Activity:
Task H8 will also support Integrated Research Assessments for the following program elements:

- IRA1: Automatic Understanding and Semantic Labeling of Environment
- IRA9: Cognitive Common Ground for Human-Robot Teams
- I3: Combining Cognitive and Probabilistic Reasoning
- I10: Maximizing Performance with Minimal Resources
- P5: Dynamic Scene Understanding

Task H8 can support Integrated Research Assessments and the subsequent scenario-based assessment:

- IRA1: Automatic Understanding and Semantic Labeling of Environment
- IRA9: Cognitive Common Ground for Human-Robot Teams

Subtasks H8-1 and H8-2 provide inputs to the following:

- HRI Comms efforts on HRI Interface for Tactical Communication, particularly in developing a vocabulary for communication.
- P5: Dynamic Scene Understanding, by identifying those spatio-temporal elements that are necessary for supporting tactical-level, team-based interactions.

Subtask H8-3 provides potential inputs and/or outputs to the following program elements:

- I3: Combining Cognitive and Probabilistic Reasoning
- H1-3: Create Initial ACT-R Representations of Selected Human Team Member Mental Models

Subtask H8-4 is closely linked to other H8 subtasks, Subtask H9-5, as well as to Intelligence and Perception:

- Experiments carried-out as part of the H9-5 task will leverage the simulation framework being developed in H8-4.
- Simulating interactions between a virtual populace and a team of human-robot Soldiers also provides invaluable insight into individual as well shared mental models, which can be leveraged by the intelligence thrust (IRA1 and IRA9).
- The unification language developed as part of the HRI Bi-directional and Team Soldier-Robot Communication Thrust will also leverage the Parameterized Behavior Tree (PBT) paradigm used for authoring the virtual populace.
- It must coordinate with and possibly exchange data with I3.
- It provides support for Integrated Research Assessments IRA1 and IRA9.
- Simulating interactions between a virtual populace and a team of human-robot Soldiers provides invaluable insight into individual as well shared mental models, which will be leveraged by the intelligence thrust (IRA1 and IRA9).
- The Parameterized Behavior Trees (PBTs) used to author the virtual populace in H8-4 will be used as the basis for designing and defining individual and shared mental models and specifying preprocessed communication data. The applied research being done as

part of the HRI Bi-directional and Team Soldier-Robot Communication Thrust will directly benefit from this task.

Subtask 1: Design and Analysis of Information-based Behavior Management in Soldier-Robot Teamwork (BYU)

Objective and Benefits:
The objectives of this work are (a) to identify spatio-temporal information needs and representations that are suitable, as an initial vignette, for coordinating behaviors between squad-level Soldiers, two robots, and other remote Soldiers who are stakeholders in the outcome of a Cordon and Search task, and (b) to use the air and ground resources to identify milling or escape attempts by nearby civilians.

1) Organize information needs within the Soldier-robot team.
    a) Identify spatio-temporal information needs to remote operators and remote information consumers. Incorporate Soldier-robot teaming findings from 2011 tasks. Identify information needs in the context of squad-level organization, roles, and Soldier equipment (e.g., GPS, radio) expected in the 2015 through 2025 time period.
    b) Develop GUI-based methods for using spatio-temporal representations to ground communications between remote and proximate interactants (Soldiers remotely interacting with unmanned vehicles and Soldiers proximately interacting with those vehicles).
    c) Support HRI Communications tasks on developing HRI vocabulary by identifying essential spatio-temporal representations for communicating relevant awareness between remote and proximate interactants.
2) Organize decisions and tasking within the Soldier-robot team.
    a) Develop GUI concepts that will allow remote interactants to task robots to regularly monitor key observation points.
    b) Develop autonomy and decision-support concepts that assist remote interactants in optimally distributing air and ground sensors (restricted to one air platform and one ground platform).
    c) Support work on developing HRI vocabulary by identifying essential spatio-temporal representations for tasking robots, thereby supporting communication between remote and proximate interactants.
3) Support subtask H9-5 with human/robot squad recognition of civilian behavior in Cordon and Search.
    a) Develop metrics for classifying milling or escape attempts by civilians during a Cordon and Search task.
    b) Support military social environment modeling by using metrics to classify milling or escape attempts from imagery obtained by air and ground robots.
    c) Explore active sensing of milling or escape attempts by sending an air or ground robot near civilians exhibiting milling or escaping behaviors.

(4) Define and conduct task assessments, and incorporate 1), 2), and 3) into the RCTA IRAs.

State of the Art:
Much work has been done on GUI-based tasking single robots by single operators, but much is unknown about how to coordinate supervisory control of semi-autonomous vehicles by a mix of proximate and remote Soldiers, each of whom is a stakeholder in the problem and some of whom are merely information consumers of robot sensor information. With respect to milling and escape attempts, much vision-based work for classifying group behaviors exists, but this work could benefit from recent advances in using metrics derived from algebraic graph theory as applied to biological systems.

Key Barriers and Approach:
As noted above, most prior work has focused on isolated interactions between one human and one vehicle. These may be appropriate for operation-level military interactions, but tactical-level and strategic-level considerations suggest that there will be multiple stakeholders in what the robot senses and what tasks it performs. To support effective teaming, all "interactants" will need a common vocabulary for interpreting spatio-temporal behaviors and objects. These spatio-temporal behaviors include immediate things like tasking a robot so that its sensor footprint samples areas of interest at desired resolutions and regular intervals, and they also include more far-reaching things like interpreting or modifying the behavior of bystanders who may exhibit differential behavior depending on whether they are civilians or enemy combatants.

We are confident that we can contribute to applied research in this area for the following reasons. First, we have extensive experience in developing and analyzing robot systems that include multiple team members and air/ground coordination [Goodrich et al., 2009]. Second, we have designed multiple user interfaces that allow operators to manage UAVs and UGVs by managing sensors or information [Lin & Goodrich, 2010]. Third, we have done much prior research on organizational-level constraints on human-robot teaming, including work in traditional fan-out models and in human interaction with bio-inspired teams [Crandall et al., 2006; Whetten & Goodrich, 2009; Goodrich et al., 2011].

Metrics for Evaluation:
- Improved taxonomy of spatio-temporal information needs for tasking robots and interpreting sensor information.
- Increased task performance of shared proximate/remote Soldier teams over either type of team alone.
- Soldier role task load levels and role accession
- Soldier equipment capabilities to support the defined information exchanges

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Identification of spatio-temporal needs for understanding context from robot sensors.
Q2	Identification of methods for tasking robots to satisfy spatio-temporal needs.
Q3	Technical report and preliminary evaluation of those needs.
Q4	Prototype GUI software for understanding and tasking air and ground robots in Cordon and Search task. Metrics for milling and escape attempts. Task experiments and IRA support.

Subtask 2: Soldier-Robot Organization (BYU and GDRS)

Objective and Benefits:
BYU and GDRS will collaborate to characterize organizations in existing and planned military HR teams and to define how the organizational findings from Subtask H8-1 could fit into and inform evolving U.S. Army operations with realistic robots. The results will be used to support the definition and execution of RCTA experiments at the task and subtask levels and integrated research assessments. This involves identification of existing and future military tasks which can be used for scenario development and model testing. Tasks would be categorized along scenario-relevant parameters such as heterogeneous teams with multiple human and robot team members. The time period for evaluation coincides with the current Army Capstone Concept for 2016-2028 (TRADOC Pam 525-3-0). The integration goal for this is to assess the theories of system stability and system robustness in terms of a proposed human-robot teaming.

State of the Art:
This subtask exists to ensure our research is grounded in the appropriate operational context so SoA is not necessarily relevant. SoA in this context means that research results have the potential to impact real military operations within a five- to ten-year horizon.

Key Barriers and Approach:
This subtask exists to ensure our research is grounded in the appropriate operational context so it requires timely input from GDRS with regard to military organizational structure.

Metrics for Evaluation:
We will use the following metrics to measure success:

- We will create a report that identifies an operationally feasible implementation of human-robot teaming in a Cordon and Search task. This will consist of a detailed vignette of how humans and robots fulfill roles and exchange information over a set of reasonable contingencies.
- We will create a technical report identifying required autonomy, mobility, and communication capabilities of unmanned assets that could realistically be used in an

integrated assessment activity that demonstrates human-robot teaming in a Cordon and Search task.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Preliminary vignette of human-robot teaming in Cordon and Search task.
Q2	Report on human-robot teaming requirements and implementation.
Q3	Identification of realistic robots capable of participating in Integrated Research Assessment activity.
Q4	Simulation of Cordon and Search for human-robot teaming that supports the Integrated Research Assessment objectives. Support IRAs organizational and Soldier-robot interaction experiment design and execution.

Subtask 3: ACT-R for Modeling Dimensions of Social Intelligence in Human-Robot Teams (UCF)

Objective and Benefits:
This subtask seeks to provide RCTA research teams ACT-R/machine learning based guidelines, algorithms, data structures, and analyses for addressing social dimensions, especially social cues and related context factors, of Soldier-robot, Soldier-civilian, and robot-civilian interactions. That is, we will explore the computational use of social information and how it could be skillfully integrated and managed within the processes of human and robot social intelligence in social situations (e.g., processes directly associated with communicating and coordinating with team members) and also how it is possible to use social information to make solutions to other robotics problems (e.g., task-oriented problems like identifying objects to navigate around) easier, more useful, more confident, or more tractable.

State of the Art:
While more work is starting to focus on using machine learning techniques for perception and modeling of social dynamics (e.g., behavior or facial recognition, social signal processing), there is no coherent approach which attempts to combine the guidance of an overarching model of social intelligence (developed in H9 for 2011), the constraints of a cognitive architecture, and the inference and learning power of machine learning techniques. Further, this approach seeks to manage and combine social variables not yet addressed in the RCTA cognitive models

Key Barriers and Approach:
- Explore machine learning/ACT-R modeling paradigms for use in successful socio-cognitive models of social intelligence (interpret, manage, signal).
- Explore levels of abstraction and representation for learning, decision-making, and communicating using socio-cognitive models (e.g., choosing and developing scenarios, cues, and levels of abstraction like decision/inference trees).
- Explore/learn ways of combining socio-cognitive cues to make effective social decisions and/or inferences that might not otherwise be effective taken separately.

- Implement domain-based, theoretically-derived models of social intelligence, interaction, and social signal processing as developed in H9 for 2011 as well as consider how computational modeling might alter or contribute to those frameworks depending on their learning performance (tested e.g., via classification accuracy in supervised learning involving social cues) in human-Soldier-civilian scenarios.
- Try to have the model learn shortcuts.
- Explore Integration of ACT-R and social intelligence modeling with other RCTA groups for collaborations, for example, on graphical simulations using ADAPT and Unity, on Shared Mental Model performance, or on Perception problems.

This modeling effort could potentially influence every thrust involved with Human-Robot Interaction as it directly draws from and plays into shared mental model and communication parts of Soldier-robot teamwork, both for robot design and human training. It might also play into (supposedly) less socially-oriented robot problems such as Perception and Intelligence by (socially) informing computational problems on which these groups work. Thus, this could conceivably inform capabilities 1 through 11 of the capstone. The approach also follows the approach begun by collaboration between Task I3: Combining Probabilistic and Cognitive Modeling and the HRI work on Shared Mental Modeling, applying it specifically to social cues and social intelligence.

Metrics for Evaluation:

- Completion of preliminary ACT-R socio-cognitive framework
- Testing of ACT-R social model in simulation
- Data structures, algorithms, scenarios, and data used for cognitive model development
- Publication of experimental results

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Exploratory ACT-R coding. Literature review.
Q2	Exploration of representations, social intelligence framework, and modeling scenario possibilities.
Q3	Code, train, and experiment with initial models.
Q4	Data collection, analysis, and complete/explore additional model representations, scenarios, and techniques (e.g., learning expertise short cuts, use ACT-R results to analyze H9 model, use social information to inform nonsocial robotics problems).

Subtask 4: Simulating Dynamic Environments with Social, Cultural Parameters (UCF)

Objective and Benefits:

The objective of this subtask is to provide a model of interaction among a robot and a group of Soldiers and civilians. The Soldiers and the robot are forming a search party participating in a Cordon and Search operation. In the ongoing scenario, we assume a group of two to four

Soldiers joined by a robot. We assume the physical capabilities of the robot being roughly equivalent to BigDog. We will consider both the cases of a Cordon and Kick (where the patrol unit is performing search without previous contact with the civilians) and the Cordon and Knock/Ask type (where the patrol is telling the civilians about the search it will conduct or asking permission to conduct a search). We assume that the searched area is an urban area of streets, possibly a busy market with a large local population.

We shall develop different behavior models of the robot, which needs to recognize the movement formation of the Soldiers and identify whether it needs to participate or stay out of way. The robot must consider the cultural sensitivities of the civilians.

- Robot supports a fixed formation of Soldiers by staying close but avoiding positions where it will interfere with the movement of the Soldiers.
- Robot recognizes the fixed formation taken by Soldiers and identifies its own position in it.
- Robot supports/participates in the movement of the Soldiers in the presence of civilians in a Cordon and Ask scenario (minimum interference).
- Robot supports/participates in the movement of the Soldiers in the presence of civilians in scenarios of various level of force (Cordon and Kick/Knock/Ask).

We shall focus on real-time interactive simulation which allows any of the Soldiers, civilians, and the robot to be controlled by either an autonomous agent or a human. We will work with Norm Badler's group at UPenn for the simulation framework and authoring tools.

This effort will contribute to the following capabilities in the capstone vision:

1. Understand command
    - We are looking for implicit understanding, identifying by team activity recognition what activity the Soldiers execute, and what the robot's role in that is (which can be as simple as stay out of the way or as complex as taking the place of a human participant)
2. Repeat back commands
    - As the model of the robot of the task it executes is based on cultural norms, this can be communicated.
4. Moving autonomously
    - This task mainly focuses on the interaction of the robot with civilians, not only who are moving but also how their movement is affected by the presence of the robot. The movement of the robots with respect to the civilians must reflect the nature of the mission (Cordon and Kick/Knock/Ask), as well as the balancing of cultural sensitivities with the objectives.

Existing research is used to, at best, recognize teamwork patterns. This work will extend the existing work for the robot to:

- Participate in the Soldier's teamwork
- Complement it with a task that adapts to the success of the human teamwork
- Take into consideration the civilians and their cultural sensitivities

State of the Art:
Human-robot coordination is in early stages, and in general, it assumes the coordination between the robot and its human controller. Most current work ignores the fact that other human players are present at the scene of action. There is significant work in team activity recognition (including by the PI), but it has not yet been applied to models where the robot participates in the recognized action.

Key Barriers and Approach:
The robot needs to perform team activity recognition in order to recognize the activity performed by the team. It must have a model of the activity, and it can infer the role of the individual Soldiers. Alternatively, the Soldiers can simply communicate the nature of the activity and their role in it to the robot.

- This activity builds on our previous team activity recognition work, funded by ARL and NSF. See: http://hector.cs.ucf.edu/teamwork/.

We have extensive previous experience with team activity recognition, including recognizing roles:

- L.J. Luotsinen and L. Bölöni. Role-Based Teamwork Activity Recognition in Observations of Embodied Agent Actions. Proc. of the Seventh Intl. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS-08), pp. 567–574, Estoril, Portugal, May 2008.

The prior work on Task H10 (cultural /social models) is good preliminary work for the interaction with the civilians. A significant part of the preliminary work for the simulation had been demonstrated through the contributions at UCF and UPenn.

Metrics for Evaluation:

- Human factors experiments with users
- Expert testimony about the realism of the simulation

Measurement of progress toward ultimate goal: Cultural models of interaction. We plan to measure intermediate progress and determine whether the research is on track to enable the capstone capability:

- Provide open-source framework to other teams for evaluation and integration
- Conduct human factors experiments in collaboration with H4 and H9, and collect metrics for evaluation
- Publish and present important contributions at regular checkpoints
- Identify and communicate requirements and/or recommendations for Perception, Intelligence, and Integration

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Robot supports a fixed formation of Soldiers by staying close but avoiding positions where it will interfere with the movement of the Soldiers.
Q2	Robot recognizes the fixed formation taken by Soldiers and identifies its own position in it.
Q3	Robot supports/participates in the movement of the Soldiers in the presence of civilians in a Cordon and Ask scenario (minimum interference).
Q4	Robot supports/participates in the movement of the Soldiers in the presence of civilians in scenarios of various level of force (Cordon and Kick/Knock/Ask).

H9-2012 – Social Dynamics Experimentation

Level: 6.1 Basic Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Stephen Fiore	UCF	H9-1, H9-2, H9-3, H9-4
Additional Investigators	Shaun Gallagher	UCF	H9-1, H9-2
	Peter Hancock	UCF	H9-4
	Deborah Billings	UCF	H9-4
	Patricia Bockelman Morrow	UCF	H9-1, H9-2, H9-3
	Jonathan Streater	UCF	H9-2
	Kristin Schaefer	UCF	H9-4
	Norm Badler	UPenn	H9-5
	Susan Hill	ARL	H9-6

Objective:
This task consists of an integrated set of experiments which seeks to empirically examine the social dynamics emerging within human-robot teamwork. In this context, we address the following needs of the RCTA:

- Identification of data parameters for anticipated human-robot behaviors (Soldier and civilian)
- Articulated means for integrating perception and context through behavior identification
- Values and metrics for social processes which can be used within models for robotics architectures
- Creation of a real-time, interactive simulation of a socially and culturally accurate middle-eastern marketplace scenario that is populated with event-driven personnel who exhibit rich interactions and the use of this simulation environment as a testbed to perform human factors experiments

Background:
This task will organize its experimental subtasks around the framework generated by H9 in the 2012 APP. This initial research identified that a key component of human-robot interaction can be conceptualized via juxtaposition and integration of two complementary ideas within the cognitive sciences – direct perception with reflective perception:

- *Direct Perception cues: Gesture, facial expression, language, ecological context. Embodied.*
- *Reflective Perception cues: Narrative, socio-political context, reasoning processes which result in interpreting another's motives or thoughts. Embodied.*

Both forms of perception influence the judgments and actions of humans. Consequently, they become points of research for understanding the Soldiers in human-robot military teams, and they also can provide input to the development of robot cognitive algorithms and architectures. Based upon our initial theorizing, our framework has integrated the need of programmers to rely upon quantifiable metrics for their architectures. This has been accomplished through the use and analysis of perceptive units. These have the following benefits:

1) Perceptive units are in contrast to elements; here, the key issue is that "elements" require static conditions and become obsolete with evolving contexts.
2) Perceptive units build upon the notion that meaning is an evolving and often contextually determined process (Vygotsky, 1986).
3) Perceptive units are characteristics which may not be further divided but retain properties of the whole; that is, units are irreducibly composite (McNeill, 1992).
4) Perceptive units provide quantifiable variables (e.g., interpersonal distance, volume, gait, etc.; ratios of multiple signals such as roles, context, etc.).
5) The approach requires norms against which standard units may be established; this allows for ratios and thresholds.

Our framework also takes into account the following additional important background considerations:

1) The social behaviors of the robot and how the robot is perceived by others (e.g., perceived form, appearance, attributes, and perceived function and affordances) can impact the civilian's subsequent decisions, actions, and behavior.
2) Research has shown that the size and speed of a robot can impact how a bystander interacts with and trusts the robot (Tsui, Desai, & Yanco, 2010).
3) Research indicates that the social behaviors of humans can, in fact, reflect the amount of trust a human has in another entity.
4) Social distance between humans and robots: Bickmore and Cassell (2001) suggest the possibility that as trust builds between partners, the acceptable social distance will change.
5) Attentional cues: A behavioral marker of trust may include attentional cues, or the extent to which a human attends to the robot (e.g., eye gaze).

Research Description:

This task will expand and develop the existing framework into a functional representational model of social signal cognition for HR teams. Towards this end, we have a multi-pronged experimental approach for shaping and evaluating the model using the following subtasks:

H9-1: Social Dynamics Model – Experimental Framework Development (UCF-Fiore)
H9-2: Social Dynamics Model Experimentation (UCF-Fiore)

H9-3: Experimentation with Soldier-Robot and Civilian-Robot Interactions: Integrating Context and Cues to Determine Intention (UCF-Fiore)
H9-4: Empirical Investigation of Social Dynamics and Trust in Soldier-Robot Teams (UCF-Hancock/Billings)
H9-5: Simulating Dynamic Environments with Social, Cultural Parameters (UPenn-Badler)
H9-6: Social and Cultural Impact on HRI (ARL)

In addition to theory evaluation, the results of these experiments will provide data for algorithms that can be implemented in robot programming and simulation environments. Additionally, the research findings can be used to help Soldiers predict human behavior (Soldier and civilian) so that safety and success can be better supported.

Metrics:
Each experiment has its own metrics (see below), but they are unified in the expectation that the research teams follow protocols in the design and approval processes for human experimentation, maintain appropriate records of data, and record finding with accuracy. Additionally, successful progress will be measured by the communication of results, submission of findings for peer review, and implementation of outcomes in the RCTA project.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
There will be substantial collaboration between UCF researchers and ARL given the interaction required for this work. While no staff rotations are anticipated, there will likely be site visits to foster an understanding of particular research environments. Additionally, these visits will provide an opportunity for guest seminar lectures that foster an increased understanding of each team's research efforts.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Host collaborative meeting to receive feedback regarding model, functionality, and considerations; make changes to model to address points from meeting (H9-1). Methodology and deployment plan for experiments (H9-2, H9-3, H9-4.1). Develop real-time agent simulation framework (v0.5) with programmatic authoring of scenarios and events (H9-5).
Q2	Distribute conceptual model for researchers in Social Dynamics; continue structuring model (H9-1). Experimentation (H9-2, H9-3, H9-4.1). Methodology and deployment plan for experiment (H9-4.2). Extend and refine simulation framework based on feedback from initial experiments (v1.0) (H9-5). Design and author a simple set of scenarios that can be used to perform initial human factor experiments with human-robot Solider teams in collaboration with UCF (H9-5). Develop a graphical user interface for agent and event authoring (H9-5).
Q3	Incorporate preliminary findings from experiments into model (H9-1). Data analysis from experiments (H9-2, H9-3, H9-4.1). Experimentation (H9-4.2). Refine authoring interface based on feedback from initial experiments (H9-5). Extend agent simulator to include multi-modal communication capabilities (H9-5).
Q4	Distribute updated model with experimental findings (H9-1). Completion of remaining analysis (H9-4). Perform human factor experiments (H9-5). Design specific scenarios for exercising communication and coordination between human-robot Soldier teams (H9-5). Communication of robot design recommendations based on empirical research findings to related thrusts. Technical report on findings to date for experimental research to the Consortium; submissions of full reports for peer review (all H9 subtasks).

Related Research:

- Perception: Behavior Recognition and Understanding
- Perception: Scene Context Recognition and Understanding
- Intelligence: Behavior Confabulation and Communication from Robot to Human and other Robots
- Task H8: Generating Adaptable Tactical Behaviors
- Work from Task H8 to analyze use of human team metaphor for Soldier-robot military teams and organizations by providing metrics and applying topographical overviews
- Support analysis of narratives used in H8 as they apply to dynamic social environments
- Simulation environment of Task H8 as a potential testbed for social signal model
- Advance work in gesture recognition (H4, P6) to provide framework for meaning

- Provide categorization for communication methods (H6)
- H9-5 receives input from P5, and it can provide output to the same task with respect to the interpretation of the scene with the mental models of participating humans

Integration and Assessment Activity:
This task provides input to Perception in support of how to direct visual attention to areas that are used for non-verbal communication, such as the hands and face (e.g., task P3). This task provides input to Intelligence in support of how to use cultural context to determine meaning behind non-verbal symbols (e.g., task I2). Additionally, in identifying gestures and social cues relevant to human-robot teaming in military contexts, this task works in collaboration with Task H5 and later with Task H6. Finally, this task provides input and/or output with I2: Data Mapping for Inference and Focus, I3: Combining Cognitive and Probabilistic Reasoning, I10: Maximizing Performance with Minimal Resources, and P5: Dynamic Scene Understanding. Finally, H9 can support Integrated Research Assessments and the subsequent scenario-based assessment with IRA1: Automatic Understanding and Semantic Labeling of Environment and IRA9: Cognitive Common Ground for Human-Robot Teams.

Subtask 1: Social Dynamics Model – Experimental Framework Development (UCF)

Objective and Benefits:
- Develop a rudimentary social intelligence framework that can support a practical model of social cognition for support of Soldier-robot interactions during team missions
- Develop a model that will provide information for more accurate decision and learning algorithms used in artificial agents in military contexts
- Expand FY11 H9 framework involving perception, language, and decisions to inform algorithmic development for situation awareness
- Articulate the critical social aspects of Soldier-robot interaction which would support team collaboration and coordination
- This effort primarily supports capabilities 1, 2, 3, 4, 5, 6, 7, 9, and secondarily all remaining aspects of the Capstone Assessment Vision Document:
    - The model uniquely examines the social and embodied aspects of meta-cognition and narrative which are key aspects of 1, 2, 6, 7, and 9
    - Semantic perception, learning, and adaptive behavior generation affect both the cognitive aspects of the robot and the practical mobility of the machine of 4, 5, and 8
- The model will be useful beyond the Capstone since it includes the social components needed for other missions, such as Counter-IED, Urban ISR patrol, and Logistics

State of the Art:
A key component of this model is the integration of direct perception with reflective perception:
- Direct Perception cues: Gesture, facial expression, language, ecological context. Embodied.

- Reflective Perception cues: Narrative, socio-political context, reasoning processes which result in interpreting another's motives or thoughts. Embodied.

Both modalities of perception influence the judgments and actions of humans. Consequently, they become points of research for understanding the Soldiers in human-robot military teams, and they can also provide development direction for robot cognitive platforms and algorithms.

Furthermore, the model assists in the design direction for moving from pure interface with a robot team member to interaction with the robot team member.

- The role these factors play in a given phenomenon depends on how one describes and observes
- These factors impact the "experience of engagement" (DeJaeger et al., 2010), which is essential for the robot's separate engagement with civilians and Soldiers
- Follows a scale of specificity, which our experiments will seek to define so that parameters and quantifiers may be distributed to the RCTA teams
- The model recognizes that interaction (versus interface) demands autonomy "to take on a life of its own" – if one agent takes on the sole responsibility of regulating the engagement, then it is no longer a social interaction

Key Barriers and Approach:
Design and deliver a high-level model for social interaction necessary in SR tactical engagement:

- Present model in visual form
- Provide insight regarding the ratios and hierarchies intrinsic to social understanding
- Communicate with RCTA researchers the qualitative and quantitative implications derived from the model
- Provide information for algorithms to support modeling of social-cognitive factors in either tabular form or in form of decision trees

The approach moves beyond the state of the art by:

- Providing a model which can be integrated into various programming aspects of the robotic team member, including decision trees for existing robotic cognitive architectures
- The end product will include ratios and hierarchies of human perception which can inform robotic world models
- The model uniquely addresses the challenges of integrating perception and social awareness, which is needed to identify threats in field

Indicators of success include:

- We have assembled a collaborative team of human factors researchers, computer engineers, and robot technology specialists who have the experience to execute well-designed experiments with the best technology available
- Incorporation of Army literature and vocabulary to assure that concepts are communicated in a manner which will apply directly to military contexts

- The model is based on substantial body of empirical literature in cognitive psychology, social sciences, and artificial-agent behavior
- The model answers requests from RCTA engineers for guides which address socio-communicative challenges, such as the integration of attended objects and behavior responses

Our approach ties into additional research:

- This task is developed through ongoing collaborative efforts with the other RCTA thrusts
- This task directly addresses attention, perception, narrative, and culture
- This task also includes relevant social considerations for Army culture during dangerous situations

Metrics for Evaluation:

- Completion of model
- Distribution of model and supporting literature to related RCTA teams, including colleagues working in perception and intelligence
- Submission of model for peer review
- Coordination with simulation teams to test components of model
- Distribution of information for algorithms informed by model predictions

The following progress has been made toward the goals:

- Based upon our year one efforts, a framework has been developed which accommodates the integration of direct perception via sensory paths and reflective perception (i.e., narrative and culture) and will serve as the basis for the proposed model
- Collection of relative literature
- Conceptual collaboration from experts in social and cognitive sciences

We plan to measure intermediate progress and determine whether the research is on track to enable the capstone capability by:

- Ongoing checks with other researchers to maintain a closed feedback loop
- Development and execution of timeline for completion of model
    - Periodic assessment of utility of information generated for algorithms to support modeling of social-cognitive factors

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Host collaborative meeting to receive feedback regarding model, functionality, and considerations. Make changes to model to address points from meeting.
Q2	Distribute conceptual model for researchers in Social Dynamics. Continue structuring model.
Q3	Incorporate preliminary findings from experiments into model.
Q4	Distribute updated model with experimental findings. Submit article for peer review.

Subtask 2: Social Dynamics Model Experimentation (UCF)

Objective and Benefits:
This subtask is divided into two experimental parts:

Part 1 – An evaluation of inference and interpretation associated with observing action in humans and robots
Objectives:

- Support the development of social perception in Soldier-robot teams
- Test components of social models which directly impact SR interactions for tactical executions
- Assess behavioral responses to robot activity
    - Participant responses provide data ranges which can be directly used to:
        - Guide robot programming for eliciting responses associated with matching observation of others with understanding of behavior of others
            - This study will identify robot signals which, even if subtle, will prime the human observer for certain behaviors or understandings
        - Test the validity of simulation testbeds for human-robot interactions. We hypothesize that elicitation of human interaction responses in the simulation which are normal in the real world will lead to higher reliability of findings in the simulated HRI studies
- Provide an experimental outlet for testing of Social Signal frameworks developed in FY11 under H9
- Provide a comparison of laboratory, simulation, and live HRI environments
    - An important aspect of this study is to see if intention-reading activity is similar in simulated and live contexts

Part 2 – Experimentation with SR teams aimed at identifying and anticipating human responses to shifts in context
Objectives:

- Manipulate aspects of the Capstone scenario to identify context factors which will elicit responses in Soldiers, robot, or civilians

- Use a factorial design experiment to control for contextual shifts which are likely to occur during a real-world example of the scenario
    - Compare Soldier response to information from robot and human team members
    - Compare social signals for team action if the "building" is a residence versus a business
        - We hypothesize that behaviors will shift due to contextual changes
        - The data gathered from these shifts can support predictive models
        - The findings can also be used to support shared mental model construction and maintenance
- Provide an empirically supported and statistically parameterized set of context factors
- Support the ongoing efforts of robot programming and shared mental model development by collaborating with these research areas
- Test social perception framework developed in FY11 under H9 to refine accuracy of the social perception model
- Articulate the critical social aspects of Soldier-robot interaction which would support team collaboration and coordination

State of the Art:
Based upon our year one efforts, a framework has been developed which accommodates the integration of direct perception via sensory paths and reflective perception (i.e., narrative and culture). The framework includes behaviors which, even if unacknowledged by the observer, can lead to intuitive interaction with a robot.

Research needs to more fully examine the understanding of intention and the human capacity to match perception and action.

Key Barriers and Approach:
This approach moves beyond the state of the art by:

- This approach considers the effect of robotic versus human movement on the perception systems using multiple assessment methods
- This approach will allow for the collection of quantifiable parameters which can be leveraged in programming the world models
- These experiments can test the interaction of direct perception (via immediate, embodied, sensory channels) and reflective perception (via narrative, socio-political norms, memory)
- The end product will include ratios and hierarchies of human perception which can sustain Shared Mental Models during situational/contextual changes
- The experiment provides straight-forward data classifications which will help robot developers meet and surpass capstone goals

Indicators of likely success include:

- We can move beyond state of the art to identify movement ranges which provide the best social signals.

    - This research also can be applied to simulations to compare the understanding of intention and any associated inferences during social interaction in a virtual environment
- We have assembled a collaborative team of human factors researchers, computer engineers, and robot technology specialists who have the experience to execute well-designed experiments with the best technology available.
    - UCF has the simulation and multi-media equipment available to execute the experiment successfully
- The experiments will expand and modify successful research to address the specific problems of Army mission relevancy.
    - Includes incorporation of Army literature and vocabulary to assure that concepts are communicated in a manner which will apply directly to military contexts
    - The experiment builds on context information from military engagements, sociological research, and studies in HRI
    - The data directly responds to requests from RCTA engineers for guides which address socio-communicative challenges, such as the accurate perception of situations

The approach ties into other research activities by:

- This subtask is developed through ongoing collaborative efforts with the other RCTA thrusts
- This subtask directly addresses organizational (Shared Mental Models) and communication (vocabulary, comprehension) concerns
- This subtask also includes relevant social considerations for Army culture during dangerous situations

Metrics for Evaluation:

- Design of experiment and approval from IRB and Army
- Completion of experiment
- Submission of findings to RCTA colleagues

The following progress has been made toward achieving our goal:

- Development of framework which identifies the interconnected roles of direct and reflective perception
- Collection of pertinent literature
- Conceptual collaboration from experts in social and cognitive sciences
- Acquisition of needed materials to conduct experiment

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Methodology and deployment plan for experiments.
Q2	Experimentation.
Q3	Data analysis from experiments.
Q4	Robot design recommendations based on empirical research findings. Technical report on findings to date for experimental research. Submission of full report for peer review.

Subtask 3: Experimentation with Soldier-Robot and Civilian-Robot Interactions: Integrating Context and Cues to Determine Intention (UCF)

Objective and Benefits:
We hypothesize that social and safety distances will influence the outcome of semantic perception as described in the Capstone Assessment.

- Social distance metrics will support clear communication and shared mental models
- Safety distance metrics will support both the squad's safety as a team and anticipate the behaviors of civilians
    - This experiment seeks to contribute data that will help predict civilian behavior in a Cordon and Search scenario, as the threat of the robot is likely to elicit responses which human Soldiers would not
- These measurable distances result from the interaction and engagement of two distinct yet complimentary systems:
    - Direct perception – from sensors or sensory systems
    - Reflective perception – from socio-cultural norms, i.e., narrative, culture, and geo-politics

The successful assessment of another's intention will depend on input from both modes of perception.

Research Description:
We will use a factorial design experiment to control for contextual shifts which are likely to occur during a real-world example of the scenario.

- Compare Soldier response to information from robot and human team members
- Compare social signals for team action if the "building" is a residence versus a business
    - We hypothesize that behaviors will shift due to contextual changes
    - The data gathered from these shifts can support predictive models
    - The findings can also be used to support shared mental model construction and maintenance

State of the Art:

- Social distance metrics will support clear communication and shared mental models

- Safety distance metrics will support both the squad's safety as a team and anticipate the behaviors of civilians
    - This experiment seeks to contribute data that will help predict civilian behavior in a Cordon and Search scenario, as the threat of the robot is likely to elicit responses which human Soldiers would not
- These measurable distances result from the interaction and engagement of two distinct yet complimentary systems:
    - Direct perception – from sensors or sensory systems
    - Reflective perception – from socio-cultural norms, i.e., narrative, culture, and geo-politics
- Robot adaptive behavior (as required for Capstone Assessment item 7) will depend in part on the assessment of the human's actions.
    - These actions are influenced by the presence of the robot itself
    - This experiment provides data which may be used to interpret and predict human movement by both Soldiers and civilians in response to the robot
        - We will compile and interpret measurements for the robot's physical distance from humans which support Cordon and Search tactics
- The results from this effort will support Intelligence, Surveillance, and Reconnaissance (ISR) missions, routing to a back door and pursuing an enemy exiting a building.

Key Barriers and Approach:

Research Questions:

- How do direct perception and reflective perception support understanding of intention between humans and robots?
- How does the dependence on these modes of perception differ when the Soldier observes a robot team member, as opposed to another human?
- How does proximity alter the dependence on direct versus reflective perception?

Hypotheses:

- We hypothesize that social and safety distances will influence the outcome of semantic perception as described in the Capstone Assessment
- The successful assessment of another's intention will depend on input from both modes of perception

Independent Variables:

- Narrative provided/not provided for reflective perception
- Modality of social signal: audio, visual, gestural for direct perception
- Soldier/civilian/robot observed
- Proximity: safety/social

Dependent Variables:

- Accuracy in determination of intention (is the agent fleeing, milling, dangerous)
- Distance at which intention may be accurately predicted

Outcomes for RCTA Consortium use:

- Identification of data parameters for prediction of intent
- Information for integration of implicit social signals with verbal directives
- Contribution of data requested by engineers at RCTA workshop for measurements which reflect perception

This approach moves beyond the state-of-the-art by:

- This approach is unique in its capacity to simultaneously explore the multiple levels of social interactions (Soldier, robot, and civilian) in a Cordon and Search scenario
- This approach will provide data ranges for modeling aspects of HRI which support the advancement of autonomous agency

Indicators of likely success include:

- We are building the experiment on a generous corpus of literature in social behavior relevant to the Cordon and Search scenario
- The robotic and virtual reality/simulation resources are accessible at UCF, as is a large pool of potential participants
- We have assembled a collaborative team of human factors researchers, cognitive scientists, computer engineers, and robot technology specialists who have the experience to execute well-designed experiments with the best technology available

Ties to other research activities include:

- Ties to the accurate representation and interpretations of intention in humans and robots
- Ties to perception by providing specific foci for attention and saliency
- Ties to implicit communication in the support of shared mental models
- Ties to trust in SR teams development and maintenance

Metrics for Evaluation:

- Completion of experimental design
- Empirically-based requirements and recommendations for Perception and Intelligence tasks
- Publication of experimental results
- Collect and describe parametric data relevant to the design and execution of robots in Capstone Assessment

The following progress has been made toward achieving our goal:

- We have already established a collection of social cues and signals which will provide a starting point in the experiment design
- The trust metrics are clearly articulated and ready for testing
- We have acquired all materials and resources necessary to conduct the experiments

We plan to measure intermediate progress and determine whether the research is on track to enable the capstone capability by:

- Regular updates with RCTA colleagues to inform on progress of experiment
- Collaborative publication/presentation for experiment results
- Support documentation and communication of requirements and recommendations for Perception and Mobility
- Provide data for further examination in the cognitive models being used in the broader RCTA project

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Methodology and deployment plan for experiments.
Q2	Experimentation.
Q3	Data analysis from experiments.
Q4	Robot design recommendations based on empirical research findings. Technical report on findings to date for experimental research, including suggested modifications to the H9 social model. Submission of full report for peer review.

Subtask 4: Empirical Investigation of Social Dynamics and Trust in Soldier-Robot Teams (UCF)

Objective and Benefits:
Our research explores the impact of social dynamics on trust in Soldier-robot teams and in civilian-robot encounters.

The two main objectives for this subtask are:
1) to provide empirical evidence for the relationship between social dynamics and human trust in a robot; and
2) to develop a set of robot design recommendations to improve Soldier and bystander trust in the robot, based upon data collected through human-in-the-loop experimentation.

Our approach will allow us to determine appropriate robot system social behaviors and cues (i.e., deliverables in the form of design specifications for Perception and Intelligence) that can foster trust within teams and with bystanders, thereby contributing to the success of missions. In particular, this research will contribute to the success of the Capstone Vision by ensuring that interactions with teammates and civilians will be successful.

State of the Art:
Trust is a critical component of HRI because it directly impacts the outcomes, safety, and overall success of dynamic human-robot military operations (e.g., Cordon and Search, counter-IED, urban ISR patrol, and logistics mission contexts).

- If trust erodes or is absent in human-robot teams, the likelihood that the robot will be used (or used properly) decreases. Calibration of trust is essential to the successful functioning and outcomes of a human-robot team (Lee & See, 2004).
- Trust can be compromised by many factors, including inappropriate or misinterpreted social behaviors exhibited by the robot.
- Our unique contribution to the goal of the RCTA is to examine the social dynamics of HRI and how these impact trust (and consequently the success and safety of the mission).
- Collaborating with human teammates and interacting with civilians (directly and indirectly) are important components of the Capstone Vision; in order to meet these requirements successfully (e.g., navigating among civilians), trust must be present in the interaction.

Robot Design:

- The social behaviors of the robot and how the robot is perceived by others (e.g., perceived form, appearance, attributes, and perceived function and affordances) can impact the civilian's subsequent decisions, actions, and behavior.
- Research has shown that the size and speed of a robot can impact how a bystander interacts with and trusts the robot (Tsui, Desai, & Yanco, 2010).

HRI Training:

- Civilians do not get formalized training for how to interact with the robot, making obvious social cues a key element in interacting with civilians and fostering their trust.
- A civilian's mental models of the robot are, therefore, likely to be inaccurate. Leveraging social interaction dynamics (identified in our proposed research) will allow civilians to have more accurate perceptions of a robot and to know how to behave safely in those situations.

Trust Measurement:

- Research indicates that the social behaviors of humans can, in fact, reflect the amount of trust a human has in another entity.
- Social distance between humans and robots: Bickmore and Cassell (2001) suggest the possibility that as trust builds between partners, the acceptable social distance will change.
- Attentional cues: A behavioral marker of trust may include attentional cues, or the extent to which a human attends to the robot (e.g., eye gaze).

Key Barriers and Approach:

In order to examine social dynamics and trust from the varying human roles encountered in human-robot interactions (teammate versus bystander), two independent experiments will be conducted. Results from these studies can be used to inform robot design, such that the appropriate social behaviors can be incorporated into a robotic system. Thus, the robot will be able to perform its given task and also behave appropriately (i.e., encourage trust and confidence in the system) in the presence of civilians and teammates as it navigates through its operational environment, contributing to the success of the Capstone Integration Vision.

Experiment 1: The first experiment, "Social Dynamics, Trust, and Teammates," will assess the impact of various social cues on a trust in a robotic teammate.

Research Question:

- When interacting with a team consisting of a robotic team member, how do social dynamics (e.g., robot behavior and form) of the robot impact human social behavior and trust in an operational context?

Hypotheses:

- We hypothesize that the type/size/speed of the robot will influence the degree of trust a human teammate has in the robot.
- We predict that behavioral markers of trust (e.g., the perceived acceptable social distance between human and robot, and other social behaviors) will be associated with subjective measures of trust.

Independent Variables:

- Speed of the robot
- Size/type (i.e., form) of the robot

Dependent Variables:

- Behavioral markers of trust (e.g., the distance that the human stays away from the robot and how often the individual attends to the robot)
- Subjective measures of trust
- Human performance/success of collaboration

Outcomes for RCTA Consortium use:

- Results from this study can be used to inform robot design (algorithms for Perception and Intelligence), such that the appropriate social behaviors can be incorporated into a robotic system to foster trust within a human-robot team.

Experiment 2: The second experiment, "Social Dynamics, Trust, and Civilians/Bystanders," focuses primarily on civilian encounters with a robot.

Research Question:

- How should the robot interact with human civilians in order to elicit appropriate trust responses from the civilians/bystanders?

Hypotheses:

- We hypothesize that the type/size/speed of the robot, and social behaviors exhibited by the robot (e.g., social distance and how close it comes to an individual), will influence the degree of trust a human bystander has in the robot.
- We hypothesize that there is an association between gender of the human bystander and the degree of trust in the robot.

Independent Variables:

- Robot's gaze (i.e., making eye contact versus looking away from human; pending availability of a platform capable of such a manipulation)
- The social distance that the robot "allows" between itself and the human
- The type/form of the robot
- The speed of the robot
- Gender of the human bystander may also impact trust

Dependent Variables:

- Identification of the different behaviors exhibited by participants
- Human performance
- Subjective measures about trust and perceived intention of the robot

Outcomes for RCTA Consortium use:

- This research effort will identify a set of robot design heuristics (algorithms for Perception and Intelligence) for appropriate social behaviors that can facilitate a bystander's trust in a robot, which will contribute to mission success and safety when the civilian populace is involved.

Metrics for Evaluation:

- Improved understanding of how a robot's social behavior influences trust development in human-robot teams
- Empirical support for how aspects of social dynamics influence trust in robots for bystanders
- Identification of the relationship between changes in social dynamics and the development and maintenance of trust
- Set of robot design guidelines (e.g., algorithms) for the social cues and behavior to build into a robot so that it will foster an appropriate degree of trust and/or prevent trust breakdowns in human-robot interactions

We have made the following progress toward our goal:

- We have conducted extensive prior research on trust in human-robot collaborations (for example, see Hancock et al., 2011), which has provided the foundation for this experimentation.
- We are in the process of completing a trust scale to measure trust in HRI, as part of the 2011 work under Task H3.
- We have the expertise and knowledge necessary to conduct the experiments.

We plan to measure intermediate progress and determine whether the research is on track to enable the capstone capability by:

- Regular updates with RCTA colleagues to inform on progress of experiment
- Collaborative publication/presentation for experiment results

- Provide documentation of findings and communicate requirements and design recommendations for Perception, Intelligence, and Mobility

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Methodology and deployment plan for experiment 1.
Q2	Experimentation 1. Methodology and deployment plan for experiment 2.
Q3	Data analysis from experiment 1. Experimentation 2.
Q4	Data analysis from experiment 2. Robot design recommendations based on empirical research findings. Technical report on findings to date for experimental research.

Subtask 5: Simulating Dynamic Environments with Social, Cultural Parameters (UPenn)

Objective and Benefits:
The objective of this task is two-fold:

(1) To create a real-time, interactive simulation of a socially and culturally accurate middle-eastern marketplace scenario that is populated with event-driven personnel who exhibit rich interactions with the environment, other members of the virtual populace, and human-robot solider teams in a manner that conforms to social and cultural norms of the region.

(2) Use the simulation environment as a testbed to perform human factor experiments.

State of the Art:

- Crowd research [Pelechano, Allbeck, & Badler, 2008] models agents as generic pedestrians with spatial goals – focus is on macroscopic phenomena of crowd flow.
- Behavior Authoring Systems (e.g., Massive Software) require programmers and generate non-interactive simulations.
- No framework exists to model functional, purposeful agents that are affected and react to human players or robotic agents in real time.

Key Barriers and Approach:
As part of this research, we will leverage prior art (RVO, HiDAC) for agent navigation and collision avoidance as well as the research that was undertaken in H10 to create a real-time, interactive framework for simulating functional, purposeful autonomous agents whose behavior conforms to social and culturally defined parameters. These autonomous agents must exhibit complex interactions between each other as well as human agents (Soldiers or robots).

Agents in the simulation will be defined using *traits* and *capabilities*. Traits define the individual mental model of the agent, while capabilities define the affordances of the agent. Agent-specific

behaviors and scenario narratives (complex events) are authored using Parameterized Behavior Trees (PBTs) at two levels:

- Individual actions and reactions
- Events coordinating multiple agents

This hybrid approach ensures that agents are completely autonomous and can interact with human agents while simultaneously exhibiting complex multi-agent interactions.

We have a prototype framework to simulate functional, purposeful agents and have demonstrated a plausible depiction of a lively marketplace. The research undergone for this subtask has resulted in four publications.

The research for this subtask is closely related to the applied research in the communications thrust (H3: HRI Tactical Communication Protocol). The authoring tools for defining behaviors will provide a basis for defining the constructs of the unification language for communication between different thrusts. Also, the results of human factors experiments will help define shared mental models of human-robot members that can be leveraged by the Intelligence thrust.

Metrics for Evaluation:
The simulation framework will be evaluated using a variety of metrics that quantify the performance, quality, and correctness of the simulation. These include the number of agents that can be simulated at interactive frame-rates, number of simulation artifacts (e.g., collisions, incorrect animations etc), number of agents that can participate in events, and complexity of interactions between agents. Most of these metrics can be automatically computed, while some may require user input.

We will provide a graphical user interface for authoring complex multi-agent simulations. Early versions of the simulation and authoring framework will be provided to UCF for feedback which will be taken into account for subsequent releases. Next, in collaboration with UCF, we will design and author a set of complex scenarios in a middle-eastern marketplace scenario (described in H8-5) that can be used to perform human factor experiments with human-robot Soldier teams.

	Subtask 5 Quarterly Milestones and Deliverables
Q1	Develop real-time agent simulation framework (v0.5) with programmatic authoring of scenarios and events.
Q2	Extend and refine simulation framework based on feedback from initial experiments (v1.0). Design and author a simple set of scenarios that can be used to perform initial human factor experiments with human-robot Solider teams in collaboration with UCF. Develop a graphical user interface for agent and event authoring.
Q3	Refine authoring interface based on feedback from initial experiments. Extend agent simulator to include multi-modal communication capabilities.
Q4	Design specific scenarios for exercising communication and coordination between human-robot Soldier teams. Perform human factor experiments in collaboration with UCF.

Subtask 6: Social and Cultural Impact on HRI (ARL)

Objective and Benefits:
This subtask objective is to better understand the impact of social and culture context on human-robot interaction.

State of the Art:
A literature review will be conducted to assess the current state of HRI research examining the impacts of culture on human performance related to interaction with autonomous systems. There will be planning and execution of experimental investigation of social interaction and cues between humans and robots.

This work is related to other tasks within H9, particularly Subtask H9-3: Experimentation with Soldier-Robot and Civilian-Robot Interactions and Subtask 9-4: Empirical Investigation of Social Dynamics and Trust in Soldier-Robot Teams.

This work will support the ongoing Social Dynamics theory and experimentation work of Consortium members. Socially aware navigation work may be integrated into an indoor-based integration and assessment activity in the mid-term.

Key Barriers and Approach:
The use of robots is moving from single operator interaction within a specific application to human-robot teams within broader social and cultural milieus. More and more research is being conducted to address human-robot interaction (HRI) within social environments; research on cross-cultural impacts on HRI is relatively new. The subtask will examine some research involving both social and cultural aspects of HRI. As part of the cultural and social impacts, the area of *proxemics*, defined as the study of measurable distances between people as they interact, will be examined with respect to interactions between humans and robots. Just as humans give cues and have implicit understandings about distance interactions, we expect the need for robots to have similar understandings of human-robot interactions.

Metrics for Evaluation:
The research review of the social and cultural impact on HRI will be measured by the completion of a final draft technical report. The collaboration for the socially aware navigation experiment will have milestones of planning document (i.e., experimental protocol), experiment execution, and then post-data collection analysis and documentation.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Substantial interaction between the Consortium and ARL is anticipated. ARL collaboration with UCF and iRobot within the HRI area in the planning, execution, and analysis of the socially aware navigation experiment is expected.

	Subtask 6 Quarterly Milestones and Deliverables
Q1	None – task not started.
Q2	Plan socially aware navigation experiment.
Q3	Execute socially aware navigation experiment.
Q4	Complete ARL technical report on HRI and culture (final draft).

5. DEXTEROUS MANIPULATION AND UNIQUE MOBILITY OVERVIEW

Objective

Our goal is to change the course of robotics through fundamental research and advanced technology development leading to a new generation of highly dexterous, ultra mobile, autonomous robots. They will succeed by exploiting coordinated, dynamic, and multi-modal mechanisms coupled with a fundamental understanding of how robots interact with objects and terrain.

We seek to integrate this research with outcomes from Perception, Intelligence, and HRI to demonstrate quantum leaps in platform capabilities, enabling man-portable platforms with advanced manipulators, capable of performing dexterous mobile manipulation and autonomous operations in urban and mountainous environments; legged, mobile manipulators that utilize hands, feet, and whole-body mechanics to grasp, push, pull, and drag while operating in the roughest terrain; and unique mobility platforms that can reconfigure to efficiently and covertly run, climb, and crawl up the sheer faces of buildings and through culverts and tunnels where other vehicles and Soldiers cannot currently go.

In practical terms, we envision enabling robotic systems that are able to serve as capable members of a team by gaining access to complex terrain (traverse cluttered allies, climb walls and poles, wiggle into otherwise inaccessible tight spaces, etc.) where they can manipulate the environment in order to perform tasks for the team they are part of (open or close doors, clear rubble, search cluttered spaces, excavate suspicious terrain, etc.).

State of the Art and Technical Barriers

The reality is that today's fielded robots have mechanically crude manipulators with almost no dexterity, and their mobility is remarkably limited in all but the smoothest terrains. These robots are equipped with manipulators and end effectors with such limited geometry and actuation as well as rudimentary control and perception algorithms that they can only support a very limited subset of missions. What dexterity they do have comes at the expense of requiring the full attention of human operators who do all the perception, planning, and direct control of each motion and task, down to the smallest detail. The mobility of existing ground robots in rough terrain is no better; these systems are limited by poorly controlled wheel or track mechanisms that do not even have the level of automation for stability and traction found in common automobiles. Their travel is nearly paralyzed in sand, rocks, rubble, snow, and mud, all places where Soldiers routinely operate and fight. Although some robot vehicles are learning to navigate through open country and on trails, they cannot perceive rugged terrain and make their way through it as Soldiers can on foot. Lastly, no fielded robots can climb buildings or trees, leap over obstacles or from great heights, or travel in tight places such as within cluttered buildings.

Moving Beyond the State of the Art

Our efforts seek to move far beyond this current state of the art. Our fundamental research and technology development efforts will enable a revolution in robot manipulation and mobility, providing an order of magnitude more dexterity, strength, speed, agility, and autonomy. The work we propose provides robots the ability to perform complex tasks with minimal direct human intervention. Such robots sense and perceive the environments in which they operate and the objects that they manipulate. They have control algorithms based on a fundamental understanding of the mechanics and dynamics of how hands, fingers, and objects interact during grasping and manipulation. They reason and plan about the task elements that achieve desired results as well as execute fine-grained sensor-based control of their mechanisms and of objects, both autonomously and in conjunction with guidance provided by human and robot partners. Such robots have new kinds of sensors, actuators, and materials with significantly more strength, lightness, sensitivity, efficiency, and bandwidth to achieve these goals. The algorithms, planning, sensors, actuators, and controls are woven into both traditional robot systems as well as new integrated robot designs. This research enables new classes of robot platforms. One class of robot is in the 25- to 30-pound range, fits in a backpack, and has extraordinary manual dexterity, autonomy, and enhanced rough-terrain mobility. A second class of robot is somewhat larger, 175 to 250 pounds, with legs for mobility, advanced terrain sensing and reasoning for extraordinary agility in diverse terrains (rough, rugged, mountainous, urban), and with extraordinary dexterity, manipulating with one or more arms, legipulating with its legs, and coordinating both to rearrange the world around it. A third class of robot combines unusual climbing designs with snake-like articulation. It uses its many degrees-of-freedom, its fundamental understanding of mobility and grasp mechanics, its perception, and its reasoning to get places no conventional vehicle can go.

Relationship to Other Technical Areas

Successful development of robots capable of performing the dexterous manipulation tasks described here and imbued with unique mobility capabilities suitable for gaining access to these types of terrain requires careful integration of research results from across the entire Robotics CTA effort. The specific efforts described within each of the DMUM tasks depend on the Perception tasks to provide the information necessary to understand the environment in which the robot is operating. Similarly, collaboration with the Intelligence tasks will provide the robotic system with the tools necessary to reason about and learn how to robustly interact with and manipulate the environment in which the robot is operating. Finally, in order to enable the robots to function as effective team members and communicate efficiently with their human teammates, it is critical that results of HRI research are used to develop effective interfaces for the resulting systems. Practically speaking, much of the DMUM research provides the physical context through which the “rubber meets the road” for a significant portion of the RCTA research. As such, it provides a natural vehicle for integration, testing, and refinement of many technologies from the Intelligence, Perception and HRI thrust areas.

Summary of Proposed Research

Our basic research, based on our team members’ seminal work in all areas, focuses on the fundamental theory and science of articulated systems, unifying adaptive control and sensing; exploiting mechanism dynamics, grasp theory, and passive compliance; and providing high-

performance coordination of locomotion, manipulation, and learning; all in the context of operations in complex terrain. We emphasize developing an understanding of the underlying science and application of grasp mechanics, dexterous and high degree-of-freedom manipulation, multi-modal locomotion, terrain interaction, and adaptive mechanisms, and we apply that understanding to advance the practice of manipulation, mobility, vehicle-terrain-object interactions, advanced actuation, and the design and control of high-performance mechanisms, resulting in a unified effort that will both advance these key technologies as well as improve our ability to integrate and leverage each to achieve higher levels of advanced functionality and behavior. We will discover new ways to design with and exploit new algorithms, materials, actuators, and mechanisms, furthering the state of the art to enable new capabilities for all scales of robots.

Our efforts are segmented into three thrusts (Dexterous Manipulation, Unique Mobility, and Next Generation Actuation), encompassing basic research in each thrust coupled with applied research. During the upcoming project year, we will focus on transitions of our basic research into demonstrable ideas and technologies that can be realized using several scenarios that require robots to grasp, position, assemble, and disassemble natural and man-made objects; manipulate objects using whole-body mechanics; smoothly negotiate complex, dynamic, and unconventional (trees, walls, wires, sewers) terrain; learn terrain interactions; and coordinate mobility and manipulation to achieve maximum efficacy.

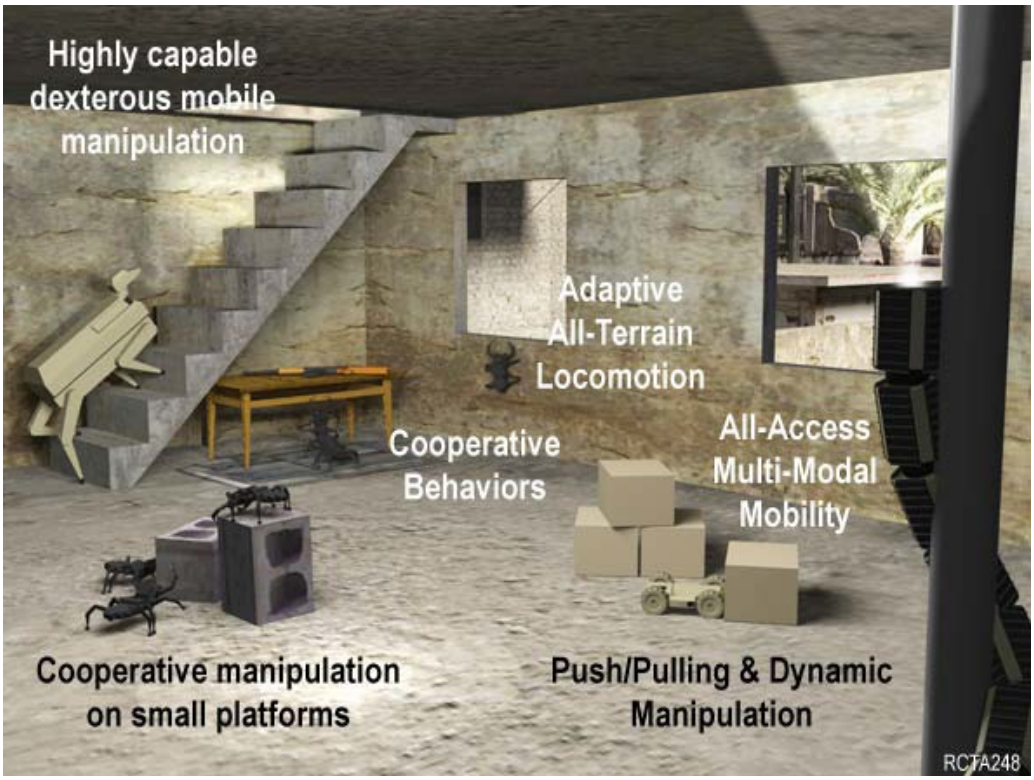

Dexterous Manipulation. The tasks in this area address the key challenges of creating robots with dexterous manipulation, dynamic mobility, sensing, and effective, easy-to-use operator interfaces. These tasks seek a high-level of dynamic performance suitable for effective operation in the field environments where Soldiers operate.

Unique Mobility. The tasks in this area focus on design and control of unique, multi-modal mechanisms that can adapt to uncertain changing environments and morph to traverse diverse 3D terrain including walls, trees, poles, and other structures.

Next Generation Actuation. In this topic area, we investigate novel actuator technology, multi-functional materials, and morphing structures to enable power-limited vehicles to manipulate their environments and achieve robust high-performance mobility.

Within the DMUM area, Tasks M1 and M2 are performing basic research on advanced manipulation algorithms and using tools/system available within their laboratories to perform exploratory experiments and develop proof-of-concept demonstration. Tasks M3 and M4 are applying the results from those basic research tasks to current and future robotic platforms, respectively. They plan to continue to use and develop appropriate hardware within each of those tasks and in collaboration with efforts being undertaken as part of the overall integration effort. Specifically, both tasks are involved in the development of manipulator systems that are more capable than the current state of the art but are specifically targeted toward integration onto traditional mobile robots (TALON and Dragon Runner platforms, M3) and more advanced robots (the BigDog platform, M4). Both M3 and M4 are highly integrated with ongoing tasks in Integration, Perception, Intelligence, and HRI. Tasks M5 and M6 are performing basic research on mobility for articulated and limbed systems, respectively, and both are using existing and new prototype hardware developed within the laboratories performing those tasks for experimentation and proof-of-concept demonstrations. Task M9 is exploring fundamental questions related to materials and actuation technologies that can apply across the entire research area but are specifically focused on robust compact multi-modal actuation and sensing for locomotion in the short term. Tasks M7 and M8 seek to apply results from M5, M6, and M9 along with results from associated Perception and Intelligence tasks on both existing and novel mobility platforms. The specifics of each task are summarized below in Table 5-1.

Task	Primary Research Outcomes	Operational Implications
M1: Theory and Principles of Mobile Manipulation	Investigate fundamental methods for robust and general automation of mobile manipulation in complex environments. Develop a formal understanding and framework to enable mobile systems to perform highly dexterous manipulation operations.	More robust tactical behaviors due to behavioral-level ability to reason about and interoperate the results of physical interaction. Improved ability to perform whole body manipulation including door opening/closing, complex environment search and inspection.

Task	Primary Research Outcomes	Operational Implications
M2: Principles of Generalized Grasp Mechanics	Formalize principles for grasping a wide variety of objects, including the special problem of grasping with all surfaces of a robot. Develop representations suitable for reasoning about for grasping a wide variety of complex objects.	Enable automated and robust manipulation of the environment and objects in the environment. Reduced computation and planning burden to plan and control physical interaction with a variety of complex objects.
M3: Sensor-Based Dexterous Manipulation	Reactive algorithms to integrate sensing, perception, and understanding of the environment with manipulation of unknown objects and terrain.	Early deployment of the capability to inspect and manipulate cluttered, unknown spaces and use specialized tools for digging, drilling, and cutting on existing robotic platforms.
M4: High Degree-of-Freedom Dynamic Manipulation	Integrated reactive control and planning algorithms capable of utilizing high-degree-of-freedom whole-body motion to effect dynamic action in cluttered and complex environments. Development of compact high-performance robust hydraulically powered manipulation systems. Integration of perception and control to develop reactive algorithms for integrated dynamic manipulation and locomotion.	Delivery of high-performance integrated mobility and manipulation capabilities on an advanced quadruped robot platform (BigDog). Demonstration of robust proof-of-concept field plausible advanced mobility and manipulation capability.
M5: Theory and Principles of Multi-Modal Locomotion Planning and Control	Integrated control and planning algorithms for robust, general locomotion of reconfigurable or highly articulated multi-modal platforms. Algorithms to enable learning of novel locomotion strategies and optimization of existing strategies.	Enable robot systems to robustly gain access to confined, cluttered and constrained spaces (such as rubble, pipes, chimneys, etc.).
M6: Principles of Locomotion Mechanics	Develop models of animal behavior to inform design methods and principles for adaptive gaits. Apply biologically inspired control strategies to and dynamic control of running, climbing, hopping and leaping in complex terrains.	Provide theoretical framework to enable robots to achieve performance comparable to biological systems.

Task	Primary Research Outcomes	Operational Implications
M7: Learning Terrain Interactions	Integrated algorithms to couple proprioceptive sensing and visual recognition with physics-based dynamic vehicle-terrain interaction models and learning to improve locomotion.	Improved semi- or fully-automated "traction control" for both traditional wheeled and tracked vehicles as well as more advanced legged platforms.
M8: Dynamic Multi-Modal and Reconfigurable Mechanisms	Develop design principles of locomotion and control that leverage novel actuators, materials, and modular designs to produce efficient, robust, and simple robots that can traverse complex environments.	Enable small and simple robots to exploit dynamic capability and novel design principles to robustly traverse rubble or debris; climb complex such as walls, poles and other structures; and reconfigure themselves to tailor capabilities to the environment in which they operate.
M9: Next Generation Actuators and Materials	Combine theoretical and computational mechanics to guide the development of actuators and mechanisms. Develop novel hybrid active materials, sensors, and actuators for powerful and efficient manipulators, both large and small, new gripper designs, and variable stiffness legs.	Small, robust, multi-functional, efficient integrated actuation and sensing capabilities. Improved robustness and adaptability of robotic systems.

Table 5-1: DMUM Task Outcome and Operational Implications table.

Dexterous Manipulation and Unique Mobility Tasks

M1-2012 – Theory and Principles of Mobile Manipulation

Level: 6.1 Basic Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Vijay Kumar	UPenn	M1-1, M1-2
Additional Investigators	Katherine Kuchenbecker	UPenn	M1-3
	Mark Yim	UPenn	M1-2
	Maxim Likhachev	CMU	M1-1

Objective:
The research objective of this task is to advance the theory and principles of mobile manipulation to enable robots to enter buildings by opening doors and avoiding or moving obstacles inside. We aim to accomplish this complex objective by pursuing coordinated research in three main areas: efficient high-dimensional motion planning, reconfigurable manipulator design, and accelerometer-based contact sensing. The three thrusts will each advance robotic interaction with doors and piles of clutter in a different yet synergistic way. After an initial probing phase (estimating required forces), the planner will generate a feasible trajectory that respects the constraints of the environment (perceived obstacles) and the robot itself (joint limits, torque limits, etc.). Rather than being constrained to a static arm design, the robot will also be able to make use of modular reconfigurable components to form shapes suited to interacting with a large range of objects. As the robot executes the chosen plan (reconfiguration and task motions), it will use high-bandwidth accelerometers distributed across its arm to monitor the progress of its actions and detect both intentional and unintentional object contacts. Capable robotic mobile manipulation requires fundamental research in all three of these areas.

Our work in motion planning aims to answer the question, “How can a mobile manipulator efficiently choose motions to accomplish a task in a real-world environment?” We will investigate dynamic constraints by extending search-based planners to incorporate dynamics on the large state spaces required for high-DOF manipulators. In collaboration with Boston Dynamics (M4, Rizzi, Murphy, and Stephens), we will explore the use of search-based planning algorithms to bootstrap optimization algorithms. Finally, we aim to develop graph searches that take advantage of the specifics of coordinated manipulation problems to achieve efficiency. This subtask will yield abstractions and algorithms for reduced dimensional planning for coordinated mobile manipulation with the secondary goal of exploring complex 3D spaces using a robotic arm equipped with 3D sensors.

Figure 5-1: PR2 opening a door [Chitta et al. 2010].

Our work in modular robotics aims to answer the question, "How can one design a mobile manipulator to enable physical interaction with a wide variety of real-world objects?" Conventional robots maintain a fixed kinematic arrangement, link morphology, and end effector, regardless of the current task and environment. The robot designer can try to meet the needs of the majority of the expected scenarios, but this process inherently necessitates compromise. Instead, a modular reconfigurable manipulator can be adaptively reshaped to create an arm, hand, and/or fingers to suit a particular obstacle arrangement and the encountered object's size and shape. This subtask will yield approaches to synthesize arms and effectors for grasping and methods for using these effectors for whole-arm grasps.

Our work in contact perception aims to answer the question, "How can a robot reliably execute its movement plans, detecting and correcting mistakes as they occur?" The traditional robotic manipulation approach of sense-plan-act breaks down in unstructured environments because neither the environment nor the robot pose can be perfectly perceived. Good mobile manipulation requires a robot to carefully monitor intended contact between its end effector and objects or tools, along with unintentional collisions between its arm and the cluttered environment; interestingly, almost all contacts between the robot and the world cause distinct measurable vibrations, but astonishingly few robots pay attention to this rich source of information. We aim to endow autonomous robots with heightened awareness of all of these types of physical contact events so they can quickly and accurately execute challenging movement plans; this sensitivity will also be useful for human operators controlling robots when fully automated execution is not possible, as we are exploring in our collaboration with QinetiQ (IR2, Cascone). This subtask will yield robust approaches for perceiving and categorizing both intentional and unintentional physical contact events using high-bandwidth accelerometers.

Background:

Current robotic systems struggle to perform mobile manipulation tasks in real-world environments because they are ill-prepared for the complexity and unpredictability that characterize such scenarios. Hence, current field-deployed robots are all controlled by human operators at a distance. It is the human who decides which joint to move at what time, always constrained to the original robot arm and gripper design, trying to judge whether and where the robot has made contact based on what can be seen from the on-board cameras. Fundamental advances are needed to offload this burden so the robot can become a capable and independent teammate rather than a cumbersome puppet.

Automating mobile manipulation requires robust execution of tasks using many coordinated degrees-of-freedom with complex constraints in uncertain and unknown environments. This

problem is high-dimensional and computationally difficult. Solving it requires principles for understanding the problem structure and intelligently exploiting and decomposing the structure without sacrificing robustness or generality.

Grasping objects requires the use of effectors beyond the simple parallel-jaw grippers that are typical in current mobile manipulation robots like the QinetiQ Dragon Runner and the Willow Garage PR2. Current grippers and arms are unable to allow robots to grasp arbitrarily shaped objects or work in arbitrary environments. We have previously used modular robots to create highly successful reconfigurable locomotion systems, and we believe they hold great promise for manipulation as well.

Finally, we want teleported and autonomous robots to be able to quickly approach, grasp, and manipulate objects in cluttered environments. Visually detecting physical contact is usually almost impossible due to occlusion. Monolithic force sensors are expensive, delicate, and difficult to integrate with many end effectors, while pressure-based broad-coverage robot skin is not yet technically feasible. Nonetheless, robots need robust ways of judging whether their actions are transpiring as planned. Our prior work has shown that high-bandwidth accelerometers can detect both intentional and unintentional contacts in robotic surgery, and we believe they will provide a similarly rich source of information for mobile manipulators in the field.

Research Description:

Progressing from intensive and challenging tele-operated control of mobile manipulation to efficient and capable autonomous task execution requires the ability to automatically plan and robustly execute tasks in high-dimensional environments that are only partially known. We are working toward this important long-term objective by developing feedback policies in high-dimensional spaces, creating novel reconfigurable manipulator designs for whole-arm grasping, and investigating an accelerometer-based approach to detecting robot contacts with the world.

Motion planning and control for mobile manipulation

Path planning in a mobile manipulator's high-dimensional state space is computationally expensive and results in individual, open-loop paths. On the other hand, developing feedback controllers that ensure global convergence in complex spaces is also difficult. Instead, we combine the two concepts: we develop feedback controllers that operate in local regions and drive the system to a local goal state, and we compose these controllers using deliberative planning to ensure convergence to the desired global state. This approach reduces the dimensionality of the planning problem while maintaining the robustness of closed-loop feedback controllers.

We propose to develop control policies that drive the system to sub-manifolds of its full state space while maintaining certain boundary properties that allow them to be composed in sequence or parallel. In our preliminary work in planning feedback control policies for mobile manipulators for simple environments, we investigate the synthesis of piecewise smooth controllers and search-based planning in the space of feedback control policies. We developed such local cell-based controllers for an obstacle-filled environment for the navigation problem

for a nonholonomic planar cart [Kallem et al. 2010]. Because we can compose controllers in parallel, the approach extends naturally to multiple manipulators or manipulation via multiple robots for simple cases (where the tasks are not spatially or dynamically coupled). This proposed work is naturally connected to the development of software architectures (Tasks I1 and I2) and the development of algorithms for planning and learning (Tasks I3 through I6). Further, we developed a receding horizon control algorithm to optimally control nonlinear systems such as mobile manipulators while satisfying the input and state constraints [Kallem and Kumar, 2010] if a nominal control law is known. The local cell-based controllers provide these nominal controllers in complex environments, thus leading to optimality in finite horizon.

While the composition of local controllers makes a system robust to local perturbations, we must still address the problem of planning when the world is not fully known. Autonomous systems must be able to discover new information and reason how to optimally explore an unknown space with the expectation of discovering new information. To address operating with partially known world states, we build on our previously developed search-based planning approaches for manipulation. These have been shown to be successful in complicated but known configuration spaces. We proposed to adapt them for planning in partially known environments with generalizations that allow them to be updated incrementally instead of replanning from scratch when new information is discovered. To enable exploration in unknown, constrained environments, we focus on combining these techniques with planning under uncertainty in low-dimensional spaces by exploiting the structure or sparseness of problems. This approach allows the problems to be solved efficiently using a series of simpler-to-solve searches, and it also provides the ability to prove and control the bounds on the sub-optimality of the solution.

Finally, we will consider trajectory generation of a mobile manipulator while considering the nonlinear dynamics of the robot in addition to its geometric and kinematic constraints. This task is linked to the work in Task M4 and will be in collaboration with Boston Dynamics.

Synthesis of arms and effectors and whole-arm grasping
We will address the problem of synthesizing bimanual, whole-arm grasps by developing the abstraction of an open chain gripper, a serial chain of rigid links contacting an object. Previous work on caging and immobilization (including work by collaborators Mason and Burdick) will be leveraged to construct an algorithm that synthesizes contact configurations for stable grasping. We will use simulations to show that our methodology can be applied to grasp a wide range of objects without relying on special-purpose end effectors. We will also conduct representative experiments with the PR2 humanoid robot to illustrate that this approach is practical.

We will also synthesize reconfigurable robot effectors using chains of hinged space-filling polyhedrons (cubes or right angle tetrahedrons). Griffith and Demaine have proven that chains of these polyhedrons are able to approximate any 3D shape. We have prototyped shape reconfiguration with such chains [White, Thorne, and Yim, 2009] and generated control algorithms to form different shapes. We believe such chains will provide the capability for gross conformational changes in size, shape, and topology that are much larger than those attainable with soft grippers, such as the Festo gripper or humanoid style high-DOF hands. These

reconfigurable robot arm and hand designs should enable improved performance at mobile manipulation tasks.

Accelerometer-based contact sensing for autonomous manipulation
To succeed at tasks like covering the back door of a building, a robot needs to be highly aware of its physical interactions with the world. Over what terrain is it driving? Are the wheels/treads slipping or sliding? Has the gripper made contact with the intended object? Is the grasp steady or slipping? Has it hit any obstacles? If so, where? Both intentional and unintentional contacts cause vibrations that easily travel through the solid structure of a robot. Humans strongly rely on this very same type of cue when controlling machines and manipulating objects, but contact vibrations have seldom been studied by roboticists. Our approach centers on the use of high-bandwidth three-axis MEMS-based accelerometers; these commercially available vibration sensors are very sensitive, tiny (a few mm per side), low cost (less than $20), and very robust (perfect for field applications). We rigidly mount one or more of these sensors to important parts of the robot like the chassis and inside the gripper. We plan to develop signal processing techniques that improve contact event detection, enable contact localization, and perhaps yield additional information about the contact, such as the stiffness of the object or the relative speed of the impact.

Our prior research in robotic surgery [McMahan et al. 2011b] has shown that humans significantly prefer robotic tele-operation interfaces that relay these vibrotactile sensations in real-time (typically played through a voice coil actuator attached to the controller). These cues strengthen the connection between the robotic tools and the operator's hands, improving his or her awareness of the actions taking place. We also have preliminary evidence that providing these cues to an expert operator makes him or her more delicate in contacts with objects [McMahan et al. 2011a]. Furthermore, our recent work in autonomous robotic manipulation [Romano et al. 2011] has shown that attending to contact vibrations makes a robot (the PR2) much more dexterous and capable with its hands. For example, a robot attempting to set down a delicate object open-loop will drop it a short distance or crash it into the table if the eye-hand calibration is not perfect. We use dynamic contact events detected through a high-bandwidth accelerometer and fingertip pressure sensors to enable the robot to identify contacts and react appropriately, in this case by releasing the object immediately following table contact. We have also shown that listening to the sounds that occur during manipulation (which are closely related to the felt contact vibrations) dramatically improves the success of autonomous manipulation [Romano et al. 2012]. This past experience gives us confidence that a similar approach will prove beneficial and enable novel extensions in fielded mobile manipulation systems.

As we move toward full autonomy for mobile manipulators, we also recognize that unknown environments will occasionally present challenges that merit human intervention. For these critical actions, we must dramatically improve the control and feedback methods available to the remote operator so that they may complete the task quickly and accurately. We are grounding these investigations in the

military context through collaboration with researchers at QinetiQ North America in Pittsburgh. We have already been working with Matthew Csencsits and his coworkers, and they are eager to continue collaborating with us on this project, supported through their own RCTA funding on the integration thrust. Specifically, we are working to adapt our approaches to mobile manipulator haptic feedback into their DR20 research platform; we will test this system's performance as part of IRA2. We are also planning a more extensive experiment on multi-modal human-robot interaction with Stephanie Lackey and colleagues at UCF as part of Task H7. We will use the results from these joint experiments in haptic tele-operation to verify that useful vibration contact signals occur in fielded mobile manipulation systems. The insights gained from these collaborations will continue informing our more basic work on physical contact perception for autonomous mobile manipulation.

Metrics:
Motion planning and control progress will be measured by comparing results to those of other kinodynamic planners (e.g., the RRT and KPIECE). Specific metrics include planning time, replanning time, and execution time for the planner.

Our reconfigurable gripper designs will be evaluated by comparing the range of objects the system can grasp versus traditional two-fingered grippers (e.g., the PR2) and other non-traditional grippers (e.g., vacuum-based granular media gripper). Specific metrics include the size, weight, and canonical shapes of objects that can be grasped.

Our autonomous contact sensing work will be evaluated by the proportion of object contacts the system can detect and localize. In collaboration with QinetiQ, our vibrotactile haptic feedback system for tele-operation will be evaluated in a formal human subject experiment as part of IRA2; specific metrics will include task completion time, error rate, and subjective assessments of mental workload, stress, and preference with and without the feedback.

Ultimately our metrics will be derived from the application of:

- New algorithms that will allow arms to open doors to offices or cars and explore in the interior of confined spaces like cars;
- Novel effectors that will allow grasping and retrieval of awkwardly shaped objects; and
- New contact sensing methods to enable autonomous and human control of mobile manipulation and grasping.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
M1-1 includes collaborations with Rizzi and Boston Dynamics (Task M4) and with Likhachev and CMU on algorithms for whole-body and mobile manipulation. M1-2 will establish collaborations with Mason and Burdick on Task M2. M1-3 represents collaboration with Csencsits and QinetiQ (Task IR2) and with Lackey and colleagues at UCF (Task H7).

Kuchenbecker, Kumar, and Yim expect to be actively involved in seminars and workshops in the context of the RCTA and in conferences such as Robotics: Science and Systems and the IEEE International Conference and Robotics Automation.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Design of algorithms for mobile manipulation, reconfigurable grasping, and contact sensing.
Q2	Implementation on 3D simulation tools and proof-of-concept prototypes.
Q3	Benchmarking of performance and demonstrations.
Q4	Plans for transitioning to M3 and M4.

Related Research:

- Task I1: Framework for Intelligence
- Task I4: Generating Adaptive Tactical Behaviors
- Task I5: Learning through Experience
- Task H5: Gestures, Posture, and Haptics in HR Communication
- Task H7: Integrating Multi-modal HR Communications in Live and Virtual Environments
- Task P1: Exploiting Novel Sensor Phenomenology
- Task M2: Novel Grasping Mechanisms and State-Estimation
- Task M3: Sensor-driven Autonomous Behaviors (JPL) and Planning (GDRS)
- Task M4: High Degree-of-Freedom Dynamic Manipulation
- Task M9: Next Generation Actuators and Materials
- Subtask IR2-3: Haptic Control of a Small Unmanned Ground Vehicle

Integration and Assessment Activity:
All the models and algorithms described here will be implemented on either the Barrett arm or the PR2 robot, testbeds to which all researchers will have access. We will also continue working with QinetiQ to adapt our tele-operation contact sensing and feedback methods to their DR20 platform. These results will be shared within and outside the RCTA community. Our results will be directly or indirectly related to IR2-3: Haptic Control of a Small Unmanned Ground Vehicle and IR3-3: Exploring Mobile Manipulation.

Subtask 1: Motion Planning and Control for Mobile Manipulation (UPenn)

Objective and Benefits:
We focus our research on the dynamics, control, and planning of agile mobile manipulators where the dynamics of the arm and the base are coupled and interaction with the external

environment is needed to accomplish tasks. We will consider the right abstractions for such tasks. These abstractions will involve decomposition of the state space into simplicial complexes and use a framework of sequential composition to generate a successful plan. In collaboration with CMU (Likhachev), we will approach mobile manipulation tasks with dynamic constraints using a search-based planning framework. We propose extending search-based planners to incorporate dynamics on the large state spaces required for manipulation tasks for high-DOF manipulators. For example, consider a robot opening a heavy door. While a dexterous wrist may not be strong enough to push open the door, it excels at unlatching the handle. The wrist motion can be supplemented by the momentum of a heavy torso or using the base to push open the door.

In addition, we will develop graph searches that target specifically mobile manipulation tasks by exploiting the properties of these planning problems. Our goal is to provide methods that generate predictable consistent motions with rigorous guarantees on completeness and bounds on sub-optimality. In particular, in the coming year, we will be developing graph search-based planning specific to dual-arm manipulation, and we will investigate how the coupling constraints of dual-arm manipulation can be exploited to make search-based planning more efficient.

State of the Art:
Mobile manipulators are typically given open-loop trajectories to accomplish tasks; these trajectories should be feasible to avoid obstacles and satisfy input bounds as well as joint limits. One approach to generate trajectories is to convert the problem to a nonlinear constrained optimization problem and solve numerically. Finding an optimal solution to the nonlinear problem is not guaranteed in general, but we can find proper states of mobile manipulator that will simplify and help build a well-conditioned nonlinear problem.

Search-based planning with dynamic motion primitives has been demonstrably effective on low-dimensional systems, such as four-wheeled vehicles. Current state of the art involves generative methods for creating families of motion primitives. Inputs are fed to a dynamic simulation and trajectories recorded. Longer motions are decomposed and represented as compositions of shorter motions, until the motions are cut down to a desired number. Other approaches, for instance those based around RRTs, require dynamic simulation of the entire system while planning. The RRT may plan in either the space of control inputs (requiring an accurate knowledge of system dynamics) or in the positions and velocities of the system.

Key Barriers and Approach:
We propose to extend search-based planners to incorporate dynamics and function on the large state spaces required for manipulation tasks for high-DOF manipulators. *The approach to maintain computational tractability is three-fold: using adaptive motion primitives appropriate to the scale of the motion (fine motion versus gross motion), pre-computing constraint queries, and developing graph searches that exploit the specific properties of the graphs that represent the problem.*

The approach in the second technique is to use optimal control techniques to synthesize open-loop controls for the motion primitives and trajectories. Consider opening a spring-loaded door [Gray et al. 2011, Gray et al. 2012]. The task involves tight coordination of the arm and base motions and has constraints on the force needed to push the door open. In addition, the robot must maintain contact with the door and actively counteract the spring to keep it from closing. Either arm or the base may be used to contact the door. We use a compact graph-based representation of the problem to keep planning times low and solve using a search-based planner with kinematic motion primitives. The motions must satisfy joint torque limits, assuming quasi-static motion for the door opening. We propose to extend this work by using dynamic and adaptive motion primitives, eliminating the quasi-static assumption, and enabling more efficient motions which make greater use of the motion of the robot base. The results will be extensively validated on the PR2 platform before being adapted for the DR20.

In our research, a mobile manipulator needs to perform tasks such as moving a heavy object, manipulating a long, heavy object, or opening a heavy door. For these tasks, it is not enough to reach a particular configuration; the mobile manipulator must be able to exert necessary forces or dynamically balance itself while performing the task. We consider trajectory generation of a mobile manipulator while considering its own nonlinear dynamics in addition to geometric and kinematic constraints. The dynamics of mobile manipulator are:

$$M(q)q + \ddot{C}(q,\dot{q})\dot{q} + G(q) = u$$

where q is the vector of joint angles and u is the control input. We can re-parameterize the states of the system to simplify the constraint of the problem and use a nonlinear program to generate the trajectory.

To demonstrate the trajectory generation by solving an optimization problem, consider a mobile manipulator required to open a heavy push-bar door [Kim et al. 2012]. The mobile manipulator needs to exert a certain amount of force on the bar to unlatch the door while also exerting force to open the door. To simplify the problem, we consider a linear spring-damper contact between the door and the mobile manipulator and formulate an optimization problem. Figure 5-2 shows snapshots of a similar experiment with a heavy box. The base moves backward to generate momentum and then pushes the box with hand and base concurrently.

Figure 5-2.

To the third point of developing graph searches that gain efficiency by exploiting the specifics of manipulation tasks, we will investigate graph searches that capitalize on the fact that planning for dual-arm manipulation corresponds to planning trajectories in two graphs with continuous constraints on the two trajectories. One possibility is to represent such planning using a different

representation in terms of the end effector positions and free degrees-of-freedom (instead of joint angles for each arm). This is a more efficient lower dimensional representation. In addition, we will derive inadmissible but highly effective for dual-arm manipulation tasks heuristic functions and will develop graph searches that will use these heuristics to find arm motions that are within the sub-optimality bounds despite the inadmissibility of the heuristic function.

The proposed techniques will provide consistent, predictable motions for mobile manipulation with dynamics constraints and dual-arm mobile manipulation with rigorous deterministic guarantees on completeness and bounds on sub-optimality. This goes well beyond the current state of the art in mobile manipulation which typically relies on sampling-based approaches that cannot provide consistent predictable motions due to their random nature.

Metrics for Evaluation:
We will measure our progress using computational speed, percentage of successes and failures, and by proofs of completeness of the algorithms we are developing.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Search space for quasi-static manipulation including discrete modes and continuous state spaces; search space for dual-arm manipulation.
Q2	Realization on the PR2 ROS simulator.
Q3	Preliminary simulation results with a mobile manipulator in a three-dimensional, complex environment. Planning for transition to M3 and M4.
Q4	Software for realizing motion plans on the hardware testbeds.

Subtask 2: Synthesis of Arms and Effectors and Whole-arm Grasping (UPenn)

Objective and Benefits:
A modular reconfigurable manipulator can be used to adaptively shape the manipulator "fingers" to suit object size and shape. Reconfigurable grippers or end effectors can form shapes suited to grasping a large range of objects both in size and shape. Our reconfigurable effectors will start with a chain of hinged space-filling polyhedron (cube or right angle tetrahedron). This chain will be folded into different 3D shapes as needed by changing the relative configuration and connectivity between adjacent polyhedron modules. The hardware design and the algorithms will allow for gross conformational changes in size, shape, and topology much larger than attainable with soft grippers, such as the Festo gripper or humanoid style high-DOF grippers.

State of the Art:
Our recent paper [Seo et al. 2010] has presented one idea about how to reconfigure chain-type modular robots which can be thought of as the modular reconfigurable gripper here. The paper addresses the reconfiguration problem by presenting control algorithms based on the Carpenter's Rule Theorem. The theorem guarantees the existence of collision-free motions between two legitimate configurations devoid of self-intersections. Our algorithm provides decentralized

computations such that the idea can be applied to a chain-type modular robots with local-scale, distributed proximity sensing, and communication.

In terms of grasping methodology, we have been developing an algorithm for planar, bimanual, whole-arm grasping by introducing the concept of a generalized contact [Seo et al. 2011]. In addition to traditional contact models addressed in fingertip point contacts, planar enveloping grasps, and parallel-jaw grasps, a generalized contact can model all types of a contact between an open chain gripper and a polygonal object. Our algorithm can synthesize kinematically feasible contact configurations for immobilizing or equilibrium grasps derived from two generalized contacts. Figure 5-3 shows that the algorithm can be applied to a wide range of planar objects without customized grippers or end effectors.

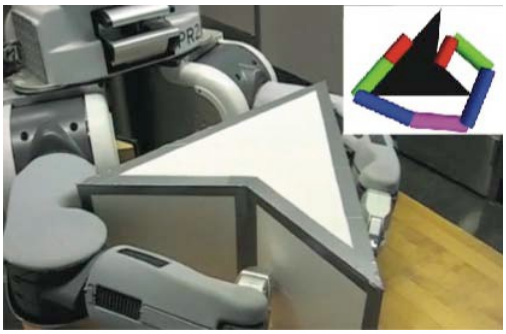

Figure 5-3.

Key Barriers and Approach:
There is extensive literature on grasping. Our focus, however, is on whole-arm or whole-body grasps where the goal is to restrain objects using multiple contacts without using hands or effectors made specifically for a few objects. We take inspiration from the simplest approach to grasping derived from caging, where multiple contacts are used to enclose an object, preventing it from escaping. While caging does not directly address the problem of completely restraining or even grasping the object, it can be a useful waypoint to grasping. For example, two-finger point grasps have an interesting property that every grasp that immobilizes an object must be preceded by a grasp configuration that cages the object. Whereas currently many approaches to grasping rely on passive conforming to objects (through compliant surfaces or joints) to achieve caging, we will explore actively formed rigid cages. This shape formation will at first be explored with object geometry entered manually with the future goal of automatically sensing the object shape. Furthermore, these ideas rely only on geometric constructs and do not require to model or assume the coefficient of friction or material stiffness which complicates analyses. Seo et al. 2011 shows the soundness of this approach in terms of synthesizing quality planar grasps. We will further investigate this strategy to synthesize general three-dimensional whole-arm grasps.

In order to implement the algorithms we are developing, we propose a modular approach to constructing arms and effectors that will leverage our experience and previous work in modular robots. We can quickly incorporate CKBot into our planning and control framework and synthesize effectors for specific tasks. There are two forms of effectors: one which uses active links in a hyper-redundant snake for fingers, and a second which uses chains with passive links that can be locked into place through externally applied forces from a robot arm [White, Thorne, and Yim, 2009].

Metrics for Evaluation:
We will measure our progress by evaluating our ability to grasp a wide class of objects. We will be specifically interested in cataloging objects that can be grasped using our approach that could not be grasped using alternative approaches. In addition, we will also measure speed of

reconfiguration, the efficiency of the planning algorithm, and proofs of completeness of the algorithms we are developing.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Algorithms for 2D whole-arm grasping using simple open chain grippers.
Q2	Developing CKBots and control strategies to synthesize open chain grippers on demand.
Q3	Evaluation of whole grasps using hardware testbeds: (a) the PR2 and (b) the reconfigurable CKBots.
Q4	Algorithms for 3D whole-arm grasping

Subtask 3: Accelerometer-based Contact Sensing for Autonomous Manipulation (UPenn)

Objective and Benefits:
The research objective of this subtask is to test the hypothesis that *high-bandwidth accelerometers can provide mobile manipulation robots with invaluable haptic perception information under control paradigms ranging from direct tele-operation to full autonomy*. As discussed above, robots need to be highly aware of their physical interactions with the world to be successful at mobile manipulation. As the robot executes a pre-planned sequence of movements, it should be able to confirm that desired contacts have been made so it can correctly proceed to the next step. Similarly, it should be able to detect when unintentional contact events occur, such as a collision with an unseen obstacle or an object slipping in the grasp. Robust physical contact sensing will enable a mobile manipulator to be much more capable in its interactions with unstructured environments, correcting any mistakes that may occur.

State of the Art:
Non-contact sensors such as cameras and laser range finders typically cannot detect robot-object contacts due to occlusion and limited viewing angle. Though correctly configured force sensors and distributed pressure sensors are more suited to this task, they are not currently feasible for field applications due to cost, complexity, and fragility. Autonomous mobile manipulators (and human operators controlling such robots through tele-operation) are, thus, left with no good method for perceiving physical contacts between the robot and its environment.

Key Barriers and Approach:
Fortuitously, almost all contacts between a robot and objects in the world cause distinct measurable vibrations, from driving over a rock to picking up a pipe. We, thus, aim to use an array of high-bandwidth three-axis MEMS-based accelerometers distributed across the robot's structure to enable it to perceive, localize, and categorize the intentional and unintentional contacts that arise between it and the environment. We have already successfully used the single high-bandwidth accelerometer mounted inside the PR2's gripper to improve its ability to manipulate objects [Romano et al. 2011]. We are working with QinetiQ to integrate a similar sensor on their Dragon Runner platform to understand the types of signals that will be measurable in field-deployed situations during both mobility and manipulation. At present, we

are focusing this thrust on haptic feedback for tele-operation to validate the usefulness of these cues, moving toward autonomy in the future.

Though clearly promising, especially for field applications, this robotic accelerometer-skin idea also has associated technical challenges. Most notably, the robot's own movements and internal mechanisms cause vibrations that are overlaid with contact signals, as illustrated in Figure 5-4. We have already developed a preliminary method of modeling and canceling out these ego-vibrations [McMahan and Kuchenbecker, 2012], but additional work will be required to make these methods robust, automatic, and extendible to high degree-of-freedom manipulators. We want a very high signal-to-noise ratio so that the operator or the autonomous robot can rely on the accelerometer signal without false alarms or missed contacts; detected transients should all indicate a contact between the robot and an item in its environment.

Figure 5-4.

Our past work has used just one sensor per robotic arm, but we envision placing eight or more MEMS-based accelerometers on a mobile manipulation platform to increase coverage. The sensors located closest to a contact yield the strongest and quickest signals, giving hope that we can deduce contact location by monitoring all signals simultaneously. For example, this would let us judge whether a grasped object was touched on the left or right fingertip, or whether a collision was occurring on the proximal or distal link of the arm. Furthermore, we have anecdotally observed that the shape of the contact acceleration transient contains information about the type of object being contacted, and the overall magnitude provides information about the relative speed of impact. We believe we can develop machine-learning-based signal processing algorithms that estimate the location of contact, object stiffness, and/or relative speed of impact with reasonable accuracy to increase the usefulness of this information.

We are planning a series of project components to moving from validation of the usefulness of this information for field-relevant tasks via tele-operation experiments with human subjects to technical refinements of the signal processing and event/pattern classification, moving toward providing a new channel of perception for fully autonomous mobile manipulation robots. 2012 will focus on refining our haptic feedback approach, understanding the effect of this type of feedback during tele-operation (as part of IRA2), refining our ego-vibration cancelation algorithms, and exploring the potential benefits of systems with multiple accelerometers.

Metrics for Evaluation:
We plan to conduct a formal experiment with the haptic Dragon Runner as part of IRA2 in March 2012. Operators will repeatedly complete a series of challenging mobility and manipulation tasks at QNA's test site. They will be situated in a nearby vehicle viewing the robot's video feed on the controller; for half of the trials, they will receive haptic vibration feedback, and for the other half they will not. We will record quantitative performance metrics such as task completion time and error rate, along with subjective assessments of mental

workload, stress, and preference. We will statistically analyze the results to determine whether and how the haptic feedback affects the operator. We expect it will increase task performance and reduce mental workload and stress. We will publish the results of this experiment in a peer-
reviewed haptics or robotics conference and/or journal. This experiment marks an important step on the path toward our ultimate goal of providing a new highly useful perceptual channel for autonomous mobile manipulation robots. If the study shows no or low benefit for the haptics, we will address any obvious deficiencies in the approach and/or reconsider this line of research.

Once we understand how these signals may benefit human operators, we will focus on analyzing the measured vibrations in real-time to help autonomous robots understand how their physical interactions with the world are progressing. We will conduct experiments with multiple accelerometers attached to the PR2 robot arm and develop algorithms for detecting contacts and localizing them on the arm. We will refine our methods for handling ego-vibration and look toward integrating them with the contact detection algorithms for improved performance when the robot is moving. These activities will be assessed through quantitative metrics on the proportion of contacts that our algorithms correctly detect and localize, along with false positive rates.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Work with QinetiQ to finalize haptic Dragon Runner for IRA2. Conduct preliminary experiments with multiple accelerometers on the PR2.
Q2	Analyze data recorded during IRA2 to determine usefulness of vibration feedback in field scenarios. Assess need for ego-vibration cancelation on Dragon Runner.
Q3	Publish results from IRA2. Refine processing algorithms for multiple accelerometers.
Q4	Test processing algorithms for multiple accelerometers, incorporating ego-vibration cancelation as needed. Publish results for contact sensing with multiple accelerometers.

References:

Sachin Chitta, Benjamin Cohen, and Maxim Likhachev. Planning for Autonomous Door Opening with a Mobile Manipulator, International Conference on Robotics and Automation, Anchorage, Alaska, 2010.

Steven Gray, Christopher Clingerman, Sachin Chitta, Vijay Kumar, and Maxim Likhachev. Search-based planning for autonomous spring-loaded door opening. Submitted to IEEE International
Conference on Robotics and Automation, 2012.

Steven Gray, Christopher Clingerman, Maxim Likhachev, and Sachin Chitta. PR2: Opening spring-loaded doors. PR2 Users Workshop, IEEE International Conference on Intelligent Robots and
Systems, 2011.

Vinutha Kallem, Adam T. Komoroski, and Vijay Kumar. Sequential composition for navigating a nonholonomic cart in the presence of obstacles. 2010. In review.

Vinutha Kallem and Vijay Kumar. Tuning vector fields for optimality in nonlinear control systems, 2010.

In preparation.

Soonkyum Kim, Vinutha Kallem, and Vijay Kumar. Motion planning for dynamic mobile manipulation. Submitted to IEEE International Conference on Robotics and Automation, 2012.

William McMahan, Karlin Bark, Jamie Gewirtz, Dorsey Standish, Paul D. Martin, Jacquelyn A. Kunkel, Magalie Lilavois, Alexei Wedmid, David I. Lee, and Katherine J. Kuchenbecker. Tool vibration
feedback may help expert robotic surgeons apply less force during manipulation tasks. In G. Z. Yang and A. Darzi, editors, Proc. Hamlyn Symposium on Medical Robotics, pages 37–38, June
2011a.

William McMahan, Jamie Gewirtz, Dorsey Standish, Paul Martin, Jacquelyn Kunkel, Magalie Lilavois, Alexei Wedmid, David I. Lee, and Katherine J. Kuchenbecker. Tool contact acceleration
feedback for telerobotic surgery. IEEE Transactions on Haptics, 4(3):210–220, July–September 2011b.

William McMahan and Katherine J. Kuchenbecker. Spectral subtraction of robot motion noise for improved vibrotactile event detection. Submitted to IEEE International Conference on Robotics and Automation, 2012.

Joseph M. Romano, Kaijen Hsiao, Günter Niemeyer, Sachin Chitta, and Katherine J. Kuchenbecker. Human-inspired robotic grasp control with tactile sensing. IEEE Transactions on Robotics, 27(6):1067– 1079, December 2011.

Joseph M. Romano, Jordan P. Brindza, and Katherine J. Kuchenbecker. Using environmental sound detection to improve the success of autonomous robotic manipulation. Submitted to IEEE
International Conference on Robotics and Automation, 2012.

Jungwon Seo, Steven Gray, Vijay Kumar, and Mark Yim. Reconfiguring Chain-type Modular Robots Based on the Carpenter's Rule Theorem. The Ninth International Workshop on the Algorithmic Foundations of Robotics. December 2010.

Jungwon Seo, Soonkyum Kim, and Vijay Kumar. Planar, Bimanual, Whole-Arm Grasping. Submitted to IEEE International Conference on Robotics and Automation, 2012.

P. J. White, C. Thorne, and M. Yim, "Right Angle Tetrahedron Chain Externally-Actuated Testbed (RATCHET): A Shape-Changing System", Proc. of ASME Intl. Design Engineering Technical Conferences & Computers and Information in Engineering Conference IDETC/CIE, San Diego, CA, Aug. 30 – Sep. 2, 2009.

M2-2012 – Principles of Generalized Grasp Mechanics

Level: 6.1 Basic Research

	Name	*Organization*	*Subtask*
Principal Investigator	Matthew Mason	CMU	M2-1
Additional Investigator	Joel Burdick	Caltech	M2-2

Objective:
This task's high-level objective is to develop new manipulation techniques suitable for use on lightweight mobile robots such as the Dragon Runner (DR20) or TALON. For the foreseeable future, these platforms will require simple, lightweight, low degree-of-freedom arms and grippers. Autonomous manipulation techniques previously developed for automated manufacturing or for anthropomorphic grippers are not applicable. We will explore and develop a different set of principles, better attuned to the complex interactions between practical effectors and objects of uncertain shape in pose in clutter.

Background:
Previous deployed robotic manipulation is of two kinds. In factory automation, autonomous manipulation usually requires a different gripper for every part. In tele-operation systems, a single general-purpose gripper can handle several different shapes but only with a human in the loop. Previous research on autonomous manipulation with general-purpose grippers most often focuses on anthropomorphic grippers with weight, cost, and complexity that are impractical for the platforms and applications of interest.

Research Description:
Autonomous manipulation of a wide range of shapes with a simple lightweight gripper raises problems not addressed in previous manipulation research. A simple gripper lacks sufficient degrees-of-freedom to drive a set of fingertips directly to a stable grasp configuration on an arbitrary shape in an arbitrary pose. When a typical simple gripper squeezes an object of uncertain shape and pose, the object shifts between the fingers, palm, and other objects in the environment until the gripper motion stalls, perhaps resulting in a stable grasp. Previous approaches based on accurate modeling, analysis, and prediction cannot be directly extended to such a complicated process. New problems arise, and old problems attain a new form:

- *Caging:* A "cage" is a loose grasp where the object is trapped by the effector, even though it is not completely immobilized. This reduces the requirement that an effector

conform to an object shape and pose and is, thus, especially well-suited to grasping a wide range of objects with a simple effector.

- *Effector design*: Previous effectors are designed either for grasping a single part or for tele-operation or are so complex and heavy as to be impracticable for anticipated scenarios. Our research focuses on novel designs better adapted to autonomous perception, planning, and control and suited to new manipulation strategies such as caging.
- *Grasp state estimation*: Since it is impractical to predict the detailed motion of a generic shape in clutter being grasped by a simple effector, grasp state estimation is a central problem. Our research explores machine learning algorithms, kinesthetic sensing, and computer vision to track the grasping process.
- *Planning*: Since grasping involves processes not readily modeled or predicted with accuracy, different approaches are required. Our research is exploring stochastic models of grasping, the use of machine learning, and planning for caging grasps.

Metrics:
Our ultimate goal is autonomous manipulation on small mobile platforms capable of a variety of missions including clearing rubble, uncovering IEDs, and supporting mobility. For that ultimate goal, the metric would be comparison with tele-operated manipulation on the targeted applications. Shorter-term work involves a variety of metrics: accuracy of state estimation, adaptability of mechanism designs, and speed and success rate of grasping behaviors.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
In September 2011, the following two papers were submitted to the IEEE International Conference on Robotics and Automation. If accepted (acceptance announcements will be made in mid-January), they will be presented at the conference in Minneapolis during May 2012.

1) R. Allen, J. Burdick, E. Rimon, "Two-Fingered Caging of Polygons via Contact-Space Graph Search," *IEEE Int. Conf. Robotics & Automation*, 2012.
2) P. Hebert, N. Hudson, J. Ma, T. Howard, T. Fuchs, M. Bajracharya, & J. Burdick, "Combined Shape, Appearance, and Silhouette for Simultaneous Manipulator and Object Tracking," IEEE *Int. Conf. Robotics & Automation*, 2012.
3) M. T. Mason. Harvard EE Seminar Series. February 17, 2012.
4) M. T. Mason. MIT workshop "Celebrating the academic work of Rod Brooks". March 9, 2012.
5) M. T. Mason. College de France colloquium "Robotics: Science and Technology". June 12-13, 2012.
6) A. Rodriguez et al., Workshop on the Algorithmic Foundations or Robotics, (assuming paper is accepted), June 13-15, 2012.
7) R. Paolini et al, International Symposium on Experimental Robotics (assuming paper is accepted), June 17-21, 2012.
8) A. Rodriguez et al., Robotics: Science and Systems, (assuming paper is accepted), July 9-13, 2012.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Integrate Kinect. Complete design of prototype P3. Online two-fingered grasp planner. Integrate tactile sensing fusion mechanism.
Q2	Stochastic model for pick-and-place. Learning to employ Kinect data. Test planner for JPL gripper in simulation. Test grasp state estimator in lab.
Q3	Fabrication of prototype P3. Learn highlighter placing behaviors. Test grasping of tools in lab. Integrate estimator with online planning.
Q4	Tests of prototype P3. Mixed objective 3D design of phalanges. Test planner and grasp observer in field and IRA2.

Related Research:
The developed planning and grasp state estimation algorithms will be combined with manipulator planners (Task M1), behavior-based control systems (Task M3), as well as high-level object segmentation, tracking, and perception information (developed in Tasks P3, P6, and P7).

Integration and Assessment Activity:
In addition to grasp planners, this task will also produce optimal hardware configurations (end effectors, kinematics, and required arm DOF) to enable dexterous manipulation tasks on DR20 platforms. As part of the Capstone Assessment, this system will manipulate rubble piles, expose and extract objects, and open doors (element 10) as well as make pathways traversable by moving obstructions and opening gates (element 5). For IRA2, we will use a DR20 to assess our approach on an IED task at the FTIG MOUT site. FY12 tasks will not only include door opening and rubble grasping behaviors but also improve upon FY11 trenching and digging behavior execution. During FY12 and beyond, we will extend the planning system to 3D objects and to whole-body manipulation strategies in preparation for the Capstone Assessment.

Subtask 1: Simple Hands (CMU)

Objective and Benefits:
The goal is to develop autonomous manipulation capabilities with simple effectors. Working with simple effectors shortens the path to deployment on practical platforms and, at the same time, focuses our work on critical problems with broad impact. Our efforts are focused in four

areas: First, learning to estimate grasp state from kinesthetic data; second, learning effective grasping and manipulation behaviors; third, designing effectors; and fourth, exploring a new formulation of caging that connects directly to grasping. In the coming year, we plan to expand our work to a broader class of shapes, to learn good grasp locations based on 3D perceptual data, to learn more effective grasp behaviors, to explore intelligent phalanx form design, and to explore the application of caging functions to enveloping grasps (in collaboration with Kumar.)

This effort directly addresses various manipulation tasks that arise in the "cover the back door" scenario, particularly clearing rubble or other obstacles from the path of the robot. The capabilities are broad enough to also support other scenarios and applications of interest, such as inspecting and uncovering IEDs or supporting building clearing behaviors.

State of the Art:
Autonomous grasping attempts usually fall into one of two approaches. (1) Design the hand, and the grasp, for the object in question. This approach is ubiquitous in manufacturing automation but not practical for our applications. (2) Use an anthropomorphic hand and plan finger placements based on detailed knowledge of object shape and pose.

More recently a third approach has been developed: offline simulation of grasping behaviors to build up a grasping table, again based on detailed knowledge of object shape and pose.

In work so far, we have been exploring a different approach called "let the fingers fall where they may" or "grab first, ask questions later." We use near-blind grasping techniques without precise prior estimates of object shape or pose. We have developed effectors with broad capture regions and employed machine learning techniques to produce haptic perception routines. Thus far, we have had excellent results demonstrating grasping and pose estimation of a single shape, specifically highlighters. We have also had some outstanding new theoretical results, suggesting interesting directions for phalanx form design and the relation of caging to grasping.

Key Barriers and Approach:
There are several barriers to deployment in practical scenarios. The greatest and most interesting challenge is to clear a pile of rubble. While the simple hand philosophy is ultimately very well suited to the task, continued progress is necessary along several dimensions.

Generality of the hand. To clear a pile of rubble, and in other scenarios, the hand must deal with a variety of shapes, both familiar and novel. So far, our work has focused on a single object: highlighters. Nonetheless, we anticipate that the hand can handle other shapes. The hand was not designed for highlighters. The design was guided primarily by ease of analysis and fabrication. In blind grasping experiments, the hand can grasp a highlighter in any of several different poses and also grasp two or three or even four highlighters. The next step is to repeat our experiments with additional shapes.

Beyond blind grasping. Blind grasping is not an essential part of our approach. Rather, it is an extreme case that proves we can grasp objects without detailed knowledge of shape and pose, which is essential to planned application scenarios. The question now is how to use vision that

provides coarse knowledge of a scene. While we do not need detailed knowledge of shape and pose, we do need information about the scene, including the approximate object position and the surrounding clutter. Most previous grasping research has ignored the issue of clutter, focusing on neatly isolated objects. We have focused instead on high-clutter, target-rich experiments, but ironically, we have not yet demonstrated grasping of isolated objects. The next step is to integrate a Kinect to provide coarse three-dimensional maps of a scene and develop heuristic behaviors for guiding the grasp, with or without clutter.

Beyond grasping. Grasping is not the whole problem. In the "cover the back door" scenario, the real problem might be manipulation in support of locomotion, such as clearing obstacles. This does not necessarily involve grasping at all, such as when pushing or kicking objects. When grasping is employed, some disposal technique follows, such as throwing, placing, or pushing. The first step is to expand our experiments beyond grasping to demonstrate an entire task. We are focusing on placing of highlighters. Preliminary experiments show that we can place highlighters, and are working on a stochastic model allowing us to explore possible advantages of end-to-end learning, rather than viewing grasping as an intermediate sub-goal.

Learning behaviors. So far, we have used ad hoc grasping behaviors based on our own intuition and experience. Our machine learning techniques are focused on perception. The next step is to apply machine learning to generate more effective behaviors. The main problem is that policy space is combinatorically larger than the space of sensations. The first step is to develop a low-dimensional set of policies that would be amenable to learning yet still rich enough to cover the variety of grasping behaviors required in like application scenarios. Initial experiments with placing behaviors are encouraging.

Effector design. Our approach is fundamentally motivated by the simple effectors used on robots currently in theater or likely to be used in an application scenario. It is unlikely that anthropomorphic hands will be practical for such robots. Nor will practical perception devices be able to provide the detailed information on shape and pose that other approaches assume. So far, we have designed, fabricated, and tested a sequence of three prototype designs, each more capable than the one before. Our current prototype was designed to eliminate inter-finger contact (to improve learning) and to be more rugged. Ruggedness is essential not just for practical application but for machine learning research. Previous prototypes were capable of only a few hundred grasps before wear made the learned information obsolete. The current prototype has performed thousands of grasps without significant wear.

We are exploring several directions related to effector design. Our next prototype will have additional sensors in the form of a force and torque sensing palm. It also will have a simpler transmission, moving us closer to a design that will be simple, rugged, and small enough for practical application.

The most interesting new direction is in finger design. Previous work on finger design has focused on kinematics of multi-joint multi-phalanx fingers. Our focus on single-phalanx fingers has led us to an interesting new direction: design of phalanx form to provide mechanical adaptation to variations in object form and pose. Single-phalanx grippers are ubiquitous on the

platforms envisioned for application scenarios. This will be the focus of Alberto Rodriguez's PhD thesis.

From caging to grasping. Caging is a type of loose grasp where the object is not immobilized but, nonetheless, cannot escape from the hand. Our main focus is on the connection of caging to grasping. Our main focus is the *pregrasping cage:* a configuration of the hand that is a suitable via point to a grasp, in that a path to a grasp exists that maintains caging. Not all cages are pregrasping cages. Our main contribution was to formulate a function similar to a Lyapunov function and to use that function to generalize previous results on two-finger cages to n-finger cages. Our immediate goal for the future is to explore connections with Kumar and the work proceeding under Subtask M2-2.

Metrics for Evaluation:
For our grasping, localization, and placing work, there are a variety of metrics: precision, recall, and accuracy, and number of trials for learning, to name a few. For our effector design work, mean number of grasps before failure is an important metric since our learning systems depend on numerous trials. We also use mean time between successful grasps. Informally, our effector design is guided by the goal of minimizing the weight, size, cost, number of motors, and number of sensors. In the future, we will measure the breadth of shapes that can be handled. Thus far, the most improvement has been in the effector's ruggedness, which can now perform thousands of grasps before breaking or wearing to the point that the learned data is obsolete. Precision, recall, number of trials required, and mean time to successful grasp have all improved dramatically from our first experiments.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Integerate Kinect. Complete design of prototype P3.
Q2	End-to-end stochastic model, applied to pick-and-place. Learn high-payoff target areas from Kinect data.
Q3	Fabrication of prototype P3. Learn highlighter placing behaviors.
Q4	Initial tests of prototype P3. Initial design concepts for prototype P4. Extension of phalanx design to 3D, mixed objective.

Subtask 2: Novel Mechanisms for Caging Manipulation (Caltech)

Objective and Benefits:
This task's high-level objective is to develop and deliver fieldable robotic manipulation strategies, planning algorithms, grasp state estimators, and novel grasping mechanisms that can be used on lightweight man-portable mobile robot platforms, such as the Dragon Runner (DR20) or TALON robotic systems.

State of the Art:
The grippers and arms on field robots are likely to have lower numbers of degrees-of-freedom (DOF) and greater inaccuracies than conventional robots for which the bulk of grasping and manipulation theory was developed. Achieving robust interaction with the wide range of objects found in field operations requires an understanding of the mechanics and geometry underlying the grasping process so that simpler robots can achieve near human-like capabilities in such tasks as lifting objects from a rubble pile (which may be hiding an IED), turning a doorknob (in support of Soldier reconnaissance), or transporting supplies (e.g., a logistics robot). Robust sensing of the manipulator posture and grasped object state are also required for low mass and lower DOF mechanisms which are inherently less accurate. Sensing, coupled with appropriate planning and task execution strategies, can compensate for lower DOF.

Key Barriers and Approach:
Our approach to practical grasping strategies is based on the notion of a *caging grasp.* A cage is a configuration where the physical presence of the finger tips, finger linkage surfaces, and/or arm surfaces guarantees the object to be grasped cannot escape – the finger/arm/body surfaces form a "jail" which traps the object. A cage can serve as a pre-grasping phase which prepares the hand (or opportunistic assembly of robot body parts and environmental leverage points) to safely grasp an object. From the caged configuration, the contact surfaces are then "squeezed" inward to secure the grasp (in preparation for object transport or object manipulation). If the hand starts in a cage, it is guaranteed not to lose the object during the grasp acquisition phase. A cage can also serve as a flexible manipulation strategy for low degree-of-freedom field grippers. The posture of the caged object can evolve within the confines of the cage while respecting constraints of the low-DOF arm/gripper. Conventional grasp and manipulation strategies assume sufficiently high degrees-of-freedom and an accurate instant-by-instant assessment of the state of contact between the object and the gripper's fingers. These constraints are unworkable for field systems.

The concept of a cage was introduced to robotics by Elon Rimon and Andrew Blake in 1996. However, until recently, the fundamental theory of caging grasps has not been well developed, particularly with respect to field applications. To date, caging algorithms are only known for two fingers or for three-finger grasps of planar polygonal objects in limited fashions. No caging algorithms have been developed for reduced degree-of-freedom hands. As described below, we have developed a new analytical framework for caging which should allow us to develop new caging algorithms that meet the demands of field deployment.

Note that while a squeezing cage is guaranteed to end up with a secure grasp, it cannot be precisely predicted where the fingers will come to rest upon the object after squeezing. Hence, to robustly execute a subsequent manipulation task, it is necessary to estimate the pose of the grasped object within the grasp and track its location during manipulation. Rapid advances in optical sensing technology have drastically reduced the cost of using stereo cameras, flash LIDAR (3D cameras), and/or RGB-D cameras (such as the Microsoft Kinect). Hence, rich visual information is potentially available at low cost to enable such tracking if appropriate sensor fusion algorithms can be developed. We have begun developing such a fusion framework which is also compatible with the efforts of task M3 and perceptual Tasks P3, P6, and P7.

Research Description: Over the next three years, this task will enable fieldable manipulator and gripper platforms to interact autonomously with the environment, by finding a set of robust grasps for identified objects, computing grasp motion strategies, and monitoring the grasp execution process. During the first program year, we made progress in two areas:

(1) new caging analyses and algorithms; and
(2) new vision based grasped object tracking and pose estimation strategies that use low-cost optical sensors.

In the area of caging, we developed a fundamentally new analytical framework for developing caging algorithms, based on a reformulation of the critical caging events in contact space. This reformulation has led to two innovations: (1) the first caging planner suited for reduced degree-of-freedom grippers, and (2) a new efficient algorithm for two-fingered whole gripper manipulation (as needed for the 2012 IRA). We have initially tested this planner offline on a complex object (a Bosch drill). Also, during this previous year, we have successfully developed an algorithm to fuse stereo vision, high resolution monocular vision, and 3D camera (either flash LIDAR or RGB-D) output with simple binary tactile sensing of the type which can be found on a field gripper. We have implemented this algorithm in a C++ package and have tested it in a laboratory setting. We have found that typical objects (e.g., tools such as screwdrivers) can be localized to within 3 to 5 millimeters.

The near-term goals of this work are to transition the offline grasp planner to an online architecture and specialize it to the geometry of the gripper to be used in the 2012 IRA. Next, we will integrate the planner, a novel gripper, and the grasp state estimator into a DR20 platform. These efforts will prepare our system to be evaluated in the next IRA. In the longer term, we seek to extend the theory to robust whole-gripper manipulation planning for more than two fingers.

Metrics for Evaluation:
The conceptual goal of this task is to enable robust autonomous grasping and manipulation behaviors on fieldable systems. This will enable more efficient and intuitive operation of such vehicles by remote operators. Hence, the most effective metric to demonstrate progress is direct comparison of task execution time between the developed autonomy system and that of a tele-operated DR20 performing the same task.

	Subtask 2 Quarterly Milestones and Deliverables: Caging Grasp Planner
Q1	Convert offline two-fingered grasp planner to online architecture.
Q2	Specialize planner for new JPL reduced DOF gripper; test in simulation.
Q3	Integrate online version of planner into DR20. Test automated grasping of tools in laboratory environment. Compare speed and robustness against tele-operation of same task.
Q4	Test and refine planner in field conditions and IRA2.

	Subtask 2 Quarterly Milestones and Deliverables: Grasp State Estimator
Q1	Integrate improved tactile sensing fusion mechanism into grasp state estimator.
Q2	Test grasp state estimator in laboratory environment (on fixed Barret Wam arm/hand) on tool acquisition task. Compare speed and robustness against tele-operation of same task.
Q3	Integrate estimator with online planning system.
Q4	Test and refine estimator in field conditions and IRA2. Test operation of integrated planner and estimator on tool acquisition, rubble removal, and door opening tasks. Evaluate metric of autonomous operation versus tele-operation on these tasks.

M3-2012 – Sensor-based Dexterous Manipulation

Level: 6.2 Applied Research

	Name	*Organization*	*Subtask*
Principal Investigator	Nicolas Hudson	JPL	M3-1
Additional Investigators	Matthew Csencsits	QNA	M3-2
	Juan Pablo Gonzalez	GDRS	M3-3

Objective:
Thi task will develop the capability to make deployable, man-portable mobile manipulation platforms such as the Dragon Runner (DR20) or TALON robot perform tasks autonomously or semi-autonomously in unstructured environments as well as approach the capabilities of high-performance, non-fieldable, state-of-the art systems such as the PR2 or Barrett WAM manipulators. Most relevant mobile manipulators (DR20, TALON, Mars Rovers) are characterized by low degrees-of-freedom (<6 DOF), precision and control bandwidth, and high gear-rations, backlash, and impedance. Three complementary approaches are taken to enhance deployable autonomy capabilities: M3-1 addresses operational and system uncertainty by developing multi-sensor feedback behaviors, which learn and adapt to the environment through estimation of the joint robot-world interaction state. M3-2 investigates and prototypes the next generation of field hardware by augmenting the current DR20 and TALON platforms with swappable task-specific end effectors, optimal kinematic configurations, and required low-level control interfaces. M3-3 develops planning algorithms for small mobile platforms in military relevant environments.

Background:
Military robots must routinely manipulate previously unseen and unmodeled objects while in uncharacterized terrain. Currently, human operators account for this high degree of uncertainty when performing manipulation tasks, but many of these tasks are tedious or difficult using current tele-operated interfaces with monocular video feedback. Most current autonomous manipulation systems operate in known environments, where the system control gains and planners have been tuned for particular configurations and situations. High-performance manipulators systems can reduce some of these issues as they are inherently compliant, through design or high-bandwidth control or have high precision perception and operation. Autonomy for deployable mobile manipulators instead requires adapting to the environment and incorporating uncertainty and limitations of the hardware systems in both planning and control.

Research Description:
To move beyond current tele-operated mobile manipulator systems, new sensing, tools, kinematic configurations, controllers, and planners must be added to existing deployable systems. To gain the maximum possible benefit from current systems while maintaining cost effective and deployable solutions, augmentation of current hardware configurations are considered. This includes both new end effector tools which are compatible with current generation hardware and kinematic and electro-mechanical modifications to enable execution of developed autonomy systems. In partnership with mechanical modifications to existing systems, new local task frame control behaviors are developed. These Generalized Compliant Motion (GCM) behaviors utilized multiple sensor modality feedback to learn and adapt to the environment, modifying task frame motion in response to new information. These behaviors allow high-level semantic commanding of tasks, such as dig, trench, open, without the need to specifically tune or apriori define models of the environment. Instead, the system estimates the robot-environment interaction state, including uncertainly in the position of the robotic manipulator, and environmental parameters by utilizing both visual and kinesthetic information. Current mobile manipulator planners are being analyzed with respect to deployment on small mobile manipulators. Adaptation of these algorithms will take into account the system constraints including limited perception capabilities and reduced open-loop precision.

Metrics:
The utility of this task can be directly measured by comparing the speed, robustness, and flexibility of tasks performed by tele-operation and developed autonomy (planning and adaptive behavior execution). New hardware development, including kinematic configurations and end effector tools for specific tasks will also be evaluated using the same metrics by comparing task execution against the baseline DR20 or TALON platforms. The speed of the task is measured as the time it takes the operator to complete a specific task. The robustness is measured as the number of times that (or the degree to which) the operator is able to successfully complete the task. The flexibility is measured as the number of environments/situations in which one task can be performed or the number of different tasks that can be performed.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Nicolas Hudson (JPL) will be present at ICRA and RSS conferences and workshops related to mobile manipulation. Matthew Csencsits replaces Scott Thayer at QNA in M3 from the FY11 APP. Juan Pablo Gonzalez at GDRS is now performing work under M3-3 rather than M7-2 due to the related nature of the tasks.

Milestones:
Initial work in FY12 Q1 will focus on improving autonomous trenching capabilities, where a DR20 will expose shallow buried wires. This work includes assessment of new trenching end effector tools and autonomous behaviors in IRA2. Also in Q1, end effector tools for effective opening of car doors will be analyzed. Q2 work will develop autonomous behaviors for car door opening, enabling operator point-and-click commanding of the robotic system. In Q3,

demonstration and assessment of car door opening will occur, utilized direct comparison of tele-operated and autonomous operations using the metrics defined below. Using developed technologies from M2 (Robust Caging) and P3 (Static Scene Understanding), hardware and behaviors for manipulation of rubble piles will be investigated. Manipulation of rubble piles will be assessed in Q4.

	Quarterly Milestones and Deliverables
Q1	Trenching for IRA2. Initial work on autonomous DR20 object grasping and dragging. Preliminary designs for end effector tools to enable easy opening of car doors. Kinematic analysis of car door opening and rubble clearing with low-DOF manipulators. Analyze existing mobile manipulation approaches.
Q2	Develop door opening behaviors using developed end effector tools. Evaluate performance and limitations of HDT manipulator. Evaluate existing manipulation planning frameworks.
Q3	Demonstration of car door opening task. Initial work for rubble pile grasping and manipulation behaviors. Develop global manipulation planner suitable for low-DOF manipulators.
Q4	Demonstration of rubble pile grasping. Work with I2 to develop initial implementation of manipulation planning framework for RCTA.

Related Research:

- Task M1: Theory and Principles of Mobile Manipulation
- Task M2: Principles of Generalized Grasp Mechanics.
- Task P3: Static Scene Understanding
- Task P6: Perception for Missions in Dynamic Environments
- Task I2: Data Mapping for Inference and Focus

Integration and Assessment Activity:
This task will provide trenching behaviors and DR20 tools and hardware for IRA2. Ongoing work on platform development, planners, and adaptive behaviors for autonomous digging, trenching, grasping, dragging, and opening will be provided for integration activities. Developed technologies and capabilities correspond to elements 5 and 10 in the RCTA Capstone Integration and Assessment Vision.

Subtask 1: Sensor-driven Adaptive Behaviors for Autonomous Manipulation (JPL)

Objective and Benefits:
The objective of this task is to develop adaptive, sensor-driven, autonomous manipulation behaviors for low degree-of-freedom (<6 DOF) arms on lightweight man-portable mobile

manipulation platforms such as the Dragon Runner (DR20) or TALON robotic systems. Over the next three years, this task will enable fieldable (highly geared, lightweight, non-backdrivable) manipulators to interact autonomously with the environment by estimating the combined (robot and environmental) state during interaction for tasks such as digging, trenching, door opening, and grasping. As part of the Capstone Assessment, this autonomy system will enable manipulation of rubble piles, expose and extract objects, and open doors (Element 10) as well as make pathways traversable by moving obstructions and opening gates (Element 5).

The developed task frame control behaviors and state estimation algorithms are designed to be compatible with future open-loop manipulator planners (M3-3 and M1), grasping algorithms (M2), as well utilizing object segmentation, tracking, and perception information (P3 and P6). This task will evolve in lockstep with M3-2, which will investigate and produce optimal hardware configurations (end effectors, kinematics, and required arm DOF) to enable complicated manipulation tasks such as door opening on DR20 platforms.

State of the Art:
Today's fielded hardware manipulation systems either consist of precise devices working in known environments, such as manufacturing, or are low precision tele-operated or scripted systems (EOD robots, Robotic Explorers). Most research manipulator systems use heavy, high bandwidth joint-torque control (industrial manipulators) or non-deployable low impedance arms (PR2, Barrett WAM) and rely on highly accurate positioning and controllable arm dynamics.

Key Barriers and Approach:
The current challenge is to approach the autonomy and performance of highly accurate, low-impedance manipulators with high-impedance, low cost-and-weight fieldable manipulation systems. These systems, such as the DR20 and TALON platforms, are low precision, low control bandwidth, low repeatability (backlash), and have minimal degrees-of-freedom (3 to 5 DOF). Instead of requiring higher performance arms, we are developing Generalized Compliant Motion (GCM) behaviors to enable these manipulator systems to interact with the environment through sensor-driven state estimation and adaptive feedback behaviors. This task uses real-time estimation of relative manipulator-environment pose as well as object properties and closed chain properties such as stiffness to adapt low-rate task frame behaviors during task execution. Manipulator "compliance" is derived from visual and kinesthetic feedback behaviors rather than being implicit in the arm hardware.

These GCM behaviors utilize multiple sensor modalities to learn and adapt to the environment and change task frame motion in response to new information. By fusing both visual information (point cloud data, visual odometer, arm tracking) and kinesthetic feedback (contact maps, forces, kinematic arm positions), the joint robot-environment state can be inferred. This state includes the position of the end effector, properties of the environment, and pose of objects in the environment. As properties of the environment become better understood, task control behaviors can be updated. GCM then provides a framework to merge concurrent task frame goals such as open-loop plans (M3-3), force control, visual tracking, and kinematic constraints. In Subtask 1, tasks such as "open" or "trench" are decomposed into a hierarchical task network of actions,

where each action is a specific set of adaptive control behaviors, which are merged and executed in the task frame.

In FY11, a cost share DR20 mobile manipulator platform was augmented at JPL with a turret-mounted stereo camera pair, a wrist-mounted 6-axis force-torque sensor, and associated computational hardware. In addition, the adaptive behaviors for estimation of system stiffness were developed as well as behaviors for trenching. Also in FY11, initial work was conducted with M3-2 to integrate developed M3-1 behaviors interfacing with the DR20 Development Kit joint control interface. Refinement of this interface is ongoing.

Initially in FY12, the developed trenching behaviors will be finalized for IRA2, where the DR20 platform will expose shallow buried wires in a semi-autonomous manner. This development will involve M3-2 by updating the interface to DR20 manipulator joint control. M3-1 and M3-2 will continue to work collaboratively to define the minimum requirements for DR20 manipulator control. The developed adaptive behavior framework will continue to be developed for the DR20 and TALON platforms. In FY12, we extend our approach to continuously estimate the pose and properties of movable objects in the environment during interaction as well as develop behaviors for door opening with low-DOF mobile manipulators. This capability will be applied to opening car doors and manipulating rubble piles. Beyond FY12, increased use of perception and modeling tasks outputs will generalize developed behaviors and will lead to the Capstone Assessment.

Metrics for Evaluation:
Task execution of the developed autonomy behaviors will be directly compared to baseline task execution by tele-operation where the speed, robustness, and flexibility of tasks is quantified. The speed of the task is measured as the time it takes the operator to complete a specific task. The robustness is measured as the number of times that (or the degree to which) the operator is able to successfully complete the task. The flexibility is measured as the number of environments/situations in which one task can be performed or the number of different tasks that can be performed.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Trenching for IRA2. Initial work on autonomous DR20 object grasping and dragging.
Q2	Door opening behaviors using tools developed in M3-2.
Q3	Demonstration of car door opening task. Initial work for rubble pile grasping and manipulation.
Q4	Demonstration of rubble pile grasping (using with M2-2 algorithms and M1 or M3-3 planning).

Subtask 2: Car Door Opening and Rubble Clearing with a Small UGV (QNA)

Objective and Benefits:
We will push the limits of what can be accomplished by a low-DOF (< 6 joints) manipulator by exploring the tasks of opening a car door and clearing debris from a rubble pile with the Dragon Runner 20 (DR20) system. We will investigate approaches to each of these tasks and assess the best combination of end effector design and kinematic arrangement to perform these tasks. The results of this analysis will then be used under QNA's IR2-3 task to modify the existing DR20 manipulator to achieve this kinematic arrangement as well as enable the necessary quality of joint control, feedback, and update rates. The successful completion of this effort will provide the capability of opening car doors to enable inspection of their interiors and the ability to clear debris from a desired/necessary path enabling inspection and traversal. These capabilities correspond to items 5 and 10 in the RCTA Capstone Integration and Assessment Vision. We anticipate performing the car door task in a 6- to 8-month timeframe. We expect the efforts of clearing debris from a rubble pile to span two years, with incremental progress in performance to be shown every other quarter starting Q4 FY12.

State of the Art:
The state of the art in fielded mobile manipulator systems consist of low-DOF manipulators with low-fidelity, joint-level control which are tele-operated.

Key Barriers and Approach:
We will work with JPL to apply generalized force control to the tasks of opening car doors and removing debris from a rubble pile. We will examine each task for ways to do them with few degrees-of-freedom and simple end effectors. We will leverage our IR2-3 work to make improvements to current SoA hardware to enable superior feedback and control and enable tighter integration of other Consortium members' work. We will also make modifications to current SoA hardware (under IR2-3) to better align with the "best" kinematic arrangement for given tasks. We will move beyond the current SoA by improving current SoA hardware and by automating tasks which would otherwise be tediously performed via tele-operation. We introduce the capability to perform higher-level autonomous manipulation functions such as force-control, policy-based manipulation planning, and coordinated autonomous mobile manipulation.

The majority of car doors are opened with either a lever-style or pull-to-open handle. Both types can be addressed with a simple hook-like end effector, which would allow for the manipulator to be less precise in its control and less capable in its dexterity but still able to perform the motions necessary for opening a car door. We believe the current kinematic arrangement of the "3-DOF" manipulator on a DR20 will be sufficient for opening car doors but may benefit from modifications in link lengths and joint offsets. This should reduce the risk for the mechanical aspects of the car door task and permit us to focus on the control and perception aspects.

The mobility of the DR20 platform should be capable of overcoming the limited dexterity of the manipulator and will allow us to begin removing debris early in the effort. Continued effort to apply grasping knowledge and force control should enable us to decrease the time it takes to

clear debris throughout the year. As every manipulator platform can be confronted by objects that exceed its capabilities, we plan to test the DR20 manipulator against debris within its capacity to lift or drag, but the approach developed will be transferable to a manipulator of any scale.

Ultimately, we will utilize results of Tasks P3, P4, P5, P6, M1, M2, and M7 to assemble a system capable of identifying car door handles and debris, planning motions to grasp and manipulate them, and leverage generalized force control to maintain contact and grasp stability. We will utilize work in M2 to determine how to manipulate car door handles and design simple end effectors for opening car doors. This task also relies on QNA's IR3 work related to the DR20 Development Kit.

Metrics for Evaluation:
We will gauge the performance of our system by measuring the time it takes to complete the following tasks along with the number of errors occurred during a task attempt (e.g., unintentionally dropping debris or letting go of the door handle).

Arm modifications have been made to support perception, force-torque sensing, and position control with SoA hardware. Requirements and preliminary design work have been started to improve the SoA hardware.

During this task, we will periodically evaluate our task metrics utilizing the rubble pile in QNA's small UGV proving ground, which was developed under task IR2 in 2011 and the Mars Yard at JPL. When appropriate, we will find other test sites like Fort Indiantown Gap, Disaster City, and the Southwest Research Institute. We will also discuss other applicable metrics with ARL.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Preliminary designs for end effector tools to enable easy opening of car doors. Kinematic analysis of car door opening and rubble clearing tasks with low-DOF manipulators.
Q2	
Q3	Demonstration of car door opening task with current DR20 manipulator kinematic design.
Q4	Initial evaluation of rubble clearing capability.

Subtask 3: Mobile Manipulation for Small Platforms (GDRS)

Objective and Benefits
The small robotic platforms targeted by the current CTA present unique challenges and opportunities for mobile manipulation. They are able to enter rooms and buildings to retrieve a variety of objects of interest, but they are limited in perception capabilities, processing power, and payload capacity. Currently, almost all mobile manipulation tasks for small platforms are achieved using tele-operation.

This subtask seeks to identify relevant mobile manipulation tasks and develop planning algorithms to perform them autonomously in small platforms.

This effort will enable the capability to manipulate objects in the environment, either to clear the way for a robot to move or to retrieve objects of interest.

This is a multi-year effort that will start by addressing manipulation of known objects and will eventually enable manipulation of arbitrary objects.

State of the Art:
Currently, almost all mobile manipulation tasks for small platforms are performed using tele-operation. There is a growing body of literature in mobile manipulation for larger platforms with high-end, high-DOF manipulators, such as the PR2, but not much in small platforms or in small platforms with low-end manipulators.

Key Barriers and Approach:
We will survey the state of the art and identify elements of existing technologies that can be used in small platforms. We will then identify the shortcomings of such approaches and propose new ways to address such shortcomings in ways that are applicable to small platforms. For example, the planners being developed in M1-3 assume a high-end manipulator that provides accurate feedback at a relatively high rate. Transitioning from this 6.1 research to a manipulator that is suitable for a small platform would likely require significant modifications or brand new algorithms.

Metrics for Evaluation:
Initially, the metrics for assessment will be the ability to successfully manipulate objects in a relevant scenario. Once this capability matures, the focus will shift to improving the op-tempo of the process and addressing more complex situations where the manipulator and the base need to coordinate its motion in order to complete a task.

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Analyze existing mobile manipulation approaches.
Q2	Evaluate performance and limitations of HDT manipulator. Evaluate existing manipulation planning frameworks.
Q3	Develop global manipulation planner suitable for low-DOF manipulators.
Q4	Work with I2 to develop initial implementation of manipulation planning framework for RCTA.

M4-2012 – High Degree-of-Freedom Dynamic Manipulation

Level: 6.2 Applied Research

	Name	*Organization*	*Subtask*
Principal Investigator	Alfred Rizzi	Boston Dynamics	N/A
Additional Investigator	Mike Murphy	Boston Dynamics	N/A

Objective:

Overall objectives:

- Develop and demonstrate techniques for dynamic mobile manipulation on a high degree-of-freedom legged platform.
- Explore full-body articulation control for dynamic manipulation that utilizes all of the available degrees-of-freedom in the manipulator and mobility systems system to accomplish tasks beyond the individual capability of either the manipulation or locomotion subsystems.
- Develop full-body articulation control to realize *legipulation* – the use of legs for opportunistic execution of dynamic manipulation tasks.
- Demonstrate high-DOF dynamic manipulation tasks such as digging, door traversal, rubble clearing/inspection, route clearance, pushing, dragging, throwing, lifting, etc.

Specific 2012 objectives:

- Transition the BigDog robot and the associated manipulator arm from laboratory testing to field operation.
- Integration of a perception subsystem to support autonomous execution of dynamic manipulation tasks.
- Perform *Dynamic Heavy Lift* and *Throw* behaviors in the field.
- Develop new behaviors and behavioral transition strategies to accomplish tasks such as *Door Transversal*, *Rubble Pile Inspection*, etc.

Background:

This task focuses on the development of control strategies for integrated manipulation and mobility systems that take advantage of the increased performance (range of motion, strength, speed, etc.) afforded by utilizing all of the available degrees-of-freedom in a coordinated manner. Anyone who has programmed a 6- or 7-DOF manipulator knows how limited their workspaces can be and how capabilities such as speed and payload can vary over that workspace. Observation of human manipulation shows how important the use of body degrees-of-freedom (e.g., shoulder, torso, and legs) can be in performing any significant manipulation task. This task

aims to significantly improve robot manipulator effectiveness by leveraging the capabilities of a highly articulated mobile base in a coordinated manner to affect whole-body manipulation.

Currently fielded robots have a range of motion and lifting capacity limitations which both increase the burden on their operator and ultimately reduce their overall effectiveness. Our approach to manipulation within this RCTA task is to use whole-body manipulation to take advantage of the additional degrees-of-freedom in the locomotion platform to effectively increase the workspace, effective strength, and velocity capability of the manipulation system.

Typically, robotic mobile manipulation has meant a manipulator system mounted to a mobile base with the two systems operated relatively independently. Recently, researchers have demonstrated a technique for performing fine manipulation from a platform which is not statically stable (Segway RMP) [Stilman, 2010]. In this example, the techniques developed demonstrate an ability to perform manipulation despite the presence of a mobile platform rather than enhancing their manipulation performance by utilizing their mobile base. There have been a few exceptions to this approach, most notably [Tooru, 2002; Deegan, 2006], which take advantage of the additional degrees-of-freedom provided by the mobility system to improve kinematic conditioning (avoid singularities) and maintain balance while resisting/producing external forces through the manipulation system.

The core difficulty in whole-body manipulation is resolving redundancy – choosing how to position and allocate forces among the entire collection of joints in order to satisfy constraints and ensure robustness to disturbances. In robotics and computer animation, researchers have explored prioritized task space control strategies to turn low-dimensional (i.e., abstract) task space commands into high-dimensional joint level controls [Abe, 2006; Sentis, 2006]. Later, these approaches were extended to employ unilateral Coulomb friction models [Abe, 2007] with mixed, rather than prioritized, task space control. Another approach to whole-body manipulation is to decouple body and limb control. Rather than explicitly form a Jacobian that includes all of the robot's degrees-of-freedom [Shkolnik, 2007], we could treat the base of the robot as if it was floating. Wrenches on the base are then converted into desired foot forces in a least-squares fashion [Nelson, 1998]. The result is that force-control can be realized using a prioritized version of virtual model control [Pratt, 1995].

Research Description:

In the initial two project years of this program, we have created tools to rapidly develop dynamic manipulation behaviors, generated several dynamic full-body behaviors, and tested those behaviors using our simulation tool (Digital Biomechanics). During the final quarter of this project year, we will be demonstrating those behaviors on our BigDog platform. As part of this task and the associated Integration task, BigDog has been modified with the addition of a custom designed hydraulic 7-degree-of-freedom manipulator that includes a gripper. By the end of the 2011 project year, we expect to demonstrate *Dynamic Heavy Lifting* and *Throw* behaviors in a laboratory, using a 17.5kg cinder block.

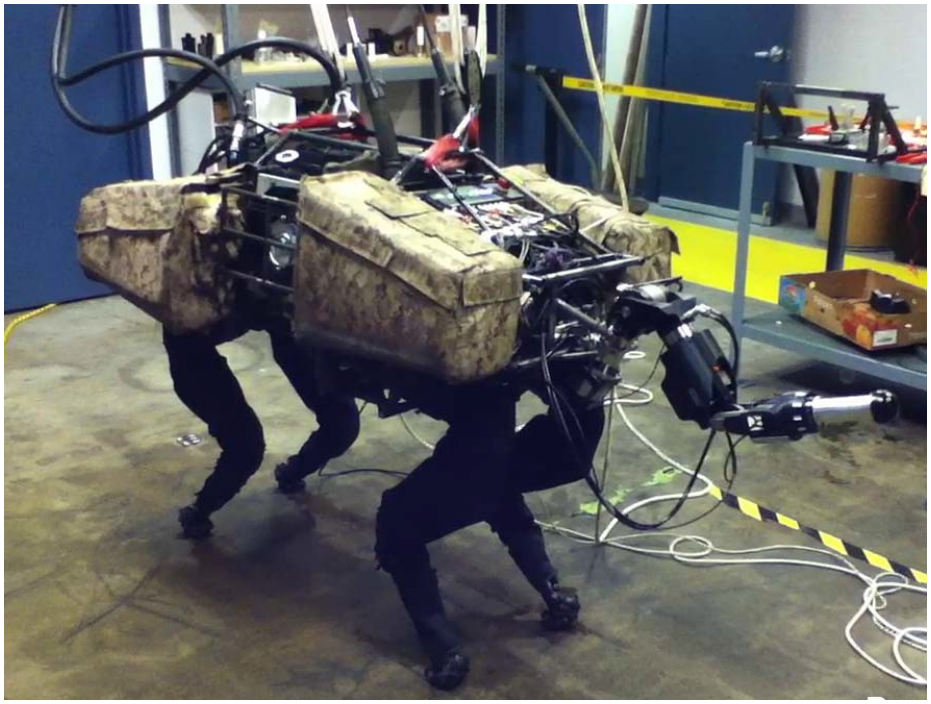

A significant goal for the 2012 project year will be to transition the platform from laboratory operation to field operation and to demonstrate the aforementioned behaviors as well as newly developed behaviors in field relevant settings and scenarios. To accomplish this, we will need to improve our algorithms as well as the physical robot in order to ensure that we respect the various power and strength constraints present in the physical robot. Specifically, we will increase the hydraulic efficiency of our manipulator so that it is suitable for use with engine power and develop hydraulic-system-aware versions of our control system. The latter of these tasks will be the primary focus of our research effort and will seek to leverage results being produced as part of Task M1.

Furthermore, in order to bring the robot out of the laboratory setting (where we are able to make use of motion capture/tracking tools to measure the location of the robot and the environmental objects with which it is interacting), we will also integrate a perception subsystem to sense the environment in close proximity to the robot and provide pose estimates (odometry) to keep track of the robot's position in the environment. The sensor deployment effort will be done in collaboration with Task P4 and P6. Our plan is to initially use an existing set of sensing hardware that was previously developed for the BigDog platform to facilitate rapid deployment of and experimentation with a perception subsystem. Executing this effort in collaboration with the JPL group will allow us to rapidly experimentally evaluate the sensing needs associated with this type of dynamic manipulation task and leverage JPL's prior experience writing software for this sensor subsystem that is currently mounted on BogDog. Also, making use of this existing hardware resource that has been successfully operated in field settings eliminates a significant risk associated with deploying a new hardware sensing suite.

Finally, we will develop new behaviors that will continue our behavior development efforts by focusing on a set of novel mission relevant capabilities such as door traversal where the

manipulation of the door handle and the door itself are required. Our target is to produce fully coordinated motion that explicitly makes use of the manipulator to interact with the door handle, using body degrees-of-freedom to "put the robot's weight" into pulling on the door and seamlessly transitioning to a *legipulation* strategy for using a foot, elbow, or the robot body to hold the door while finally passing through the doorway. This effort will be accomplished through advances in our optimization-based constrained planning tools in addition to transitioning key technical results from 6.1-level research being performed under Task M1. The development and demonstration process will follow the standard we have used throughout the RCTA effort to demonstrate viability in simplified simulations, prove viability through high-fidelity physics-based simulations, and finally demonstrate functionality on the experimental platform.

Metrics:
The primary metrics used to evaluate performance will be directly related to the behaviors being implemented.

Heavy Lift Behavior: The ratio of manipulator weight to the payload that the robot can lift to the top of its workspace. We will compare this metric to the results from laboratory results in 2011.

Throwing Behavior: The distance a 17.5kg cinder block can be thrown. Again, we will compare this metric to the results from laboratory results in 2011.

Door Traversal: The time elapsed during door traversal tasks. Also, the number of types of doors that the robot is able to successfully traverse will be used to measure performance.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Oral presentation and Manuscript, "High Degree-of-Freedom Dynamic Manipulation" SPIE Defense, Security, and Sensing, April 2012, Baltimore, Maryland, United States.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Integrate perception subsystem. Design hydraulic efficiency improvements.
Q2	Integrate hydraulic efficiency improvements. Develop and tune perception software for field use.
Q3	Develop *Door Traversal* behaviors. Perform *Dynamic Heavy Lift* behaviors in field; compare results.
Q4	Perform *Door Traversal* behaviors; record results. Perform *Throw* in field; compare results.

Related Research:
We will take input from 6.1-level research in *Door Traversal* and other dynamic manipulation from UPenn's Task M1: Theory and Principles of Mobile Manipulation.

We will also take input from JPL's perception team to provide support for our perception subsystem. This includes Task P6 (Perception for Missions in Dynamic Environments) for dynamic scene mapping and Task P4 (Perception for Missions in Complex Environments) for terrain classification (e.g., identifying blocks from tufts of grass). We also plan to eventually move away from a large complex sensor head and toward the simpler low power real-time stereo camera system that JPL proposes to develop.

Integration and Assessment Activity:
This research will enable such tasks as door traversal, interaction with piles of clutter. Capstone Assessment activities such as traversal of rough terrain, building entry, path clearing, interior vehicle inspection, interactive inspection, and IED discovery are clear applications of this research.

By the end of the 2012 year, we will be prepared to demonstrate several mission-relevant capabilities in the field as part of assessment activities.

References:
Y. Abe, M. da Silva, and J. Popovic. Multiobjective Control with Frictional Contacts. J. ACM Siggraph/Eurographics Symposium on Computer Animation 2007.

L. Sentis and O. Khatib. A Whole-Body Control Framework for Humanoid Operating in Human Environments. In Proc. of the IEEE International Conference on Robotics and Automation, pp. 2641-2648, Orlando, FL, USA, May 2006.

A. Shkolnik and R. Tedrake. Inverse kinematics for a point-foot quadruped robot with dynamic redundancy resolution. In *Proceedings of the 2007 IEEE International Conference on Robotics and Automation*, April 2007.

G.M. Nelson and R.D. Quinn. Posture control of a cockroach like robot. In IEEE Int. Conf. onRobotics and Automation. 157-162 vol 1. May 1998.

Pratt, J. Virtual Model Control of a Biped Walking Robot , M.Eng. Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1995.

M. Stilman, J.Olson, W. Gloss, Golem Krang: Dynamically Stable Humanoid Robot for Mobile Manipulation, Proceedings of the 2007 IEEE International Conference on Robotics and Automation, 2010.

Tooru et al.,Whole Body Manipulation with a Quadruped Robot. Avoidance of Increased Mobility Singularity, Journal of the Robotics Society of Japan, 2002.

Deegan et al., Designing a Self-Stabilizing Robot For Dynamic Mobile Manipulation, 2006.

M5-2012 – Geometric Mechanics for Multi-modal Highly Articulated Systems

Level: 6.1 Basic Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Howie Choset	CMU	M5-1, M5-2
Additional Investigator	Alfred Rizzi	Boston Dynamics	M5-3

Objective:
The objective of this research is to increase the ability of robotic systems to gain access to and perform surveillance and manipulation tasks in environments that present difficult locomotion challenges. The environments we target are typically inaccessible to a person or conventional robot, with wheels or tracks, or where gaining access through conventional means would risk compromising covertness. While the techniques in this section apply to a wide range of articulated mechanisms, we use snake robots as a platform to demonstrate the ideas cultivated in Task M5. Our efforts fit directly into the "unique mobility" portion of the overall program vision. As highlighted in the Capstone Assessment Vision document, this research will enable robots to navigate "…mobility challenges on the way due to both natural terrain and urban clutter."

In 2012, we will use the task of navigating a notional cluttered alleyway between close buildings as a target demonstration. The alleyway will have bumpy terrain, a small enclosure through which the robot must pass, e.g., a chain link fence, pipe, or an open staircase, and potentially a structure or rubble pile. This supports the "cover the back door" application identified as a component of the capstone assessment. We also believe that culverts and pipe networks will provide unprecedented access to key locations in urban settings. Therefore, we will use culverts and pipes, both straight segments and networks, as an additional notional challenge to demonstrate capabilities derived from this task. The objectives of this task include:

- Prescribe gait evaluation and generation for diverse serially actuated locomotion systems, both biological and man-made. (Subtask 1 and Subtask 2)
- Develop theories for learning and optimizing gaits for a variety of locomotion modes. (Subtask 1)
- Develop theories for feedback and state estimation for highly articulated locomoting systems, and determine ways in which this information can feed into a universal world model. (Subtask 2)
- Generate automated gait transitions in transitional terrain. (Subtask 3)
- Demonstration enhanced capabilities derived from this task's concepts on highly-articulated "snake-like" systems with two target regions in mind: the notional alleyway and culvert/pipe networks. (All Subtasks)

Background:
We believe the cornerstone of new locomotive capabilities lie in cyclic motion controllers which we call gaits. Gaits – repeatable cyclic inputs to the system's shape – provide a library of "default" (e.g., slithering or sidewinding) and "reflex" (e.g., tail sweeps by a falling animal to right itself in midair) actions for a system that can be executed as maneuvers in which the system controller does not need to judge the optimality of the input it applies at each time, trusting instead in the pre-calculated result that the net motion at the end of the cycle displaces the system in the desired manner.

The state of the art in generating gaits for articulated locomoting systems tends to fall into two categories: theoretic, which is rigorous but tends to be limited to low degree-of-freedom systems, and empirical, which is pragmatic for a specific system but are merely rules-of-thumb that do not generalize to a broad class of systems. There is good reason for both approaches, and yet they have their drawbacks. Dynamic models of robots are often so complicated that they can only be analyzed as black boxes: for a given input to the joints, we may be able to simulate the resulting motion of the robot through the world, but finding the "best" inputs often requires trial-and-error ("shooting method") optimization, rather than examining features of the model. *Geometric mechanics*, an area of research combining classical mechanics with *differential geometry* (an extension of vector calculus), offers tools to look inside these black boxes and identify good inputs for the system based on features of the system equations of motion. Unfortunately, from an applications perspective, most of the geometric mechanics literature to date has been developed as an internal discussion between mathematicians and is written in a language not readily accessible to engineers.

Planning, estimation, and control of a field robot needs to be performed from the perspective of that robot's body frame. Typically, a body frame is interpreted to mean any frame rigidly fixed to a point on the robot. Unfortunately, for highly articulated robots, the motion of any fixed point on the robot is rarely representative of the overall motion of the robot. As such it is more difficult to model, tele-operate, and automate articulated robots than it is for simpler wheeled robots.

Even with the best motion controllers, wheeled, legged, snake, and shape-changing robots still fall short of their true potential because they do not exploit all modes of operation when the terrain greatly varies. In other words, we need hybrid controllers to navigate through hybrid terrains. No such locomoting system truly exists that automatically selects among individual controllers to negotiate hybrid terrains. The third subtask seeks to develop a hybrid controls framework to select among gaits and the necessary transitions to facilitate such transitions. Once again, we will expand upon the proposed work to include developing novel means by which feedback can be readily incorporated to select the appropriate modes.

Research Description:
The approach taken in this task builds upon our previous work in geometric mechanics, not specifically to prescribe controllers for our high-DOF systems but rather to confidently make broad statements about the systems we care to design. This in turn guides our intuition in

understanding how both natural and man-made systems locomote as well as the design of man-made high-DOF systems. We are developing an alternate introduction to this material that draws as much as possible on engineering background knowledge, bringing-in more advanced mathematical terminology only where simpler explanations do not exist. We will take two parallel efforts of theoretical development in this task. The first line of development will focus on deriving solutions from fundamental mechanics. While such results will be quite informative, they will be limited to low-DOF systems. Therefore, the second line of research will use lessons learned from the low-DOF results and seek to generalize and/or approximate them on the high-DOF system. For both lines of research, experiments will feedback and drive theoretical pursuits.

Our new treatment of the geometric mechanics framework is also useful for describing the motion of biological systems. From a mathematical standpoint, many problems in biological and robotic locomotion are identical: the animal or robot deforms its body to propel itself through the world. In the biological context, the problem at hand is often not so much the determination of an optimal input given a cost/result metric, but rather the dual problem of identifying the cost/result metric given a presumed-optimal motion used by the organism. This problem, however, must often still be addressed though a trial-and-error approach. By bringing geometric mechanics to bear (aided by our simplified treatment), opening the black box and relating the optimization function to the system's equations of motion becomes feasible.

Typically, the systems we are considering are way too difficult to model as they have many degrees-of-freedom and interact with the environment in a seemingly stochastic manner. To address this complexity, we seek to "reduce" the locomotive challenge by approximating the locomoting system's relationship between external motion and internal shape variables. With such a relationship in-hand, then cyclic controllers, such as gaits, will start to behave like wheeled systems, and thus we can bring to bear many of the results already developed for autonomous wheeled systems. This also allows us to plan in the space of controllers for locomoting systems, thereby reducing the computational burden of planning in a high-dimensional configuration space. This also allows us to prescribe estimators, necessary for feedback, for our locomoting system. We call this approach locomotive reduction.

With well-approximated controllers in-hand, we can then prescribe a hybrid control framework in which we can transition from one controller to another as the environment which surrounds the robot varies.

Metrics

Since the ultimate goal of this task is autonomous operation, it only makes sense to measure our overall success with respect to this capability. In terms of our current, fieldable, high-degree-of-freedom system, the snake robot, the degree of autonomy achievable can be measured gradually by continuously removing human-in-the-loop dependency. Enhanced feedback and an autonomous control architecture are the objectives of Subtasks 2 and 3. Our overall progress can thus be charted through the successes of each of these subtasks.

Capability/Metric	Units	Subtasks
Steepness of slope able to traverse (various gaits).	degrees	5-1
Diameter of pipes/poles able to autonomously traverse.	range (in inches)	5-2
Accurate world map reconstruction through locomotion (below ground).	% error	5-2
Ability to reliably navigate a T-junction in pipes (various sizes and configurations).	pass/fail time to completion (sec)	5-2 and 5-3

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:

- Teaching the Geometric Mechanics Course remotely to ARL personnel.
- Development of tactile skins jointly with ARL for snake robots to improve gait performance.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Document gait design techniques for highly articulated robots (published research paper).
Q2	Initial mapping of a pipe-like environment using locomotive feedback.
Q3	Locomotion in pipe or culvert environment with transition from a main line to an off-shoot or lateral.
Q4	Locomotion for guard the back door task, overcoming at least two mobility challenges.

Related Research:

Related Research to Other Parts of the RTCA:

- We are currently working with Subtask P2-3 to incorporate a 3D sensor suite into our mobile platform. This collaboration is synergistic in that 3D knowledge of the surrounding environment increases our ability to automatically generate locomotion behaviors, and our high degree-of-freedom system provides unique capabilities in terms of the kinematic motions that are potentially necessary for 3D reconstruction of diverse environments.
- This task is related to the ongoing work in M6-3, as we have actively been pursuing work on autonomous-dynamic transitions for some time. We seek to collaborate with the investigators of M6-3 to create a meta-level theory on transitions between steady-state behaviors, such as gaits, in our efforts in Subtask M5-3.
- In collaboration with task M1, we are investigating the necessary symbolic techniques that allow a planner to assemble the controllers efficiently while simultaneously respecting their underlying dynamics and keeping the computation tractable.
- We explore using accelerometers, currently incorporated in our platform, to ultimately help contribute to the world model is directly related the work of M1-2, utilizing high-

bandwidth/low-cost accelerometers to gain knowledge of the environment; we believe that we will ultimately be able to use our platform as a distributed sensor network capable of topological reconstruction.

- In addition, we have also opened dialogue with our colleagues at ARL to explore a project involving the design of force sensing skins. The skins would be designed by ARL and fielded on our platform. This collaboration would enable the exploration of force control strategies that we believe will enhance the ability of our mechanisms to sense and traverse complex terrain.

Related Research of Others Outside the RCTA:
The proposed work builds on the body of locomotion literature which uses geometric mechanics to separate internal shape changes from the external motions they produce as well as external shape changes from reconfiguration. The application of geometric mechanics to locomotion, pioneered by Shapere and Wilczek and further developed by Murray and Sastry, Krishnaprasad and Tsakiris, and Kelly and Murray, provides a powerful mathematical framework for analyzing locomotion. A key product of this work is the reconstruction equation for locomoting systems, which relates internal shape changes to body velocity for locomoting systems. The reconfigurable latching mechanisms of Nilsson and Shen are examples of latching mechanisms that enable multiple singly connected component robots to join.

Past researchers have derived the reconstruction equation for a diverse set of systems, including those that locomote across land, swim in a variety of fluid regimes, or float freely in space. Additionally, a number of algorithms using the reconstruction equation to plan motion trajectories for these systems have been proposed, with emphases on either implementing asymptotically stable controllers or designing gaits, i.e., cyclic changes in shape, which produce desired net changes in position and/or orientation.

Related Research and How it Pertains to the RCTA:
Choset, along with Co-PI Schneider at CMU, is developing learning techniques on the high-DOF snake robot under ARL support (Biologically-Inspired Modeling and Learning of Gaits for Snake Robots, Proposal Number 54341-CI). Building upon existing gaits, which have mainly been empirically derived, this work uses a learning approach in optimizing parameters of the current gait model to improve and discover new gaits. This work develops a novel learning algorithm that allows for learning to occur on the mechanism itself by finding an optimal solution with minimal trials.

Choset, along with Co-PI Hossoi at MIT, is developing analytic tools which, although are currently limited to low-DOF systems, provide deep insight on how to design optimal gaits for locomoting systems, starting with kinematic systems and moving on to dynamic systems, both with nonholonomic constraints. We believe that this new framework will enable us to define new metrics to evaluate optimal agility in locomoting systems. This may also prove useful when seeking to vary the configuration of the robot to improve locomotion efficiency.

Integration and Assessment Activity:
In 2012, we plan to transition our research in order to contribute to the capstone vision "cover the back door" in a number of ways. Already, in Fort Indiantown Gap, we demonstrated a snake robot climbing a pole to gain a vantage point for covering the back door. We will also demonstrate a snake robot breeching a chain link fence. These and other capabilities that we are developing address elements 3, 4, and 5 of the capstone vision. Our work in Subtask 2, in which we aim to use proprioceptive sensing, as well as collaborating with Subtask P2-3 in order to incorporate a vision system into our platform, directly addresses needs sought in item 3 of the capstone vision, "Understand world through perception and world model."

We would like to note that due to the unique forms of mobility of which snake robots are capable, we can achieve a vantage point from the top of a pole or mapping a pipe or culvert network by first entering the network through a storm drain. In fact, to highlight the efficacy of snake robots, we are going to navigate a network of culverts and pipes; this will include making meaningful transitions at key junctions, which is a direct application of Subtask 3, as well as relying heavily on previous and continuing work in Subtasks 1 and 2. This task is aimed at displaying a unique capability which is not currently listed in item 5 of the capstone vision but potentially should be: "Moving covertly to the back door through a system of difficult-to-enter underground passageways." We will use the sewer pipe network in the Principal Investigator's home as a test site and will collaborate with members of the RCTA in order to properly document this deployment.

Our work in Subtasks 1, 2, and 3 currently addresses capabilities desired in item 4 of the capstone vision. In order for our robot to move in its environment, we must first determine good ways to represent the robot in its environment and basic forms of control which are addressed in Subtasks 1 and 2. In order to then truly move autonomously, we need a framework that links a series of controls for potentially varying tasks. This capability is addressed in Subtask 3. Although we are not currently working on the "tracking movers" portion of item 4 in the capstone vision, it should be noted that we are actively adding vision capabilities to our platform under the RCTA and, thus, addressing this desired capability as a possible avenue for future CTA collaboration.

Our work in this task has a large potential for future innovation in item 5 of the capstone vision, "Move through mobility challenges on the way due to both natural and urban clutter." This item is again addressed in Subtasks 1, 2, and 3. In the coming year, we will be doing work in Subtasks 1 and 2 which is aimed at robust gait generation. We will be doing work in Subtasks 2 and 3 which relates to adaptive behavior generation through the use of feedback control. We will also be doing work in Subtask 2 that is related to learning. All of the capabilities that we have just discussed will be evaluated and assessed at the Fort Indiantown Gap site.

Subtask 1: Geometric Methods for Locomotion (CMU)

Objective and Benefits:
The first subtask stems from our basic research in geometric mechanics, a new branch of mechanics that combines differential geometry and applied mechanics. It should be stressed that the study of this field will not immediately or directly impact high degree-of-freedom systems. The pursuit of geometric mechanics for the purposes of advancing locomotion will bear considerable fruit, however. The results of this research will allow us to confidently make broad statements about the systems we care to design. This in turn guides our intuition in understanding how both natural and man-made systems locomote as well as how to design man-made high-DOF systems.

State of the Art:
Snake robot locomotion in rough terrain has received surprisingly little research interest. Most snake robot literature has focused on the "slithering" motion over smooth terrain, with only a few other investigations into sidewinding [4,5], also focused on flat ground. We have developed a new gait, based on sidewinding, which allows the robot to also rotate in place [4]; the translational sidewinding, combined with the rotational gait, gives the snake robot full controllability in the plane with this locomotion mode. Some effort has been made towards using obstacles in the environment as push points [6], but this kind of locomotion is substantially different from the disturbance-rejecting motions we are considering. For bumpy terrain with features smaller than the snake robot, sidewinding offers a naturally robust means of locomotion. Previous work [3] has already established sidewinding's power efficiency, and it applies minimal shear forces on the terrain [1], leaving it additionally robust to differing frictional coefficients. Finally, with the notable exception of [9], few prior efforts considered affect of skin on a mechanism, and in cooperation with ARL personnel, we will try different skin morphologies, most notably skins which include tactile sensors.

Key Barriers and Approach:
Our fundamental geometric mechanics approach informs the gait design in two ways. First, by applying our low-DOF approaches to modal models of the snake robots (i.e., by restricting the full 16-dof snake to sine-wave-based sets of joint angles), we can get a good feeling for the forms that useful gaits will take. Second, understanding the differential geometric aspects of snake locomotion directly helps us to understand the three-dimensional shapes assumed by our snake robots. Three-dimensional kinematics are a traditionally "hard" problem, but the geometric methods we employ provide structure to mitigate this complexity.

The geometric mechanics work will inform the control designer of a "good" gait, and then the next step is to make it optimal. Gait optimization uses machine learning techniques in order to tune gait parameters, which are a product of previous work in gait generation, such as decreasing the amplitude of the sidewinding gait in order to traverse steeper slopes.

M5-1.1: Backbone Curvature-based Gaits
Most of the successful gaits that have been developed so far are built around the framework of the compound serpenoid curve. Recall that the serpenoid curve, originally designed solely for

planar motion, is a biologically inspired approach conceived by Hirose in his pioneering work in snake robots. We relate this work to our fundamental geometric analysis which allows joint angles to be prescribed in a manner that reduces the dimensionality of the controls of the full system while still making good use of whole-body motion. An additional advantage is that the gait parameters provide intuitive "knobs" for control that can be used to easily adjust the gait on the fly.

M5-1.2: Optimal Wave Fitting

While the backbone curvature-based gaits are powerful, they are disconnected from the true shape of the snake robot in the world. Outside of environments that possess symmetry (flat ground, pipes, poles, etc.), finding backbone curvature-based gait parameters that correspond to true world shapes is difficult. Noting the geometry of the snake's underlying shapes, we are applying tools from geometric mechanics to make the process of going from shapes to backbone curves automatic. The end product will be an ability to create gaits that have their basis in the true backbone shape (allowing us to prescribe shapes that are relevant to any given environment) but their controls based on backbone curvature (so that they can be smoothly and seamlessly adjusted on the fly).

M5-1.3: Tail Modeling

Our previous work with Goldman at Georgia Tech in modeling the sand-swimmer showed that a complicated system – a locomoting device maneuvering through a granular medium – can be quite accurately and succinctly approximated with a simple kinematic system that uses the same geometric mechanics described above. With this in-hand, we will extend and apply these tools to modeling and evaluating highly articulated appendages, such as the tails on geckos. We have already initiated some dialogue with Full's group to see if this approach is feasible.

Metrics for Evaluation:

Since this area deals primarily with understanding the fundamental geometry of gaits and gait design, metrics are somewhat difficult to define. However, we intend to demonstrate this work by improving on a number of gaits to make them more stable over a range of slopes.

Capability/Metric	Units	Subtasks
Steepness of slope able to traverse (various gaits).	degrees	5-1

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Develop gait optimization techniques to improve gaits derived from geometric mechanics.
Q2	Gait design for tails.
Q3	Gait generation on uniformly "bumpy" terrain and various ranges of pipe diameters and configurations.
Q4	Optimal gaits on mixed bumpy and inclined terrains.

Subtask 2: Locomotive Reduction (CMU)

Objective and Benefits:
For performing meaningful tasks with highly articulated locomoting robots (like snake robots), there are a myriad of challenges. Two of them can be summarized as follows:

1) The connection between the controls and the actual motions of the robot are generally considered to be complex and difficult to model.
2) "Closing the loop" on the robot in a meaningful way is also considered to be complex due to the many degrees-of-freedom of the robot.

We believe that if we can reduce the model of the highly articulated system to one of a "simple" wheel-like vehicle, then we can bring to bear all of the benefits of world interaction (i.e., process) and sensor models developed for wheeled vehicles.

State of the Art:
In order for a robot to estimate its position in the world, assumptions need to be made about how the robot interacts with its surroundings as well as how the state of the robot and characteristics of the environment affect those interactions. Given a well-understood system, there are a myriad of techniques that can be applied to estimate these interactions. For example, recent research [8] has used an extended Kalman filter (EKF) to estimate the slip parameters of a wheeled robot on various types of terrain. It is important to realize that this technique is inherently empirical and serves simply to adjust the parameters of a simple model rather than generate estimates based on a full understanding of wheel-ground interaction. For highly articulated systems, a key stumbling point has been that finding "simple" wheel-like models for world interaction has remained elusive.

Key Barriers and Approach:
Our technical approach to extending the unique mobility of snake robots is one that revolves around simplification and intuition. Both simplification and intuition provides benefits to the Soldiers in the field who ultimately will be using and interacting with these robots. The most significant limitation of state-of-the-art snake robots is their applicability to only simple environments. Changes to their shape that require more complex modes and variables have been elusive because the proper modes of control to add are unclear, and the way in which these modes should respond to feedback from the environment is also unclear. The recent developments with both the basic representation and state estimation of snake robots will allow us to address both these deficiencies simultaneously.

M5-2.1: Virtual Chassis for Gait Development
The planning, estimation, and control of a field robot needs to be performed from the perspective of that robot's body frame. Typically a body frame is interpreted to mean any frame rigidly fixed to a point on the robot. Unfortunately, for highly articulated robots, the motion of any fixed point on the robot is rarely representative of the overall motion of the robot. However, a body frame is by no means limited to a fixed point on the robot. In fact, any shape-dependent transform to a body frame is also a valid body frame. This is a result which we have shown to be true in the

geometric mechanics framework. This includes transforms based on shape-dependent features like center of mass, mean orientation, or principal moments of inertia.

In this task, we build upon last year's work that developed the *virtual chassis* for the robot [14]. Whereas for simple systems, tools from geometric mechanics cleanly separate the internal shape changes from external motion in the world, the virtual chassis seeks to optimally approximate this relationship. This is accomplished by determining the "optimal" body coordinate frame which measures the most meaningful motion of the system given the available feedback, which in this case is "joint angles." In doing so, we now have struck a balance between the benefits of theory and the realities of practical implementation on a multi-degree-of-freedom system.

M5-2.2: State Estimation and Feedback from Locomotion

We then apply these results to obtaining proper feedback, with respect to the external environment, for the highly articulated locomoting robot. This issue can be addressed by realizing that the same modal controls that we use to simplify the control of the robots can also be used to simplify state estimation [15]. In short, we are providing tools that allow us to treat snake robots more like wheeled robots, allowing us to leverage the wealth of prior work that has been developed and applied to that domain.

We will continue to look at our currently available proprioceptive data as well as generally increasing other sensing capabilities, such as our collaboration with Subtask P2-3, to characterize the snake's motion in the world. We will also develop a deeper understanding as to which coordinate frame best represents the internal configuration of the robot for both an operator and state estimator. This work will build upon current efforts in deriving optimal coordinate frames for locomoting systems.

Metrics for Evaluation:

Capability/Metric	**Units**	**Subtasks**
Diameter of pipes/poles able to autonomously traverse.	range (in inches)	5-2
Accurate world map reconstruction through locomotion (below ground).	% error	5-2
Ability to reliably navigate a T-junction in pipes (various sizes and configurations).	pass/fail time to completion (sec)	5-2 and 5-3

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Derive gait-based feedback controllers for a variety of gaits operating in a culvert, pipe, or otherwise structured space.
Q2	Incorporate accelerometer and other proprioceptive sensor data into estimation and controls.
Q3	Incorporate sensing data from 3D sensor developed in Subtask P2-3
Q4	Robust autonomous locomotion and estimation of straight pipes/poles, assisted tele-operation of T-junctions and bends.

Subtask 3: Hybrid Controls (Boston Dynamics)

Objective and Benefits:
With a mechanical platform that can physically change modes, the system must be able to recognize and adapt to different terrains. Currently, systems capable of complex reconfiguration and different modes of locomotion for crawling, climbing stairs, or climbing poles fall short of their true potential due to a lack of intelligent motion planning and control. This is because a single gait, as defined today, is limited to a single terrain. We, therefore, develop a hybrid control framework to invoke and blend the behavior produced by multiple locomotion strategies to traverse multiple terrains. This framework will enable nontrivial-autonomous task execution. We use a hybrid control framework that has shown to be effective in composing controllers for wheeled systems with dynamics and nonholonomic constraints. In this framework, simple feedback controllers are defined over small, simply shaped regions of free space. A planner then sequences them together so that the system can achieve a desired global behavior such as moving from a start to a goal configuration. In collaboration with Tasks M1 and M6, we will investigate the symbolic techniques that allow a planner to assemble the controllers efficiently while simultaneously respecting their underlying dynamics and keeping the computation tractable.

State of the Art:
Sequential composition is a technique for developing hybrid controller systems where individual controllers are defined over limited subsets of the robot's state space such that the stable goal state of one controller lies within the domain of attraction for another local controller. Sequencing such controllers together, one can create a nearly global policy that directs the robot from any start configuration to a goal while respecting the constraints present in the robot's work space. In this subtask, we build upon prior work in sequential composition to handle high-degree-of-freedom systems. We, once again, reduce the problem by considering gaits as controllers which casts the sequential composition as a matter of planning from one gait controller to another.

Key Barriers and Approach:
Transferring motion from one gait to another, however, is a subtle yet important problem. When changing gaits, the end state of one gait is the same as the start state of the next, which generically will not be the case. This problem has been largely overlooked in the locomotion literature, with the notable exception of Haynes and Rizzi [13] with climbing locomotion, and is now being addressed for legged systems in Task M6. We are interested in gait transitions for highly articulated locomotors. In this area, only central pattern generators (CPGs), as discussed in [6] and the references therein, implement some transition techniques but require much empirical hand-tuning to provide natural-looking motion, none of which has been demonstrated on real robots in hybrid terrains.

Metrics for Evaluation:
Transitioning between gaits includes a variety of possibilities, such as traversing a T-junction in a pipe network, locomoting up to a tree and climbing it, or even dynamically changing the robot morphology to facilitate switching between snake-like undulation to tread-like rolling. For a set

of such tasks, we will measure the time it takes to successfully execute a transition from start to finish.

Capability / Metric	Units	Subtask
Ability to reliably navigate a T-junction in pipes (various sizes and configurations).	pass/fail time to completion (sec)	5-2 and 5-3

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Provide practical examples to facilitate the development of gait-like motions required to transition between two disparate gaits.
Q2	Define specific constraints and how they impact motion during transitions.
Q3	Develop framework for evaluation of transition behaviors.
Q4	Support demonstrations of transitions in a culvert environment.

References:

[1] W. Mosauer, "A Note on the Sidewinding Locomotion of Snakes," The American Naturalist, vol. 64, no. 691, pp. 179–183, March 1930.

[2] R. L. Hatton and H. Choset. Sidewinding on Slopes. In the Proceedings of the IEEE International Conference on Robotics and Automation, May 2010.

[3] M. Tesch, K. Lipkin, I. Brown, R. Hatton, A. Peck, J. Rembisz, and H. Choset. Parameterized and Scripted Gaits for Modular Snake Robots. Advanced Robotics, 23(9):1131–1158, 2009.

[4] J. Burdick, J. Radford, and G. Chirikjian, "A "Sidewinding" Loco- motion Gait for Hyper-redundant Robots," Robotics and Automation, vol. 3, pp. 101–106, 1993.

[5] J. Gonzalez-Gomez, H. Zhang, and E. Boemo, "Locomotion Principles of 1D Topology Pitch and Pitch-Yaw-Connecting Modular Robots," in Bioinspiration and Robotics: Walking and Climbing Robots. Ad- vanced Robotics Systems International and I-Tech Education and Publishing, 2007.

[6] Transeth, A.A.; Leine, R.I.; Glocker, C.; Pettersen, K.Y.; Liljeback, P.; , "Snake Robot Obstacle-Aided Locomotion: Modeling, Simulations, and Experiments," *Robotics, IEEE Transactions on* , vol.24, no.1, pp.88-104, Feb. 2008.

[7] Ijspeert, A.J.: Central Pattern Generators for Locomotion Control in Animals and Robotics. Neural Networks **21**, 642–653 (2008).

[8] Rogers-Marcovitz, F., Kelly, A. On-line Mobile Robot Model Identification using Integrated Perturbative Dynamics, ISER 2010.

[9] D. Goldman and D. Hu, "Wiggling through the world," American Scientist, vol. 98, pp. 314–393, July 2010.

[10] Yim, M., Shirmohammadi, B., Sastra, J., Park, M., Dugan, M. and Taylor, C.J., "Robustness and self-repair in modular robots," Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), San Diego CA, Oct. 27-Nov. 2, 2007.

[11] Yim, M., Shirmohammadi, B., Sastra, J., Park, M., Dugan, M. and Taylor, C.J., "Towards Robotic Self-reassembly After Explosion," Video Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), San Diego CA, Oct. 27-Nov. 2, 2007.

[12] Shirmohammadi, B., Taylor, C.J., Yim, M., Sastra, J., and Park, M., "Using Smart Cameras to Localize Self-Assembling Modular Robots," Distributed Smart Cameras, ICDSC'07. First ACM/IEEE Intl. Conf. on, pp. 76- 80, 2007.

[13] Haynes & Rizzi "Gait Regulation and Feedback on a Robotic Climbing Hexapod: Robotics Systems and Science (RSS) August, 2006.

[14] Rollinson, D. & Choset, H., "Virtual Chassis for Snake Robots", Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), September 2011.

[15] Rollinson, D., Buchan, A., & Choset, H., "Virtual Chassis for Snake Robots", Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), September 2011.

M6-2012 - Principles of Locomotion Mechanics

Level: 6.1 Basic Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Dan Koditschek	UPenn	M6-3, M6-4
Additional Investigators	Robert Full	UC Berkeley	M6-1
	Mark Cutkosky	Stanford	M6-2→ M9-2

Note: Subtask M6-2 has been moved to M9-2.

Objective:
The objectives of this task are:

- Create principles adapting gaits to variable, non-steady terrain with focus on transitional behaviors.
- Gait generation for diverse systems with focus on seamless and increasingly automated gait transitions.
- Mixed horizontal and vertical locomotion.
- Focus on "legged" systems with actively tunable and passive compliance in relation to terrain and substrate properties.
- Explore the role of appendage inertia in behaviors and behavior transitions.

Background:
In contrast to existing robotic vehicles, animals operate almost anywhere. Robotic vehicles lack the intelligent and general control systems that can transform the vehicles' distinctive physical endowments into appropriately coordinated ground reaction forces required to traverse complex, dynamical terrain. On the other hand, animals fluidly combine aspects of tuned feed-forward, reactive, and deliberative online control mechanisms to negotiate highly varied and ever-changing environments. They use a combination of neural and mechanical (tuned musculoskeletal material properties) designs to achieve their robust, agile passage through the environment.

Research Description:
We investigate incorporating a similar combination of all these control styles into a new generation of bio-inspired, nimble-legged running machines. Furthermore, we develop the mathematical insight into their underlying principles to translate them to non-biological mobility platforms, including vehicles with traditional wheels and tracks, articulated suspensions such as those on the Crusher vehicle and the FCS MULE, and novel advanced designs. This allows robotic vehicles to perform ably in the face of the inevitable, severe, unanticipated perturbations

they suffer in natural terrain. We have previously developed the broadly cited bio-inspired controls abstraction stack that uses "programmed" attractor-repellor basins assembled sequentially or in parallel at each layer of the architecture [1,2,3]. While the theoretical foundations of these ideas and their connections to biology are well advanced, they have been developed largely around steady-state gaits and behaviors. By close observation and understanding of highly adaptive mammalian, lizard, and insect models of mobility, we extend our framework to more variable, non-steady conditions where rapid adaptation to the environment is required. Our current animal studies on highly perturbed, non-steady modes of terrain engagement have led us to hypothesize how increasingly high levels of sensory feedback may be recruited to adapt to large, unanticipated perturbations.

We extend our current work on translating these insights into non-steady controls, transitional behaviors, and mechanisms of adaptation, permitting machines to change their gaits in response to the interaction of their sensorimotor apparatus with the environment. These algorithms aim to exploit both the purely mechanical components of the system below them and the high-level planners above them. The passive compliant multi-modal mechanisms being developed under Task M8 (Dynamic Multi-modal and Reconfigurable Mechanisms) provide a preflex control system that offers passive self-stabilization of steady state gaits and robust disturbance rejection mechanisms. We adapt these to new environments with both offline tuning and online learning of feed-forward controllers. We combine these controllers with deliberative planners, leveraging work performed by Tasks I2: Data Mapping for Inference and Focus, I4: Generating Adaptive Tactical Behaviors, and M7: Learning Terrain Interactions to predict the vehicle response to the terrain.

Metrics:

- Time required to negotiate complex terrain (e.g., in the early quarters of the project, simply affecting transitions between level and vertical terrain).
- Quality of the locomotion through terrain (e.g., how many falls or stumbles) as a function of terrain complexity (e.g., surface variability or stability or slipperiness).
- Amount of human driver attention (e.g., in operator command rate) required to negotiate the terrain.
- Amount of prior information required about the nature of terrain negotiated.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
These project lines have been extremely collaboratively focused during 2011. There is a very close collaboration between UPenn and Berkeley concerning appendage inertia that has entailed a weekly telecon, culminating in a week-long visit by Full's students Chang and Libby. We anticipate a similar (or perhaps even heightened) degree of collaboration between the two labs in 2012. UPenn also hosted a week-long visit from Jonathan Clark's group at FSU/FAMU for purposes of consulting in the duplication of an XRL RHex-lookalike [4].

The RCTA M6 project line personnel will be hosting a special session in the 2012 CLAWAR[1] meeting on the role of appendage inertia in maneuverability and agility.

We anticipate that Jason Pusey, one of the ARL technical staff members under the direction of Harris Edge, will be spending roughly one week a month in the UPenn GRASP lab for much of 2012, dividing effort between M6-3 and M8-4.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Progress report on understanding control signals to leg muscles of animals on level, sloped, and vertical surfaces. (Subtask 1)
Q2	Progress report on tests of RHex dynamic transitions using appendage inertia. (Subtask 3)
Q3	Progress report on identification of complex terrains and substrates of greatest immediate Army interest and the nature of available models. (Subtask 4)
Q4	Comparison of animal and RHex control of locomotion on level, sloped, and transitions over vertical surfaces. (Subtask 1 and Subtask 3)

Related Research:

- Task M7: Learning Terrain Interactions (inputs to and outputs from)
- Task M8: Dynamic Multi-modal and Reconfigurable Mechanisms (inputs to and outputs from)
- Task I1: Framework for Intelligence (inputs to and outputs from)
- Task I2: Data Mapping for Inference and Focus (inputs to and outputs from)
- Task I4: Generating Adaptive Tactical Behaviors (inputs to and outputs from)

Integration and Assessment Activity:

Results of this research will impact the integration and assessment activity through outputs to efforts I2, M7, and M8. Specifically, M6 basic research activity will help develop the control of attention and focus algorithms of I2 and inform the adaptive tactical behavior of I4 as well as the terrain interaction algorithms of M7. It will have a major impact on the mechanical designs of M8 and M9.

Subtask 1: Principles of Animal Locomotion Mechanics (UC Berkeley)
Studies of animals operating on level, sloped, and vertical ground are starting to elucidate the roles of the musculoskeletal system in conferring agility and, particularly, in making fast transitions from running to climbing, often in part through appendage-steered ballistic flight.

[1] 15th Intl. Conf. on Climbing & Walking Robots http://www.niu.edu/clawar2012/images/cfp_2012_Final.pdf

Objective and Benefits:
This subtask pursues an investigation of animal agility and, in particular, how animals affect transitions (e.g., between level, inclines, and vertical terrain). It is increasingly apparent that along with changes in the feed-forward and feedback control provided by the neuromuscular system, systematic use of appendages plays a significant role in such capabilities. In particular, we are increasingly focused on the specific investigation of how animals use appendage inertia to increase agility.

Animals possess mobility competence far beyond any presently demonstrated in robots. This work aims to reveal for targeted species in targeted environments the nature of specified agility capabilities and strategies as a means of providing clues to the engineering of comparable capabilities. The huge advantage of animals relative to machines will surely persist well beyond the duration of the RCTA. We expect this effort to continue over a similarly extended period.

State of the Art:
Results of this research led to the first use of tail inertia in righting a falling (wheeled) robot [5].

Key Barriers and Approach:
The Berkeley Polypedal lab uses a mixture of external (e.g., still and video images, force-plate measurements, etc.) and internal (e.g., muscle action potential recordings, inter-neural trunk recordings, etc.) in free running and prepared animals, in combination with mathematically formalized hypotheses to reveal the materials, morphological, and controls origins of animal motor competence. Novel empirical methods [6] and novel data analysis methods [7], combined with the novel hypotheses [8], emerging from the Polypedal lab are widely acknowledged to represent the cutting edge of animal neuromechanics research. The Polypedal lab has a continuing track record of novel discoveries (e.g., one recent example in [9]) that we expect to continue unabated in this project. The new observations about animal transition strategies can be directly transferred to the M8 locomotion experiments. The new discoveries about appendage inertial are inspiring completely new approaches to limb design and control within the parallel M6 and corresponding M8 tasks.

Metrics for Evaluation:
Prime metrics include number of publications in top biology venues as well as the rate of transition of insights into robot design and control. Specifically, we aim for publications in JEB, PNAS, Science, Nature, as well as robot transitions within RCTA. A near-term outcome of the appendage work will be a series of tails implemented on the RHex platform and evaluated for increased agility. We will compare the capabilities of the same legged machine, RHex, with and without a tail.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Study measurements of control signals to leg muscles (EMGs) of animals on level, sloped, and vertical surfaces. Prepare for measurement of dynamics in tail and compliant backbone use in animal locomotion.
Q2	Compare measurements of control signals to leg muscles (EMGs) of animals on level, sloped, and vertical surfaces with alternative template controls hypotheses [10]. Begin measurement of dynamics in tail and compliant backbone use in animal locomotion.
Q3	Initiate paper on animal template control of locomotion on level, sloped, and vertical surfaces. Continue measurement of dynamics in tail and compliant backbone use in animal locomotion.
Q4	Continue paper on animal template control of locomotion on level, sloped, and vertical surfaces. Complete measurement of dynamics in tail and compliant backbone use in animal locomotion.

Subtask 2: Principles of Machine Locomotion Mechanics (Stanford)
This project, subcontracted by UPenn to Stanford University's Mark Cutkosky, has been rebranded to better align the nature of work and intellectual effort with the larger RCTA effort. It will henceforth appear as a subtask within task M9, led by Jonathan Clark at FSU/FAMU, but to assure continuity of funding, UPenn will continue to cover the cost of the subcontract from M6 task lines for the year 2012.

Subtask 3: Integrating Principles of Animal and Machine Locomotion for Unique Robot Mobility (UPenn)
This subtask aims to increase legged robot capability beyond quasi-static transitions [11] to dynamical transitions through complex terrain. We are particularly focused on incorporating the use of appendage inertia to increase agility.

Objective and Benefits:
Mobility over uneven, broken, and poorly structured terrain plays a central role in any of the outdoor RCTA scenarios. The combination of fundamental questions about dynamical transitions and commitment to provably correct methods imply a multi-year effort, likely exceeding the duration of the RCTA.

State of the Art:
We advance our past traditions of comparing details of animal behavior with engineering best practice [12,13,14]. In particular, existing robots specialized for running or climbing are reaching a level of sophistication that transitions between such modes of operation can be studied – theoretically and empirically – and compared with the strategies found in nature.

Key Barriers and Approach:
We seek to generalize such transition planners and execution methods previously articulated for quasi-static settings [15]. Whereas legged machines have achieved increasingly impressive performance in steady state operation on level horizontal and vertical terrain, transitions remain a difficult and largely unexplored domain of operation, and performance on unfamiliar highly irregular surfaces remains to be achieved at all. The method of dynamical task specification and control [1,2] has a long and still vigorously growing track record of success in theory [3] and practice [13]. We hope to show the applicability of a very new theoretical foundation for transitions in general by recourse to Conley Invariant theory [16] recently rendered in algorithmic form via the computational homology methods of [17].

We are working very closely with the Polypedal lab in M6-1 to integrate animal strategies into our designs and algorithms. We seek the foundations of legged locomotion capabilities that can serve as vehicles for Intelligence and Perception tasks within the RCTA, such as Subtask I2-3. This 6.1-level effort will find its tech transition in the 6.2-level effort of Subtask M8-4.

Metrics for Evaluation:
We use measures of efficiency such as specific resistance [18] as well as standard measures of environmental complexity [19] and normalized speed or acceleration. Following on the "front-flip" maneuver previewed in last quarter's progress report, we are now very close to a breakthrough in dynamical maneuvering that would allow transitions across broken ground more than twice the machine hip height. We continue to measure components (e.g., specific resistance, normalized speed, and acceleration) as well as target specific maneuvers (e.g., the front flip just mentioned).

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Continue dynamical self-righting test using appendage inertia on RHex; add closed-loop control on body angle.
Q2	Test dynamic transitions using appendage inertia. Begin to work on a formalism for dynamic transitions.
Q3	Begin analysis on tail effects while running under various gaits. Continue work on a formalism for dynamic transitions using appendage inertia.
Q4	Continue analysis and begin experiments of tail effects and generalized appendage inertia while running under various gaits.

Subtask 4: Environment Mechanics Representation for World Model (UPenn)
This subtask is a more mechanics-focused "rebranded" version of the former 2011 Subtask I1-2. We are hoping to use physics-based and simplicially represented phenomenological models of complex terrain and substrates to increase locomotion performance.

Objective and Benefits:
We would like our robots to develop robust, parsimonious models of the environment enabling better locomotion. Locomotion over broken, unstable, slippery, complex terrain is an essential feature of any of the RCTA outdoor scenarios. The combination of fundamental questions about

environmental mechanics alongside of the open questions about how best to represent and then deploy such models for control purposes implies a multi-year effort, likely exceeding the duration of the RCTA.

Such insights would provide immediate application to IRA activities. At the recent IRA1 session in FTIG, we presented RHex with one of the various gravel piles on site, and it immediately became clear how important substrate mechanics will likely be.

State of the Art:
This is a rebranded version of I1-2 (Simplicial World Model), intended to leverage recent advances in phenomenological modeling and simulation of complex terrain [20] yet with the hope of maintaining the simplicial decomposition of the underlying environment. We propose pursuing phenomenological representations of terrain conditions to circumvent sensitive dependence of a kind explored in [21] to achieve better locomotion over more varied substrates.

A good feeling for the present empirical SoA in the field is conveyed by considering our experience at the end of the IRA1 assessment when the RHex robot attempted to climb a very challenging gravel pile. It was able to climb slopes up to 35 degrees but had trouble on the steeper parts of this loose, changing terrain. We believe that Boston Dynamics has had some significant success in instrumenting their "Rugged" RHex with a "force allocated" climbing gait based upon work on the RiSE controller reported in [22], and we would like to pursue, formalize, and extend this line of effort with particular emphasis on coming to understand how more careful models of the environment's geometric shape and substrate mechanics might be brought to bear.

Key Barriers and Approach:
We seek to generalize such transition planners and execution methods previously articulated for quasi-static settings [15]. Whereas legged machines have achieved increasingly impressive performance in steady state operation on level horizontal and vertical terrain, transitions remain a difficult and largely unexplored domain of operation, and performance on unfamiliar highly irregular surfaces remains to be achieved at all. The method of dynamical task specification and control [1,2] has a long and still vigorously growing track record of success in theory [3] and practice [13].

We are working very closely with the Polypedal lab in M6-1 to integrate animal strategies into our designs and algorithms. We provide empirical legged locomotion capabilities achieved in this project for higher-level activities within the RCTA, such as Subttask I2-3. The 6.1-level effort progressing within this project (i.e., Subtask M6-4) will find its tech transition in the 6.2-level effort of Subtask M8-4. At the same time, we have initiated collaboration with a FSU/FAMU researcher in Task M7, starting to make a connection between their past and continuing terrain identification techniques (proven in various wheeled and tracked vehicles) and seeking to integrate UPenn's disturbance detection with FSU's terrain identification to perform proprioceptive terrain identification on the XRL robot.

Metrics for Evaluation:
Metrics are largely based upon empirical agility experiments over various terrains. We use measures of efficiency such as specific resistance [18] as well as advanced measures of substrate complexity [21] and normalized speed or acceleration.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Initiate literature search and review of terramechanics models as well as physical models of media.
Q2	Conclude literature search and review of terramechanics models as well as physical models of media. Initiate discussion with ARL sponsors concerning complex terrains and substrates of greatest immediate interest.
Q3	Initiate empirical locomotion studies with RHex in substrates and terrains of interest.
Q4	Review data from experiments and compare with terramechanics models to identity dimensions of RHex performance that might be enhanced by better world mechanics model.

References:
[1] D. Koditschek, "Task encoding: toward a scientific paradigm for robot planning and control," *Robotics and autonomous systems*, vol. 9, no. 1-2, pp. 5-39, 1992.
[2] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek, "Sequential Composition of Dynamically Dexterous Robot Behaviors," *The International Journal of Robotics Research*, vol. 18, no. 6, pp. 534-555, 1999.
[3] P. Holmes, R. J. Full, D. E. Koditschek, and J. Guckenheimer, "The dynamics of legged locomotion: Models, analyses, and challenges," *SIAM Review*, vol. 48, no. 2, pp. 207-304, 2006.
[4] K. C. Galloway et al., "X-RHex: A Highly Mobile Hexapedal Robot for Sensorimotor Tasks," University of Pennsylvania School of Engineering and Applied Science, ESE Technical Report, 2010.
[5] E. Chang-Siu, T. Libby, M. Tomizuka, and R. J. Full, "A lizard-inspired active tail enables rapid maneuvers and dynamic stabilization in a terrestrial robot," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, 2011, pp. 1887–1894.
[6] K. Autumn et al., "Evidence for van der Waals adhesion in gecko setae," *Proceedings of the National Academy of Sciences*, vol. 99, no. 19, pp. 12252-12256, 2002.
[7] S. Revzen and J. M. Guckenheimer, "Estimating the phase of synchronized oscillators," *Physical Review E*, vol. 78, no. 5, p. 51907, 2008.
[8] R. Full and D. Koditschek, "Templates and Anchors: Neuromechanical hypotheses of legged locomotion on land," *J. of Experimental Biology*, vol. 202, pp. 3325-3332, 1999.
[9] E. Pennisi, "Racing Crash-Happy Cockroaches," *Science*, vol. 327, no. 5967, p. 776, Feb. 2010.
[10] S. Revzen, D. E. Koditschek, and R. J. Full, "Towards Testable Neuromechanical Control Architectures for Running," in *Motor Control: A Multidisciplinary Perspective*, vol. 629, D. Sternad and M. Turvey, Eds. NY: Springer, 2009, pp. 25-55.
[11] G. C. Haynes, F. R. Cohen, and D. E. Koditschek, "Gait Transitions for Quasi-StaticHexapedal Locomotion on Level Ground," in *Proc. International Symposium of Robotics Research*, Lucerne, Switzerland, 2009, p. (in press).
[12] R. Altendorfer et al., "RHex: A Biologically Inspired Hexapod Runner," *Autonomous Robots*, vol. 11, no. 3, pp. 207-213, 2001.

[13] D. E. Koditschek, R. J. Full, and M. Buehler, "Mechanical aspects of legged locomotion control," *Arthropod Structure and Development*, vol. 33, no. 3, pp. 251-272, 2004.

[14] M. J. Spenko et al., "Biologically inspired climbing with a hexapedal robot," *Journal of Field Robotics*, vol. 25, no. 4-5, pp. 223-242, 2008.

[15] G. Haynes, F. Cohen, and D. Koditschek, "Gait transitions for quasi-static hexapedal locomotion on level ground," in *Robotics Research*, Springer-Verlag, New York, Heidelberg, Berlin: , 2011, pp. 105–121.

[16] C. C. Conley, *Isolated invariant sets and the Morse index*. Amer Mathematical Society, 1978.

[17] W. D. Kalies, K. Mischaikow, and R. C. A. M. VanderVorst, "An algorithmic approach to chain recurrence," *Foundations of Computational Mathematics*, vol. 5, no. 4, pp. 409-449, 2005.

[18] G. C. Haynes and D. E. Koditschek, "On the Comparative Analysis of Locomotory Systems with Vertical Travel," in *ISER'10*, New Delhi & Agra, India, 2010.

[19] U. Saranli, M. Buehler, and D. E. Koditschek, "RHex: A Simple and Highly Mobile Hexapod Robot," *The International Journal of Robotics Research*, vol. 20, no. 7, p. 616, 2001.

[20] A. Slatton et al., "Integrating a Hierarchy of Simulation Tools for Legged Robot Locomotion," in *Proceedings of Workshop on Robot Simulators, IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2008, p. 474.

[21] C. Li, P. B. Umbanhowar, H. Komsuoglu, D. E. Koditschek, and D. I. Goldman, "Sensitive dependence of the motion of a legged robot on granular media," *Proceedings of the National Academy of Sciences*, vol. 106, no. 9, pp. 3029-3034, 2009.

[22] G. C. Haynes, "Gait regulation control techniques for robust legged locomotion," Carnegie Mellon University, Pittsburgh, PA, 2007.

M7-2012 – Learning Terrain Interactions

Level: 6.2 Applied Research

	Name	*Organization*	*Subtask*
Principal Investigator	Emmanuel Collins	FSU	M7-1
Additional Investigator	Jonathan Clark	FSU	M7-1

Objective:
The objective of this research is to enable a wheeled, tracked, or legged autonomous ground vehicle to locomote safely and efficiently on a variety of terrains by developing terrain-dependent control and planning algorithms.

Background:
Development of terrain-dependent control rules for wheeled and tracked vehicles is still at its infancy. It has been approached by quantifying and automating expert off-road driving rules using measures, such as slip and power consumption, with perhaps the most complete result being the rut detection and following algorithm that was developed by the PI and his coworkers. The development of terrain-dependent control rules for legged vehicles is also at its infancy, although limited research has been conducted in the development of specialized gaits for sand. Development of control rules for legged vehicles may rely on a combination of intelligent experimentation and optimization based on dynamic models of these vehicles.

Terrain-dependent motion planning is also needed for safe and efficient maneuvers of ground vehicles in environments with difficult terrains. This type of motion planning can, in principle, be accomplished with the aid of dynamic models since these models capture the interaction of the vehicle with a given terrain type. However, modeling is underdeveloped for most classes of vehicles apart from Ackerman steered vehicles, and the vast majority of planning algorithms are not capable of using dynamic models. This research builds upon recent work by the PI and his coworkers in terrain-dependent modeling of skid-steered vehicles, which is relevant to many wheeled outdoor vehicles and unique mobility robots such as the RHex. It also builds upon planning algorithms designed to plan trajectories based on a model of the vehicle dynamics such as Sampling Based Model Predictive Optimization (SBMPO), developed by the PI and his coworkers. SBMPO is also expected to be useful in developing terrain-dependent gaits for legged vehicles as discussed above.

Research Description:
We push the development of terrain-dependent control and planning algorithms by 1) developing dynamic models of vehicles, 2) using these models with offline optimization to develop control

laws, and 3) developing online planning algorithms that can directly use vehicle models or constraints developed offline by vehicle models. Motivated by both control and planning, we develop kinematic and low-order dynamic models of vehicles and terra-mechanical models of vehicle-terrain interaction and use these to predict trafficability for given terrain types. The terrain models include terrain elevation profiles and elements of soil and vegetation mechanics that impact trafficability. We learn parameters of the terra-mechanics models (both offline and online) from sensor feedback, including slip, motor currents, wheel sinkage, and pressure and moisture sensors. The critical challenges are how to (1) recognize terrain types accurately ahead of the vehicle, (2) model complex vehicle-terrain interactions and make parameters of those models observable, (3) deal with mixed terrain types, (4) learn new locomotion behaviors for new terrain types, and (5) keep the computational load manageable for all of this.

Metrics:
The metrics for this research are: 1) the correspondence between the dynamic models and experimental measurements, 2) the energy savings associated with using these dynamic models, and 3) the improvement in the ability of the robot to traverse difficult mobility challenges.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Collaboration with JPL on Dynamic Modeling
Presently, JPL's visual odometry (VO) system is employed by FAMU/FSU as a key sensor to determine terrain dependent model parameters for skid-steer vehicles. However, the current VO setup is relatively heavy for small-scale legged robots such as the X-RHex Light (XRL), built in conjunction with UPenn.

Current work at JPL involves the development of a VO solution using the Gumstix boards which are in the 100 g weight range. This will allow the extension of the current skid-steered research to platforms like the XRL.

Collaboration with UPenn on Terrain ID
Previous work at FAMU/FSU has led to the development of a terrain identification framework, which has been successfully applied to different wheeled vehicles. Parallel to this work, UPenn has concentrated on the development of software contact-event sensors for RHex-type platforms, which have proven to be effective at detecting expected ground contact, missing ground contact, and unexpected contact.

This collaboration seeks to integrate UPenn's disturbance detection with FSU's terrain identification to perform proprioceptive terrain identification on the XRL robot.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Demonstration of online learning of reduced order dynamic model for a skid-steered vehicle; there must be a close correspondence between the model and experimental data. Simulation of longitudinal momentum-based planning for skid-steered vehicles that improves the ability of the robot to traverse difficult mobility challenges such as steep hills, mud patches, and high, stiff vegetation.
Q2	Experimental demonstration of longitudinal momentum-based planning on different terrains that improves the ability of the robot to traverse difficult mobility challenges. Development and experimental verification of an empirical power model for a walking gait of the XRL (X-RHex Lite) robot.
Q3	Experimental verification of 3D model for skid-steered vehicles on outdoor surfaces; there is expected to be a close correspondence between the model and experimental data. Demonstration of proprioceptive terrain classification for the XRL robot (in a joint effort with UPenn); this will allow the robot to sense difficult surfaces and change its control system for that surface.
Q4	Outdoor online implementation of energy efficient path planning on skid-steered vehicles; these results are expected to demonstrate substantial energy savings associated with minimum energy planning over minimum distance planning. Demonstration of energy efficient motion planning for the XRL robot; these results are expected to demonstrate substantial energy savings associated with minimum energy planning over minimum distance planning.

Related Research:
This research requires identification of the surface characteristics and slopes of terrain. Hence, it is highly related to all of the research in terrain classification, in particular Task P1: Exploiting Novel Sensor Phenomenology, Task P4: Perception for Missions in Complex Environments, and Subtask I6-4: Heightened Proprioceptive Awareness and Safe Exploration. A part of this task is *learning using dynamic models*. It is therefore related to Intelligence tasks in learning but appears to be distinct since the Intelligence tasks do not appear to be based on using dynamic models. It is expected that the use of dynamic models makes the learning much easier. This subtask relies on receiving information from the world model on the surface and slopes of the terrain.

Integration and Assessment Activity:
Using the PennBot, we will demonstrate the use of momentum-based planning to overcome obstacles that cannot be traversed quasi-statically, i.e., if the vehicles are moving too slowly, they will actually be immobilized due to the inability of the vehicle motors to provide sufficient torque. In particular, we will demonstrate steep hill climbing and traversal of artificial terrains that emulate high vegetation and/or mud. This demonstration is part of Integrated Research Assessment 12: Manipulation and Locomotion in Challenging 3D Terrain.

Subtask 1: Planning with Dynamic Models (FSU)
This subtask involves the development of terrain-dependent models for a skid-steered vehicle and a RHex-type robot (particularly the X-RHex Lite) and demonstration of their use in planning. The particular focus during this fiscal year is demonstration of online learning of skid-steered vehicle models, demonstration of momentum-based motion planning for skid-steered vehicles, development of a power model for the XRL, and demonstration of energy efficient motion planning for both a skid-steered vehicle and the XRL.

Objective and Benefits:
This task seeks to develop computationally efficient physics-based vehicle-terrain models for wheeled, tracked, and legged robots. The developed models will be combined with learning and integrated with motion planners to enable traversability prediction, adaptation to different terrain surfaces, and exploitation of vehicle dynamics to overcome challenging terrains (e.g., momentum-based planning).

The outcome of this task plays a key role in capability element 5 of the Capstone Integration and Assessment Vision, which requires a robot to navigate mobility challenges due to natural terrain and urban clutter. In particular, it will provide robots with the necessary models to overcome these mobility challenges. For example, on the way to a desired location, the robot may encounter high vegetation, mud patches, steep hills, sand, gravel, piles of obstacles, etc.

This research is a multi-year effort that is expected to extend for the five-year duration of the RCTA. It involves development in the short term for wheeled and tracked vehicles with a transition to systems exhibiting more complex dynamics, in particular legged robots, in the longer term.

State of the Art:
This task seeks to develop control and planning strategies that enable an AGV to traverse difficult surfaces (e.g., sand, mud, loose rocks, and steep hills) safely and efficiently. Very little research has been done in this area. Some work has been recently accomplished at CMU (using the Crusher vehicle) on pure learning approaches to difficult environment navigation. However, learning by itself is somewhat limiting since it requires a huge number of experiments. What could dramatically reduce the amount of experiments is the use of dynamic models. For example, if one had to produce a curve of the function $f(x) = a_2x^2 + a_1x + a_0$ for a given range of x, then a purely experimental approach would require one to discretize the range of x, perhaps by dividing it into dozens or hundreds of equally spaced points, then performing an experiment to produce $f(x)$ for each of these points. On the other hand, if one knows the model, which has only three parameters, then only three experiments are needed. Using models in planning for these difficult surfaces is still in its infancy.

A related issue is momentum-based motion planning which can be accomplished using dynamic models. This is clearly needed for climbing steep hills, i.e., hills that have grades so steep that the vehicle can only decelerate due to the torque limitations of the vehicle motors or engine. Momentum-based motion planning is also needed for traversing mud puddles and high, stiff

vegetation. However, prior to the RCTA research, no one had developed momentum-based planning algorithms.

Key Barriers and Approach:
Current vehicle-terrain models are mostly empirical and have limited generalization. On the other hand, the few existing analytical models, assume Coulomb friction models, which result in poor predictions of motor torques, actuator saturation, and power consumption.

In this subtask, we propose to use appropriate friction models to develop terrain-dependent low-order dynamic models for wheeled and legged platforms. Online learning techniques will then be employed to identify key model parameters (e.g., coefficient of friction and shear deformation modulus). The developed models will then be integrated with our Sampling Based Model Predictive Optimization (SBMPO) motion planning framework. This research needs to be integrated in the longer term with the perception research in order to couple proprioceptive sensing and visual recognition with the developed dynamic models.

Metrics for Evaluation:
The metrics for this research are: 1) the correspondence between the dynamic models and experimental measurements, 2) the energy savings associated with using these dynamic models, and 3) the improvement in the ability of the robot to traverse difficult mobility challenges. As illustrated in Figure 5-5, current progress has led to the development of a dynamic model for skid-steered robots, which has shown good correspondence between analytical and experimental motor torques for vinyl, asphalt, and grass surfaces. The models have been extended to inclined surfaces and have been integrated in simulation with SBMPO to generate energy efficient paths such as the one illustrated in Figure 5-6.

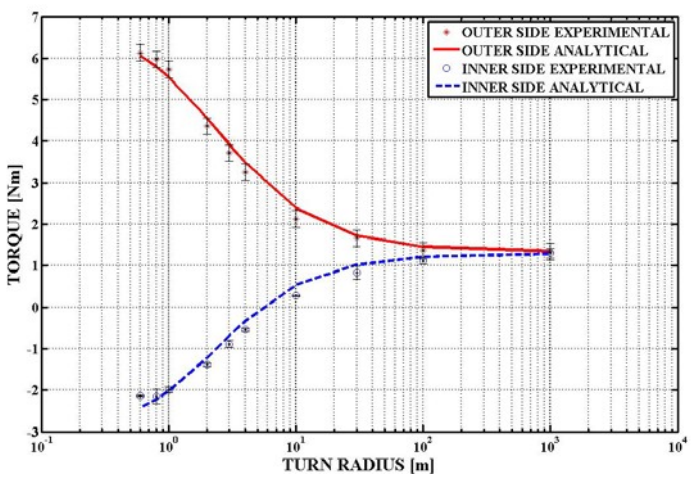

Figure 5-5: Experimental versus analytical torques.

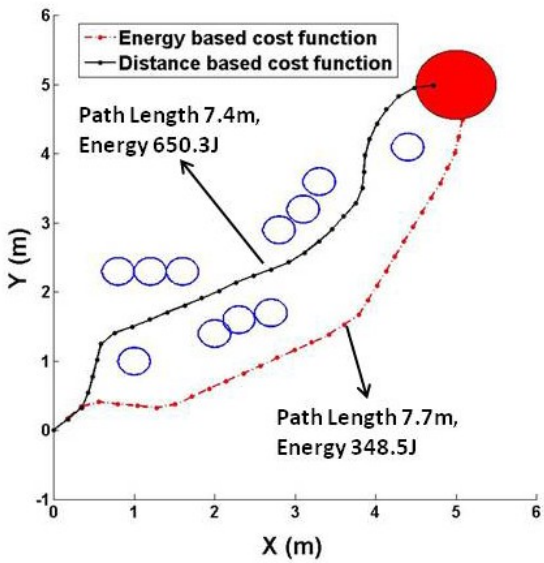

Figure 5-6: Energy optimal versus distance optimal paths via SBMPO.

In order to stay on track with the Capstone Integration and Assessment Vision, current work includes the integration of the model with SBMPO to demonstrate momentum-based planning for mobility challenges like high, stiff vegetation, steep hills, and mud patches.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Demonstration of online learning of reduced order dynamic model for a skid-steered vehicle; there must be a close correspondence between the model and experimental data. Simulation of longitudinal momentum-based planning for skid-steered vehicles that improves the ability of the robot to traverse difficult mobility challenges such as steep hills, mud patches and high, stiff vegetation.
Q2	Experimental demonstration of longitudinal momentum-based planning on different terrains that improves the ability of the robot to traverse difficult mobility challenges. Development and experimental verification of an empirical power model for a walking gait of the XRL (X-RHex Lite) robot.
Q3	Experimental verification of 3D model for skid-steered vehicles on outdoor surfaces; there is expected to be a close correspondence between the model and experimental data. Demonstration of proprioceptive terrain classification for the XRL robot (in a joint effort with UPenn); this will allow the robot to sense difficult surfaces and change its control system for that surface.
Q4	Outdoor online implementation of energy efficient path planning on skid-steered vehicles; these results are expected to demonstrate substantial energy savings associated with minimum energy planning over minimum distance planning. Demonstration of energy efficient motion planning for the XRL robot; these results are expected to demonstrate substantial energy savings associated with minimum energy planning over minimum distance planning.

M8-2012 - Dynamic Multi-modal and Reconfigurable Mechanisms

Level: 6.2 Applied Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Emmanuel Collins	FSU	M8-1
Additional Investigators	Dan Koditschek	UPenn	M8-4, M8-5
	Mark Cutkosky	Stanford	M8-6
	Raymond VonWahlde	ARL	M8-7, M8-8
	Geoffrey Slipher	ARL	M8-7
	Justin Shumaker	ARL	M8-7
	Jason Pusey	ARL	M8-8

Objective:
The objective of this task is to use bio-inspiration and engineering best practice to develop:

- Multi-modal limbs for running and climbing.
- Tuned and tunable compliance.
- Novel actuation for multi-modal mobility.
- Isolated shape/parameter changers and prime-movers.

Background:
In this task, we propose to investigate novel locomotion systems for climbing, crawling, hopping, and leaping. We also explore improving locomotion through reconfigurable mechanisms. The goal is to realize robots that combine two or more locomotion systems to achieve mechanical traversability and dexterity in dramatically varied terrains. The science and technology we propose will enable systems that can perform combinations of vertical climbing, negotiation of tight spaces, running, hopping, and leaping over or from obstacles. This task presents extremely difficult research challenges because the robotics community has only begun to understand how the form and function (i.e., the mechanical structure, actuation and sensing, and the controls and behavior) combine to achieve dynamical dexterity. Moreover, power density constraints require passive mechanical elements and novel materials to achieve fast, strong, well- timed interaction of limbs with complex environments which adds to the complexity of the problem.

Research Description:
To overcome these challenges we propose to integrate algorithmic and mechanical design insights. We merge structural innovations such as tuned and tunable compliant structures with passive mechanical hierarchies that have fine structure. We also take advantage of new kinds of actuators, as described in Task M9: Next Generation Actuators and Materials, and deeper understanding of locomotion mechanics and control from Tasks M5: Theory and Principles of Multi-modal Locomotion Planning and Control and M6: Principles of Locomotion Mechanics. This research strategy yields more general platforms offering combined capabilities in a single platform that can be reconfigured at execution time using an interchangeable family of rationally designed and rapidly prototyped modular components. This approach to designing dynamic robots with unique task-specific capabilities offers great promise for achieving innovative robots with advanced capabilities.

Specifically, we combine variable transmission elements and energy storage materials that can redirect the application of force between the different directions of effort needed for running and climbing. We generalize the theory and design methodology for constructing limbs with tunable three-dimensional stiffness properties. We introduce small, distributed "shape-changing" actuators into the mechanical design so that the larger "prime mover" actuators and the compliant elements can be altered, focusing them for the next task.

In addition to using variable compliance legs for more efficient running, we propose to explore design optimization for different running conditions. The work focuses on configuring limb and body compliance to account for nonlinear dynamic coupling and transients during mode transitions. Specifically, we begin to explore more complex morphologies – backbones and tails as well as more elaborated limbs with appendages – all with bio-inspired insights. Combining these efforts, we quantitatively evaluate the efficiency and generality of our locomotion capability with experiments running and climbing on a variety of terrains.

Metrics:

- Platform speed (bodylengths per second)
- Endurance (specific resistance) over breadth of operational regime (range of speeds, terrains, and substrate geometry)
- Failure rate in transitions between task domains (such as from level to vertical ground)
- Capability to support sensory payloads (onboard sensory modality diversity and bit rate)

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
Jonathan Clark will be visiting and giving a lecture about M8-related developments at UPenn and CMU in February and March, respectively. UPenn researchers are planning to hold a special session or workshop about using appendage inertia in running at the Climbing and Walking Robots (CLAWAR) conference in July. Stanford and FSU will be collaborating on the development of adhesive foot technologies for dynamic vertical climbing. JPL and FSU will be collaborating on integrating novel piezoelectric actuators being developed as part of M9 into the dynamic quadruped, subtask M8-1.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Develop familiarity with and port CANID and X-RHex to DB environment. Repair CANID. Perform tests on QuadClimber.
Q2	Experiment with CANID. Test XRHex on gravel hills. Demonstrate level surface running of QuadClimber.
Q3	Compare robot performance to DB simulation. Improve controllers. Prepare for IRA3.
Q4	Design review and appraisal of gait planning and development. Begin integration of smart materials from Task M9.

Related Research:
Inputs from:

- Task M5: Theory and Principles of Multi-modal Locomotion Planning and Control
- Task M6: Principles of Locomotion Mechanics
- Task M9: Next Generation Actuators and Materials

Outputs to:

- Task I1: Framework for Intelligence
- Task I2: Data Mapping for Inference and Focus
- Task I5: Learning through Experience

Integration and Assessment Activity:

- Exploiting mechanisms, dynamics, and whole-body interactions with objects and terrain.
- Manipulation and locomotion in challenging 3D terrain.
- New materials, sensors, actuators, and mechanisms.

Subtask 1: Mechanisms for Versatile Multi-legged Locomotion (FSU)

Objective and Benefits:
What are you trying to do?
Develop robotic platforms that can effectively exploit their dynamics to rapidly locomote in multiple domains. In particular, we are interested in the combination of running over horizontal and sheer vertical surfaces.

How does the effort fit into the overall program vision?
Developing robots with unique mobility characteristics that allow them to run up to walls and then climb them will dramatically improve the speed and robustness of robots moving in

cluttered environments and extend the range of missions that small mobile robotics can be asked to perform. Improving the range of domains that dynamical robots can achieve is an essential feature for RCTA outdoor scenarios.

If this is a multi-year effort, how long do you expect it to continue?
This is a multi-year effort, and it is expected to continue throughout and beyond the life of the project.

Technical Approach to Overcome Barriers and Advance SoA:
What is the essence of your tech approach?
Utilize insights and models based on biological studies of dynamical motion of animals in varied environments. By embedding self-stabilizing models of vertical and horizontal locomotion into the same robotic platform, highly efficient and stable running should be achievable, even over vastly different types of terrain.

How does your approach move beyond the state of the art?
Our present work in dynamical climbing represents the state of the art in rapid vertical locomotion. This effort extends that work to combine it with well-established principles of horizontal (level ground) locomotion. This combination has not yet been achieved in a single robotic platform.

Why do you think you will succeed?
We have a proven track record in the development of dynamical running and climbing robots that serve as the basis for this work. One of the primary challenges to dynamical multi-modal locomotion is the difficulty of adapting the tuned dynamics of running to climbing. The developments in variable impedance mechanisms in Subtasks M8-6, M9-1 and M9-6 are addressing this very challenge.

How does this tie into other research activities?
In addition to the tasks listed above, this task builds upon developments in understanding animal locomotion (Subtask M6-1) and will utilize the transition framework being developed in M6-3.

Metrics and Progress Assessment:
How are you measuring your progress?
Through empirical tests of robot performance, such as speed, efficiency, rate of failure (falling), etc.

What specific metrics are you using?
We use measures of specific resistance, non-dimensionalized speed, ground reaction forces, etc.

What progress has been made toward achieving your ultimate goal?
In the past year, we have constructed a quadrupedal robotic climber that has limbs capable of generating the appropriate forces and stiffnesses for both climbing vertical surfaces and horizontal running.

How do you plan to measure intermediate progress?
We plan to measure the performance of the robotic test platform in controlled tests.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Demonstrate robust climbing on vertical surfaces in a dynamically similar manner to the Full-Goldman (FG) template, and quantify the behavior.
Q2	Demonstrate robust running on level surfaces in a dynamically similar manner to the Lateral Leg Spring (LLS) model. Additionally, investigate the use of FG-like climbing and LLS-like running on sloped surfaces.
Q3	Develop rudimentary heading control for running and climbing on vertical and level surfaces. Furthermore, feet from Stanford's Stickybot will be integrated to allow for climbing on a more varied selection of surface types.
Q4	Initial integration of smart materials for use in aiding transitions between modes of locomotion.

Subtask 4: Integrating Mechanisms of Animal and Machine Locomotion for Unique Robot Mobility (UPenn)
This is the continuation of the CANID project that aims to design and build a quadrupedal robot with an actuated, articulated compliant spine.

Objective and Benefits:
We hope to achieve the first actuated compliant spine-based quadrapedal bounding machine. This is a 6.2-level implementation opportunity for some of the ideas arising from the 6.1-level effort of Subtasks M6-1 and M6-3. The potential for achieving much higher efficiencies and speeds in locomotion has clear impact on any locomotion application within RCTA.

This is a continuation of the CANID effort started last year. Our speed of progress, and hence likely duration, are strongly linked to the level of sponsor support we can achieve. Last year, we moved ahead rapidly because the project enjoyed the participation of an IC Supported Postdoctoral fellow who moved on to his next position at the end of the summer. We are presently looking for additional funding to move this project ahead. Discussions with ARL DMUM Lead Harris Edge suggest significant internal interest within his organization in this project, so we are hopeful that sufficient funds can be found to keep the project going despite the cuts in 6.2-level RCTA funding for FY12.

This is largely a 6.2 effort (because of the focus on applying existing notions of modular design and dynamical task encoding), but it represents a very high-risk/high-reward project with a very uncertain payoff time because of the novelty of the design and goals. External assessments of progress will be based upon scientific documentation of performance in refereed technical publications as well as visits to UPenn GRASP Lab by RCTA collaborators and ARL and other independent Government personnel to review and assess progress. Even in the best scenario, given the necessarily limited support possible within the RCTA for any single project activity,

we should not imagine that the experimental hardware developed in this project will ever get out of the laboratory for any kind of external assessment during the five-year lifetime of this project. We expect that prototype hardware based upon and benefiting from lessons learned on this project will indeed be ready for assessment in the field over a time frame of a few years past the nominal end date of the present RCTA.

State of the Art:
At present, to the best of our knowledge (and certainly as documented in the technical literature), there is no functional compliant spine quadrapedal machine. We think the traditional steady state analysis introduced to robotics in [1] and [2] (and presented in the context of "spinal" bounding say in [3]) might be accelerated by introduction of (appropriately modified) computational homology methods [4] that we will be attempting to develop in the M6-3 investigation on a more fundamental level; demonstrating a CANID application would represent an important benchmark instance of the utility of the abstract idea.

Key Barriers and Approach:
We have a long track record in pushing out a series of firsts in dynamical robotics, including the first dynamical hexapod [5] and the first dynamical climbers [6,7]. There is strong and growing competition from excellent researchers in this important domain; thus, whether we can continue the record of "firsts" given the inevitable funding constraints within RCTA is unclear, but surely we have a place at the table of research in the area.

More efficient and faster steady state locomotion has relevance to all outdoor RCTA tasks. Specifically, this is a 6.2-level tech transition opportunity for some of the ideas arising from the 6.1-level effort of Subtasks M6-1 and M6-3 along the lines indicated above as well as, potentially, for other tasks within DMUM.

Metrics for Evaluation:
We use measures of efficiency such as specific resistance [8] as well as standard measures of environmental complexity [5] and normalized speed or acceleration. Last summer, we completed the assembly of the prototype machine and achieved a three-stride benchmark. Most immediately, we seek a stable steady state bounding gait. We believe a ~10 stride hands-off experiment would justify a claim of stable steady state gait.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Complete repairs to CANID limbs from last quarter's experiments. Study and verify the anti-buckling effect of the "vertebral" cable guides installed in the spine last quarter. Learn to use Boston Dynamics' CANID DB simulation environment.
Q2	Pursue further empirical tests of gait stability on the CANID test platform. Develop parallel behavioral controllers for physical CANID and CANID-DB simulation.
Q3	Detailed comparisons between physical CANID experiments and CANID-DB simulation.
Q4	Initial appraisal of potential Conley Index application to CANID steady state gait planning.

Subtask 5: Using Environment Mechanics Representation in Mobility Tasks (UPenn) This is the 6.2-level empirical characterization of the models, algorithms, and new controllers based upon the 6.1-level effort arising from Subtask M6-4. The target will be legged locomotion platforms (principally RHex), and the line of effort will be largely devoted to translating the M6-4 abstract ideas into specific robot controllers and preparing the robot for empirical assessments at RCTA IRA events.

Objective and Benefits:
Locomotion over broken, unstable, slippery, complex terrain is an essential feature of any of the RCTA outdoor scenarios. This is a new project startup intended to leverage recent advances in phenomenological modeling and simulation of complex terrain [9] as well as our ideas about "foveated" (cellular decomposition) representations arising from moving the abstract world model effort (formerly Subtask I1-2, now proposed as Subtask M6-4). We have cited in the M6-4 project proposal our experience at the recent IRA1 session in FTIG where we presented RHex with one of the various gravel piles on site, and it immediately became clear how important substrate mechanics will likely be.

State of the Art:
Whereas legged machines have achieved increasingly impressive performance in steady state operation on level horizontal and vertical terrain with simple uniform substrate features, locomotion over slippery, fragile, or broken substrates remains a difficult and largely unexplored domain of operation. We have presented early work on feedback-based gait adjustment for RHex [10], and Rizzi and Haynes have extended and improved those ideas in the context of vertical climbing [11]. These efforts did not try to incorporate explicit terrain or substrate models in the design of the controllers, yet it is clear empirically that small changes in substrate mechanics can have catastrophic effects on theretofore robust gaits [12]. Important recent applications of resistive force theory to snake mechanisms operating in homogeneous granular media suggest that explicit substrate models can be used effectively for locomotion [13], but most substrates present dramatically inhomogeneous and changing features that would require multiple high-dimensional models with great uncertainty as to which might best be applied at any instant.

Key Barriers and Approach:
The method of dynamical task specification and control [14,15] has a long and still vigorously growing track record of success in theory [16] and practice [17]. We hope to extend the methods of [10,11] into a more model-based framework arising from the abstract results of Subtask M6-4.

We are working very closely with the Polypedal lab in M6-1 to integrate animal strategies into our designs and algorithms. We hope to connect as well with work ongoing in Tasks M5 and M7. Most specifically, this is the 6.2-level tech transition opportunity for the models, algorithms, and new controllers based upon the 6.1-level effort arising from Subtask M6-4.

Metrics for Evaluation:
We use standard measures of efficiency such as specific resistance [8] as well as advanced measures of substrate complexity [12] and normalized speed or acceleration.

	Subtask 5 Quarterly Milestones and Deliverables
Q1	Port to X-RHex empirical platform and Boston Dynamics X-RHex-DB environment suitably adapted versions of the controllers reported in [10,11].
Q2	Initiate experiments with X-RHex on gravel hills. Port to X-RHex-DB environment an appropriately abstracted simulation model of granular media [9].
Q3	Prepare X-RHex for and participate in IRA3 with focus on climbing gravel hills as baseline for future advances.
Q4	Design review: port M6-4 world model representation of complex terrain and granular media to actual use in X-RHex and X-RHex-DB controller.

Subtask 6: Utilization of Variable Impedance Mechanisms (Stanford)

Objective and Benefits:
What are you trying to do?
Apply directional dry adhesive technologies to develop feet for the dynamic quadrupedal climbing robotic platform being developed in Subtask M8-1.

How does the effort fit into the overall program vision?
Adaptive dynamical locomotion over varying terrain requires the alteration of climbing strategy and foot attachment as substrates change. For effective climbing on smooth surfaces, such as buildings, directional dry adhesives are an effective approach. The integration of these feet onto dynamic climbing robots will greatly expand the range of surfaces they can traverse.

If this is a multi-year effort, how long do you expect it to continue?
While directional dry adhesives has shown to be effective on small climbing robots, the challenges involved in incorporating them onto dynamic platforms with masses of 1 kg and greater are unknown. The development of robust versions of these mechanisms may take a few years.

Technical Approach to Overcome Barriers and Advance SoA:
What is the essence of your tech approach?
Utilize directional dry adhesives as animals like geckos and insects use when climbing to improve attachment for dynamically climbing robotic systems. Micro-fabrication techniques, including angled-exposure lithography, have been used to create hierarchical structures that adhere when loaded and then self-detach when unloaded. By incorporating the proper degrees-of-freedom and hierarchies of compliance, practical levels of surface adhesion can be achieved on relatively smooth surfaces for areas of contact large enough to support dynamic robot platforms with masses of 1 kg and greater.

How does your approach move beyond the state of the art?
Over the past few years, we have developed the most effective gecko and insect-inspired dry adhesives and applied them to climbing systems. In this subtask, we will be applying these

techniques and structures to dynamical climbing systems with relatively shorter and larger loading conditions than found in slower climbing systems.

Why do you think you will succeed?
Our recent developments in micro-fabrication of polymer materials and successes on a range of both small (Stickybot, Spinybot) and large (RiSE) climbing platforms give us confidence that this task will be successful.

How does this tie into other research activities?
This effort feeds directly into Subtask M8-1, the development of mulit-modal dynamical robots.

Metrics and Progress Assessment:
How are you measuring your progress?
Through empirical tests of climbing behavior, adhesive force, range of motion, and speed. Successful integration onto robotic platforms of increasing size and with increasing payloads.

What specific metrics are you using?
We use measures such as range of motion, speed of attachment and detachment, strength of attachment, and effectiveness and efficiency of the climbing robot to gauge progress.

What progress has been made toward achieving your ultimate goal?
This new subtask is just starting.

How do you plan to measure intermediate progress?
We plan to measure the rate at which these devices improve performance and are adopted onto the quadrupedal climber being developed in Subtask M8-1.

	Subtask 6 Quarterly Milestones and Deliverables
Q1	Initial benchtop tests on directional adhesive arrays 3x3cm and larger, sufficient for supporting a climbing robot with a mass of at least 4 kg. Quantification of pull-off angles and maximum misalignment that can be accommodated.
Q2	Design of new foot and ankle for quadrupedal climber in collaboration with FSU. Adaptation of adhesive and microspine climbing technology to quadrupedal climber.
Q3	Delivery of first generation feet and ankles for quadrupedal climber to FSU.
Q4	Development of refined foot and ankle design in collaboration with FSU for further testing. Publication of manufacturing recipe and design specifications.

Subtask 7: Tactile Sensing for Control of Hyper-redundant Mechanisms (ARL)

Objective and Benefits:
The objective of this subtask is to contribute to research in unconventional mobility and manipulation by employing skin tactile sensors on hyper-redundant mechanisms, e.g., snake robots or tentacle-like manipulators. Additionally, we will demonstrate in simulation the benefit

to control algorithms; for example, how such data could aid in gait transition for path planning and autonomous navigation.

Tactile sensors can provide useful feedback for determining the interaction of hyper-redundant mechanisms with the environment. Such data could assist in maneuvering in constrained space or selecting an appropriate gait over the ground. Novel sensing approaches, such as the tactile sensing being pursued in this subtask, can contribute directly to the implementation of behavior generation.

State of the Art:
A number of commercially available sensors have been developed for application in robotic manipulators and prosthetic devices. Interlink Electronics holds a patent on Force Sensing Resistors (FSR). FlexiForce® sensors are piezoresistive force sensors from Tekscan. They have products that can detect and measure a relative change in force or applied load; however, neither sensor offers information concerning the position of the applied load. The BioTac® is a tactile sensor system from SynTouch LLC. The design consists of a rigid core surrounded by an elastic skin filled with a fluid similar to a human fingertip. Among other sensory information, it can detect forces using impedance within the core in the range of 0.01 – 50N with a variability of <5% within range.

Control of hyper-redundant mechanisms using tactile sensor feedback dates back to at least the mid 1970s with work done by Hirose (1). Research has been done in utilizing dense sensor arrays as artificial skins (2). Tactile sensing has been proposed as an interface between a human and robot to both control the robot and allow safe interaction (3). Most recently, researches at the Healthcare Robotics Lab at Georgia Tech, propose drobots that can control their contact forces during goal-directed motion (4). Tactile sensing would be used in a novel controller using model predictive control.

Key Barriers and Approach:
The research will follow two parallel tracks. One track will leverage ongoing research by Slipher and Pekala. It will explore the feasibility of incorporating a pressure sensitive tactile sensor into the “skin” of a hyper-redundant robot (e.g., Howie Choset’s Unified Modular Snake Robot). An appropriate sensing method will be chosen, a design will be developed, and a prototype will be constructed. The goal will be to demonstrate at least one functioning module equipped with a tactile sensing skin.

The second track will model a hyper-redundant mechanism in simulation with virtual tactile sensors. Generated feedback will be scrutinized to aid the development of control algorithms that make optimal use of the data. Dynamic simulation software like LMS Virtual Lab will be utilized for this purpose.

Metrics for Evaluation:
Autonomous contact sensing work will be evaluated by the percentage of repeatability for object contacts the system can detect and localize within the load range of the sensor. A preliminary

estimate for the load range is between 1.5 and 30 N with a variability in load measurement of less than 5%.

Related Research:
This proposed collaboration supports Tasks M3: Sensor-based Dexterous Manipulation as well as Task M5: Theory and Principles of Multi-modal Locomotion Planning and Control.

	Subtask 7 Quarterly Milestones and Deliverables
Q1	Identify most promising tactile sensing method (e.g., resistive versus capacitive). Determine method of manufacturing. Obtain material, electronics, and miscellaneous components. Select simulation software.
Q2	Construct prototype sensor. Create simulation.
Q3	Test and gather data. Verify simulation.
Q4	Integrate sensor into module.

References:
(1) Hirose S. Umetani Y. "Kinematic control of active cord mechanism with tactile sensors." In *2nd International CISM-IFTMM Symposium on Theory and Practice of Robots and Manipulators*, pages 241-252, 1976.
(2) P'erez G.B. (2006) S.N.A.K.E.: A Dynamically Reconfigurable Artificial Sensate Skin (Master's thesis), Massachusetts Institute of Technology, Cambridge, MA.
(3) Wösch, T., Feiten, W. Prassler, E., Lawitzky, G., Stopp, A., Grunwald, G., et al. (2005) Tactile Interaction between Human and Robot, *Advances in Human-Robot Interaction, Springer Tracts in Advanced Robotics* Vol. 14. (pp 91-92). Springer Berlin/Heidelberg.
(4) Jain, A., Killpack, M.D., Edsinger, A., and Kemp C. C. (2011) Manipulation in Clutter with Whole-Arm Tactice sensing, Submitted for Review.

Subtask 8: Dynamic Modeling of the CANID Platform (ARL)

Objective and Benefits:
The objective of this subtask is to examine the impact of a central flexion joint upon the overall speed, efficiency, and agility of a dynamic running legged robot. It will develop a 3D physics-based model of the CANID legged robotic platform and use this model to understand the dynamics of the system and how it can be used to form an efficient gait using a central flexion joint for a spine. The effort will focus upon dynamic modeling using Leuven Measurement Systems (LMS) Virtual-Motion and the Boston Dynamics Digital Biomechanics (DB) simulation software environments. The dynamic models will be validated and verified using data gathered from experiments with a prototype CANID robot.

The effort will be closely coordinated with Subtask M8-4: Integrating Mechanisms of Animal and Machine Locomotion for Unique Robot Mobility (UPenn) and is part of a multi-year effort

where within each year the model will be expanded in fidelity and complexity to meet the needs and understanding of the CANID platform.

State of the Art:
The state of the art for modeling legged robotic systems focuses upon examination of rigid spines for quadrupeds and simplified models such as the spring loaded inverted pendulum (SLIP) models. Boston Dynamics has developed a modeling environment called Digital Biomechanics (DB) where the BigDog quadruped is modeled. Boston Dynamics has provided a simplified model of CANID with a spine modeled as three revolute joints connected by rigid links. Each joint can be modeled as any combination of spring, damper, and actuator.

Key Barriers and Approach:
The approach being taken is to create a dynamics model of the CANID platform and use this model to understand the system dynamics as well as to aid in the development and creation of different gaits. The model will be used to support design changes and design evolution. Several design ideas are being generated as research progresses, and the development of virtual dynamic simulation models can provide insight to guide the directions chosen. Virtual models do depend on physical systems so a balance between the virtual and physical prototype will have to be made based on how far a model can estimate the physical world.

Different levels of fidelity of the model are currently being pursued to enable greater results and to verify the needed complexity for increased fidelity. Different levels of modeling fidelity supply both an understanding of the dynamics of the systems and the tools to enable exploration of different system aspects. For example, one can employ a lower fidelity model to explore basic control designs for the system and then examine specific designs implementing higher fidelity models. This stepping of the level of complexity will allow isolation of basic concepts to be roughed out then evolved as complexity is added.

Most models of legged platforms have not implemented the complexity of flexible elements working together within a high fidelity simulation environment to obtain a gait. Boston Dynamics uses their in-house developed DB dynamic simulation environment to obtain an initial understanding of the design and has provided the RCTA with their DB tool with a model CANID. The model uses a slightly simplified representation of the spine to reduce the complexity of the system while still maintaining the system's general dynamics. The LMS model captures the finite element model of the spine, which increases the fidelity and the complexity of the system. Based on the success of the development of the physical prototype and the initial development of the kinematic models, the evolution to the dynamics model is an in-depth extension of the current system.

As part of this research effort, a successful CANID prototype was designed, manufactured, assembled, and initially tested in FY11 through collaborative efforts between ARL and UPenn. ARL researcher Jason Pusey works part-time at UPenn with the Kodlab to maintain the CANID design, assist with experiments, and collaboratively develop experimental plans for the physical prototype.

Metrics for Evaluation:

- Level of completion for the model; percentage of geometry imported into the model and dynamically represented in the model
- Percentage of mass properties incorporated for all parts
- Accuracy of simulation motion compared to CANID platform

Currently, Boston Dynamics has supplied a model of CANID. The DB model of CANID incorporates the basic interconnected components and their basic inertial properties with some low-level controls to actuate the model. The LMS model of CANID has modeled most of the interconnected components and the finite element model of the spine with basic actuation schemes currently fixed in a 2D plane. Initial understanding of how to model the actuation in both models is currently being tested.

Intermediate progress will be measured by the increased degree of fidelity reached that represents the physical CANID prototype. This multi-year effort will evolve the fidelity until the dynamic motion of CANID modeled in the simulation environment portrays the physical motion of the CANID prototype. Intermediate steps of matching different isolated parts of the CANID platform between the physical and modeled worlds will be executed to concede accuracy.

	Subtask 8 Quarterly Milestones and Deliverables
Q1	Develop familiarity with and port CANID to Digital Biomechanics and LMS environment. Repair CANID.
Q2	Continue building CANID models in DB and LMS. Experiment with CANID.
Q3	Develop simulation experiments to verify CANID model, and continue to develop skills with the model environments. Develop control methodology to actuate CANID model.
Q4	Develop a plan to implement control methodology on model and physical CANID robot. Develop a plan to validate CANID model.

References:

[1] M. Buehler, D. E. Koditschek, and P. J. Kindlmann, "A family of robot control strategies for intermittent dynamical environments," *IEEE Control Systems Magazine*, vol. 10, pp. 16-22, 1990.

[2] D. E. Koditschek and M. Buhler, "Analysis of a simplified hopping robot," *International Journal of Robotics Research*, vol. 10, no. 6, pp. 587-605, 1991.

[3] U. Culha and U. Saranli, "Quadrupedal bounding with an actuated spinal joint," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, 2011, pp. 1392–1397.

[4] W. D. Kalies, K. Mischaikow, and R. C. A. M. VanderVorst, "An algorithmic approach to chain recurrence," *Foundations of Computational Mathematics*, vol. 5, no. 4, pp. 409-449, 2005.

[5] U. Saranli, M. Buehler, and D. E. Koditschek, "RHex: A Simple and Highly Mobile Hexapod Robot," *The International Journal of Robotics Research*, vol. 20, no. 7, p. 616, 2001.

[6] G. C. Haynes et al., "Rapid Pole Climbing with a Quadrupedal Robot," in *Proc. IEEE International Conference on Robotics and Automation*, 2009, p. (to appear).

[7] J. Clark et al., "Design of a Bio-inspired Dynamical Vertical Climbing Robot," in *Robotics: Science and Systems III Atlanta, Georgia*, 2007.

[8] G. C. Haynes and D. E. Koditschek, "On the Comparative Analysis of Locomotory Systems with Vertical Travel," in *ISER'10*, New Delhi & Agra, India, 2010.

[9] A. Slatton et al., "Integrating a Hierarchy of Simulation Tools for Legged Robot Locomotion," in *Proceedings of Workshop on Robot Simulators, IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2008, p. 474.

[10] J. D. Weingarten, R. E. Groff, and D. E. Koditschek, "A framework for the coordination of legged robot gaits," in *Robotics, Automation and Mechatronics, 2004 IEEE Conference on*, 2004, vol. 2.

[11] G. C. Haynes, "Gait regulation control techniques for robust legged locomotion," Carnegie Mellon University, Pittsburgh, PA, 2007.

[12] C. Li, P. B. Umbanhowar, H. Komsuoglu, D. E. Koditschek, and D. I. Goldman, "Sensitive dependence of the motion of a legged robot on granular media," *Proceedings of the National Academy of Sciences*, vol. 106, no. 9, pp. 3029-3034, 2009.

[13] R. D. Maladen, Y. Ding, P. B. Umbanhowar, and D. I. Goldman, "Undulatory swimming in sand: experimental and simulation studies of a robotic sandfish," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 793 -805, Jun. 2011.

[14] D. Koditschek, "Task encoding: toward a scientific paradigm for robot planning and control," *Robotics and autonomous systems*, vol. 9, no. 1-2, pp. 5-39, 1992.

[15] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek, "Sequential Composition of Dynamically Dexterous Robot Behaviors," *The International Journal of Robotics Research*, vol. 18, no. 6, pp. 534-555, 1999.

[16] P. Holmes, R. J. Full, D. E. Koditschek, and J. Guckenheimer, "The dynamics of legged locomotion: Models, analyses, and challenges," *SIAM Review*, vol. 48, no. 2, pp. 207-304, 2006.

[17] D. E. Koditschek, R. J. Full, and M. Buehler, "Mechanical aspects of legged locomotion control," *Arthropod Structure and Development*, vol. 33, no. 3, pp. 251-272, 2004.

M9-2010 – Next Generation Actuators and Materials

Level: 6.1 Basic Research

	Name	*Organization*	*Subtask*
Principal Investigator	Jonathan Clark	FSU	M9-1
Additional Investigators	William Oates	FSU	M9-1
	Yoseph Bar-Cohen	JPL	M9-4
	Mark Cutkosky	Stanford	M9-6
	Geoffrey Slipher	ARL	M9-7

Objective:
The objective of this task is to:

- Characterize and model novel solid state actuators and active materials; and
- Identify and exploit opportunities for implementation on manipulation and mobility platforms.

The goal of this task is to change the course of robotics through fundamental research and advanced technology development leading to a new generation of actuation systems for highly dexterous, ultra mobile, autonomous robots. The use of effective and compact actuators is critical to performing robotic missions, and it is in line with the overall program vision as articulated in the FY11 APP and the Capstone Assessment Vision document. In particular, it is focused on enhancing the capability in several capstone assessment areas including miniaturized pointing devices to enable surveillance (3a, 6a) and mobility (4a, 5a). The objectives of Subtasks 1, 2, and 6 are to support Task M9 through the development of adaptive structures and integration into legged robots that rapidly change their elastic modulus, viscoelastic properties, and shape from an applied electric field. Subtask 4 investigates new, super thin smart material-driven motors for integration directly into the robot's frame. These advancements are critical to provide transformative capabilities for robotic limbs that can supply advanced maneuverability for running and climbing over a variety of terrains.

Background:
Actuators based on smart materials have great potential to transform robotic systems by improving the strength-to-weight ratio, speed, range of motion, compactness, efficiency, and controllability of manipulators, legs, sensors, and other robotic components. The most mature of technologies with respect to robotics include conventional smart materials: ferroelectric (piezoelectric) materials, magnetostrictive compounds, and shape memory alloys. Piezoelectric materials still provide robust opportunities to provide unique step-and-repeat actuators (for example), which have not been exploited to their full potential within robotic manipulators. Some of these promising opportunities have begun to mature through research at JPL. Alternatively, multifunctional polymers have begun to emerge as viable, light-weight alternatives

that not only provide actuation and sensing but also can be exploited to create real-time stiffness and damping changes which are highly relevant to robotic manipulation and control.

Whereas multifunctional polymers provide a number of unique attributes that cannot be obtained with "conventional" smart materials, they are limited by electrode performance, fatigue, and environmental sensitivity. Both dielectric elastomers and shape memory polymers require electrical input or heat generation to achieve changes in shape and stiffness. Flexible electrodes are critical to fully utilize these materials by effectively delivering electric fields as well as heat. Thus, collaborations with Drs. Lenhart and Slipher at ARL will continue to be an integral part of this research to support materials development. Transition of materials developed at ARL to FSU is expected to begin early 2012. (This has been primarily limited to ARL patent submissions that began summer 2011).

Research Description:

To overcome these challenges we propose to develop active materials for legged robotic platforms by formulating new constitutive models with complimentary experimental characterization for integration into legged robotic structures. Material modeling and characterization research will focus on electroactive and thermal responsive elastomers and polymers to enhance control and eliminate the need for additional DC motors and complex transmission designs. We will also focus on smart structures that utilize shape-changing actuation and real-time material property control for dynamic adaptation and superior maneuverability. We will investigate effective techniques of leveraging forces that are induced by piezoelectric stacks and seek the development of structurally integrated motors. Such motors are expected to offer significant improvement in operation efficiency as well as reduced mass and size.

Electroactive dielectric films will be explored to develop systems and structures that can be used for actuation and electric field properties changes (i.e., stiffness and potentially damping control). During this phase, a combination of electro-mechanical characterization and material/structure modeling will be done to identify key performance attributes over a range of deformation states and loading rates.

We plan to extend prior development of large deformation electroactive membrane finite element models and benchtop experiments to robotic platforms to achieve this goal. Thus, far a maximum stiffness percent change of 80% has been achieved for the typical electric field range of 0 to 6.5 kV. This will include modeling and characterization of time-dependent behavior to identify bandwidth issues relative to robotic dynamics. This will also include scalability challenges. For example, we have identified new adaptive structure designs to overcome displacement ranges needed for the dynamic climbing platform being developed in Task M8. Benchmark tests in simulation and on robotic platforms will be conducted to evaluate the real-time affect of stiffness and damping variation on running performance. Material modeling and characterization research will be extended from electroactive large deformation coupling to include thermal effects. This will be done to identify temperature ranges for field-deployed environmental limits of existing electroactive materials. Further, it will provide ARL colloborators (Lenhart and Slipher) with

benchmark constitutive relationships for future comparisons with new polymer compositions that they are currently developing for this application.

This research will include collaborations with ARL (Drs. Joe Lenhart and Geoffrey Slipher) on enhanced electrode designs. One of the key challenges with dielectric elastomers is the application of robust and flexible electrodes that do not impede robotic systems integration. We will work with ARL on integrating flexible electrodes onto dielectric elastomer surfaces and compare their electro-mechanical properties to conventional carbon grease electrodes. Model development will focus on expanding stochastic homogenized energy modeling to finite deformation of the DEAs. This approach is ideal for fast computations of underlying material physics that can be integrated into finite element models and model-based nonlinear control designs based on prior research conducted by the co-PI Oates.

In Subtask 4, the focus is on developing actuators that utilize electroactive materials with leveraging techniques, simultaneous sensing, and real-time material property control for dynamic adaptation and superior maneuverability. Methods of displacement leveraging piezo-stacks are investigated in order to develop rotary and linear actuators. To optimize the design of effective actuators and sensors, various novel configurations are analytically modeled and simulated a multiphysics finite element-modeling package that allows examining physical phenomena at multiple time and space scales. Complementing the developed technologies, the PI of Subtask 4 is investigating potential applications of electroactive polymers and will coordinate his efforts with the efforts of the other subtasks' PIs as well as the activity in labs worldwide using his international role in this field.

Metrics:
The electroactive constitutive response of dielectric elastomers and thermal response of shape memory polymers will be assessed to identify the feasibility for integration into unique mobility platforms. The property attributes will be compared with conventional servo actuators or passive materials, when appropriate, to determine ideal mobility platforms that will advance dynamic capabilities.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:

1) Y. Bar-Cohen, Guest Editor of a Special Issue on Biomimetics and Bioinspiration, International Journal of Smart and Nano Materials, 2012
2) Y. Bar-Cohen, "Biologically Inspired Technologies for Aeronautics," Chapter 4 in T. M. Young (Ed.), Innovation in Aeronautics, Woodhead Publishing Limited, Cambridge, UK (Submitted for publication). Expected in the summer of 2012.
3) Y. Bar-Cohen, Chair, 14th "Electroactive Polymer Actuators and Devices (EAPAD)," SPIE's 19th Annual International Symposium on Smart Structures and Materials, held in San Diego, California, March 11-15, 2012.
4) Y. Bar-Cohen, Member of the Symposium Executive Committee, 19th SPIE SSM/NDE Symposia, San Diego, CA, March 11 – 15, 2012.

5) Y. Bar-Cohen, Member of the Program Organization Committee for the Conference "Health Monitoring of Structural and Biological Systems," SPIE SSM/NDE Symposia, San Diego, CA, March 11 – 15, 2012.
6) Y. Bar-Cohen, Member of the Program Organization Committee for the Conference "Nondestructive Characterization for Composite Materials, Aerospace Engineering, Civil Infrastructure, and Homeland Security V," SPIE SSM/NDE Symposia, San Diego, CA, March 11 – 15, 2012.
7) Y. Bar-Cohen, Member of the Program Organization Committee for the Conference "Bioinspiration, Biomimetics, and Bioreplication," SPIE SSM/NDE Symposia, San Diego, CA, March 11 – 15, 2012.
8) Y. Bar-Cohen, Advisory committee, International conference on Advanced Polymeric Material 2013, Agra, India expected to be scheduled in Jan or Feb, 2013.
9) Y. Bar-Cohen, Y. Bar-Cohen, Member of the International Advisory Committee, 5th International Conference on Electroactive Polymers: Materials and Devices (ICEP-2012), November 4-10, 2012 at Banaras Hindu University, Varanasi, India.
10) Y. Bar-Cohen, International Journal of Advanced Robotic Systems Editorial Advisory Board.
11) Y. Bar-Cohen, Editorial Board of the "Bioinspiration & Biomimetics: learning from nature" Journal, The Institute of Physics, England
12) Y. Bar-Cohen, Advisory Board of the Asia Pacific Committee on Smart and Nano Materials (APCSNM)
13) Y. Bar-Cohen, Editorial Board of the Journal of Bionic Engineering (JBE)
14) Y. Bar-Cohen, Editorial board of the Advances in Materials Research
15) Y. Bar-Cohen, Associate Editor of the International Journal of Smart and Nano Materials, Taylor & Francis Group,
16) S. Sherrit, M. Badescu, Y. Bar-Cohen, and X. Bao, "Miniaturization of planar horn motors," SPIE Smart Structures and Materials/NDE Symposium, held in San Diego, CA, March 12-16, 2012

Milestones:

	Quarterly Milestones and Deliverables
Q1	Validate and optimize field-driven material stiffness for robotic implementation. Complete the fabrication of the miniature piezo-Barth actuator.
Q2	Implement membrane dielectric elastomers on a robotic platform. Start testing the performance of the miniature piezo-Barth actuator.
Q3	Corroborate the model of scaling the developed miniature piezo-Barth actuator, and fine-tune the analytical model towards optimized design. Characterize thermal and fatigue affects on electro-mechanical performance.
Q4	Test and characterize the performance of smart-material mechanisms on robotic platforms, and investigate alternative designs for larger robotic platforms. Implement the developed miniature piezo-Barth motor into the selected robotic application, and perform tests.

Related Research:
We work closely with other tasks to identify requirements and jointly formulate control algorithms for the technologies we develop. We work with Task M2: Principles of Generalized Grasp Mechanics to develop variable compliance, under actuated grippers, and with Task M8: Dynamic Multi-modal and Reconfigurable Mechanisms to develop multi-function limbs for running and climbing. In the long term, the actuators and materials can be used to develop larger manipulators with fingers/toes for legipulation as well as user interface gloves that change stiffness for tactile and force feedback.

In the future, we will also integrate pressure sensors being developed by Task P1: Exploiting Novel Sensor Phenomenology into gripper fingers and appendages for tactile sensing, and we will exploit the strategic alliance between GDRS and 3M to manufacture custom materials to advance this research.

Integration and Assessment Activity:
For each material investigated, we will quantify the relevant materials' properties and their response to stimulation, and we will compare their capabilities to standard actuation techniques. We will also assess difficulties in and progress toward integration into robotic structures.

Subtask 1: Electro-Adaptive Structures for Legged Robots (FSU)
Electric field control structure stiffness will be developed for changing the stiffness of robotic legs for enhanced mobility and manipulation.

Objective and Benefits:
What are you trying to do?
The objective of this effort is to achieve large and fast stiffness changes of adaptive structures that will ultimately enhance the mobility of legged robots. This will be achieved using electric field controlled stiffness in contrast to thermal approaches (e.g., shape memory polymers).

How does the effort fit into the overall program vision, particularly as articulated in the 2011 APP and the Capstone Assessment Vision document? What particular capability(ies)of the Capstone Assessment will this research enable?
Legged robots afford new capabilities in mobility over complex terrains not easily accessible to wheeled robots. However, stable and agile motion introduces challenges in adapting to these complex terrains (running as well as transition to climbing). The proposed field-driven stiffness control provides a route towards rapid adaptive structures that can overcome this limitation.

If this is a multi-year effort, how long do you expect it to continue?
We expect this task to continue over the course of the project. This is partially based on conversations with ARL materials researchers who are interested in testing low viscoelastic dielectric elastomer materials on our robotic platforms to achieve higher efficiency and faster response. In addition to dielectric elastomers, Oates is involved with glassy azobenzene liquid crystal polymer networks for polarized light control of polymer deformation through a NSF

CAREER award. Opportunities exist to exploit these materials in future robotic manipulators. Similar materials are under investigation within Lenhart's group which will be explored as this program matures.

State of the Art:
Prior work in electroactive elastomer membranes has focused primarily on theoretical balance laws, constitutive relations, and experimental characterization at room temperature. This has included a detailed description of nonlinear continuum mechanics and field theory necessary to describe the complex constitutive behavior governing the material under moderate to large deformation and electric fields. Experimentally, these materials have been demonstrated as actuators as step-and-repeat systems and some work on capacitive sensing for feedback control. More recently, field-driven stiffness control has been considered by Mark Cutkosky's group at Stanford with an emphasis on linearized reduced-order modeling for future robotic control.

Key Barriers and Approach:
What is the essence of your tech approach?
The essence of our technical approach is to quantify the electro-*thermo*-mechanical performance of soft dielectric materials and structures using modeling, experimental material characterization, and structure testing to develop electric field controlled properties changes (i.e., stiffness and potentially damping control).

How does your approach move beyond the state of the art?
This work will build upon prior theoretical and computational electro-mechanics research to exploit large deformation mechanics necessary in predicting dynamic motion on a robotic platform. Moreover, thermal coupling will be theoretically, computationally, and experimentally characterized as the stiffness may significantly change over the typical temperature range experienced in the field. Lastly, we will develop new structure designs that will be critical for larger scale robots relevant to the RCTA – diaphragm designs have been identified to be limited in stiffness to small robots (<1 kg).

Why do you think you will succeed?
We have extensive research experience in modeling and experimental solid mechanics of electroactive materials. This includes nonlinear field theory and continuum mechanics, experimental materials characterization, and adaptive structure implementation on other platforms. For example, Oates has worked at Eglin AFRL for the past three summers through a Summer Faculty Award Program where these materials have been implemented for flight control of micro air vehicles. Significant control authority has been achieved through wind tunnel testing. We have now developed and calibrated a thermal chamber capable of electro-mechanical measurements with a temperature range of up to 90°C. This will be used to test the limits of these materials in high temperature environments.

How does this tie into other research activities?
Oates' group is working with Clark on implementing this material on legged robotic platforms to test the feasibility for rapid field-driven stiffness control and, ultimately, enhanced mobility. We

are also working with Dr. Joe Lenhart's group on characterization and implementation of new polymer formulations that exhibit enhanced electrode performance and reduced damping.

Metrics for Evaluation:
How are you measuring your progress? What specific metrics are you using?
We are quantifying the amount of stiffness change driven by a field that can be implemented on a robotic platform in Dr. Jonathan Clark's lab. This is based on a given desired stiffness change needed to transition from running to climbing. We are also monitoring and measuring the amount of displacement needed for a given robotic kinematic motion which requires special adaptive structure designs. Lastly, we measure the transient stiffness changes and power necessary for this process. This will be used to determine how many strides a robot may undergo before complete stiffness change.

What progress has been made toward achieving your ultimate goal?
We have developed and modeled membrane materials that can undergo large electro-mechanical deformation and have experimentally demonstrated significant stiffness changes well above the requirements for the given robotic platform. A thermal, electro-mechanical testing chamber has also been developed for characterizing thermal effects that may influence the constitutive behavior and structural dynamic performance of these devices.

How do you plan to measure intermediate progress and determine whether the research is on track to enable the capstone capability?
We intend to implement these adaptive structures on a robotic platform in Clark's lab and directly measure key performance attributes necessary to achieve the desired mobility enhancement.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Validate and optimize field-driven material stiffness for robotic implementation.
Q2	Implement the membrane dielectric elastomers on a robotic platform. Upgrade to larger scale robotic platforms.
Q3	Characterize thermal effects on electro-mechanical performance. Quantify electro-mechanical performance of flexible electrodes developed at ARL.
Q4	Test and model dielectric elastomer structure designs applicable to larger robotic platforms.

Subtask 4: Next Generation Actuators and Materials (JPL)

Objective and Benefits:
The objective of this subtask is to develop effective actuators and sensors that can enable robotic locomotion and manipulation capabilities that are superior to current technology. For this purpose, we will continue investigating and producing novel displacement leveraging techniques and use piezo-stacks to develop effective rotary and linear actuators. To optimize the design of effective actuators and sensors configurations, we are analytically modeling and simulating the configurations using a multiphysics finite element-modeling package that models physical

phenomena at multiple time and space scales. Complementing the developed technologies, we will also investigate the use of electroactive polymers while coordinating the efforts with those of the other subtasks co-PIs as well as the activity in labs worldwide using the international role that the JPL Principal Investigator has in this field [1].

State of the Art:
Actuators based on smart materials are being developed towards revolutionizing robotic systems by improving the strength-to-weight ratio, speed, range of motion, compactness, efficiency, controllability, and reliability of manipulators, legs, sensors, and other robotic components. The most promising technologies include ferroelectric materials, magnetostrictive compounds, shape memory alloys, and electroactive polymers. Each of these materials offers actuation capabilities that are faster and lighter than conventional actuators. However, these materials have constraints that require special attention during the design process. For example, piezoelectric actuators offer high power densities but have very small displacements, resulting in the need for effective mechanical amplification techniques. On the other hand, while electroactive polymers are emerging with high flexibility, high fracture toughness, and large displacements that cannot be matched by other actuators, the current ones generate low forces.

Key Barriers and Approach:
The lack of knowledge of the underlying constitutive behavior of electroactive materials at large drive levels makes it difficult to design smart structures with desirable performance characteristics. These materials also have constraints that require special attention during the design process. For example, piezoelectric actuators offer high power densities but have very small displacements, resulting in the need for effective mechanical amplification techniques. On the other hand, while electroactive polymers are emerging with high flexibility, high fracture toughness, and large displacements that cannot be matched by other actuators, they are still limited to generating low-level forces. To overcome these limitations we are investigating synergistic actuator and sensor material systems with multifunctional characteristics and employing new methodologies for adaptive structural control. We are investigating next generation actuators for enhanced manipulation and mobility with improved performance and power density, as well as integrated sensing, that eliminate the need for DC motors and complex gear and transmission designs. Specifically, we are investigating shape-changing actuation, simultaneous sensing, and real-time material property control for dynamic adaptation and superior maneuverability.

In the first two years of this subtask, we developed novel actuators that leverage the induced forces from piezoelectric stacks, and we filed a patent [2,3,4]. Once we demonstrate the capability of the miniature piezo-Barth actuator that we developed, we intend to integrate it into a robotic mechanism that is part of this robotics program. For this purpose, we will work with our team partners from Florida State Univeristy to choose the most effective robotic demonstration. We have accomplished the following thus far:

- We established the parameters for the design of the testbed for testing developed novel actuators and produced the first breadboard.
- We designed and started fabricating key components of promising actuation mechanisms.

- We investigated horn designs, flextensional configurations, and ratcheting methods to identify piezo-actuation options of novel displacement leveraging for effective solid state actuators.

Metrics for Evaluation:
Novel designs that employ electroactive materials as miniature, lightweight actuators that generate large forces and torques with large displacement using low power are being developed and demonstrated to surpass the performance of existing actuators while using compact, low mass, and low power mechanisms. The performance of the developed actuators will be compared with conventional servo actuators to demonstrate performance significant enhancement of the capability in torque, force density, electro-mechanical conversion efficiency as well as the displacement of electroactive material actuators over conventional motors of similar mass. The developed actuators are documented in New Technology Reports towards potential filing of Patents, and we are publishing our results in Proceedings of Symposia such as the SPIE Smart Structure and NDE Symposium. We are also planning to publish completed results in peer-reviewed journals.

We have already developed a novel actuator design [4] for which we filed a patent, and we are making an order of magnitude smaller version that we expected to demonstrate soon. The developed actuator has a superior advantage of making multi-dimension motors that are integrated into the robotic structure, allowing multi-functionality, low mass while operating as efficient low power actuators. We are documenting the performance of our developed novel actuators and comparing the results with the prior capability that we recorded as well as the capability of existing technology.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Complete the fabrication of the miniature piezo-Barth actuator. Develop a testbed to enable testing the miniature actuator. Continue investigating novel configurations of high-performance actuators that are lightweight, compact, and highly power efficient.
Q2	Start testing the performance of the miniature piezo-Barth actuator. Investigate other novel configurations of effective piezoelectric actuation horn designs and flextensional configurations with ratcheting.
Q3	Corroborate the model of scaling the developed miniature piezo-Barth actuator, and fine-tune the analytical model towards optimized design. Jointly with our with our task team partners from the University of Florida, we will investigate potential robotic application for implementation of the developed motor. Continue investigating other novel configurations of effective piezoelectric actuation horn designs and flextensional configurations with ratcheting.
Q4	Implement the developed miniature piezo-Barth motor into the selected robotic application, and perform tests. Continue to explore leveraging techniques for the development of compact high-torque density motors as well as explore possibilities of using EAP actuators. Fabricate a novel piezo-actuator that operates at extreme environments using acoustic-mechanical feedthrough mechanism.

Subtask 6: Novel Actuation for Variable Impedance Mechanisms (Stanford)
Develop electro-active polymer-based variable impedance mechanisms that can be utilized on robotic platforms. These devices can be actuated by simple electrical fields and are capable of rapid variation of passive properties.

Objective and Benefits:
What are you trying to do?
Develop electro-active polymer-based variable impedance mechanisms that can be utilized on dynamic robotic platforms. These devices can be actuated by varying an applied voltage and are capable of rapid variation of passive properties.

How does the effort fit into the overall program vision?
Efficient dynamic locomotion over varying terrain requires the alteration of dynamic properties as environmental conditions change. A second goal is to modify transient behavior, for example, to land stably at the end of a fall or jump. To implement these behaviors without the addition of undue mass, compact and light mechanisms and materials are essential. Improving the range of domains that dynamical robots can achieve is an essential feature for RCTA outdoor scenarios.

If this is a multi-year effort, how long do you expect it to continue?
The integration of non-standard actuation schemes into robotic platforms has taken many years, and the development of robust versions of these mechanisms is expected to take many more.

State of the Art:

What is the essence of your tech approach?

Utilize the dielectric properties of VHB and other electro-active polymers to rapidly vary the impedance properties (stiffness and energy dissipation) of small discrete elements including stacks of deformable membranes.

How does your approach move beyond the state of the art?

A large number of researchers have looked at using EAPs as artificial muscles or actuators to replace DC motors for robotic applications. However, it remains difficult for EAPs to compete with electromagnetic motors, which have been under development for over a century, as prime movers. In contrast, this effort employs the active material properties to modulate the inherent passive properties to alter the resulting dynamics while minimizing energy consumption and size.

Why do you think you will succeed?

We have recently developed variable impedance modules that have controlled deflection and robustness properties that exceed the state of the art. We have also developed new high-voltage switching circuits and fabrication techniques that allow for faster changes in material properties, reduced mass, and more robust specimens.

How does this tie into other research activities?

This task builds directly upon developments made on Subtask M9-1. The results of this effort will feed into Subtasks M8-6 and M8-1, the development of multi-modal dynamical robots.

Metrics for Evaluation:

How are you measuring your progress?

Through empirical tests of performance and through integration of successful designs onto test platforms and robotic platforms.

What specific metrics are you using?

Performance metrics include: maximum deflection, change in stiffness (range and response time), change in damping (range and response time), life (cycles before failure), and manufacturing yield (number of robustly functioning units per batch).

What progress has been made toward achieving your ultimate goal?

Individual membrane elements have been made and tested. Four-fold variations in stiffness and damping have been demonstrated [5,6].

How do you plan to measure intermediate progress?

We plan to measure the rate at which these devices improve performance and are adopted onto robotic platforms.

	Subtask 6 Quarterly Milestones and Deliverables
Q1	Implement new batch manufacturing process; test multi-actuator stacks in benchtop tests. In parallel, compare predicted behavior with results of nonlinear finite element model to provide guidelines for future designs.
Q2	Test new actuators on robot hopper, and measure variations in periodic and transient response.
Q3	Refine manufacturing process; quantify yield and robustness. Publish manufacturing recipe and design specifications for actuator.
Q4	Test multi-actuator elements in robotic hopper or other application (to be determined through discussion with FSU).

Subtask 7: Engineering Large Displacement Actuators for Enabling Biologically Inspired Modes of Mobility on Small Robotic Platforms (ARL)

Objective and Benefits:
The basic thrust of this research task is to create enabling technologies for the practical implementation of solid state soft material-based actuation and sensing for control of centimeter to meter scale robotic systems. We are primarily focusing on overcoming the current limitations of dielectric elastomer actuator (DEA) systems, which operate essentially as high-voltage compliant capacitors. This actuation technology is attractive for biologically inspired robotic systems since its theoretical performance is similar to biological muscle: it typically can change its length by up to 30% actively and tolerate much higher deformations, is theoretically capable of similar force outputs as biological muscle, and is capable of providing a DC/steady-state response as well as large displacement actuation up to 10Hz. If large displacements are not required, DEAs are capable of actuation exceeding 1kHz. As a secondary actuator, it holds promise for mechanisms with tunable mechanical impedance. In addition to actuation, the same device can be employed as a sensor for self-sensing feedback by switching the duty cycle between sensing and actuation modes. Such actuation capabilities will enable expansion far beyond the current modes of mobility achievable with electric motor-based actuation schemes.

State of the Art:
Current state of the art for soft polymeric electric field-based actuation is primarily at the lab bench-top demonstration stage of development. There is one known commercial company (Artificial Muscle, Inc.) that has a commercialized product based on the dielectric elastomer actuator technology. In the strain ranges of interest under this project, 10% and greater, stretchable sensor packages are still in the lab demo phase as well.

Key Barriers and Approach:
There are two key barriers to implementing dielectric elastomer technology on a small-scale robotic platform that we are focused on addressing the following two principal areas:

Understanding basic physics behind actuator failure modes. In order to maximize efficiency and force output of dielectric elastomer actuators, operation very near to dielectric breakdown of the dielectric material is desirable. When operating close to the dielectric breakdown point, small

spatially distributed variations in material properties, such as dielectric constant and stiffness, and in geometric parameters, such as dielectric thickness and micro-voids, can lead to large variations in electric field strength. These spatial variations in electric field strength can manifest as electro-mechanical instabilities, such as pull-in instability or dielectric breakdown of the material. Such instabilities lead to device failures. Currently, we do not have an adequate understanding of these phenomena to develop detection and mitigation strategies to prevent catastrophic failure during device operation near dielectric breakdown. Our research is focused on experimentally characterizing and developing predictive modeling of these failures which will be used to feed back into the actuator design process in order to maximize device efficiency and performance.

Conductor flexibility and fatigue tolerance. DEA technology relies on introducing large temporal shifts in electric field strength, which requires a decent conductor that can also stretch and maintain conductivity while subject to the same deformations as the actuator. The stretchable conductor must be able to survive deformation cycling without fracture. It is also important that the stretchable conductor be of similar or lower stiffness than the dielectric material so as not to hinder the displacement of the actuator. There is currently no available COTS material that meets all of these requirements. We are working closely with materials scientists and polymer chemists to develop specifications for the stretchable conductor, to engineer a material to meet these developed specifications, and to characterize and validate the performance of the engineered material.

Metrics for Evaluation:

Key performance metrics include:

1) Reduction in operating power of 20% which will primarily be achieved through improved conductive materials.
2) Reduction of stretchable conductor stiffness to an order of 1 MPa.
3) Development of a method for predicting device failures by utilizing *in situ* measurements of spatio-temporal electric field abnormalities and current-voltage hysteresis curves.

	Subtask 7 Quarterly Milestones and Deliverables
Q1	Complete experimental technique validation. Validation report.
Q2	Acquire data for assessing performance of actuators constructed from COTS materials. Performance data.
Q3	Acquire data for engineered stretchable conductor material. Performance data.
Q4	Establish performance criteria and form factor for the technology demonstrator. Technology demonstrator design and method of manufacture.

References:

C. Jean-Mistral, A. Sylvestre, S. Basrour, and J.-J. Chaillout, “Dielectric Properties of Polyacrylate Thick Films Used in Sensors and Actuators,” *Smart Materials and Structures*, **19** (2010) 075019 (9pp).

Mrozek, R.A., *et al*, "Highly conductive, melt processable polymer composites based on nickel and low melting eutectic metal," *Polymer*, 51, 2954-2958, 2010.

Slipher, G.A. and Hubbard, J. E. Jr., "Exploitation of Higher Order Membrane Modes for Improved Synthetic Jet Performance," *American Institute of Aeronautics and Astronautics Journal*, Vol. 47, No. 6, June, 2009.

References:

[1] Y. Bar-Cohen (Ed.), "Electroactive Polymer (EAP) Actuators as Artificial Muscles - Reality, Potential and Challenges," 2nd Edition, ISBN 0-8194-5297-1, SPIE Press, Bellingham, WA, Vol. PM136, (March 2004), pp. 1-765

[2] Y. Bar-Cohen (Book Editor and author of 5 chapters out of 20), *Biomimetics: Nature-Based Innovation*, ISBN: 9781439834763, ISBN 10: 1439834768, CRC Press, Taylor & Francis Group, Boca Raton, Florida (Sept. 2011), pp. 1-788.

[3] S. Sherrit, X Bao, M. Badescu, Y. Bar-Cohen, and P. Allen, "Monolithic Flexure Pre-stressed Ultrasonic Horns," NTR Docket No. **47610**, patent was submitted on July 7, 2011.

[4] S. Sherrit, Y. Bar-Cohen, and X. Bao, "Acoustic Mechanical Feed-throughs (AMF): for producing work across a structure," NTR Docket No. **48422**, submitted on September 22, 2011.

[5] S. Dastoor and M. Cutkosky, "Variable Impedance due to Electromechanical Coupling in Electroactive Polymer Actuators," Proceedings: IEEE IROS 2011.

[6] S. Dastoor and M. Cutkosky, "Design of Dielectric Electroactive Polymers for a Compact and Scalable Variable Stiffness Device," submitted to IEEE ICRA 2012.

6. INTEGRATED RESEARCH OVERVIEW

Robotics, by its nature, is a diverse research field that incorporates concepts from many traditional disciplines. While each of the four research areas stands on its own as a scientific field, the tight integration of this research is paramount to achieving breakthroughs in robotics research. Given this, maintaining an integrated focus as our research evolves is crucial to planning, managing, and evaluating the success of the Robotics CTA.

Integrated Research Philosophy

The principal way the Robotics CTA shows progress is through the Integrated Research Assessment (IRA). The target is *Integrated Research* in which more than one RCTA technology working together is necessary to achieve component behaviors of an operationally relevant mission task. Also, the target is *Formal Assessment*, where statistical rigor and sound experimentation practices are preferred over demonstrations that sometimes have neither. The question is not, "Can it do it?" in a specific circumstance, but rather, "How well does it do it?" over a relevant space of circumstances. A key purpose of the IRA is to objectively measure the current capability in light of some operational context and to assist the developers in documenting their progress. An assessment is merely a data point in the development cycle, an opportunity to stress the technology in a system and to identify what it does well and what it could do better. It is not pass/fail. When advancements are made and a follow-on IRA occurs, the cycle continues with the bar a little higher.

This approach to robotics advancement differs from traditional system development against target performance specifications. The RCTA strikes a balance between advances in envisioned systems and advances in the component level technology that would make envisioned systems possible. The IRA provides a snapshot of what a system or subsystem might deliver in an operationally relevant mission task, while encouraging researchers to take certain risks to uncover concepts that may ultimately result in breakthrough technologies paying substantial dividends to the Army. Instead of managing development to meet pre-defined goals that may limit basic research, we assess our research against performance benchmarks to evaluate how well that research stands to enhance or even revolutionize robotics and related disciplines. Assessment allows the RCTA to measure current research progress as well as guide future research direction to enhance our research products' ultimate value to the technical and user communities.

To support field assessment, our Consortium has direct access to some of the most capable and widely deployed robotic platforms as well as experience integrating on platforms from several manufacturers. This again provides flexibility to match appropriate platforms to research and to minimize platform development and maintenance costs. Our existing development programs provide knowledge and access to other Government-furnished platforms, whose access we coordinate with the relevant agency, if their use furthers our research goals.

Underlying this unified approach is our Robotic Framework (RFrame). RFrame is the set of simulation engines, software middleware, interfaces, development tools, and robotic platforms that enables a geographically distributed Alliance to work toward common goals in a single

environment. RFrame is more than an arrangement of standard interfaces; it is a comprehensive suite of tools with which the entire Alliance can perform groundbreaking research in robotics using an environment that facilitates distributed software development.

Integration and Assessment Process

Integration and Assessment (I&A) activities begin at the inception of the research, usually during the annual program planning process. Not all the work in the RCTA will be ready in the near term for an IRA. Some technology under the RCTA is too immature to be integrated. Other assessment methods may be used to verify progress on individual technology capability in accomplishing lower-level tasks. However, technology integration for an IRA should be the goal for every task in the RCTA. The more rigorous we make the task-level assessment, the smoother the transition to an IRA. The I&A Team works with Principal Investigators (PIs) to ensure definition of a notional plan for experimentation and measurement. Since our process operates at multiple levels, this could include one of three approaches, depending on the nature and maturity of the research: *Task-based Assessment, Data Collection,* or *Integrated Research Assessment.* At each level, it is important to note that the integrator has many avenues at his disposal to facilitate integration and assessment of the research, ranging from data collection and analysis tools, simulation platforms, and robotics test facilities owned by the Consortium or the Government.

Task-based Assessment. At the early stages, we assess research tasks as standalone entities without pairing them with other research. This is similar to unit testing in a system engineering process and intends to demonstrate the utility of the research at the task level. The goal of this type of assessment is often to collect data that either directly exercises the research or assists the Principal Investigator to further his research. We conduct numerous independent assessments each year, usually in simulation or laboratory environments.

Data Collection Approach. Like many independent assessments of basic or applied research**,** the case above involves data collection. When such datasets feed machine-learning algorithms, we separate data into training and test sets so researchers can properly assess algorithms against the test set. Other data collections may involve the sequestering of a portion of data to challenge researchers to solve problems in the general case. We have adopted the standardized HDF5 file format for collection of multiple different types of data into a single file. Open-source readers and writers of this format are available and are used internally by the Matlab calculation package.

Integrated Research Assessment. Integrated research assessment (IRA) occurs at the stage where research starts to show its utility as part of a larger robotic entity. During this stage, the I&A Team works with multiple Principal Investigators to pair complementary research elements to evaluate the value of the combined research. This is the stage where the true meaning of *integrated research* becomes apparent. Because we experience the boundaries and interactions among those areas in the most practical sense, this stage is the most technically challenging aspect of performing integrated research, and this is the focus of most of our integration and assessment activities. The scientific value at this stage is profound: not only do we start to validate system architectures, but we also start to show that complementary research can lead to robotic capabilities that are greater than the sum of their parts.

For technology that is ready for an IRA, devising a good assessment can be challenging. It requires open communication between those overseeing integration and assessment and developers whose work is being assessed. No one knows the technology as well as the researcher. It is important for the I&A Team to know aspects of the technology that are mature and reliable as well as those that are still developing. The I&A Team has been tasked with staying current on operationally relevant tasks that would exercise the technology consistent with the RCTA vision. They can also marshal resources, design the experiment, and conduct or participate in the analysis.

The key to a successful IRA is an open dialogue. To help facilitate the dialogue, we have prepared a start-up guide with the type of information we generally need at the beginning of the IRA planning process. The guide encourages researchers whose technology is ready for an IRA to envision a notional experiment and consider questions pertaining to background and motivation, details of the physical set-up, test protocol, and measurement.

Motivation

1. What is the objective of the experiment? What are the core capabilities you would like to show?
2. What military capability could this task enable or support? Answer in terms of the RCTA think-look-move-talk-work spectrum.
3. What past data has been collected, including simulation, and what was the performance of the technology either integrated or singularly? Please provide documentation.
4. Within the experimental region proposed, what do you expect the result to be? What's easy? What's hard?

Physical Set-Up

5. What are the details of the physical set-up? What is the physical environment in which a trial/run would occur? What is the human landscape within the physical environment?
6. Does the experiment require a human operator? If so, technician or subject?
7. Does the experiment require more than one robot? If so, how many?
8. Does the experiment require independent ground truth?
9. Do you have a specific facility or facility features (e.g., CACTF, stairs) in mind?

Protocol

10. Describe a trial/run. What would it look like?
11. Are there software flags/settings that would have to be adjusted or physical changes to the platform that would be required as the test conditions are varied?
12. What physical test conditions relevant to the outcome would you vary?
13. What physical test conditions relevant to the outcome would you control?
14. What physical test conditions relevant to the outcome are not controllable? Can they be measured?

15. Is there a natural progression of stages in the experiment, perhaps defined by different behaviors to be explored?
16. For two or more runs, what steps would have to occur to successfully move within the run schedule from run to run?

Measurement

17. What responses would be measured? What is the metric for performance?
18. What is the anticipated complexity of the relationship between experimental variables and the response variables? Would you expect an interaction between variables? (For example, the effect of increasing incline is greater over a hard flat surface than over a vegetated surface.)
19. How can data be captured for post processing and analysis? (The intent is to capture run data that will both allow performance to be assessed and provide insight on how the run developed.)

20. Are there other special considerations not addressed?

We plan to conduct two integrated research assessments in the coming year. The first will explore advances made in autonomous mobility research since the baseline experiment in 2010. This will combine research in semantic perception, behavior generation, and human-robot interaction to advance the state of the art in intelligent navigation. The second IRA will occur in April 2013, the fourth quarter of the program. This assessment will combine technologies from perception, intelligence, manipulation, and mobility to determine the improvement in capability over the baseline DMUM assessment occurring in April 2012.

Taken as a whole, this tiered approach provides a mechanism to build and validate new concepts in robotics, starting with basic research and integrating it to form advanced capabilities. This enhances the probability for successful research and forms a powerful case for transition of the technology.

Lessons Learned. Our experience in this area emphasizes the importance of involving the research PIs throughout the process, not only by contributing technologies but also by understanding test sites, experimental designs, and insights into the Army's needs. We recognize the difficulties inherent in working at multiple sites and on multiple platform types. Our I&A Team is distributed among the industrial members, who are experts in the platforms used. This group expects to travel frequently to Government and academic sites in support of integration efforts. Some sites may not be able to support a platform for integration due to space, security, or other constraints. In these cases, we support the equipment in a local industrial facility, such as the GDRS or QinetiQ facilities in Pittsburgh or GDRS facilities in Westminster, Maryland. Finally, we have received great value by having independent institutions such as NIST participate in our assessment process, providing unbiased assessments, experimental designs, technical reports, and suggestions.

Integration and Assessment Team

I&A is a collaborative effort in that it involves the Principal Investigator, supporting researchers, program management, ARL, other Army decision makers, and a dedicated I&A Team. Within the I&A Team, we identify overlapping subgroups: research integrators and support staff, the assessment team, and I&A management.

Our research integrators are experts in design of the mechanical, electrical, and software systems needed for robotic operations. Our experienced technicians understand how to deploy the many systems required, including computer and operating system setup, software module build and installation, configuration of a myriad of electronic devices such as gyroscopic navigation sensors, GPS receivers, wireless communications, sensor alignment and calibration, power electronics, and control systems for small and large robots. This expertise frees the engineering staff to address more complex issues and substantially reduces the cost of the integration effort.

Task Breakdown of Integrated Research

Integration, Experimental Design, and Assessment. Building on our regular integrated research assessments, we stage scenario-based assessments of integrated Robotics CTA research that build upon more frequent integrated research assessments. Complementing the research assessments that measure the performance of the research, the scenario-based assessments are a significant forcing function for our applied research program in two respects. First, they drive integrated performance of the research, which is paramount to a diverse field like robotics. Second, they ensure the relevance of the applied research to the Army's vision, providing results that quantify the *effectiveness* of the research to transitional stakeholders. In addition to the research itself, an important output of these integrated assessments is a publishable and scientifically defensible capability report to highlight the results of the experiment. This task includes 6.1 elements, such as data collections and experimental design support of 6.1 research, as well as 6.2 elements involving the integration and assessment of robotic technologies.

Fielded systems currently use tele-operation for the bulk of their tasks. The Soldier receives limited feedback from cameras on the platform and uses these to complete the task while directly operating the mobility and manipulation systems of the platform. In the Integration task, we examine the effects of enhanced feedback mechanisms for tele-operation by assessing the performance of the Soldier given haptic (tactile) feedback while performing his task.

Platforms and Testbeds. We provide robotic platforms to researchers under the platforms task. In addition to the platforms themselves, we provide computing systems, sensors, and navigation systems for these platforms. In 2012, the navigation systems for the non-ITAR K-bot platforms are a priority, along with support for the GDRS microLADAR and Velodyne's small LADAR system. In addition, we are upgrading the DR20 platforms with a Development Kit (DevKit) computing and communications solution as well as providing maintenance and support for all the RCTA platforms. On the BigDog platform, we will work on transitioning the research from M4 (High Degree-of-Freedom Dynamic Manipulation) from the laboratory to the robot and support assessment of that research.

RIVET. Our research is supported by a simulation environment that allows full hardware-in-the-loop operation of robots. Simulation has a number of benefits, including significant cost savings as duplication of the simulation is easier, fewer resources are required for testing, complex interactions can be verified prior to field exercises, and many other well-known benefits of simulated operation. The RIVET simulation environment will be upgraded with physics-based rendering to support the new research areas in the RCTA. Better object interactions are needed for research in manipulation and in mobility on complex, non-rigid terrain. In order to accomplish these things, we are defining the Task IR4-2012.

RFrame Tools. A large number of software tools are needed to support the research in robotics. These include messaging systems such as ROS and RCSLIB, computers and interfaces for interacting with the robots while the natural interfaces are developed, build tools and code repositories, and the Shared Data Environment (SDE). Development and support of these tools are described in the Task IR5-2012.

Integrated Research Tasks

IR1-2012 – Integration, Experimental Design, and Assessment

Level: 6.1 Basic Research

	Name	*Organization*	*Subtask*
Principal Investigator	Brad Stuart	GDRS	N/A

Objective:
Applied research outcomes are integrated and assessed as outlined in the IR2 task description that follows this task. However, that process does not directly address the issue of assessing basic (6.1) research. It is important to have a process that does assess basic research because it provides a basis for better managing the overall research portfolio. Basic research ideally produces capabilities that can be transitioned to the applied level, but the timescale for that transition can vary considerably, and some basic research paths may not prove fruitful. Thus, we briefly describe here our process for monitoring basic research tasks so that the Alliance can make informed decisions each year on which tasks to continue, redirect, discontinue, or transition to 6.2.

Background:
The task of integrating and assessing RCTA fundamental research applies to both the 6.1 and 6.2 levels of research. As quoted in the RCTA Program Announcement, the DoD Financial Management Regulation, Volume 2B, Chapter 5 (June 2006) provides the following definitions of basic and applied research:

> Budget Activity 6.1 – Basic research is systematic study directed toward greater knowledge or understanding of the fundamental aspects of phenomena and of observable facts without specific applications towards processes or products in mind. It includes all scientific study and experimentation directed toward increasing fundamental knowledge and understanding in those fields of the physical, engineering, environmental, and life sciences related to long-term national security needs. It is farsighted high payoff research that provides the basis for technological progress. Basic research may lead to: (a) subsequent applied research and advanced technology developments in Defense-related technologies, and (b) new and improved military functional capabilities in areas such as communications, detection, tracking, surveillance, propulsion, mobility, guidance and control, navigation, energy conversion, materials and structures, and personnel support.
>
> Budget Activity 6.2 – Applied research is systematic study to understand the means to meet a recognized and specific need. It is a systematic expansion and application of knowledge to develop useful materials, devices, and systems or methods. It may be

oriented, ultimately, toward the design, development, and improvement of prototypes and new processes to meet general mission area requirements. Applied research may translate promising basic research into solutions for broadly defined military needs, short of system development. This type of effort may vary from systematic mission-directed research beyond that in Budget Activity 6.1 to sophisticated breadboard hardware, study, programming and planning efforts that establish the initial feasibility and practicality of proposed solutions to technological challenges. It includes studies, investigations, and non-system specific technology efforts. The dominant characteristic is that applied research is directed toward general military needs with a view toward developing and evaluating the feasibility and practicality of proposed solutions and determining their parameters. Applied Research precedes system specific technology investigations or development.

Research Description:

Our basic research assessment process focuses on task- and/or subtask-level assessments of standalone research outcomes. These assessments are typically conducted by the researchers, but they are reported to and monitored by the I&A team. Task and subtask assessments are provided in quarterly reports and at other times as requested. An example of an individual assessment is the measurement of precision/recall performance for object or activity detection on a given dataset. Results of such individual task assessments will help determine which basic research efforts are ready for transition to the applied level. Basic research that is producing demonstrated results beyond the state of the art creates a “push” to be expanded or transitioned. We also define integrated capability goals each year that reflect expected outcomes. Thus, we create a “pull” to set an expectation for research outcomes and form a basis for prioritizing the basic research portfolio.

In addition to (sub)task-level assessments, we can begin to assess integrated capability even at the basic research level. We can do this by taking representative basic research results and feeding them into a simulation or into an integrated applied research subsystem. For example, perception results from basic research could be fed into a real-time planner to assess system effects. This early partial integration can help determine whether to invest in transitioning the research to the applied level so it can be included in a future integrated research assessment. As an example of basic research assessment during 2011, we monitored the testing by UPenn researchers at Ft. Indiantown Gap of a version the RHex robot that had been enhanced with rudimentary sensing and planning capabilities. That testing revealed both impressive capabilities and the need for integration of intelligence and perception from other RCTA research, which has been incorporated into the 2012 APP.

For FY 2012, we are considering several candidate assessments of basic research; one is the use of predictions from shared mental models to alter simulated robot behaviors. Another candidate is the assessed improvement in perception due to contextual guidance from the cognitive component of the world model. We also expect results from the Intelligent Navigation IRA to highlight salient basic research needs that may warrant both standalone testing and early integration.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Task-based assessment support. Planning for DMUM and Intelligence tasks.
Q2	Data collections for Perception tasks using RIVET and real sensors.
Q3	Data analysis tools and methods.
Q4	Data collections for DMUM, Intelligence, and HRI tasks.

IR2-2012 – Integration, Experimental Design, and Assessment

Level: 6.2 Applied Research

	Name	*Organization*	*Subtasks*
Principal Investigator	Brad Stuart	GDRS	IR2-1
Additional Investigators	David Duggins	QNA	IR2-2, IR2-3
	Alberto Lacaze	Robotic Research	IR2-4
	Marshall Childers	ARL	IR2-5
	Barry Bodt	ARL	IR2-5

Objective:
This task is to provide for the assessment of candidate technologies as integrated into a robotic system. It assesses the application of technology generated in other areas of this program as well as technology that is developed outside the program, either research grade or commercial grade. We integrate these technologies on the RCTA platforms or other platforms available to the program. Once integrated, we define metrics that can be used to evaluate the technologies, thresholds for these metrics if applicable, and assessments that will reveal the values of these metrics on the integrated system. We work closely with RCTA researchers and Government partners to ensure that integration is done as well as possible and that assessments are fair to all parties.

Subtask 1: Integrated Research Assessments (GDRS)

Integrated Research Assessment 3 – Intelligent Navigation: Our vision for robotics is one of a partner that can act in concert with its teammates; autonomous in simple actions yet part of a larger team. As part of this goal, we need to have a capability for intelligent navigation; movement through the world that takes account of the various properties of objects as they pertain to movement. This means an understanding beyond “floor” and “blockage;” we need the robot to understand mobility-related features, including walls, doors, buildings, stairs, movable objects, and independent movers such as people and vehicles.

This IRA will assess three different areas pertaining to intelligent navigation as we envision it. First, the robot needs the ability to understand the world through its perception and world modeling systems. This is a combined Intelligence/Perception capability. Second, the robot needs to use this knowledge of the environment and its features to determine where it needs to go and how to get there. Third, the robot needs to actually move to where it needs to be, using its mobility systems and manipulation if necessary. It should be communicating with its teammates

either implicitly or directly so that they have a clear understanding of what it is doing and what it is planning to do. We envision that part of this assessment will include determining the efficacy of such communications.

We anticipate that this assessment will be performed on a TALON-class platform, though it will be using the RCTA Architecture computing system, allowing portions to be developed on other platforms and then ported over as necessary.

Integrated Research Assessment 4 – Autonomous ISR: The second assessment of the year will determine the performance of the RCTA research as a platform for autonomous surveillance and reconnaissance. We will build upon the results of the behavior classification and human detection systems and incorporate the ability to track and follow persons of interest. The robotic system will need to recognize possible paths, including doors, stairs, fire escapes, and the like. It will need to move at human-like speeds to keep up with the natural motion of the target and its teammates. We will integrate and assess systems to keep our teammates up-to-date with the robot and target actions, and inform when the target is lost or when target loss is predicted.

RCTA Architecture for Intelligence and Perception on Small Platforms: This will select appropriate computing hardware and perception sensors to facilitate research and assessment. Currently, we are using the Mac Mini on a large number of our platforms. These computers are several years old now, and better performance is available on the market for a reasonable cost. A newer Mac Mini is one possible selection; however, we will look at other platforms for ruggedness, performance, size, weight, cost, and power requirements.

Online Strategies to Manage Integrated Components: In addition to the computing hardware, we will develop the software architecture to support the multitude of systems that are being developed for perception, route planning, manipulation planning, subsystem control, learning, and modeling. This framework module needs to have the knowledge of what the various modules do and control, what they need, and what they provide. Then, it combines the needed modules dynamically to solve the problem at hand. This module is being developed as part of the world model framework and in conjunction with the Intelligence researchers.

Early Preparation for IRA 5 – Physical Interaction: A vital capability that robots need to operate in the world is the ability to manipulate objects for mobility purposes. That is to say, the robot itself decides what objects to move and where to move them, in order to achieve a larger purpose such as clearing a route through a crowded room. Current robotic path planners take the world as given and not subject to change by the robot; either the world is static and immovable, or it includes movers on particular paths but who cannot be influenced by the robot.

This assessment will provide an opportunity for various groups of researchers to explore alternative approaches to the problems of perception, intelligence, and controlled full-body manipulation of the environment. We envision a robot the size of a Dragon Runner navigating through a cluttered environment, pushing and pulling objects out of the way to clear its path. Tasks to prepare for the assessment include decisions regarding the base platform, perception systems to be made available, and the basic outline of the assessment activities. Making these

available to the various researchers will enable them to pursue different approaches and provide a level field for assessing the different technologies that are developed.

Background:
While task-level assessment is necessary to maturing the technology and a critical element of our approach, the integrated nature of robotics calls for higher levels of assessment in which one pairs complementary research and existing technology to form an integrated testbed. At this level, quantitative assessment of the research can be performed on the integrated technology, and it provides scientific value of its own. For instance, integrated research helps discover the somewhat blurry boundaries and interactions between perception, intelligence, platform and manipulator control, and human interaction. This is important to subsequent technology development efforts as it strengthens and validates system architectures and identifies boundary issues which are typically the most difficult issues in development efforts. Integrated research also helps identify applications for the research, which sometimes differ from those identified at the outset.

Research Description:
Integrated Research Assessments: Assessing technologies in the robotics context serves a variety of purposes. It provides a goal for technology producers to make their systems reliable and real-time so that they may function in real-world settings. It documents the readiness of robotic capabilities for transition to other programs, such as 6.3 development programs. It highlights new goals for robotics to tackle as weaknesses of the current approaches are revealed. Also, it stimulates the imagination for new applications of the robotic technology that we are developing so that it can provide tangible benefit to our forces.

RCTA Capstone Integration and Assessment Vision: Now we describe an integration/assessment vision for what the RCTA program could perform about three years from now, i.e., late 2014. This would be a “capstone” assessment that spans the new autonomy capabilities the program will have created by that time. This vision has been discussed throughout the RCTA and assesses research elements from all areas of the program.

The vision is decomposed into constituent capabilities, and each capability is associated with one or more of the “think-look-move-talk-work” items (a). Each constituent capability is also identified with one or more of the technical foundations of autonomy (b). Also, a short narrative describing the capability is given, indicating some of the complex challenges that might arise (c). A short description is given for how such a capability might be assessed (d). Finally, other types of missions that require similar capabilities are listed (e).

This is still far from the descriptive assessment plan we will produce together.; that document will explain not only more details of the capstone assessment but also the incremental steps we will take in research, integration, and assessment over the next several years to achieve our shared vision.

A Notional Capstone Assessment Vision: "Cover the back door!"
This section describes one instance where a set of general capabilities is required for an autonomous robot to collaborate with Soldiers to perform elements of a mission. Other examples would occur in other mission contexts:

- Counter-IED: the EOD robot is to go downrange to investigate a rubble pile on the side of a road next to a large tree, look for objects of interest, send back information about them, and possibly disarm or detonate the device.
- Urban ISR patrol: the reconnaissance robot is to check an upcoming intersection for a platoon that is advancing through an urban environment and warn the platoon of suspected danger with backup information, such as annotated photos or video.
- Logistics for dismounted Soldiers: the logistics robot is to retrieve ammo and water supplies and carry them to a platoon inside a nearby village.

Thus, it is useful to think of the following description as an iconic example, not a prescription for an assessment nor a call for a point solution to this particular set of problems.

The narrative begins with a team of perhaps four Soldiers and a robot moving in an urban environment. An event occurs such that the unit leader believes an insurgent is in a certain building they are facing. He believes the insurgent is likely to escape out the back but doesn't want to risk splitting-up his already small unit. So he wants the robot to move to the back of the building and determine whether someone does actually try to escape that way. The commander instructs the robot something like: "Go to the back of this building, look for people leaving, get imagery of them, report back, call for help…."

This is a very difficult task, far beyond the capabilities of today's robots. We have broken-down the elements of this task into eleven distinct items, ordered roughly chronologically in terms of the narrative description. Following are the capability elements needed for a robot to perform in such a manner:

1. Understand commands

Task Areas – Think; Intelligence and Human-Robot Interaction
Technical Foundations – World Model and Meta-cognition
Narrative – A Soldier in the platoon indicates to a robot that it should "cover the back of the building." The Soldier does this either by a hand signal, a voice command, a text message, or some other effective method. The Soldier has to communicate in some fashion which building – or group of buildings – is being designated. There may also be some direction about what specific activities are of interest, and what actions should be taken upon observing given activity at the back of the building.
Assessment – This needs to include assessment for multiple commands, such as "Check the next intersection", or "Follow that Guy", etc. A possible metric could be the number of distinct terms understood in a simple "verb-descriptor-object" imperative sentence structure.
Other missions – Similar capabilities are needed to command a logistics robot (see Figure 6-1) or direct an EOD robot to assess a suspected area downrange.

Figure 6-1: The RCTA program is performing research to achieve semantic communication with robots for other missions, in this case logistics.

2. Repeat back commands to demonstrate understanding

Task Areas: Talk; Intelligence and HRI

Technical Foundations: Meta-cognition

Narrative – The robot needs to demonstrate its understanding of the command in some reasonable way. It has to manifest having a shared mental model with the Soldier about what they are doing together. This may, for example, be a text message to the Soldier such as "I'm going to cover the back of this building" while it designates the building with a laser. It would probably be good to say, "I will send you a message when I'm in position." The message could be a text or maybe a haptic communication.

Assessment – same comment as for 1. Also, understanding may be demonstrated through answering questions from Soldier. The Soldier should also be asked if the explanation from the robot is clear, and whether the robot correctly understood.

Other missions – Similar capabilities to manifest task and mission understanding are needed for an EOD robot or for a robot providing logistics support to dismounted Soldiers.

3. Understand world through perception and WM
Task Areas – Look; Perception and Intelligence
Technical Foundations – Semantic Perception and WM
Narrative – The robot needs to have been building a semantic world model of the environment in order to make sense of the command from items 1 and 2 – it at least has to know what buildings are nearby. But beyond that, it needs to plan an initial path with an understanding of what it means to get to the back of a building, which may be in an alley, on another street, next to an open area, behind a fence, etc.
Assessment – key to this capability is the semantic labeling of major types of entities in the urban landscape, such as buildings, streets, sidewalks, alleys, doors, stairs, porches, and windows.
Other missions – Similar capabilities are needed to command a logistics robot to associate perceived objects to "tire pallet" and "water" and "ammo" or for an EOD robot to understand "curb" or "rubble pile".

4. Figure out where to go in the world and go there.
Task Areas – (I/P) Move; Intelligence, Perception
Technical Foundations – Adaptive Behavior Generation, Semantic Perception
Narrative – The robot needs to begin moving autonomously through the urban setting with some initial plan. As it goes, it learns new things about the environment, such as the position of stationary vehicles and other urban clutter – road signs, guard rails, dumpsters, gas pumps, traffic cones, garbage cans, etc. It may also encounter movers – vehicles, people, and animals – and needs to move without hitting them or being unduly distracted. Assessment – key to this capability is robust tracking of all movers. This is a great opportunity for comparison to previous assessments where vehicles and/or people were detected and tracked. In particular, we should be able to demonstrate marked improvement in our ability to track movers through occlusions, distinguish vehicles from people from animals, and avoid false detections due to urban clutter.
Other missions – Similar capabilities are needed for an EOD robot to proceed downrange to a suspected IED location or for an ISR robot to move to a suitable observation position.

Figure 6-2: Scene from 2009 RCTA Capstone Experiment. People are moving in a moderately cluttered setting and an autonomous vehicle is detecting and avoiding them. We want to put this capability on smaller bots, while also dealing with substantially higher levels of urban clutter.

5. Move through mobility challenges on the way
   Task Areas – Move; Perception, Intelligence, Unique Mobility
   Technical Foundations – Learning, Adaptive Behavior Generation, WM, Semantic Perception
   On the way to its observation position, the robot may encounter many mobility challenges. For example, it may need to go down an unpaved alley with various kinds of debris (as at the FTIG MOUT site). It may encounter a gate or stairs. It may encounter mud, water, sand, gravel, high grass, etc. Robots with unique mobility capabilities will find vantage points to support situation awareness. This includes climbing poles, moving through underground drains, climbing walls, etc. Robots with dynamic manipulation capabilities will help clear rubble and other obstacles out of the way.
   Assessment – this is a rich area for assessment across many parameters. We can use FTIG, the QNA test site that was recently completed, and probably other facilities to assess how well our systems can learn from experience, recognize mobility challenges, and traverse them successfully.
   Other missions – Similar capabilities are needed across all missions that require significant movement.

6. Watch for people leaving the building
   Task Areas – Look; Perception, Intelligence

Technical Foundations – Semantic Perception, Meta-cognition
Narrative – Now the robot is conducting its primary mission – observing the back of the building. Probably the first thing it needs to do is announce somehow that it is now in position. It also needs to assess what the "back of the building" really means – the underlying question is "what are the points of egress." So it needs to find and label doors, stairs, windows, porches, balconies and any other likely means of escape. And it needs to monitor activity. This could be relatively simple if a single fugitive has entered an unoccupied building. It could be extremely complex – even for a human – if several insurgents have entered a building that many uninvolved civilians are already occupying. Another compounding factor would be the prior presence of people behind the building.
Assessment – the assessments here will depend heavily of the degree of difficulty mentioned above. In the simpler cases, the primary metric is the timely detection of any person leaving the building. In the case of many people being involved, an important metric becomes identifying a particular person as the individual of interest.
Other missions – Similar capabilities are needed by a robot at a checkpoint to watch for people leaving their vehicles or for vehicles leaving the queue.

Figure 6-3: This is a photo of the Fort Indiantown Gap MOUT site (or "CACTF"), showing a back view of the Police Station which is the reddish building in the left foreground. It exemplifies the complexity implied by the phrase "back of the building." In this case, there is a back porch with a double door, as well as several low windows with shutters and a narrow alley to the right of the building. People could emerge from any egress point.

7. Respond appropriately if/when people do leave
   Task Areas – Think, Work, Talk; Intelligence, HRI
   Technical Foundations – Adaptive Behavior Generation, Meta-cognition
   Narrative – Up to now, the robot is only observing what is going on: has anyone left the building from the back or not? That is challenging in itself, but now it has to move to a higher cognitive level of understanding and decide what action to take in response. In the simple case of a single individual emerging from an otherwise empty building, it needs to respond in some useful way. This may be simply sending a message that someone has emerged – this could be via text, or haptics, or some other means. In a more complex situation, it may need to somehow realize that a person emerging is just escaping from

the insurgent, or it may transmit a photo or video snippet for human interpretation. It may track the person and report where he/she goes next.
Assessment – assessments here occur on two levels: one is the correct interpretation of the salience of activity detected in (6), and the other is generation of an appropriate set of responses.
Other missions – Similar capabilities are needed to respond to suspicious activity at a checkpoint.

8. Maintain situation awareness and respond appropriately to other events
   Task Areas – (I/HRI) Work
   Technical Foundations – Adaptive Behavior Generation, WM, Meta-cognition
   Narrative – While performing its primary surveillance mission behind the Police Station, the robot should maintain SA in a broader context. So, for example, if a person emerges from behind the wall of the cemetery in the bottom right of Figure 6-3, it should be aware of that too. Also, it should report if someone is seen in an upper floor window or on the roof. Being able to do all this should have been taken into account when it chose its position from which to monitor the back of the building. In the case of any of these observations, it should take some course of action to alert the unit of salient events.
   Assessment – the metrics here will be both recognition of activity and classification of the salience of that activity.
   Other missions – Similar capabilities are needed across all missions that require maintenance of SA.

9. Interact with user/teammate to learn calibrated importance of observations
   Task Areas –Talk; Intelligence, HRI
   Technical Foundations – Learning, Meta-cognition
   Narrative – Because there a so many observations that could potentially be reported, the robot should have some way to calibrate itself on what to report. If it reports every single moving entity in a complex situation, it will become a distraction and annoyance to the Soldiers in its unit. If it reports only the most obvious things, it may miss something important. Hopefully it can both be programmed as well as trained to make these distinctions.
   Assessment – one assessment could consist of letting the robot report exhaustively on its observations and then get feedback on which were and were not actually salient. We can test on its ability to learn about salience, and to generalize that to other situations.
   Other missions – Similar capabilities are needed across all missions that require significant interaction with human teammates.

10. Enter building by opening door and avoiding or moving obstacles in doorway or building
    Task Areas – Move; DMUM, Intelligence, Perception
    Technical Foundations – Adaptive Behavior Generation, WM, Semantic Perception
    Narrative – A really complex portion of the action can occur when the robot enters the building and begins to interact directly with the environment. There are many possibilities here. The robot may need to open doors, which could be spring-loaded. It may find obstacles such as furniture or rubble in the way, and need to maneuver around

them or move them out of the way. It may need to go up or down steps and maneuver in tight quarters. The robot may be asked to look for and fetch objects of interest such as UXOs, weapons, etc.

Assessment – there are many individual assessments that can be done for the types of capabilities listed above. What types of doors can the robot deal with? How long does it take? How quickly can it ascend or descend stairs? Are these capabilities combined and integrated appropriately into complex sequences of actions? What objects can the robot find? What objects can the robot grasp and fetch? What level of clutter can the robot handle?

Other missions – Similar capabilities are needed across all missions that require physical interaction with the world, such as inspecting a vehicle, removing rocks in a rubble pile that might conceal an IED, or picking up supply items for a logistics mission.

Figure 6-4: Robots may need to enter a building through doors or windows, deal with furniture and other obstacles and ascend or descend stairs in order to execute a complete surveillance mission.

11. Return to user/commander

    Task Areas – Move; Intelligence, Perception

    Technical Foundations – Adaptive Behavior Generation, Semantic Perception

    Narrative – Once the action is complete, the robot needs to return to its unit in an efficient and safe manner. During the mission it may have moved a substantial distance from its unit, or the unit may have moved. It needs to be able to return in a flexible, adaptable manner – it may simply return to where it started, or it may need to rejoin the unit in a new location. It should be able to return to a semantically described location such as "the next intersection beyond the building we just cleared."

Assessment – measure ability of robot to return to the desired location. Is it able to return to its starting point, or go to a semantically described location? Is it ready to resume its previous mission or perform a new one? Does it behave as the Soldiers expect?
Other missions – Similar capabilities are needed across all missions that require significant travel and then return to the unit.

Integrated Research Assessments for 2012
The above section describes a target for an assessment to occur in 2014; for 2012 we plan two assessments that will address portions of this vision. Our third integrated research assessment is planned for October. In this assessment, we will address the capability of the robot to understand the world through perception and world modeling, figure out where to go and go there, and move through mobility challenges on the way; sections 3, 4, and 5 above. Specific RCTA task products that we are assessing come from the Semantic Perception task (P), World Model (I2), Adaptive Behavior Generation (I3/I4), and Learning (I5/I6). These technologies will be integrated onto the DR20 platform. We will develop the needed interfaces and data collection systems to conduct the assessment efficiently. We intend to integrate systems in pieces from June to August, conducting integration testing working closely with the technology developers through the summer. More specifically we will work to the following milestones:

- May: requirements gathering from Semantic Perception, WM, Learning, and Adaptive Behavior investigators. Define site and participants in the IRA, perform travel planning.
- June: complete any needed hardware purchases and begin integrating WM and Semantic Perception.
- July: Begin integrating Learning and Adaptive Behavior modules; provide data collection for Semantic Perception investigators, if needed.
- August: Finish integrating learning and Adaptive Behavior modules; provide data collection for these modules if needed. Define metrics for the assessment and possible approaches to discovering them.
- September: begin testing complete system; apply updates as provided from investigators. Complete planning for the assessment, including list of test runs.
- October: perform assessment as planned; collect and archive data, begin report.
- November: finalize assessment report, including lessons learned about assessment process.

The fourth IRA is planned to occur in February 2013. It revisits the topic of autonomous manipulation that was initially covered in IRA 2 (autonomous trenching). In this IRA, we intend to assess progress in autonomous grasping, human interaction to identify target objects, combined manipulation and mobility, and manipulation for the purpose of mobility. Autonomous grasping involves the use of perception and intelligence software to identify targeted objects, effective grasp locations, and motion sequences to accomplish the grasping and lifting. Human interaction for grasping involves the matching of instructions to objects in the world and actions that can be performed on them. Combined manipulation and mobility can include actions such as door opening and closing, dragging heavy objects, and moving objects with a teammate. Manipulation for mobility is the act of moving objects out of the way so that mobility can continue. Progress toward this assessment is envisioned to be as follows:

- September: requirements gathering from Semantic Perception, HRI, WM, Learning, and DMUM investigators. Define site and participants in the IRA, perform travel planning; select robotic platform.
- October: complete any needed hardware purchases and begin integrating WM and perception and manipulation/mobility results.
- November: Begin integrating HRI and Intelligence modules; provide data collection for Perception investigators, if needed.
- December: Finish integrating HRI and Intelligence modules; provide data collection for these modules if needed. Define metrics for the assessment and possible approaches to discovering them.
- January: begin testing complete system; apply updates as provided from investigators. Complete planning for the assessment, including list of test runs.
- February: perform assessment as planned; collect and archive data, begin report.
- March: finalize assessment report, including lessons learned about assessment process.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations: The RCTA IRA team includes ARL participation at its core. With our Alliance partners and the basic researchers, we design the assessment metrics and randomized trials that will provide insight into the capabilities and weaknesses of each technology. In particular, we are working with the line striper in P2, the manipulation tasks M3 and M5, and the intelligence framework developed in I2. For the autonomous ISR assessment, we will include even more work from perception and intelligence researchers. In addition, this assessment will introduce some of the work from the HRI area.

As much as possible, we will attend the workshops and monthly meetings of the other research areas in RCTA, such as the HRI workshop, the DMUM planning meetings, and the Intelligence workshops.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Framework and detailed description of IRA 5: Autonomous ISR. Definition of component technologies for IRA3; detailed plan for IRA 3. Computing platform for small robots selected and ordered.
Q2	Integration of computing platform and small platform sensing solution. Integrated Research Assessment 3 in August 2012. Initial software framework for resource allocation, integration with sensors.
Q3	Integrated Research Assessment 4: Intelligent Manipulation. Detailed plan. Integration of JPL Stereo system with computing system and software architecture. Pre-test related to IRA 5 demonstrating control of platform to move where commanded.
Q4	Integrated Research Assessment 4. Improved architectural framework delivered, with ability to allocate computing to perception, intelligence/learning, HRI, manipulation, and mobility tasks.

Subtask 2: Integration Support (QNA)

Objective and Benefits:
In this subtask, we are working to support integration of various Consortium technologies on small platforms. This task directly supports the core RCTA concept of integrating technologies developed by researchers across the Consortium into a capability which can then be assessed by the assessment team. We can then iteratively improve the work throughout the life of the program. This work is applicable across a number of tasks; specifically, we are focusing our efforts here to support proprioceptive slip sensing for interactive navigation with vegetation (Task P4), improve manipulator controllability to support autonomous trenching (Task M3), and the design of a 360-degree stereo head, in coordination with JPL, for improved small UGV situational awareness. These efforts support various capabilities identified in the APP; specifically, we look to support Intelligent Navigation (IRA3) and Physical Interaction (IRA5) by supporting Task P4 which focuses on proprioceptive sensing of the environment and interacting with the world. Autonomous Trenching (IRA2: Physical Interaction Baseline and IRA5: Physical Interaction) is supported through an improved manipulator hardware for the DR20, and finally Autonomous ISR (IRA4) functionality is enabled through the design of a 360-degree stereo head, in close coordination with JPL. We are also bringing cutting-edge haptic feedback technology to modern, field-deployed systems, and we are assessing the utility of haptic feedback for mobility and manipulation tasks. In 2012, we plan to extend and improve the leading-edge work already developed and integrated onto the DR20 Hand Controller by detecting and filtering-out the signal due to ego-motion for manipulation tasks. We are transitioning 6.1 technology developed on M1 into a 6.2 task and will assess the overall utility of such a capability. As related to the Capstone Assessment Vision, this technology would apply to items 4, 5, 8, 10, and 11. We anticipate active participation in an assessment near the end of the 2012 program year and hope to identify other potential avenues of applicability that could extend this beyond the 2012 year. We envision this task continuing throughout the life of the program to directly support these and other research tasks that are being conducted throughout the program. Finally, we are working to improve the state of the art in manipulation. We will make use of new high-DOF, high-strength robotic arms, such as the HDT MK1, for algorithmic research and improving the control of manipulators.

State of the Art:
Each of these efforts under IR2-3 work to improve the state of the art for individual research tasks. Deployed small UGVs have little to no proprioceptive sensing, specifically torque sensing for slip. Additionally, deployed manipulators neither support the control mechanisms for adequate closed-loop joint control nor provide the computational resources to coordinate that into higher-level (FK/IK, Force, Torque) control modalities. Perception modalities on small UGVs are also limited, typically a ground-mounted camera or mast-mounted camera that allows a 45-degree field-of-view or smaller. This greatly limits the applicability of autonomous navigation or automated intelligence, surveillance, or reconnaissance. The current state of the art in haptic feedback has been pioneered in the medical domain using the Verro-Touch system. The RCTA program has adapted and extended techniques used for that system and applied it to the small UGV domain. Other haptic feedback systems do exist for manipulation (Harris, etc.), but haptic capability is presently emerging in the DoD space. UPenn and QNA have collaborated to

integrate a low-cost haptic feedback system onto the Dragon Runner 20 platform, which presently defines the state of the art for tele-operated vehicles. Current state of the art for manipulation includes so-called "high-end" manipulators (e.g., HDT) and deployed systems, which have different requirements and, therefore, different capabilities. We hope to bridge the gap between these spaces with the parallel efforts to be completed here.

Key Barriers and Approach:
In 2012, this task will directly support research Task P4 by integrating torque sensors on the wheel motors of a DR20 to better support proprioceptive sensing of the environment for interaction. The torque sensors to be integrated will be specified by the P4 researchers whose plan is to enable better interaction of small UGVs with the environment, specifically vegetation. If the P4 researchers determine that an alternative sensing modality would better suit the interaction task (e.g., a force-sensor as a virtual-bumper), those alternative sensors will be integrated as an appropriate alternative. Present state-of-the-art platforms have a no need to detect slip, as they are 100% tele-operated. This perceptual enhancement will enable higher levels of both supervised (e.g., tele-operated traction control) and autonomous capabilities. This subtask will leverage use of the DR20 Development Kit as developed during the 2011 program year.

This task will also support the development of a high-fidelity manipulator controller for the DR20 manipulator in support of Task M3. In 2011, M3 researchers established that the present controller neither provides adequate computing resources to control the individual joint motors at sufficient rates for closed-loop position control, nor does it provide velocity-control and a method for jointly controlling the joints as is typical in the manipulation space. We will enable these higher-level behaviors and coordination with perception by creating a computing platform suitable for controlling the DR20 manipulator at sufficiently high rates. We envision a two-tiered system, consisting of low-level controllers for each individual joint and a high-level controller to coordinate the joint actions. The high-level controller will also support appropriate I/O to enable inputs from various cameras to support perception-in-the-loop for manipulation. These new capabilities all greatly enhance the state of the art where simple manipulators are controlled on a joint-by-joint basis in open-loop (effort-based) form.

We will coordinate with JPL to design a 360-degree stereo head to support increased situational awareness for the Soldier as well as any autonomous capability that would be used on a small UGV. We will work with JPL to design the appropriate hardware, electronics, and software that can be realized in the future to support the various Integrated Research Assessments. We will incorporate lessons learned from the existing stereo-head (integrated with an OMAP processor) and provide a sufficiently capable system toward Autonomous ISR (IRA4), Physical Interaction (IRA5), and Tactical Team Movement (IRA6). This improves the state of the art by permitting the first-known 360-degree stereo head in a compact, power-efficient platform enabling increased situational awareness for the Soldier. This task notably supports both P4 and P6.

We plan to extend the current haptic feedback system to incorporate the ego-motion detection that was developed under Task M1 at UPenn during the 2011 program year. This will allow the system to filter-out the signal due to self-motion and then convey sensations that are based solely

on the manipulator's contact with the environment to the operator. This filtering will be performed, in real-time, by comparing measured accelerations against models that are developed offline. We will also apply this technique of developing models for filtering-out accelerations due to self-motion of the tele-operated platform in order to classify the terrain being traversed by a mobile platform equipped with the haptic feedback system. RCTA has developed the current state of the art in haptic feedback for military EOD robots. The technical approach described here will extend that SoA and enable the Soldier to have increased situational awareness beyond the current haptic feedback system. We have developed the only known-working haptic system for EOD robots (both manipulation and mobility) and are working with the leaders in the field of haptic feedback (Katherine Kuchenbecker, UPenn) on Task M1. The haptic feedback system provides perceptual clues which are applicable to increased situational awareness for the operator during tele-operation, and the extensions we are pursuing in 2012 will enhance semi-autonomous operations via terrain classification. As such, the work relates to semantic mapping in P4 and interactive navigation in P2. The haptic feedback interface is applicable towards multi-modal control research in the HRI domain as well as a new sensory input to situational awareness through terrain classification.

For our manipulation support work, we plan to provide an application programming interface to the HDT arm to enable the features of the various joints (position control, etc.) as well as integrate an improved forearm section to the DR20 manipulator, leveraging QNA internal research and development.

Metrics for Evaluation:
Progress for each of the individual tasks is measured using a combination of objective metrics and an appropriate level of project management acumen to breakdown larger efforts into smaller tasks. Metrics for these efforts include appropriate design documentation (preliminary/detailed), hardware procurement schedules, and realization of actual hardware integrated with the platforms in question. Progress related to the P4 effort was reflected in the P2 work related to the "interactive navigation" task, which focused on various simulation capabilities for understanding how robots interacted with a wide variety of objects. Progress to date for the M3 supportive effort is reflected in the 2011 progress of the M3 subtask, where a large understanding of how far the current hardware can be pushed and identifying the requirements for a more capable system have been completed. Work on the 360-degree stereo system has been reflected in the JPL task for visual odometry on an OMAP processor.

For haptic feedback, we measure progress using time to complete various tasks (mobility, manipulation) in addition to number of dropped objects (manipulation) for various tasks, as planned for IRA2, which focuses on comparing the baseline (non-haptic) system to the haptic system. In 2012, we will further compare the baseline (non-haptic) system against the unfiltered haptic system and ego-motion enhanced haptic systems. Tasks to perform at the IRA are as follows:

1) *Tele-operated Driving Scenarios:*
    a) Out and Back: Direct the operator to drive the robot to a designated location, over which the surface textures change. Enable the robot to be within line-of-sight at

all times but out of earshot at the maximum travelled distance. Measure time to target, time to return, and closest distance to target.

b) Obstructed View: Same as item 1a, but obstruct a portion or all portions of the route to be traversed.

2) *Pick and Place Manipulation:* Direct the operator to pick-up various objects and place them at a designated location, all within the reach of the manipulator. Measure the number of incorrectly placed or "dropped" objects.
3) *Pick and Place Mobile Manipulation:* Same as item 2, but require the operator to conduct the driving task between the picking and placing operations.
4) *Grip strength on a fragile object:* Direct the operator to pick and place fragile objects where the objects could be damaged by the gripper itself. Measure the number of broken objects with and without the haptic system to determine what improvement there may be.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Specification of torque sensors; manipulator controller specification; operational specification (use cases and system spec.) for 360-degree stereo head. Fabrication of one or more haptic feedback devices. System requirements for ego-motion and terrain classification. Analysis of results from IRA2.
Q2	Torque sensor integration design; manipulator controller preliminary design; preliminary design of 360-degree stereo head. Revision of program plan as appropriate based on IRA2 results. Preliminary design of ego-motion and terrain classification systems. Assembled forearm for DR20 manipulator.
Q3	Torque sensor procurement and installation. Manipulator controller detailed design and fabrication. Hardware and software specification for 360-degree stereo head. Software development of ego-motion and terrain classification systems. Assembled and tested forearm.
Q4	Torque sensor bench and field testing; manipulator controller integration and testing. Detailed design/component selection of 360-degree stereo head. Assessment and results of ego-motion and terrain classification systems as compared to the baseline system and basic haptic feedback system. Integrated forearm with fabricated manipulator controller.

Subtask 3: Haptic Feedback (QNA)

Objective and Benefits:
The objective of this subtask is to bring cutting-edge haptic feedback technology to modern, field-deployed systems and assess the utility of haptic feedback for mobility and manipulation tasks. In 2012, we plan to extend and improve the leading-edge work that was explored in 2011 with the stock DR20 Hand Controller by 1) fabricating a unique haptic interface designed for an optimal user experience (based on the DR20 Hand Controller), 2) detecting and filtering-out the signal due to ego-motion for manipulation tasks, and 3) classifying terrain based on the haptic

feedback signal for tele-operated or autonomous mobility. We are transitioning 6.1 technology developed on M1 into a 6.2 task and will assess the overall utility of such a capability as part of IRA2. As related to the Capstone Assessment Vision, this technology would apply to items 4, 5, 8, 10, and 11. The system also relates to the P4 task for 2012 to focus on proprioceptive detection of and interaction with vegetation. We plan to identify other applications and design improvements as a result of our IRA2 assessment to extend this beyond the 2012 year.

State of the Art:
The current state of the art in haptic feedback has been pioneered in the medical domain using the Verro-Touch system. The RCTA program has adapted and extended techniques used for that system and applied it to the small UGV domain. Other haptic feedback systems do exist for manipulation (Harris, etc.), but haptic capability is presently emerging in the DoD space. As part of the RCTA program, UPenn and QNA have collaborated to integrate a low-cost haptic feedback system onto the Dragon Runner 20 system which, to our knowledge, presently defines the state of the art for tele-operated vehicles.

Key Barriers and Approach:
In the 2011 program year, QNA and UPenn collaborated to integrate the M1 haptic feedback system onto the DR20. This led to a number of iterative engineering evaluations that led us toward conceiving various electromechanical system configurations for effectively communicating haptic feedback to the operator. In 2012, we plan to implement one or more of these candidate designs in hardware to provide the optimal user experience. We also plan to extend the current haptic feedback system to incorporate the ego-motion detection that was developed under Task M1 at UPenn during the 2011 program year. This capability permits the system to filter-out the signal due to self-motion and then convey sensations that are based solely on the manipulator's contact with the environment to the operator. This filtering will be performed, in real-time, by comparing measured accelerations against models that are developed offline. We also plan to apply this technique of developing models for filtering-out accelerations due to self-motion of the tele-operated platform in order to classify the terrain being traversed by a mobile platform equipped with the haptic feedback system. This new capability will extend the current system, one which defines the state of the art, and enable the Soldier to have increased situational awareness beyond a baseline system without haptic feedback or a basic system that communicates all sensed vibrations. We have developed the only known working haptic system for EOD robots (both manipulation and mobility), and are working with the leaders in the field of haptic feedback (Katherine Kuchenbecker, UPenn) on Task M1. The haptic feedback system provides perceptual clues which are applicable to increased situational awareness for the operator during tele-operation, and the extensions we are pursuing in 2012 will enhance semi-autonomous operations via terrain classification. As such, the work relates to semantic mapping and interactive navigation in Task P4. The haptic feedback interface is applicable towards multi-modal control research in the HRI domain as well as a new sensory input to situational awareness for terrain classification.

Metrics for Evaluation:
We measure progress using time to complete various tasks (mobility, manipulation) in addition to number of dropped objects (manipulation) for various tasks, as planned for IRA2, which

focuses on comparing the baseline (non-haptic) system to the haptic system. In 2012, we will further compare the baseline (non-haptic) system against the unfiltered haptic system and ego-motion enhanced haptic systems. Tasks to perform at the IRA and subsequent system improvements are as follows.

1) *Tele-operated Driving Scenarios:*
    a) Out and Back: Direct the operator to drive the robot to a designated location, over which the surface textures change. Enable the robot to be within line-of-sight at all times but out of earshot at the maximum travelled distance. Measure time to target, time to return, and closest distance to target.
    b) Obstructed View: Same as item 1a, but obstruct a portion or all portions of the route to be traversed.
2) *Pick and Place Manipulation:* Direct the operator to pick-up various objects and place them at a designated location, all within the reach of the manipulator. Measure the number of incorrectly placed or "dropped" objects.
3) *Pick and Place Mobile Manipulation:* Same as item 2, but require the operator to conduct the driving task between the picking and placing operations.
4) *Grip strength on a fragile object:* Direct the operator to pick and place fragile objects where the objects could be damaged by the gripper itself. Measure the number of broken objects with and without the haptic system to determine what improvement there may be.
5) *Terrain classification:* Evaluation of the system over various terrains, and the classification system results compared to ground truth.

Our work to date shows the applicability and utility of the haptic feedback system for both tele-operated mobility and tele-operated manipulation in small, researcher-led trials. We are presently designing our third iteration of haptic hand controller hardware with a goal of optimizing the haptic feedback experience for the user. As part of Integrated Research Assessment 2, which at the time of this writing is scheduled for March 2012, we will more fully assess the utility of the technology. We will continue with our methodology of iteratively testing and evaluating on a small scale as we make incremental improvements to the haptic feedback system. We plan to utilize the RCTA Small UGV proving grounds to continually evaluate different materials for the driving task and more structured environments for manipulation tasks. We envision using the rubble pile to see the applicability of the haptic feedback system toward other research areas (e.g., Task M3).

	Subtask 3 Quarterly Milestones and Deliverables
Q1	Fabrication of one or more haptic feedback devices. System requirements for ego-motion and terrain classification. Analysis of results from IRA2.
Q2	Revision of program plan as appropriate based on IRA2 results. Preliminary design of ego-motion and terrain classification systems.
Q3	Software development of ego-motion and terrain classification systems.
Q4	Assessment and results of ego-motion and terrain classification systems as compared to the baseline system and basic haptic feedback system.

Subtask 4: Intelligent Navigation Assessment (Robotic Research)

Objective and Benefits:
The baseline assessment conducted in 2011 showed the capability of small autonomous systems to navigate through an urban environment. The robots performed well at tasks such as obstacle avoidance, maze navigation, and recognition of passable and impassable steps. Difficulties were exposed as well, including navigation among moving people, moving through tall grass and narrow corridors, and the need for operator oversight through complex operator control software.

In 2012 and 2013, we will assist in evaluating the progress made by the RCTA in the field of intelligent navigation. This will include the integration and assessment of novel planning technologies, sensor processing systems, and human interfaces to perform navigation in complex environments.

State of the Art:
Skid-steer platforms are becoming a standard for small robots. Their main advantage is in reduced complexity over Ackermann steering or legged locomotion. The RCTA includes three different skid-steered platforms: the K-bot, TALON, and Dragon Runner. We have developed autonomous navigation, sensing, mapping, and localization solutions for the TALON as part of the RASR package. As an NML-based system, ours is well-suited for integration with RFrame and the rest of the RCTA architecture to support intelligent navigation and human-robot interaction systems.

Key Barriers and Approach:
The current Hokuyo sensing system suffers from a low pixel update rate and, therefore, low numbers of pixels on obstacles. This severely limits the opportunity to classify objects into different types; the data only supports the inference that there is an object at all.

The current computing solution (Mac Mini) is somewhat power hungry for the smaller platforms in our inventory, such as the DR20. We will be porting more of the software into the smaller DevKit computers and looking for low-power (and smaller) computing systems to do the sensor processing and other intensive tasks.

We will work to integrate the intelligence systems developed in the other areas of the RCTA onto these small platforms, including short- and long-range planning, manipulation planning, and planning to change the environment.

	Subtask 4 Quarterly Milestones and Deliverables
Q1	Integration of navigation sensors and software for DR20 and K-bot platforms. Integration of TALON platform with JPL stereo and Velodyne LADAR.
Q2	Integration of TALON systems with GDRS micro-LADAR and RCTA intelligent navigation planners.
Q3	Participation and assistance in Intelligent Navigation IRA.
Q4	Data analysis and reports regarding the Intelligent Navigation IRA.

Subtask 5: Integrated Research Assessments (ARL)

Objective and Benefits:
This subtask will conduct formal technology assessments of integrated research capabilities. It will apply fundamentals of experimental design and applied statistics to enable assessment events that provide quantitative performance metrics and data. This data will be used to track incremental progress of the RCTA, to provide useful feedback to the researchers, and to communicate to stakeholders the integrated performance that was achieved. The experimental design and the data collection at assessment events are cooperative efforts led by the ARL assessment team with help from the RCTA investigators being in the forms of providing iterative feedback on the experimental design as well as data capture and analysis tools.

These integrated technology assessments are preparatory for both researcher and technology to facilitate participation in the Capstone Assessment. Conducting these assessments is a periodic activity that will continue throughout the course of the RCTA. The experimental methods that are refined and/or developed to evaluate and to stress the new capabilities will be required to obtain performance data at the Capstone. This Capstone Assessment data will be useful in evaluating the overall accomplishments of the RCTA.

In FY 2012, we will conduct two integrated research assessments of RCTA technology. In August 2012, we will assess integrated capabilities in intelligent navigation. In October 2012, we will conduct an assessment of technologies that enable autonomous surveillance and reconnaissance.

Key Barriers and Approach:
The purpose of our technology assessment efforts is to conduct quantitative measurement of autonomous unmanned vehicle technology performance. Two critical elements for effective evaluation of the technology are a relevant environment and effective integration of technology on a suitable testbed. Technology assessment is an iterative process that requires a minimum of five steps in order to function properly.

Steps in the Integrated Research Assessment Process:

1. Laboratory testing and parametric simulation performed by researchers.
2. Researchers perform task-level field evaluations with input/guidance from ARL.
3. Formal structured experimentation of integrated capabilities in relevant environments.
4. Detailed data analysis performed by ARL.
5. Provide feedback and recommendations to researchers.

Steps 1 and 2 are threshold activities for participation in an Integrated Research Assessment. Ideally, this is when the technical bugs in software, hardware, interfaces, and data collections are worked-out in preparation for the IRA. During Step 1, the functionality of the technology will be verified with respect to the design intent and the envisioned application. During Step 2, useful metrics should become apparent. Step 3 involves designing the experiment, planning the assessment, and creating a test plan/protocol. In order to design the experiment, a significant amount of dialogue must take place between the ARL assessment team and the researchers to ensure an appropriate assessment plan in keeping with design intent and functionality of the technology. Step 4 requires that raw and/or post processed data from the formal assessment be provided to ARL in a format that is readily ported to data analysis software for further reduction. Step 5 provides specific feedback on performance and recommendations for improvement.

Metrics for Evaluation:
Ultimately, our success will be reflected in the health of the feedback loop to developers and advancement of their research, but an indicator is the number of technology assessments that is conducted and documented. These can be ARL technical reports, journal articles, and conference papers. Each planned technology assessment will have a corresponding report that describes the experimental design to include performance metrics, the assessment event, the reduction and analysis of data, an assessment of performance, and recommendations on further technology development.

	Subtask 5 Quarterly Milestones and Deliverables
Q1	Plan IRA 3: Intelligent Navigation. (The schedule for planned Integrated Research Assessments appears in Figure 1-8.) Report on IRA 2: Behavior Recognition.
Q2	Conduct IRA 3; reduce and analyze data from IRA 3. Plan IRA 4: Autonomous ISR.
Q3	Finish IRA 4 plan; conduct IRA 4; reduce and analyze data from IRA 4. Report on IRA 3.
Q4	Plan IRA 5: Physical Interaction. Report on IRA 4.

IR3-2011 – Platforms and Testbeds for Integrated Research

Level: 6.2 Applied Research

	Name	*Organization*	*Subtask*
Principal Investigator	Mark Del Giorno	GDRS	IR3-1
Additional Investigator	David Duggins	QNA	IR3-2

Objective:
In this task, we provide robotic platforms as a basis for integrated research in all four of the main RCTA research areas. These platforms will include wheeled vehicles, small tracked ground robots, and small legged robots. This task upgrades existing platforms with the necessary computing, sensing, control, communications, and power systems needed to operate autonomously and support RCTA research.

Subtask 1: Robotic Testbeds (GDRS)

Background:
As detailed in our proposal, we have access to a rich variety of robotic platforms to support our research plan, and we pair research tasks with the platform most appropriate for the research. The primary robotic platforms for our integrated research are Dragon Runner, TALON, BigDog, RHex, RiSE, and Escape. This set of platforms supports the vast majority of our integrated research plan. We have chosen proven, mature platforms so our focus is on research, not platform development and maintenance. In 2011, we added the K-bot platform developed by the University of Pennsylvania. This basic platform is not a military robot, so it does not fall under export control restrictions that are difficult for University partners to work around.

Research Description:
TALON. The platform will be provided with upgrades to support research in areas of Intelligence, Perception, and Human-Robot Interaction. These upgrades include a computing system based on the Apple Mac Mini, a low-cost navigation solution based on a small MEMS IMU with GPS, and various cameras or small LADARs as needed. The control system will be upgraded from the current system which relies on the operator for feedback and control to a system that is more controllable and able to perform waypoint navigation indoors or outside.

Dragon Runner. The Dragon Runner family of man-portable robots range from 15 to 45 pounds and support many payloads, including RSTA sensors, CBRNE sensors, weapon illuminators, and manipulators. Dragon Runner provides for fine mobility control on the order of 0.1 cm. A particular advantage of the man-portable Dragon Runner is that the USMC and an allied nation both have purchased dozens of these systems and deployed them in daily use, supporting current

EOD operations in theater. Dragon Runner is the first fielded robot of its class. The DR20 with its manipulator is used in IRA 2 – autonomous trenching and searching for buried wires – and also for research on haptic control of manipulators.

K-bot. We added the new K-bot platform to our menagerie to support the research of university partners who cannot work with the ITAR-controlled platforms such as the TALON and Dragon-Runner. The platform was built by Alex Kushlyev at our partner UPenn; the design was previously used in their entry to the MAGIC 2010 competition. It is a low-cost platform; only $8,000 per unit. In 2012, we will upgrade these platforms with additional sensors, including navigation and perception systems. We will improve its weather resistance and assess its mobility characteristics on rough terrains. We have developed control interfaces that integrate with RFrame so that the same software can control this platform or another RCTA platform. We used the K-bot as an integration platform for IRA1, the classification of pedestrian behaviors.

Computing Platforms
We are largely satisfied with the Apple Mac Mini as a computing platform on the TALON and the K-bot. For smaller robots such as the DR20, DR10, and RHex, a more compact system may be needed. New processors developed for laptops and other mobile applications have become mass-market products and can provide a lot of computing power in a very small package. We are exploring packages like smartphones and tablets that can run Linux and Android systems. These provide the possibility to be used as either on-board computing and display devices or off-board interfaces that can provide bi-directional communication to the robot using voice, text, and graphics.

Sensors
We are working with JPL to bring their OMAP-based stereo camera system to our robots. These systems will provide three-dimensional data to the platform, with the demanding computation of stereo correspondences done on the OMAP chip rather than the CPU or power-hungry GPU system. In addition to this work, we are looking at low-cost alternatives that can further other parts of the research. We have purchased Point Grey Bumblebee stereo camera systems and the Microsoft Kinect structured light range sensor. These have a similar ranging geometry to the JPL stereo system and can be used as cheaper surrogates for that system in less demanding applications. We will develop interfaces to the world model for all these sensors as well as others that may be needed through the course of the year.

The GDRS-developed microLADAR is due to be delivered in early 2012. This laser ranging device is perfectly suited to small platforms; we will be incorporating its data interface into the world model for use by whatever perception systems wish to use it. In addition, we expect delivery of the small Velodyne laser we ordered last year, and we will be including the interfaces to that device as well.

Navigation Systems
We will work with QNA in exploring the capabilities of the latest generation of MEMS-based inertial navigation systems. There have been recent advances in this area, and the devices are getting more accurate and less expensive. We will examine whether any of these devices are

viable replacements for the FOG gyros that we are using. We will work with QNA and Robotic Research to develop a navigation package that is suitable for indoor (GPS-denied) or outdoor use and that requires a minimum of intervention in the calibration process.

Manipulators
We intend to supplement the current DR20 arm that we have purchased for one of our five systems with a second arm. We will work with QNA to determine what additional feedback devices can be provided to assist in the manipulation tasks that are needed. We will meet with the DMUM researchers who intend to use the platform, determine their needs, and create a workable solution. The main complaint with the EOD arms is that they lack feedback in the wrist joint; typically, such feedback is not needed by the operator during tele-operation. However, for autonomous robotic control, we will certainly need such a feedback signal.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
The testbeds team will include members from QNA, GDRS, Robotic Research, and UPenn. We will be collaborating with all the tasks involved in the various IRAs as well as any task that needs to use the K-bot or the DR20 to conduct testing of their systems. We will collaborate with the IRA team to provide the platforms needed to conduct the assessments.

As much as possible, we will attend the workshops and monthly meetings of the other research areas in RCTA, such as the HRI workshop, the DMUM planning meetings, and the Intelligence workshops.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	INS/GPS systems for small platforms; JPL stereo; platforms for IRA 3.
Q2	Interfaces for small LADARs (Velodyne and uLADAR); Kinect.
Q3	Platforms for IRA4; computing systems.
Q4	Improved manipulator system in preparation for future DMUM tasks.

Subtask 2: Testbeds (QNA)

Objective and Benefits:
This subtask provides enabling technology for Consortium members to rapidly integrate their individual Perception, Intelligence, DMUM, and HRI research with field-deployed systems. To date, we have worked to provide interfaces on the Dragon Runner 20 (DR20) and TALON platforms. Here, we bring the latest generation small unmanned ground vehicles – the Dragon Runner 10 (DR10) platform – to the RCTA program. The platforms provided by QNA enable the "think," "move," and "work" capability building blocks applicable to DMUM, Perception, and Intelligence tasks. We will also modify the existing DevKit to support the DR10 platform. These vehicle and computing platforms can permeate all aspects of the Capstone Assessment Vision and are applicable to nearly all areas of research for the RCTA program. We anticipate that use

of these testbeds will occur throughout the life of the program and will necessitate iteratively improving the available hardware and software as the program continues.

State of the Art:
The current state of the art for small UGVs are tele-operated systems that do not readily support open interfaces or the computational hardware to support higher-level autonomous navigation, manipulation, or intelligent behaviors. This subtask is viewed as an enabler for these types of capabilities on small UGVs.

Key Barriers and Approach:
In 2012, we plan to introduce the DR10 as a testbed to the program. The DR10 is representative of the class of micro UGVs (MUGVs) – the first class of UGVs that are truly man-packable. To accommodate the goals of the RCTA program, we will extend the DR20 DevKit to the DR10 platform by adapting the existing hardware and software, creating a payload which will enable higher-level perception, planning, and control methodologies. We also plan to integrate mesh network radios as part of the payload, enabling more effective team-based communication between the Soldier and the robot.

The DR10 is a significant step forward for the small UGV space. It is one of the first man-packable platforms, being significantly lighter than the DR20 platform (50% weight reduction). By enabling autonomous navigation, perception, and surveillance operations on the DR10, it will become one of the first field-ready MUGV systems to be included in Army S&T experimentation, providing the opportunity to explore the challenges and limitations of current perception, navigation, intelligence, and manipulation algorithmic approaches on a platform of this size.

QNA has a long history of developing and deploying robotic platforms. We have fielded over 4,000 systems around the world. We have an intimate knowledge of our platforms that enables us to provide appropriate APIs in a well-structured manner, empowering researchers to rapidly integrate their technologies and enable transition to the Solider. This work has application in all areas of research and in the Capstone Assessment Vision. In particular, the platform work applies to Tasks M1, M2, M3, P2, and I10, while mesh networking applies toward Tasks I4, I9, P8, and H1.

Metrics for Evaluation:
The DR20 DevKit, consisting of both a hardware and software component, is currently on its third software release, supporting both the corresponding DR20 platform and its 3-DOF manipulator. The hardware component came online at the end of the 2011 program year. We have two use cases, with JPL and UPenn, in which these integration pieces have been successfully applied to research related to manipulator force control and haptic feedback. To ensure progress in this ongoing effort, we will continually monitor technical activities to ensure that quarterly milestones are reached. This can be measured by the design documentation that will be developed for the payload as we develop the hardware and software components. This design documentation will be provided in the quarterly report format for the program.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Preliminary design of DevKit changes for DR10. Mesh network radio survey.
Q2	Detailed design of DevKit changes for DR10. Mesh network bench and/or field trials.
Q3	DR10 available to the program. Release 1.4 of DevKit to include support of the DR10. Design of mesh network radio integration with DR10 DevKit.
Q4	Release 1.5 of DR10 DevKit. Mesh network field trials.

IR4-2012 – Modeling and Simulation Environment

Level: 6.2 Applied Research

	Name	*Organization*	*Subtask*
Principal Investigator	Dave Wagner	GDRS	IR4-1
Additional Investigators	Lou Bonbrest	GDRS	IR4-1
	Ed Weller	GDRS	IR4-1
	Ralph Brewer	ARL	IR4-2

Subtask 1: RIVET Graphics and Simulation Engine Upgrade (GDRS)

Objective:
Over the next year, the RIVET team proposes to upgrade the current graphics engine that underlies RIVET 1.0 which was developed during the last RCTA. Simulating small robotic assets during the current RCTA will require enhanced simulation and modern rendering techniques. These improvements/upgrades will create RIVET ***2.0*** which will enable order-of-magnitude advances in physics, scene interaction, lighting, material properties, and visual fidelity.

The RIVET 2.0 upgrades are critical to a state-of-the-art integration and assessment platform that will allow the RCTA to maintain an integrated focus in each of the four key research task groups. Although each research group stands on their own as a scientific field, the tight meaningful integration of this research is paramount to achieving breakthroughs in robotics research.

RIVET intends to provide Consortium members the capabilities that allow:

- Basic research technologies to pass exit criteria and peer-evaluations prior to maturation and other assessments in a simulation environment resulting in a more efficient integration effort.
- A means to exercise research in a repeatable fashion, subjecting research to far more test cases than one can evaluate in the real world and without the high cost of field tests.
- Simulation environments to support the majority of our integration and assessment activities, serving as a necessary gate to field assessment.

The scientific value of integrated research is profound; not only do we start to validate system architectures, we also start to show that complementary research can lead to robotic capabilities that are greater than the sum of their parts.

Next generation robotic technologies should be supported with leading-edge simulation technologies. Our upgrade approach guarantees that we are applying leading-edge industry techniques using the latest hardware acceleration paradigms.

The RIVET team will continue to develop RIVET 2.0 as well as support the four RCTA key task areas over the lifespan of the program. This also includes meeting the ongoing needs and requirements of the researchers by adding new test areas, robots, and sensors as they are needed.

Research Description:
The new graphics engine for RIVET 2.0 will provide opportunities to further exploit the strengths of the RIVET tool as well as extend its capabilities to meet the needs of future RCTA objectives with industry leading-edge technology enhancements. The two most significant areas of upgrade that will allow RIVET 2.0 to impact the broadest research within the program are rendering and physics.

We will upgrade RIVET's rendering capability to next-generation programmable pipelines using GPU hardware accelerated parallel processing. This paradigm uses the concept of programmable "shaders" to calculate sophisticated rendering effects on graphics hardware with a high degree of flexibility and parallelism. Shaders are simple programs that describe the traits of either a vertex or a pixel as they travel through the pipeline from internal 3D object representation to the screen buffer. Vertex shaders describe the traits (position, texture coordinates, colors, etc.) of a vertex, while pixel shaders describe the traits (color, z depth, and alpha value) of a pixel. A vertex shader is called for each vertex in a primitive – one vertex in, one (updated) vertex out. Each vertex is then rendered as a series of pixels onto a surface (block of memory) that will eventually be sent to the screen.

These graphics techniques allow us to display photo-realistic synthetic scenes and, most importantly, be able to "reach back" into the various GPU buffers to "touch" the scene on a per-pixel basis. This will be extremely profound for simulated sensor imagery and data because we can process material, normal vectors, range, and more on a per-pixel basis. This will allow for extremely realistic LADAR processing as well as the generation of realistic video frame data that mimics real-world lighting. Because of the shift to high performance GPU processing, we will be upgrading our modeling tool chain and editors and exploring the benefits of defining object- and scene-based custom materials.

Another area of significant enhancement within RIVET 2.0 will be the use of a GPU/PPU hardware-based physics solution. RIVET 1.0 used a second-order Euler method to compute physics exclusively on the CPU. CPU processing limits the amount of realism that can be created between objects in a scene and significantly reduces the ability for disparate objects to be able to affect each other.

RIVET 2.0 is integrated with Nvidia's PhysX, a state-of-the-art, real-time platform for physics simulations. Hardware-based physics solutions (such as PhysX) are based on the forth-order Runge-Kutta (RK4) numerical solutions method for high-order accuracy. In addition to a rigid body physics simulation, a few of the PhysX highlights implemented in RIVET 2.0 include:

- Cloth dynamics
- Rigid body dynamics
- Destructible objects

- Destroyable joints
- Fluid buoyancy

The primary impetus to upgrade the RIVET is to improve the fidelity of the physics simulation. RIVET 1.0 used a simple solid body physics system with enhancements made for simulating vehicle suspensions and water for USVs. With RIVET 2.0 upgrade, we will be using the PhysX engine which will provide a much higher fidelity physics engine with the ability to do both solid and soft body physics simulation in a real-time environment. We will still have to provide our own enhancements to the physics engine for simulating complex physics systems, such as motors and engines, and for extremely high fidelity situations, such as manipulating objects with the gripper on the TALON's robotic arm.

To validate the new physics engine, we plan on starting with the K-bot since this is an asset available to all RCTA members and, therefore, will provide the most benefit from testing and integration in RIVET 2.0. Our approach is to first gather real-world data on the K-bot running in the NIST obstacle course setup in Gaithersburg, Maryland. This will provide a controlled environment in which we will gather specific information, including acceleration rate, turn rate, max velocity, and stopping distance. Once the data is collected, we will recreate the obstacle course in RIVET 2.0 and fine-tune our physics simulation so that we can validate the measured data against the data from the simulation. This approach will be used validate all assets as they are added to the simulation.

During our research of all the available real-time physics engines, we discovered that none were designed to simulate increasing friction from the pressure applied to objects, such as the TALON's gripper squeezing an object, which enables it to pick-up a glass or rotate a door knob. In order to solve this problem, we will use a feature of the PhysX engine that allows a function to be called when a collision occurs between specifically tagged physics materials; in this case, the gripper colliding with another object. The callback function will allow us to insert our own physics model into this situation. Here again, the simulated values will be validated against real-world measurement of the TALON's gripper.

Metrics:
The primary methods for evaluating capability and realism upgrades from RIVET 1.0 to 2.0 will be based upon differential analysis of the current and new rendering systems performance parameters as well as the subjective (yet reliable) visual and cognitive assessment of increased "realism."

We intend to verify upgraded performance of the new software-based rendering and simulation engine by running side-by-side comparisons of the old and new rendering pipelines in a standardized 3D demo world. In addition, the team will perform a series of real-world versus simulated-world physics experiments that will show the realistic capability of the PhysX RK4 system.

During development, we will measure CPU/GPU utilization using Visual Studio 10's DirectX profiler as well as nVidia's PerfHUD. PerfHUD will allow us to quantify API call lists, pipeline

dependencies, and CPU/GPU timing graphs. These performance-based metrics will show the revolutionary increase in performance over the RIVET 1.0 fixed-function pipeline. The metrics will be comprised of:

- Automated Performance Analysis
- API call lists
- Per-draw call GPU graphs
- CPU and GPU Timings graphs
- Instruction Count Ratio graphs

Our primary focus is to create graphical and simulation platform enhancements that allow for higher realism of simulated robotic technologies.

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
The RIVET team will continue to support RIVET 1.0 during the time RIVET 2.0 is under development. During the previous RCTA, the RIVET team created an extensive one-week training program that introduces students to broad-based capabilities as well as specific programming examples and support interface documentation of RIVET 1.0. Transfer of knowledge is accomplished by a 435-page Developer's Guide and an extensive PowerPoint presentation of core concepts. These learning tools are available on an as-needed basis to support all RCTA members who have a need to interface or develop using RIVET.

Related Research:
Compelling industry descriptions, research, and development using GPU based solutions:

nVidia's CUDA Community Showcases an extensive selection of some of the top GPU computing applications and the power that GPU computing has brought to the world of simulation: http://www.nvidia.com/object/cuda_apps_flash_new.html#

See nVidia's PhysX web page for in-depth discussion of detailed capability and feature set: http://developer.nvidia.com/object/physx_features.html

Paper discussing our new deferred rendering and lighting approach which will allow for greater realism and the ability to "touch" a scene on a per-pixel basis:
http://www710.univ-lyon1.fr/~jciehl/Public/educ/GAMA/2007/Deferred_Shading_Tutorial_SBGAMES2005.pdf

Integration and Assessment Activity:
RIVET will play a critical role within the Integration and Assessment task area. It is expected that RIVET 2.0 will be a significant core technology platform for integration and assessment activities of the four key task areas within the RCTA. The RIVET team will work closely with Ralph Brewer at ARL to ensure that all simulation efforts are Soldier-focused and to meet the expectations of the ARL Enterprise as well as the RCTA collaborative alliance.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Finish port of front end from RIVET 1.0 over to RIVET 2.0. Develop art pipeline for using PureLight for interior static lighting. Re-design and skin the front end UI. Creation of installation scripts. K-bot modeled, rigged, and implemented with physics in the simulation. New trees and vegetation modeled using SpeedTree. BOLT control interface to K-bot created. Ability to mount sensors on vehicles ported to the new engine. Creation of EO video sensor. Release of RIVET 2.0 with controllable K-bot using BOLT interface and video feed from sensor mounted to the K-bot.
Q2	Data collection of K-bot on small robot obstacle course. Re-creation of small robot obstacle course in RIVET to validate and tune K-bot physics. TALON, Dragon Runner, and HMMVW modeled and rigged. Creation of interior physic objects and material properties. Hardware-in-the-loop communication interface for controlling vehicles with a Mac Mini. Release of RIVET 2.0 with hardware-in-the-loop using K-bot.
Q3	Physics rigging for TALON, Dragon Runner, and HMMWV. Data collection of TALON and Dragon Runner on small robot obstacle course. Validate and tune TALON and Dragon Runner physics. Creation of new people and animations for crowds. Creation of custom in engine crowd editor. Creation of specialized AI/crowd behaviors. Soldier controlled through Kinect interface. Release of RIVET 2.0 with new vehicles, crowd edit tools, and AI.
Q4	Ability to output data recording of *a priori* datasets. Ability to DAQ selected set of people, vehicles, and robots in a mission to a database. Map creation, integration, and usage. Creation of forest mission site. Creation of Middle Eastern market site. Creation of an urban mission site. Creation of RADAR sensor. Release of RIVET 2.0 with new missions, sensors, and data capture capability.

Subtask 2: RIVET/ROS Interface (ARL)

Objective and Benefits:
There are two objectives of this research subtask: 1) to increase the utility of the RIVET tool by providing an interface to an external robotic operating system, and 2) to provide a suite of study design and analysis tools. A RIVET/ROS interface with the supporting user tools will make RIVET more accessible to members of the Consortium, especially academic partners, who are using Willow Garage's Robotic Operating System (ROS) in their robotics research. A suite of

study design and analysis tools will allow users to conduct Monte Carlo studies to analyze algorithm performance.

State of the Art:
RIVET 1.0 was created by the 2001-2009 Robotics CTA Consortium to support algorithmic development by providing a virtual environment to support frequent and early evaluation of algorithms from RCTA partners. RIVET 1.0 provides a high fidelity simulation of the RCTA LADAR sensors, a simulation of the dynamics of the physical platform, the XUV, and an interface for the Soldier Machine Interface. Since it is a hardware-in-the-loop system that utilizes the computational hardware of the XUV, the loss of fidelity in testing algorithms in the virtual environment is minimized.

Originally, RIVET only supported evaluation of algorithms designed for the XUV (4DRCS-compliant using NML messaging protocol). It provided hardware-in-the-loop support for the autonomous mobility computer from the XUV and other robots using the same control system. At this point in time, there have been few attempts to interface other robotic systems to RIVET. However, the introduction of the ESPIAL and SAMURAI components enables the RIVET tool to interface with other robot architectures and communication protocols.

One architecture of particular interest for this project is ROS – an open source software framework for robot software development, providing operating system-like functionality on a heterogeneous computer cluster. In addition to the operating system services, it is a code repository for many state-of-the-art algorithms used by the robotics community. It has been adopted by some of the Alliance members including UPenn, CMU, and ARL. By developing a RIVET/ROS interface, we will enable ROS-based robots to be tested in the RIVET environment.

While ESPIAL and SAMURAI provide a means to communicate with arbitrary robot architectures, simulating new robots in RIVET also requires a three-dimensional model of the vehicle chassis and a dynamics model for the chassis and any important components. As the RCTA explores new concepts, particularly in mobility and manipulation, we need to develop models for specific vehicles. In this project, we plan to develop models for the common CTA platform, but there is no restriction on the type of vehicle; air, sea, and ground vehicles can be simulated in RIVET.

RIVET 1.0 did provide some tools to run simulation studies; however, the process still required significant user participation, limiting the number of simulated trials. We intend to develop a tool that enables a user to specify the elements of a simulation study, such as the number of trials, robot paths, stationary objects, and dynamic elements. This tool will also allow users to vary study parameters deterministically or stochastically. The result of the efforts will be a tool that enables a deeper understanding of algorithmic performance.

Based on a 2011 initial trade study conducted by the RCTA, RIVET 2.0 will utilize the Torque game engine and the PhysX physics engine. One of the findings from the study was that the Torque game engine offered many of the same capabilities at significantly less cost. We require access to the source code to create the appropriate sensor models, to modify the game to support

experimental trials, and to create data collection and analysis tools. We considered the Unity game engine which had a higher fidelity physics model, but the cost of the source code was approximately $100,000, compared to Torque's cost of approximately $150. Nvidia's PhysX engine is a state-of-the-art, real-time physics simulator that utilizes 4-order Runge-Kutta (RK4) numerical solutions method for high-order accuracy. In addition to a rigid body physics simulation, a few of the PhysX highlights implemented in RIVET 2.0 include:

- Cloth dynamics
- Rigid body dynamics
- Destructible objects
- Destroyable joints
- Fluid buoyancy

Key Barriers and Approach:
In 2012, we plan to develop a framework for seamless communication between the RIVET simulation and the autonomous algorithms onboard the RCTA common platform, the K-bot. This interface will enable virtual testing and evaluation of smaller platforms in interior and exterior spaces. We also plan to develop a Unit Interface (UI) to allow researchers to design parameter studies to evaluate robotic systems. Through the UI, users will be able to specify changes to the design of experiments. This approach, along with the ability to conduct Monte Carlo and batch experiments, will allow for an increase of data collected on a number of levels.

Increasing the functionality and ease of use with RIVET allows more users to conduct virtual tests of their algorithms. Additional communication protocols will allow Consortium partners to utilize the tool for evaluating robotic technology as they develop it. The RIVET 2.0 includes simulation support for not only autonomous mobility but increasing levels of autonomy within perception, human-robot interaction, and dexterous manipulation. This increase will allow most members of the RCTA to use the tool for testing and evaluation prior to field studies.

The Torque game engine has its own scripting language called Torquescript along with a UI creator which will enable custom screens to be built in order to for the user to develop custom testing environments. The goal is to have a generic-enough interface that it will enable a myriad of interactions with the user and the gaming environment.

The RIVET/ROS interface will be designed using C/C++ coding on a Linux development system. Using the Eclipse IDE will allow for quicker test and evaluation of the code base. The simulation will require a minimum of three high-end computers, two for RIVET (one hosting the main RIVET simulation, the other hosting the SUMURAI tool) and one for the ROS robot. A minimum set of messages passed between the simulation environment and the robot will include point cloud, camera, and text messages. The point cloud message, in particular, is extremely large so one of our goals is to develop custom sensor models within RIVET to enable quicker communication between systems with less backend processing. Each robotic system interacting with RIVET has separate and unique communication protocols which will increase overhead when interfacing with them. It is our hope to get a unified interface that can determine which system is asking for information and send the correct message based on that set of protocols.

The game engine currently does not make use of the data management systems and techniques that would be useful to conduct large simulated studies. Data collection was hand-coded for each experiment, and there was no central repository of data gathered from simulation studies. It has been suggested to implement a SQL-type database management system. We intend to explore this idea as we develop the UI.

Metrics for Evaluation:
The value of a modeling and simulation tool to the overall RCTA effort can be very difficult to quantify. Some measures of the tool's overall utility include the number of RIVET users, the number of different environments, or vehicles built within RIVET. For this work, we will concentrate on metrics related to the ROS/RIVET interface and the UI for setup and design of experimental studies.

Potential metrics include:

- Environmental Complexity specified with the UI
- Mission Complexity that specified with the UI
- Experimental complexity that can be specified using the UI tool
- Trueness of the point cloud being passed to ROS systems
- K-bot physics versus modeled physics

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
This activity is being conducted in collaboration with Subtask IR4-1.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Develop initial RIVET/ROS interface. Develop initial simulation study interface.
Q2	Develop models of the RCTA common robotic platform, the K-bot.
Q3	Complete user interface to enable batch and Monte Carlo testing within RIVET.
Q4	Document simulation study tools and the RIVET/ROS interface.

IR5-2012 – Collaboration Software Integration

Level: 6.2 Applied Research

	Name	*Organization*	*Subtask*
Principal Investigator	Bob Dean	GDRS	IR5-1
Additional Investigator	Mike Wellfare	GDRS	IR5-2

Objective:
The RCTA Integration Toolkit is the set of software middleware, interfaces, and development tools which, when combined with RIVET and common robotic platforms, provides the infrastructure to enable a geographically distributed Alliance to work collaboratively toward common goals in a single environment. The RCTA Integration Toolkit provides facilities for distributed software development and integration and pursues research whose goal is to bring together results from the other task areas. The RCTA Integration Toolkit goes beyond defining an Application Programming Interface (API) to interact with robots, to include an instantiation of the integrated concepts.

The output of this task is a software system that provides the methods to instantiate research results from the other task areas in an integrated manner. Additionally, this task develops general interfaces, software, and analysis capabilities to simplify integration for Integrated Research Assessments.

Background:
While task-level assessment is necessary to maturing the technology and a critical element of our approach, the integrated nature of robotics calls for higher levels of assessment in which one pairs complementary research and existing technology to form an integrated testbed. At this level, quantitative assessment of the research can be performed of the integrated technology, and it provides scientific value of its own. For instance, integrated research helps discover the somewhat blurry boundaries and interactions between perception, intelligence, platform and manipulator control, and human interaction. This is important to subsequent technology development efforts as it strengthens and validates system architectures and identifies boundary issues which are typically the most difficult issues in development efforts. Integrated research also helps identify applications for the research, which sometimes differ from those identified at the outset.

Web-hosted collaboration tool suites bring together common team collaboration capabilities in a common environment. Wikis, document control, calendars, and blogs are common amongst systems such as [1][2][3]. Given the commonality of features, selection of such a system essentially comes down to the ability of an IT team to support the installation.

Though we have built the RCTA Integration Toolkit on top of many open-source or commercial tools, some of our domain-level services are Government-owned or proprietary to Consortium members. Given the advantages of collaborative research, we make the proprietary elements (including hardware, software, and documentation) available throughout the Alliance for research purposes without licensing costs or restrictions. These items will include RIVET simulation software, the RFrame integration framework, and GUI engineering tools.

Whereas we provide a robotic system architecture in the Intelligence research, we reduce this to practice using a layering scheme that facilitates integration of software into a larger environment. In a layered architecture, modules in a given layer can refer to modules at layers lower in the architecture but cannot refer to or make assumptions about modules above them in the architecture. A layered architecture is modular in that integrators can replace entire layers while maintaining the same interfaces. The layers we develop in this task refer to the broader software development environment, including hardware, operating systems, middleware, robot control systems, data visualization and analysis tools, and collaborative development environments.

Research Description:
During FY11, we developed the RFrame Integration Toolkit which was designed to move beyond the state of the art in existing robotic frameworks used within the Consortium, while still maintaining backwards compatibility to legacy systems. RFrame defines a modular middleware framework supporting messaging and a runtime environment. During FY12, we seek to use RFrame to instantiate and assess the Intelligence Architecture with a focus on the needs of IRA3.

To support the future capability needs of the program, we build upon outputs from other task areas to define Subtask 2: Moving Object Descriptions for the Common World Model. This task seeks to provide a versatile world modeling capability for moving objects that not only incorporates information from multiple sensor types, integrates properly with the common world model representation, allows combining of evidence over time and multiple views, and performs tracking of object position over time but also expands that capability to allow the incorporation of object shape, size, articulation, and gestures as well as important mobility characteristics like wheelbase or turning radius into the recorded data. This extended effort supplies the necessary raw data for a new level of intent and behavior understanding, dynamic object characterization, and future path prediction.

Metrics:
- Successful delivery of software to meet IRA requirements
- Software meets computing resource requirements and simplifies integration process
- Successful implementation of Intelligence Architecture on a robot
- For moving objects: latency of track initiation, false alarm rejection, centroid tracking and articulating assembly state error in angle, position, and velocity, and increased levels of track quality (looking at precision, recall, and F1 metric) due to evidence combining over time

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
This task is involved with instantiation of the Intelligence Architecture and support of IRA3. As such, we anticipate tight collaboration with Tasks I1, I2, I3, and I4 in addition to Tasks IR1 and IR2. Through Subtask 2: Moving Object Descriptions for the Common World Model, we foresee collaborative activities with Perception and HRI to deliver gesture support. As the architecture matures, we plan to discuss with ARL Collaborators (Young and Fields) how to integrate their research into this system.

Milestones:

	Quarterly Milestones and Deliverables
Q1	Extension of RFrame code generation to support Matlab. Definition of Intelligence Architecture components to support FY12 IRA efforts. Define detailed objectives, software architecture, primary data structures, and simple prototypes to examine tradeoffs. Review with Intelligence and Perception PIs.
Q2	Investigate engineering tools, and evaluate those available in Matlab. Complete octave model of dynamic objects.
Q3	Implementation of Intelligence Architecture components which apply to Integrated Research. Complete C++ version of dynamic object library.
Q4	Investigate building a simulation world built from sensed data from real robots. Collect data using real moving platforms and obstacles to demonstrate the baseline dynamic object model.

Subtask 1: RFrame Middleware Toolkit (GDRS)

Objective and Benefits:
The RFrame Integration Toolkit is common software middleware, interfaces, and development tools which simplify integration and assessment of task research results. RFrame is more than an arrangement of standard interfaces; it is a comprehensive suite of tools with which the entire Alliance can perform groundbreaking research in robotics using an environment that facilitates distributed software development. Through RFrame, we are able to leverage new, legacy, and third-party capabilities in an efficient manner. Additionally, RFrame provides the basis to instantiate the Intelligence Architecture.

State of the Art:
There have been attempts at building common robotic platforms, yet each falls short of a complete system. The Robotic Operating System (ROS), for example, has yielded significant acceptance in the research community by providing infrastructure, such as easy startup/shutdown of a multi-node system, and enabling the development of algorithm ecosystems around common data formats (such as the PointCloud and OpenCV libraries). Yet ROS stops short of becoming a full system by leaving a hole between message transport and application levels, where each adopter must “recreate the wheel.” Lack of real-time support prompts questions as to its ability to

support already demonstrated military capabilities. The JAUS standard implementations such as JAUS Toolset and OpenJAUS provide tools to auto-generate JAUS-compliant messaging and high-level component/service state machines allowing different subsystems to communicate reliably. However, these implementations to date stop short of defining methods above "on the wire" messaging and simple state machines, requiring additional robotic infrastructure from each vendor. Likewise, it may require significant integration effort to integrate technologies developed using different middleware frameworks.

Key Barriers and Approach:
There are a number of integration barriers, including differences in computing environment, middleware, and software-level architecture. For example, taking a manipulator developed using ROS and integrating it onto a NML-based system developed under the previous RCTA may require significant code updates.

During FY11, we developed the RFrame Integration Toolkit which was designed to move beyond the state of the art in existing robotic frameworks used within the Consortium, while still maintaining backwards compatibility to legacy systems. RFrame defines a modular middleware framework supporting messaging and a runtime environment. Applications for RFrame are implemented as dynamically loaded libraries using a common state machine. This simplifies integration in a number of ways. Use of a common module startup and shutdown procedure enables understanding of how a module is expected to operate. Dynamic libraries design allows modules to be easily moved between hardware as well as in/out of common process spaces for efficiency of communication. They also reduce build dependencies, speeding the development process and reducing overhead. The runtime environment provides common utility functionality such as data logging and ID assignment.

RFrame also provides a messaging and communication abstraction layer. Communications are implemented by a Connection Manager, allowing the system to use multiple communication protocols if desired. Currently, we have implemented a proof of concept using the NML protocol developed by NIST and used on the previous RCTA.

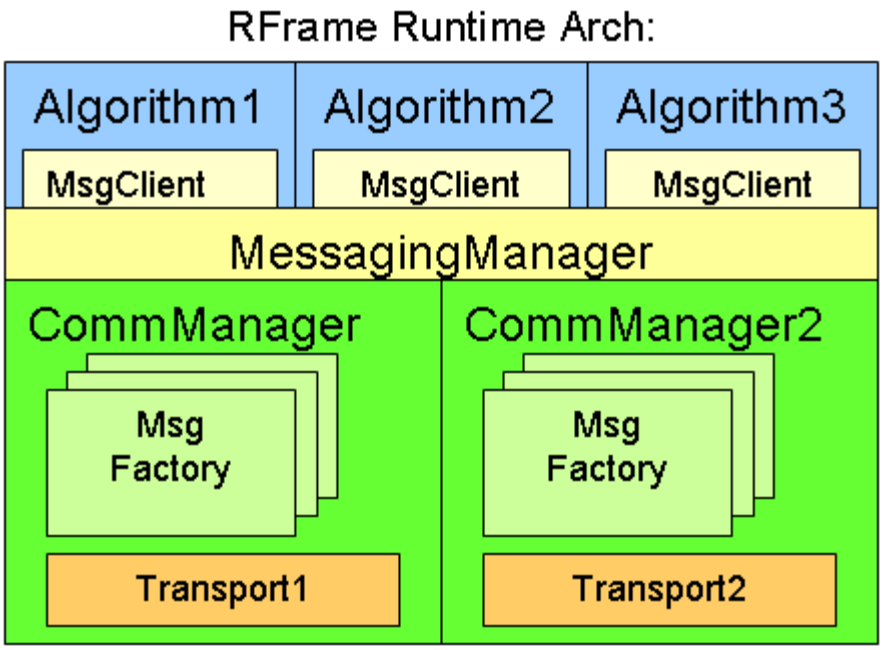

Key to the RFrame concept is a new Code Generator. This code generator differs from other implementations in that it treats each data class as data, where systems such as ROS and NML treat it as a Message. This enables the code generator to be used within algorithms itself. For example, data structures for the world model developed by Task I2 will be defined in this manner, with an expectation of great savings in development time. For example, adding data members to a class which could result in 30 to 60 minutes of work in updating data structures, serialization code, and unit tests is now automatic.

During FY11, a ROS interface was built on top of RFrame, and command line scripts were developed to update ROS Stack application source code to build against the RFrame core. As a result, we can simply integrate ROS applications and NML applications. An example of this is shown here:

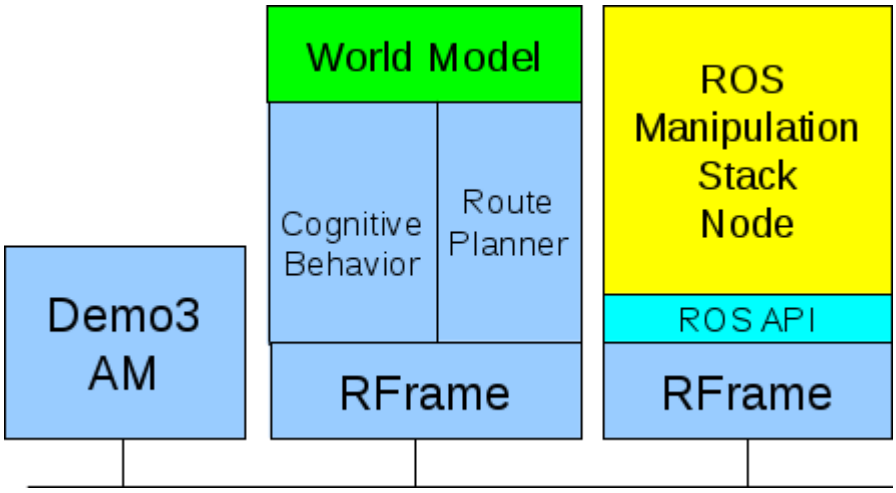

During FY12, we will extend the RFrame in a number of ways.

First, we extend the RFrame framework. The code generator will be extended to output serialization code capable of generating Matlab readable data, which will greatly improve the ability to analyze results. We plan to use the HDF5 scientific data format as it is well established in the scientific community, self documenting, and the data format used to implement the native .mat format used by Matlab. We will investigate the potential of live streaming of robot data to Matlab to allow in-line analysis of system performance.

Communications is a thorn in the side of many research efforts. To assist with this problem, we will update the RFrame NML communications manager to track network statistics and develop a Wireshark plugin able to assist with diagnosis of network issues.

Second, we will provide engineering GUI tools to control Integrated Research Platforms, integrating command and control concepts from HRI when possible. Part of this investigation will be to evaluate tools provided with ROS as the RFrame ROS interface allows for these tools to be integrated with the overall platform architecture.

Investigate a common merged DAQ database tool. Each portion of sensed data from a robot provides more conformation about the world. We should be able to load all data from test runs into a common database, crunch the numbers, and result with a world model for simulation and testing, with quality that increases over time while reducing uncertainty.

Finally, implementation of the Intelligence Architecture defined by Task I1 is a critical effort for the RCTA program as a whole for FY12. We will assist efforts in I2 to implement the Intelligence Architecture on top of the RFrame environment as the design develops. The elements mentioned above were chosen to support the implementation of the Intelligence Architecture. We anticipate that this work will include integration of the I4 Task Executor with cognitive aspects from I1 and I2, determining base module definitions for common components such as planners as well as addressing system-level integration concerns like configuration discovery.

Metrics:

- Successful delivery of software to meet IRA requirements.
- Software meets computing resource requirements and simplifies integration process
- Successful implementation of Intelligence Architecture on a robot

Anticipated Lectures, Workshops, Seminars, Collaborative Activities, and Staff Rotations:
This task implements middleware functionality necessary to instantiate the Intelligence Architecture defined by Tasks I1 and I2, with a primary focus on supporting IRA3: Intelligent Navigation. In preparation for this event, we anticipate source code delivery to the researchers involved and trips to pre-integration sites to support use of the code base.

Related Research:
The fundamental related research for this area is with Tasks I1 and I2 and assisting with implementing the Intelligence Architecture.

Integration and Assessment Activity:
This task implements middleware functionality necessary to instantiate the Intelligence Architecture defined by Tasks I1 and I2, with a primary focus on supporting IRA3: Intelligent Navigation. In preparation for this event, we anticipate source code delivery to the researchers involved and trips to pre-integration sites to support use of the code base.

	Subtask 1 Quarterly Milestones and Deliverables
Q1	Extension of RFrame code generation to support Matlab. Definition of Intelligence Architecture components to support FY12 IRA efforts.
Q2	Investigate engineering tools, and evaluate those available in Matlab.
Q3	Implementation of Intelligence Architecture components which apply to Integrated Research.
Q4	Investigate building a simulation world built from sensed data from real robots.

Subtask 2: Moving Object Descriptions for the Common World Model (GDRS)

Objective and Benefits:
This subtask seeks to add descriptive abilities unique to moving objects into the common world model to allow the UGV's Intelligence processing to: (1) better verify the existence of an object detection through repeated sensing updates and compare those detections to prior data in order to remove false detections from the world model; (2) build-up diverse sensor information that will allow an intelligent system to better classify the object through repeated looks and accumulation of evidence; and (3) process object position, detailed shape, and velocity data to understand potential UGV collision threats and test against hypothesized avoidance maneuvers. These efforts will provide the ground work for future efforts that focus on predicting future object behavior based on prior observed or trained behavior and the current local environment as well as provide position and articulation information in a framework that will allow an intelligent system to semantically classify the different behaviors observed. Providing the above will require the ability to consolidate multiple sensing updates over time, from multiple perception sensors, into a coherent object motion and object articulation description over time. It is also critical to ensure that moving objects do not interfere with or confuse static portions of the common world model.

Detection and Avoidance of Moving Obstacles is a key capability in every part of the Capstone document. It is critical to lay the proper groundwork for that capability which can be built upon in future efforts. Although much has been achieved in terms of perceiving moving obstacles, the prevailing state of the art does little to address fundamental issues relating to the representation of essential information about moving obstacles or the use of the world model to answer questions about what sort of obstacle or platform movements would cause potential conflict in

the future. Some work towards this end has been accomplished in prior robotics programs, but this effort will more fully integrate the different approaches.

State of the Art:
Numerous past and present robotic systems, including several developed by RCTA team members, implement methods to deal with dynamic moving objects external to the robot. Research in planning or object tracking in support of such capabilities includes the work of Brummit and Stenz on "Dynamic Mission Planning for Multiple Mobile Robots," the work of Phillipsen in "…Growable Costmaps," the work of Kushleyev and Likhachev on "Time-Bounded Lattice Planning…," and the work of Almeida et al. on moving object tracking in robotics. However, the listed methods do not fully address the correlation of motion with associated shape extraction, object pose determination including hierarchical articulation, and extraction of dynamic features such as gestures and gait that are useful for intuiting intent.

Current approaches typically only look at tracking of objects within the local environment or concentrate on understanding discrete behavior sets of one or two people at a time. This proposed effort moves beyond that to correlate object base motion with the extraction of current shape and pose information, including hierarchical articulation, and extraction of gestures and repetitive motion patterns that are useful for intuiting identity and intent. This effort also addresses several important UGV missions, such as moving object collision avoidance, while looking ahead to future needs like understanding the local environment to predict possible interactions. For example, it is important that a UGV knows if a car door opens, a person may soon exit. Similarly, a UGV should be aware that if a car's hood is propped open, it is likely that the car is non-functional. A sitting person is much less likely to move quickly than is a running person. There are many possible benefits to understanding that can be facilitated by this improved object-level information.

Key Barriers and Approach:
Current approaches separate the understanding of an object's position (detection and localization) from its classification (labeling) and from its behavior (tracking and intent analysis). Simultaneous representation of object shape, characteristics, and behavior is necessary to enable better categorization and prediction of the object within a full scene understanding and intentional context and will be provided under this subtask.

The achievements of the ANS/SOURCE program made it apparent that a solid baseline had been developed in this technical area. Although not all problems were solved, a clear route to problem solution exists along with a wealth of associated field test data. Research done in past Robotics CTA tasks has also developed portions of this capability (for example, in the vehicle tracking research) that will be fully leveraged under this effort. We will also carefully follow other related research that is making progress towards behavior prediction through competitions such as the ICPR Contest on Semantic Description of Human Activities (SDHA) to provide guidance on the best path to improving prediction and behavior understanding while looking for ways to combine with our own experience on prior RCTA human detection activities.

Metrics for Evaluation:
The first year metrics would focus on latency of track initiation, false alarm rejection, centroid

tracking and articulating assembly state error in angle, position, and velocity, and increased levels
of track quality (looking at precision, recall, and F1 metric) due to evidence combining over time.

	Subtask 2 Quarterly Milestones and Deliverables
Q1	Define detailed objectives, software architecture, primary data structures, and simple prototypes to examine tradeoffs. Review with Intelligence and Perception PIs.
Q2	Develop complete octave model of basic processing, and present initial processing results on previously collected data. Use results to refine final code detailed design.
Q3	Complete a baseline research platform library in C++, and prepare tools and integrated Perception capabilities from collaborators for data collection and ground truthing. Demonstrate integration with simulated sensors and platforms.
Q4	Collect data using real moving platforms and obstacles, process data, and demonstrate that the baseline model can provide the desired metric and semantic information needed for effective Intelligence decision-making according to formulated objectives while maintaining the feedback needed for directed and focused Perception processing.

APPENDIX 1. COMMON LANAGUAGE TABLE

Term	Definition	Related/Associated
Mental Model	An internal representation of some domain, system, procedure, or situation that supports description, explanation, and prediction (cf. understanding, reasoning). The representation is assumed to be primarily in the form of declarative structures to enable sharing, introspection, and explanation, but it can also be proceduralized for efficient execution. Having an appropriate and/or accurate mental model allows users to generate descriptions of the domain, system, procedure, or situation; to explain behaviors and system states; and to make predictions about its future states (based on Gentner, 2002; Rouse, 1989).	Related: Schema, script, knowledge structure, structural knowledge, teammate model, team interaction model, task model Associated: Declarative knowledge, procedural knowledge, strategic knowledge
Mental Model Representation	An external representation of a mental model in the form of text, graphics, computational code, sequences of actions, or numerical values. Example methods to create mental model representations are the application of (cognitive) task analysis, talk-aloud protocols, pairwise relatedness ratings, card sorts, concept maps, and cue-strategy associations.	Related: Externalized knowledge, knowledge elicitation
Shared Mental Model	The degree to which mental models of interacting agents (e.g., human and robotic team members) overlap, are similar, and/or are compatible. Overlap can be indicated by agreement among elements (e.g., elements are the same). Similarity can be indicated by correlation of elements (e.g., patterns among elements are similar). Compatibility can be indicated by the degree to which outcomes based on the mental model(s) are the same (e.g., based on the mental models, the same understanding, reasoning, and prediction takes place).	Related: Transactive memory system, shared knowledge structures
Situation Awareness	A resulting state of knowledge about relevant elements within a domain, procedure, or situation that supports planning, decision making, and behavior. An example method to describe the relevant elements is goal-directed task analysis (see Endsley, 1994) whereby information needs to support goal-directed behaviors are cataloged.	Related: Situational understanding, situational model

Term	Definition	Related/Associated
Situation Assessment	A series of cognitive and perceptual processes through which an individual develops individual situation awareness. Situation assessment is performed through individual information acquisition (albeit possibly from team members) and processing. Goal-directed processes include the perception of elements, comprehension of interrelationships amongst elements, and the prediction of future states of the elements. These processes are constrained by cognitive limitations (e.g., attentional focus, working memory) and by the state knowledge (e.g., interpretation schemas). Acquisition of simple perceptual information supports integration of elements and predictions, but this process is not necessarily linear or strictly hierarchical. In teams, the situation assessment conducted by individual team members can be supported by other team members through the timely sharing of relevant information about the situation and the team members' (perceived) level of situation awareness.	Related: Sense-making, cognitive function, knowledge acquisition
Team-level Situation Awareness	The degree to which the team has common situation awareness ("team situation awareness") can be indexed by aggregation (the sum of individuals' SA), sharedness (agreement, overlap, or correlation/similar relational patterns), and/or complementariness (individuals' unique and compatible SA). Based on our FY10 and FY11 work, initial results suggest that a model of shared situation understanding in teams that combines compatible situation knowledge (described by Endsley et al. as team SA) with complementary situation knowledge as an emergent property of the team (described by Stanton et al. as distributed SA) is most promising.	Related: Team SA, distributed SA, shared SA, team cognition, common ground
Natural/Intuitive Communication	Methods of communication that human beings understand innately or use without additional training outside of what they would experience in normal day-to-day life.	
Multi-modal Communication	The exchange of information through a flexible selection of explicit and implicit modalities that enables interactions as well as influences behavior, thoughts, and emotions.	
Explicit Communication	The purposeful conveyance of information through multiple modalities (i.e., audio, visual, tactile) that has a defined meaning.	
Implicit Communication	The inadvertent conveyance of information about emotional and contextual state that will affect interpretation, thoughts, and behaviors.	
Social Signal	A communicated unit; signals include explicit communications such as gestures and implicit communications such as attire (which may indicate social rank, gender, etc.). Multiple modalities of cues work in dynamic processes to indicate mental state of other agents, such as intention or affect.	Related: Cues, communication, expression Associated: Intersubjectivity, context, information consumption

Term	Definition	Related/Associated
Joint Attention	Describes capacity to at once attend to an object and recognize that another is attending to the same object by means of a reciprocal monitoring of gaze in which each person checks on the attention of the other while acknowledging that the other is doing the same.	Associated: Collaboration, coordination, shared mental models
Direct Perception	Modes of perception which allow for more immediate access to intentions or inferences in the world.	Related: Perception Associated: Situational awareness, context
Reflective Perception	Memory-bound perceptive qualities which allow for the evaluation of contextual and narrative cues of another in the world. Reflective perception often uses information from cultural expectations and goal direction to inform decision making.	Related: Memory, narrative Associated: Situational awareness, context
Intention	Generally taken to mean understanding the goal(s) pursued by that person through his/her actions and the related mental contents as it pertains to that person's situation.	Related: Mental model, intersubjectivity Associated: Situational awareness
Intersubjectivity	The interaction and engagement of autonomous agents. Intersubjective behaviors compose the dynamic social exchanges in humans. Requires mutual engagement and bidirectional feedback.	Related: Social signals, understanding others, communication Associated: Collaboration, coordination
Team Activity Recognition (TAR)	The ability to recognize from a visual input or geometrical scene description a specific team pattern being enacted by team of humans and/or robots. Usually, Team Activity Recognition (TAR) assumes a certain, well-defined movement pattern, such as those found in tactical maneuvers. In many cases, TAR is possible based on a single static snapshot of a scene.	Related: Direct perception Associated: Situation awareness
Team Activity Analysis (TAA)	A deeper understanding of an ongoing team action (compared to TAR), Team Activity Analysis (TAA) involves the identification of the roles played by the individual players, the transitions between different team behaviors, as well as behaviors which had been executed incorrectly or incompletely. By necessity, TAA must by dynamic and allows the prediction of the next move of the individual players.	Related: Intention Associated: Situation assessment
Team Activity Synthesis	Allows a player to dynamically generate the activities necessary to enact a specific role in a team. Beyond simple enactment of a predefined movement pattern, this requires ongoing TAA and appropriate reactions to the actions of the other team members and external circumstances (even if these are incorrect in the context of an idealized team action).	Related: Problem solving Associated: Situation awareness

Term	Definition	Related/Associated
Crowd Simulation	Research in crowd simulation refers to the use of computer models using various algorithms to simulate large populations of typically homogeneous agents, usually with a focus on properties of macroscopic crowd flow, such as evacuation rates or traffic flow capacities.	
Agent Traits and Capabilities	Traits define the state space of an agent, while capabilities define the affordances of the agent. For cognitive models, traits define biases and tendencies that modulate behavior, while capabilities define both available functionalities and their capacity parameters.	
Behavior Authoring	Behavior authoring concerns itself with providing a specification interface whereby people (e.g., subject matter experts) can define agent behaviors and a compatible software generation system with an optional visualization engine that simulates agents to enact the specified behaviors.	
Parameterized Action Representation (PAR)	The Parameterized Action Representation (PAR) bridges the gap between natural language and simulation by giving an expressible, computer-understandable description of an action. PAR uses a hierarchical action representation model that enables a parser to map the components of the instruction directly for agent control.	
Parameterized Behavior Trees (PBT)	Parameterized Behavior Trees (PBT) is an extension to the PAR formulation which uses the behavior tree graphical programming paradigm to author both individual and group behaviors.	