Army Research Laboratory

ROBOTICS COLLABORATIVE TECHNOLOGY ALLIANCE (RCTA)

Proposed 2017-18 Biennial Program Plan

Members

Army Research Lab (ARL)
(Alliance Lead)

Carnegie Mellon University (CMU)

Florida State University (FSU)

General Dynamics Land Systems (GDLS)
(Integration Lead Organization,
Consortium Manager)

Jet Propulsion Laboratory/
California Institute of Technology (JPL/Caltech)

Massachusetts Institute of Technology (MIT)

QinetiQ North America (QNA)

University of Central Florida (UCF)

University of Pennsylvania (UPenn)

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

For any purpose other than to evaluate the *Robotics CTA 2017-18 Biennial 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.

REVISION HISTORY			
REV	**DATE**	**BY**	**DESCRIPTION**
1	11/01/2016	D. Patel	Original Draft
2	12/13/2016	D. Patel	Added subtasks
3	01/24/17	D. Patel	Major revision of subtasks, added intro
4	01/31/17	D. Patel	Final for signature and release – R1
5	2/8/2017	D. Patel	Added subtask T3C1S1E

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

The Robotics CTA 2017-18 Biennial Program Plan presented in this document has been reviewed and approved by both the Consortium, represented by General Dynamics Land Systems, the Integration Lead Organization (ILO); and the Government, represented by Dr. Stuart Young, the Collaborative Alliance Manager (CAM).

Dilip G. Patel
RCTA Program Director
General Dynamics Land Systems
Integration Lead Organization

Dr. Stuart Young
Chief, Asset Control & Behavior Branch
US Army Research Laboratory
Collaborative Alliance Manager

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Table of Contents

1.0 INTRODUCTION 1
2.0 Thrust 1: Maneuvers in Unstructured Environments 9
2.1 Overview 9
2.2 Thrust 1/Capability 1: Mobility in dynamic scenes 9
2.3 Thrust 1/Capability 1 Subtasks 15
2.3.1 T1C1S1A - Rapid learning of recognition models 15
2.3.2 T1C1S1B - Understanding of environment at contextual and more generalized levels 17
2.3.3 T1C1S1C - Scene Understanding at the Tactical Edge 20
2.3.4 T1C1S1D - Multimodal, multiscale reasoning for scene understanding 22
2.3.5 T1C1S2A - Modeling interactions and action 24
2.3.6 T1C1S2C - Human-robot teaming in a social context 26
2.3.7 T1C1S3A - Anytime neural networks for semantic understanding on the move 28
2.3.8 T1C1S4B - Navigation among dynamic, reactive agents 30
2.3.9 T1C1S4D - Modeling trajectories for robust perception 32
2.3.10 T1C1S4E – Learning traversability cost from unsupervised semantics 34
2.4 Thrust 1/Capability 2: Op-tempo mobility in rough terrain 36
2.5 Thrust 1/Capability 2 Subtasks 41
2.5.1 T1C2S1A – Motion control policies for planning and high-speed running in varying terrain 41
2.5.2 T1C2S2A – Perception for maneuver in complex terrain 43
2.5.3 T1C2S2D – Perception for optempo mobility in rough terrain 46
2.5.4 T1C2S2G - Perception for OpTempo Mobility 48
2.5.5 T1C2S3A - Terrain-adaptive planning for autonomous movement of legged and skid-steered robots 50
2.5.6 T1C2S4C - Practical human-robot interaction for navigation 53
2.5.7 T1C2S5A - Next-gen legged platforms 55
2.5.8 T1C2S5B - Data-driven Viability state network (VSN) framework for dynamic locomotion control 57
2.5.9 T1C2S5C – Analysis of Legged Morphological Design and Control for Traversal of Complex Terrain at OPTempo for Different Scales 60
2.6 Thrust 1/Capability 3: Optempo mobility in confined spaces 61
3.0 Thrust 2: Human-Robot Execution of Complex Missions 65
3.1 Overview 65
3.2 Thrust 2/Capability 1: Situation awareness in unstructured environments 65

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.3 Thrust 2/Capability 1 Subtasks 69
3.3.1 T2C1S2A - Long term learning and prediction of 'action objects' interaction 69
3.3.2 T2C1S2B – Dynamic fine-grained understanding of human scenes 72
3.3.3 T2C1S2D – Predicting human intent and activity possibilities 74
3.3.4 T2C1S3A - Two-way translation between vision and language 76
3.3.5 T2C1S3B - Generating Natural Language Scene Descriptions for Shared Situational Awareness 78
3.3.6 T2C1S3C - Investigating Tactical Multi-modal Soldier-Robot Exchanges 81
3.3.7 T2C1S4A - A user-centered design (UCD) approach to creating usable naturalistic communication interfaces for Soldier-robot teaming 85
3.3.8 T2C1S4B - Adaptive Multimodal Communication 88
3.4 Thrust 2/Capability 2: Distributed Mission Execution 90
3.5 Thrust 2/Capability 2 Subtasks 95
3.5.1 T2C2S1A - Interactive Multi-Agent Planning 95
3.5.2 T2C2S1C – Learning Symbolic Planning Models 98
3.5.3 T2C2S1D – Adaptive Models for Human-Robot Interaction in Diverse Decentralized Human-Robot Teams 100
3.5.4 T2C2S1E - Mixed Optimization Approaches to Route Planning in Humans and Intelligent Agents 102
3.5.5 T2C2S1G - Adaptive Provable Missions over Unknown Environments 104
3.5.6 T2C2S1J - Distributed Situation Awareness and Communication through Collaborative Mission Planning and Execution 107
3.5.7 T2C2S2A - Learning to explain route choices 110
3.5.8 T2C2S2B – Grounding Natural Language Instructions with Learnt Visual-linguistic Knowledge Representation 112
3.5.9 T2C2S2C – Collaborative Perception 114
3.5.10 T2C2S2D - Transparency of Agent Reasoning and Operator Trust 116
3.5.11 T2C2S2E – Semantic Framework for Concepts in Context 119
3.5.12 T2C2S3B - Establishing Appropriate Human-Robot Trust through Empirically Validated Training Techniques: Investigating Mental Models, Situational Risk, and Trust 121
3.5.13 T2C2S3C - Investigating Multi-Modal Soldier-Robot Bidirectional Exchanges 124
3.5.14 T2C2S3D - Appropriate calibrations of trust for supporting soldier-robot teaming. 126
3.6 Thrust 2/Capability 3: Trusted Execution of Verified Missions 129
4.0 Thrust 3: Mobile Manipulation 134
4.1 Overview 134

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.2 Thrust 3/Capability 1: Manipulation in cluttered human habitable spaces 135
4.3 Thrust 3/Capability 1 Subtasks 141
4.3.1 T3C1S1A - Active tactile learning of complex objects 141
4.3.2 T3C1S1B - Planning with Uncertainty in the Model and Perception 143
4.3.3 T3C1S1C – Discovering Task-Related Implications of Object Arrangements 147
4.3.4 T3C1S1D - Lifelong Improvement of Robot Capabilities and Its Robustness 149
4.3.5 T3C1S2A - Geometry learning of rigid and deformable objects 152
4.3.6 T3C1S2B – Transferring Manipulation Strategies to New Objects 157
4.3.7 T3C1S2C - Learning Mechanical and Geometric Models of Unknown Objects Online 160
4.3.8 T3C1S3A - Learning manipulation actions from demonstration 165
4.3.9 T3C1S3C– Dual Arm Contact Driven Mobile Manipulation Behaviors on the Roman Platform 167
4.3.10 T3C1S3D – Mobility-aided manipulation on the Roman platform 170
4.3.11 T3C1S3F - Learning novel force-based mobile manipulation – Navigation in Confined Spaces 172
4.4 Thrust 3/Capability 2: Manipulation for unknown heavy objects 174
5.0 Integrated Research 177
5.1 Overview 177
5.2 Thrust 4 Subtasks 184
5.2.1 T4C1S1A - Maintain current platforms – Husky, Roman, Robosimian 184
5.2.2 T4C1S1B - Maintain current platforms – RoboSimian 186
5.2.3 T4C1S1C – Support and Maintain TL1/ROMAN 188
5.2.4 T4C1S1D – Maintain Current Platforms 189
5.2.5 T4C2S1A – RCTA Collaborative Environment (RCE) and SW Management 191
5.2.6 T4C2S1B - RCE Maintenance and Integration Support 193
5.2.7 T4C3S1A - World Model Maintenance and Enhancements 194
5.2.8 T4C3S1B – Perception Sensor and Computing Interfaces 195
5.2.9 T4C3S1C – Capability Integration 196
5.2.10 T4C3S1D – Human Location and Platform Interfaces 198
5.2.11 T4C3S1E - HRI Driven Enhancements 199
5.2.12 T4C3S1F – Capability Integration Support 201
5.2.13 T4C3S1G – Capability Integration Support 203
5.2.14 T4C3S1H - Capability Integration Support 204
5.2.15 T4C4S1A – Task Based Assessment (Local and Field) 205

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.16 T4C4S1B – Integrated Assessment (Local and Field) 206
5.2.17 T4C4S1D – Integrated Assessment (Local and Field) 207
Appendix A Quarterly Reports 209

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

1.0 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), and 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. Current members of the consortium are: General Dynamics Land Systems (GDLS), Carnegie Mellon University (CMU), Jet Propulsion Lab/California Institute of Technology (JPL/Caltech), Massachusetts Institute of Technology (MIT), University of Pennsylvania (UPenn), University of Central Florida (UCF), Florida State University (FSU), and QinetiQ North America (QNA). The consortium is also augmented with additional institutions on an as needed and yearly basis.

The PA calls for the preparation of a proposed Program Plan for the research in each year of the RCTA program. It provides a detailed plan of research activities, down to the task and even subtask levels. Each year's plan is to be presented to the Research Management Board (RMB) for comments and suggestions. Beginning with 2015 plan, RMB asked for a Biennial Program Plan (BPP) instead of Annual Program Plan (APP).

This document is the proposed 2017-18 BPP, which consists of five 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. Section 2 through Section 4 of the BPP describe in detail the research by thrust/capability 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 5 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.

[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

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

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 operated 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.

Synopsis of Progress in Base Program

To achieve the robotic team member skill set and move toward our operational vision, we identified five technical thrust areas for emphasis during the base phase of RCTA. These thrust areas have been highlighted in the 2011 through 2014 APPs and presented at the 2012 through 2014 RMBs:

1. Semantic perception – perception that understands a basic set of object types important to robotics, moving beyond just what is or is not an obstacle.
2. Adaptive behavior generation – combining previously developed robotic navigation planning together with semantic understanding in a cognitive architecture that supports context.
3. Meta-cognition – enabling the use of intuitive, human-level commands for soldier-robot communication, and creating shared mental models and situation awareness.
4. Machine learning – leveraging new learning techniques to achieve better and faster training of perception and planning algorithms.
5. Hybrid cognitive/metric world model – spanning the range from traditional metric data to associated semantic understanding to support a cognitive level of reasoning.

Key Remaining Challenges during the Extension Years

As we moved into the extension phase of RCTA, we went through a process of asking ourselves what we have accomplished so far and what the most salient remaining technical challenges are that we need to address. That process led to the identification of four technical thrusts:

1. Peer-to-peer tactical teaming (human-robot, robot-robot, etc.)
2. Proactive, fine-grained contextual understanding
3. Fast, adaptive on-line and on-the-fly learning
4. Interaction with complex 3-D environments

The 2015-16 BPP was based on these thrusts, however, it became apparent, at the July 2016 Bootcamp at Fort Indiantown Gap, PA, that we needed to enhance these thrusts so that it could lead to capabilities that were more relevant to army field applications. It also became apparent that the enhanced thrusts and capabilities would lead to significant cross-cutting research collaboration from the original four research areas: Perception, Intelligence, HRI, and DMUM.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

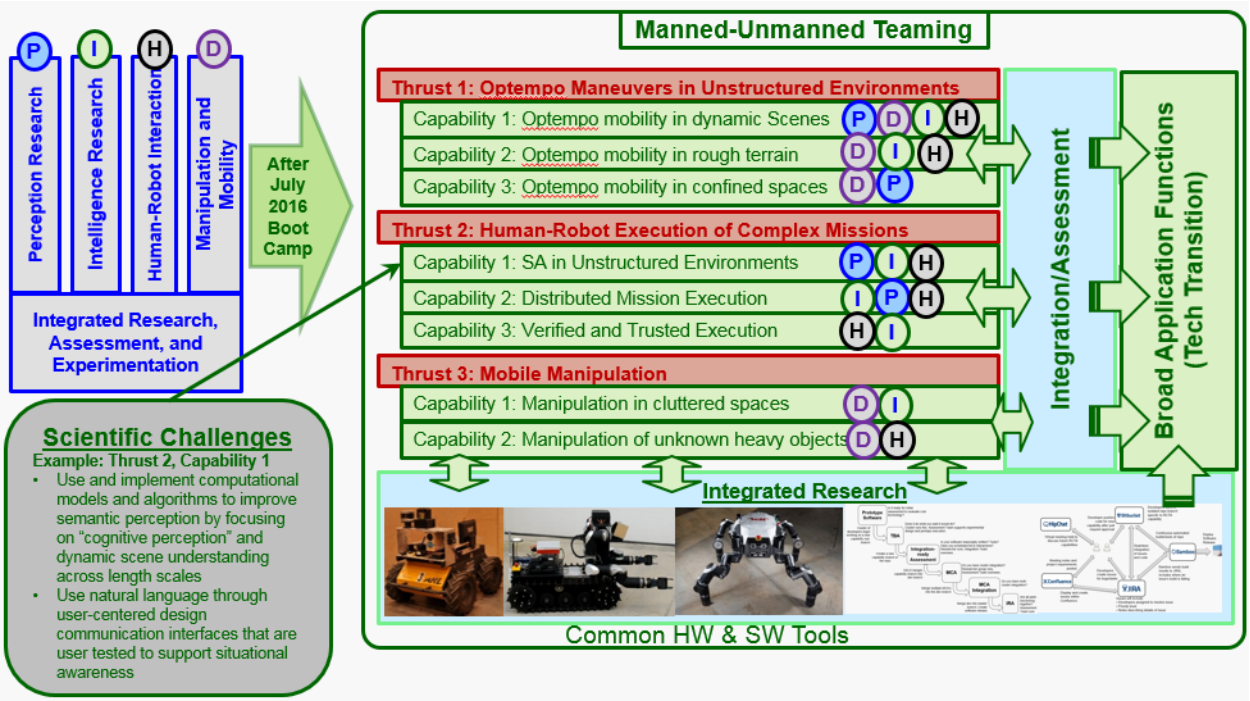

Figure 1. The 2017-18 BPP is based on the formulation of the three new thrusts, which will lead to army relevant capabilities. The main idea is to move away from technical functional area research to capability based research.

For the 2017-18 BPP, the three thrusts and very brief description of capabilities are:

Thrust 1 - Optempo maneuvers in unstructured environments: In order to be an autonomous team member, the robot must have mobility to move at an operational speed and in a relevant environment. There are three targeted capabilities:

> **Capability 1 - Optempo mobility in dynamic scenes:** The interactions between robot and moving agents in the relevant environment. This capability is crucial to any scenario in which the robot must maneuver through a populated and dynamic environment.
>
> **Capability 2 - Optempo mobility in rough terrain:** The ability to move on difficult, *rough terrain*, in particular terrain with difficult 3D geometry, requiring novel approaches to robot configuration, planning, and control. The ability to Fuse visual, point cloud, and proprioceptive sensing into a terrain model that captures geometric and terramechanical characteristics, and use near-to-far learning/labeling techniques to extrapolate trafficability inferences from underfoot to far ahead to facilitate higher speed mobility
>
> **Capability 3 - Optempo mobility in confined spaces**: The ability to move in confined spaces such as dense urban structures, caves, subterranean structures where mobility path is obstructed and poor lighting conditions for perception sensor. (*This capability is not funded at this time.)*

Thrust 2 – Human-Robot execution of complex missions: The ability to carry out complex missions as part of a human-robot team. Enable robots that can do multi-stage missions that do not have a specific

location as an endpoint. Robots should be designed to increase Soldier situational awareness, to maintain security, and to react to changing environmental conditions.

Capability 1 – SA in unstructured environments: To perform a team mission at optempo in unstructured, unfamiliar environments, it is crucial for the team to develop accurate and cohesive situation awareness. This capability requires advances in models and algorithms that allow focused "cognitive perception" and dynamic scene understanding across length scales. Also requires advances in natural language and gesture understanding, integrated with user-centered design communication interfaces that are user tested to support situational awareness

Capability 2 – Distributed mission execution: Collaborative interactions between humans and robots involves collaborative task planning and task execution. Advanced structure prediction, inference, new graphical models and algorithms, and new computational models using RNN.

Capability 3 – Trusted execution of verified missions: The underlying scientific and technological advances necessary for the verification of robot behavior and for the development of Soldier trust in their robotic teammates (*This capability is not funded at this time.)*

Thrust 3 – Mobile Manipulation: The ability to carry out physical human like manipulation tasks.

Capability 1 – Manipulation in cluttered spaces: Developing methods that cope with increasing uncertainty in the geometry of objects and their spatial arrangement.

Capability 2 – Manipulation of unknown heavy objects: Manipulation of unknown, heavy objects, with uncertain inertial properties, which requires more force control that the first capability (*This capability is not funded at this time.)*

The three thrusts, the associated capabilities, and the proposed research are described in sections 2, 3, and 4 respectively. We have used a simple coding system to link each research topic to the appropriate Thrust (T), Capability (C), Summary task (S) with numerical number (e.g. S1), and Subtask withsequential lettering, (e.g. S1A). For example, subtask number T1C1S1A refers to Thrust 1, Capability 1, Task S1, and subtask A. So the research description is always at subtask level. One or more subtasks support a task, and a collection of tasks provide a capability.

Section 5 describes the Integrated Research activities and to keep the same nomenclature flow, we have adapted the same numbering system.

Appendix A provides a template for the quarterly status report, a requirement for all funded researchers.

The following four pages shows the selected research subtasks, organization, and researcher's name.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Thrust/Cap/Task/ Subtask	Task/Subtask	ORG	Researcher
T1	T1 - Optempo Maneuvers in Unstructured Environments		
T1C1	T1C1 - Optempo mobility in dynamic scenes		
T1C1S1	Understand Environment Statically		
T1C1S1A	Rapid learning of recognition models	CMU	Hebert
T1C1S1B	Understanding of environment at contextual and more ge	UCF	Fiore
T1C1S1C	Scene understanding at the Tactical Edge	ARL	Rao
T1C1S1D	Multi-modal, multi-scale reasoning for scene understandi	CMU	Pires/Hebert
T1C1S2	Recognize and anticipate motion and activities		
T1C1S2A	Modeling interactions and actions	CMU	Navarro-Serment
T1C1S2C	Human-robot teaming in a social context	UCF	Boloni
T1C1S3	Perform perception functions robustly with limited co		
T1C1S3A	Anytime neural networks for semantic understanding on t	CMU	Muelling
T1C1S4	Planning based on dynamic representation		
T1C1S4B	Navigation among dynamic, reactive agents	CMU	Oh
T1C1S4D	Modeling trajectories for robust perception	CMU	Navarro-Serment
T1C1S4E	Learning traversability cost from unsupervised semantics	ARL	Wigness/Rogers
T1C2	T1C2 - Optempo mobility in rough terrain		
T1C2S1	High-Speed motion in changing terrain		
T1C2S1A	Motion control policies for planning and high-speed runni	FSU	Clark
T1C2S2	Perception for optempo mobility in rough terrain		
T1C2S2A	Perception for Maneuver in Complex Terrain	JPL/Caltech	Matthies
T1C2S2D	Perception for optempo mobility in rough terrain	GDLS	Henke
T1C2S2G	Perception for optempo mobility in rough terrain	UPenn	Taylor
T1C2S3	Environment and terrain-adaptive planning		
T1C2S3A	Terrain-adaptive planning for autonomous movement of le	FSU	Collins
T1C2S4	Integrated human assistance in path planning		
T1C2S4C	Practical human-robot interaction for navigation and mani	BYU	Goodrich
T1C2S5	Design and control of a dynamic, autonomous quadru		
T1C2S5A	Next-gen legged platforms	JPL/Caltech	Kennedy
T1C2S5B	Data-driven Viability state network (VSN) framework for d	MIT	Kim
T1C2S5C	Analysis of Legged Morphological Design and Control for	ARL	Pusey
T1C3	T1C3 - Optempo mobility in confined spaces		
T1C3S1	Robust perception in confined, unstructured environm		
T1C3S2	Planning in confined, dynamic environments, includi		
T1C3S3	Interacting with teammates and other agents		

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Thrust/Cap/Task/ Subtask	Task/Subtask	ORG	Researcher
T2	T2 - Human-Robot Execution of complex missions		
T2C1	T2C1 - SA in unstructured environments		
T2C1S1	Semantic Scene Understanding in Unstructured Enviro		
T2C1S2	Dynamic Scene Understanding in Unstructured Enviro		
T2C1S2A	Long term learning and prediction of 'action objects' intera	UPenn	Shi
T2C1S2B	Dynamic fine-grained parsing of human scenes	JPL/Caltech	Rothrock
T2C1S2D	Predicting human intent and activity possibilities	Indiana	Ryoo
T2C1S3	Synthesizing Semantic Content into Human-understan		
T2C1S3A	Two-way translation between vision and language	CMU	Oh
T2C1S3B	Generating Natural Language Scene Descriptions for Sha	TTIC	Walter
T2C1S3C	Investigating Tactical Multi-modal Soldier-Robot Exchang	BGU	Oron-Gilad
T2C1S4	Design of System for Effective Teaming		
T2C1S4A	A user-centered design (UCD) approach to creating usab	UCF	Jentsch
T2C1S4B	Adaptive multimodal communication	UCF	Barber
T2C2	T2C2 - Distributed mission execution		
T2C2S1	Collaborative Mission Planning and Execution		
T2C2S1A	Interactive Multi-Agent Planning	CMU	Lebiere
T2C2S1C	Learning Symbolic Planning Models	MIT	Roy
T2C2S1D	Adaptive Models for Human-Robot Interaction in Diverse D	Rochester	Howard
T2C2S1E	Mixed Optimization Approaches to Route Planning in Hu	ARL	Perelman
T2C2S1G	Adaptive, Provable Missions over Unknown Environments	UPenn	Pappas
T2C2S1J	Distributed Situation Awareness and Communication thro	UPenn	Kumar
T2C2S2	Distributed Situational Awareness and Communicatio		
T2C2S2A	Learning to explain route choices	CMU	Oh
T2C2S2B	Grounding Natural Language Instructions with Learnt Visu	MIT	Roy
T2C2S2C	Collaborative Perception	Brown	Tellex
T2C2S2D	Transparency of Agent Reasoning and Operator Trust	ARL	Chen
T2C2S2E	Semantic Framework for Concepts in Context	ARL	Summers-Stay
T2C2S3	Shared Multi-modal Dialogue		
T2C2S3B	Establishing Appropriate Human-Robot Trust through Em	UCF	Jentsch
T2C2S3C	Investigating Multi-Modal Soldier-Robot Bidirectional Exc	ARL	Elliott
T2C2S3D	Appropriate calibrations of trust for supporting soldier-robo	Brown	Phillips
T2C3	T2C3 - Verified and trusted execution		
T2C3S1	Verifiable Planning		
T2C3S2	Maintaining Correctness and Trust through Communic		
T2C3S3	System Design to Maintain Trust		

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Thrust/Cap/Task/ Subtask	Task/Subtask	ORG	Researcher
T3	**T3 - Mobile manipulation**		
T3C1	**T3C1 - Manipulation in cluttered spaces**		
T3C1S1	**Object Manipulation "Where objects are?"**		
T3C1S1A	Active tactile learning of complex objects	UPenn	Daniilidis
T3C1S1B	Planning with Uncertainty in the Model and Perception	CMU	Likhachev
T3C1S1C	Discovering task-related implications of Object Arrangem	ARL	Fields
T3C1S1D	Lifelong Improvement of Robot Capabilities and Its Robus	CMU	Likhachev
T3C1S2	**Object Grasping "How to grasp?"**		
T3C1S2A	Geometry learning of rigid and deformable objects	UPenn	Daniilidis
T3C1S2B	Transferring Manipulation Strategies to New Objects	JPL/Caltech	Detry-Pappon
T3C1S2C	Learning Mechanical and Geometric Models of Unknown	Rutgers	Boularias
T3C1S3	**Human Like Manipulation "How to manipulate?"**		
T3C1S3A	Learning manipulation actions from demonstration	CMU	Muelling
T3C1S3C	Dual-arm contact-driven mobile manipulation behaviors or	JPL/Caltech	Karumanchi
T3C1S3D	Mobility-aided manipulation on the Roman platform	JPL/Caltech	Burdick
T3C1S3F	Learning novel force-based mobile manipulation	UPenn	Yim
T3C2	**T3C2 - Manipulation of unknown heavy objects**		
T3C2S1	**"Inertial Uncertainty"**		
T3C2S2	**"Grasping to perform work"**		
T3C2S3	**How to enable physical interactions via force feedbac**		
T3C2S4	**Autonomous 3-D Limbed Mobility**		

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Thrust/Cap/Task/ Subtask	Task/Subtask	ORG	Researcher
T4	**Integrated Research**		
T4C1	**T4C1 - Platforms**		
T4C1S1	**Development and Maintenance**		
T4C1S1A	Maintain current platforms - Husky, RoMan, Robosimian	GDLS	Gonzalez
T4C1S1B	Maintain current platforms - Robosimian	JPL/Caltech	Karumanchi
T4C1S1C	Maintain current platforms - TL1 - RoMan	QNA	Taylor
T4C1S1D	Platform maintenance support	ARL	Westrich
T4C2	**T4C2 - RCTA Collaborative Environment (RCE)**		
T4C2S1	**RCTA Collaborative Environment (RCE)**		
T4C2S1A	RCE and SW Management	GDLS	Henke (A)
T4C2S1B	RCE Maintenance and Integration Support	Upenn	Daniilidis
T4C3	**T4C3 - RCTA SW Infrastructure and Integration**		
T4C3S1	**RCTA SW Infrastructure and Integration**		
T4C3S1A	World Model Maintenance and Enhancements	GDLS	PatelJ
T4C3S1B	Perception Sensor and Computing Interfaces	GDLS	Henke
T4C3S1C	Capability Integration	GDLS	Weller
T4C3S1D	Human Location and Platform Interfaces	Robotic Res	Sapronov
T4C3S1E	HRI Driven Enhancements	UCF	Barber
T4C3S1F	Capability Integration Support	ARL	Childers
T4C3S1G	Capability Integration Support	MIT	Roy
T4C3S1H	Capability Integration Support	CMU	Navarro-Serment
T4C4	**T4C4 - Assessment**		
T4C4S1	**Assessment**		
T4C4S1A	Task Based Assessments (Local and Field)	ARL	Lennon
T4C4S1B	Integrated Assessment (Local and Field)	ARL	Lennon
T4C4S1D	Integrated Assessment (Local and Field)	GDLS	Cullop
T4C5	**T4C5 - Modeling and Simulation**		
T4C5S1	**Modeling and Simulation**		
T4C6	**T4C6 - Technology Transition**		
T4C6S1	**Technology Transition**		
P1	**Program Management**		
P1M1	**P1M1 - Program Management**		
P1M1S1	**Program Management**		
P1M1S1A	Program Management	GDLS	Patel

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.0 Thrust 1: Maneuvers in Unstructured Environments

2.1 Overview

A fundamental requirement for RCTA systems to work with soldiers in the field is to move at speeds consistent with the operational requirements through unknown environments. This is a cross-cutting requirement in that mobility is necessary in all of the scenarios involving autonomy and human-robot teaming. In Thrust1 we focus on the research elements necessary to achieve the mobility requirements. Systems may operate in a variety of environments: static environments with level ground for relatively easy mobility, dynamic scenes with varying degrees of complex interactions with agents acting in the environment, non-level rough terrain precluding simple mobility mechanisms, and indoor, cluttered environments. Accordingly, we investigate two fundamental capabilities intended to provide a complete coverage of the anticipated mobility conditions.

The first capability focuses on mobility in *dynamic scenes*, that is, the interactions between robot and moving agents in the environment. This capability is crucial to any scenario in which the robot must maneuver through a populated and dynamic environment. For this capability, static understanding of the environment is not sufficient and deeper understanding of motion, activities and interaction is needed. Similarly, stop-and-go planning and action is not viable and online planning reacting to external activities is required. This first capability assumes mild terrain conditions that can be handled with a conventional platform. Our second capability focuses on the ability to move on difficult, *rough terrain*, in particular terrain with difficult 3D geometry, requiring novel approaches to robot configuration, planning, and control in particular. Finally, while these first two capabilities address scenarios in which outdoor environments are involved, the third capability focuses on mobility in indoor, confined spaces, which pose unique challenges.

While the state of the art addresses portions of these the capabilities, the proposed research in each capability addresses three cross-cutting technical gaps in the state of the art. First, state of the art techniques cannot be used directly in the unstructured, uncontrolled world envisioned in our scenarios. For example, mobility over rough ground reduces perception performance by corrupting sensor data. Thus we need to directly address *robustness* issues that are not addressed in more controlled environments with more favorable observation conditions. Second, the state of the art assumes off-line computational and data resources, and substantial online resources. For example, learning models (in particular, using deep learning) offline from a large amount of annotated samples and cannot be used in online, on the move scenarios with limited computation. More generally, *speed* for optempo operation requires all of our approaches need to be operational in an online fashion with limited computational resources, and have a platform that explicitly considers the dynamics of its motion in its control and planning. Finally, much of the state of art operates in scenarios in which the environment is constrained, for example in terms of geometry, and motions/activities. In contrast, our approaches must operate in *unstructured* environments in which the system must react to unmodeled geometric and motion configurations, e.g., in moving through a marketplace, across rubble, or through a set of rooms.

2.2 Thrust 1/Capability 1: Mobility in dynamic scenes

2.2.1 Overview

In order to serve effectively as part of a team, robots must be able to keep up with their Soldier teammates as they move toward their objective. Of particular interest are situations in which the team needs to traverse areas where other people or moving objects are present. For example, consider a scenario involving a robot that moves fast, in relative terms at speeds comparable to those of the other agents, through the market, where four or more people are present. For such circumstances, robots must be provided with a set of

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

abilities that allow them to navigate crowded environments without reducing the speed at which the team carries out its mission.

2.2.2 Relevance to the RCTA and the Army

This capability is needed to fulfill the requirements of tactical teaming at high speeds, where there is a critical need for each team member to advance the completion of the mission without extensive communication between them, while maintaining a minimum level of awareness of what each other is or should be doing. Robots must maintain suitable situational awareness without explicit indications from or interactions with Soldiers. These requirements are relevant to the Army because military operations involving robots frequently exhibit the following characteristics:

- The robot needs to operate in populated environments with no control on activities, observation conditions, or scene structure.
- Prior information is unavailable: there are no models of what is happening, who is moving, or where. Moreover, the team has to operate in an unstructured environment without a map or knowledge of area landmarks. Finally, the robot's sensors and perception systems have to deal with wildly varying environmental conditions, like large changes in illumination, without a model to adapt to prevailing conditions.

The research required to meet the challenges imposed by these characteristics is addressed by this capability, and described in the following section.

2.2.3 Required research areas

Previous research has focused on structured environments, containing mostly static or only a few dynamic agents, and where these agents interact with the robot in simple ways—or not at all. Similarly, previous work has centered on approaches based on static planning, where dynamic agents are treated as sequences of static entities. The research proposed for the remaining of the program advances the state of the art in two fundamental ways: first, increase the ability to deal with more complex environments by addressing the challenges imposed by unstructured, dynamic environments; by enhancing approaches to work with multiple agents; and by recognizing and understanding more complex interactions and social cues. Second, increase the complexity and the speed of the decision-making process by moving towards online planning, capable of accounting for complex interactions, for multiple agents, and of performing in an unknown environment. Our research plan builds upon work done in 2015-2016, and extends it to fill the corresponding gaps. We have identified the set of necessary technical elements, which are summarized in the groups described below. The first two groups address the problem of understanding both the static and the dynamic aspects of the environments. The third research area addresses the issue of performing these perception functions with limited computational resources, while retaining online, optempo performance. Finally the third research involves taking the representations generated by perception output to plan the robot's motions, again emphasizing online and time aspects.

2.2.3.1 Understanding environment statically

Team navigation across crowded spaces requires that robotic teammates understand the context of stationary objects at optempo speed. Previous work has addressed the construction of models of the environment at different levels of detail. However, these models were constructed and evaluated separately, so more recent efforts focused on providing mode complete semantic descriptions by integrating different levels of descriptions, and more detailed descriptions by increasing the resolution of semantic representations. In the upcoming part of the program, we build on these capabilities with the goal of developing fine-grained semantic interpretation of the environment in terms of objects, features, and

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

regions from sensor data. To this end, the research focuses on two principal abilities: *recognition* and *learning*.

Recognition: The state of the art on semantic labeling and recognition of coarse-grained classes assumes *good quality imagery*; unfortunately this assumption is unrealistic for the scenarios considered in the RCTA. To meet this challenge, research is needed in three different directions: 1) *Online operation with limited computation and under difficult observation conditions*; 2) *Fine-grained representations of unstructured, dynamic environments*; and 3) *Incorporation and use of domain and task information*. The body of research of the RCTA has developed and demonstrated approaches for semantic labeling based mainly on inference machines and fully convolutional networks. Additionally, we have made progress on inferring scene structure by extending the initial RCTA techniques for detecting local features (e.g. doors, windows) using reinforcement learning and grammars representing the rules governing the relative placement of these structural elements. Finally, we have developed approaches that incorporate other sources of information—different from on-board sensors—by fusing information from prior maps in inference machines and convolutional neural networks. Leveraging on this work, the research plan focuses on the development of a unique multi-modal, multi-cue approach to scene interpretation through the integration of richer cues into the semantic labeling algorithm. The expected outcome is a more accurate, more complete representation of the environment. Our plan pursues this development by fusing data from three general domains: the first one is to use *geometric* information by fusing features from 3D data; this work is based on the 3D interpretation approaches developed during earlier phases of the program. The second one is through the use of *symbolic* information by fusing features that encode interactions of agents with the environment. The third one capitalizes on *environmental* knowledge to fuse features representing the salience of the different parts of the environment. This research is cross-cutting through all three Thrusts since understanding the static environment is the most basic input to most of the capabilities, and it is incorporate in the Thrust 2 research plan.

Learning: The state of the art assumes availability of *large volumes of supervised data* and of *unlimited, off-line computational* resources—e.g. as required by deep learning techniques. However, the unique needs of the RCTA program require weakly or semi-supervised data and domain adaptation techniques to provide two important abilities: *fast* training with *limited supervised* data and *fast* adaptation with *unsupervised* data and *limited* training. The research work carried out in previous stages of the RCTA has developed learning models from weakly supervised data using fully convolutional networks and similarity graphs. Similarly, work has been done involving learning from small-samples along two basic lines: a) by creating a regression model from small-sample models to large sample-models for classification tasks, and b) by encoding of the visual space with classifiers learned from unsupervised data integrated in deep learning. The current research plan will investigate a unique approach to fast training and adaptation of state of the art techniques for operation with limited computation and supervision, divided in two principal fronts: a) *small-sample learning*, where we extend techniques from classification tasks to detection and semantic labeling tasks (fully convolutional networks), and b) *Domain adaptation*, where the Unsupervised transfer learning of CNNs between source domain (traditional imagery) and target domain (tactical scenario). The novel contribution of this work to the state of the art is a unique approach to fast training and adaptation of SoA techniques for operation with limited computation and supervision.

2.2.3.2 Recognizing and anticipating motion and activities

Teams navigating through crowded environments require for each team member to have effective models of what the other members of the team are doing at any given time. For robotic team members, the need for real-time awareness of human activity is addressed by studying approaches to recognize and anticipate the motion of humans, and the activities they are engaged in. For this extension of the program, the goal is to

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

develop approaches to recognize actions from visual data; to predict actions and motions of people and other dynamic objects. These capabilities can serve as input for planning; to provide situational awareness; and to detect threats in dynamic environments. To this end, the research focuses on three primary abilities: *recognize*, *track*, and *predict*. The first two elements are also fundamental technical elements of Thrust 2 for deep semantic understanding of the environment and are part of the Thrust 2 research plan. The Thrust 1 research plan focuses on the prediction element which is most important for being able to plan through dynamic environments.

Recognize: unlike state of the art approaches on video analysis, the *online* recognition capability for planning and decision making specified by the RCTA requires early detection of activities. The research done has developed concepts of sub-events and differentiable attention filters for activity recognition; an approach for end-to-end learning of sub-events using a new architecture involving CNN/LTSM and temporal attention filters; and a real-time implementation of multi-type activity recognition. Leveraging on these elements, the proposed research extends the CNN/LTSM model to 1) *detect* the onset of events, and 2) *anticipate* events from scene context.

Track: the early detection of activities also imposes a challenge for the timely tracking of movers in video sequences. Previous research in the RCTA has developed algorithms for frame-based pose estimation for each video frame using a residual fully convolutional network architecture. Here, joint motion trajectories are learned through recurrent LSTM elements to propagate pose state from frame to frame, using semi-supervised training. In the proposed research, we build on these approaches in several ways: a) *extend the CNN/LSTM architecture* for high-resolution estimation, including gestures and signals; b) *integrate features from context*, using object detectors; and c) *estimate 3D pose cues*.

Predict: an important challenge imposed by the task of online planning of robot actions through multiple agents is the requirement of *online, causal* generation of expected motions, including modeling interactions with the environment from sensors. Moreover, models of social interactions and interactions with the environments are crucial to predict more accurately the trajectories of people on the scene. Contributions from the RCTA program include a multi-agent prediction framework using Iterated Maximum Entropy Inverse Optimal Control—including a ROS-based implementation. This framework has also been applied--in collaboration between ARL and CMU--to transfer predictive behavior models from humans to robots. These models capture the way humans traverse a terrain in the presence of certain features. By predicting how a human would negotiate a path, robots can plan trajectories to traverse an environment in the same way a human would.

2.2.3.3 Robust perception with limited computational budget

The ability to keep up with Soldiers as the team moves through multiple scenarios imposes demands real-time and latency-sensitive perception systems that have not been fully addressed during the RCTA baseline work. Decisions must be made on the fly, allowing enough time to respond and adapt to real-time events. Predictions made by perception systems must be accurate without significantly increasing latency. In this case, it is difficult to decide whether to use a larger, more complex predictor with higher accuracy, or a less accurate but significantly faster predictor. Critical perception results must be generated within the computational budget allocated. Unfortunately, state of the art tools, e.g. deep learning, are incapable of generating predictions that adapt to different computational budgets. This part of the research addresses this challenge: the goal is to develop robust prediction algorithms that are capable of automatically adjusting to any computational budget available. These algorithms—known as *anytime algorithms*—exhibit the desirable trait of providing increasingly better results given more computation time. The body of research of the RCTA has introduced prediction algorithms based on gradient boosting and extended to work on streaming settings. Work has also been done on the derivation of theoretical anytime performance

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

guarantees. Finally, we have developed the first version ever of anytime neural networks. Leveraging on this work, the research plan focuses on developing a unique approach to designing anytime versions of the critical perception components: a) *anytime design of fully-convolutional networks* with application to semantic labeling, event detection, pose, etc., and b) *anytime prediction for high number of classes*: identify common classes quickly through simple computation, and rare and complicated classes through further detailed computation. The key contribution of this part of the work to the state of the art is a unique approach to designing anytime versions of the critical perception components.

2.2.3.4 Planning in dynamic environment

Navigation across crowded spaces presents significant challenges to robotic teammates. They must operate in dynamic—and within the RCTA, unstructured—environments; this requires online planning of robot actions using plausible models of social interactions with elements of the environment. To address this difficult task, this part of the research focuses on two basic goals: to *generate plans based on paths predictions*, and to *assemble and generate* high-level behaviors. The research done in the RCTA program has produced approaches to robot planning using a velocity-based prediction model; a dynamic local planner, which closely follows planned routes; and models of social interaction using simulation tools (crowd and bystanders scenarios), tested on "VIP scenarios" using implicit communication. The collaboration between ARL and CMU on transference of predictive behavior models described before provides online planning capabilities for local plans. These approaches do not take into account the context of the environment—social or scene features—and consequently plans can fail to deal with the actual scene. Consequently, further research is needed to produce effective plans that perform well in crowded, dynamic, and unstructured environments. To this end, the research plan proposes developing a unique approach to planning which takes into account probabilistic predictions, social interactions, and interactions with environment. This objective is pursued in three principal ways: 1) *integrate context representation* in behavior generation and planning, e.g., social and physical interaction models; 2) *refine models of social behaviors* using Social-Long Short Term Memory (LSTM) and Inverse Reinforcement Learning techniques; and 3) *develop cognitive-logical and reactive models* for the decisions of team members: train in simulation, perform in real world. The key contribution to the state of the art is a unique approach to planning taking into account probabilistic predictions, social interactions, and interactions with environment, in contrast to state of the art approaches which focus on planning against moving point representations.

2.2.4 Demonstration plan

The technical elements that comprise this capability will be continuously integrated as part of the software development process during the course of the program. Upon transitioning as mature research from the laboratory to the RCTA platforms, these elements will be incorporated into the RCTA IRA testing configuration and demonstrated as an integrated system for mobility through dynamic environments at the end of the program. We envision three milestones for integrating the technical elements for capability 1:

- *Perception with moving objects (CY17*): This milestone will demonstrate the ability to detect, track, and predict the motion of moving objectsin a representation of the static environment and to plan within a field of moving objects. Since we focus here on the dynamic aspects, with the issues related to complex terrain geometry investigated in Capabilities 2 and 3, this milestones can be executed on the existing Husky platform.
- *Interaction with humans (CY18):* The first milestone does not incorporate any semantics about actions and motions. The second milestone adds explicit recognition of human actions, interactions, and predicted behaviors into the system, including online planning based on reinforcement learning.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

- *Optempo operation (CY19):* The first two milestones focus on dealing with robustness and unstructured aspects of the problem. The third milestone focuses on the speed aspects, including in particular the issue of online, anytime computation with limited resources. This milestone includes evaluating the performance of the system in terms of number of moving agents.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.3 Thrust 1/Capability 1 Subtasks

2.3.1 T1C1S1A - Rapid learning of recognition models

PI: Hebert, Martial (CMU), hebert@cs.cmu.edu, (412) 268-5704.

Description

A major limitation of all of our approaches is the need for extensive labeled data for training. This classical limitation of all statistical learning-based approaches has two consequences. First, it limits the ability to rapidly add new models to the collection of models used in the system. In particular, this limitation prevents the users from defining new models that they want to find in the data based on the current environment, forcing them instead to a static collection. Ideally, we would like the user to be able to generate a new feature model by showing only one example ("one-shot" learning) or a small number ("few-shot" learning). In our scenario, this is important to enable the system to be tuned to the needs of users looking for specific visual concepts that emerge during the mission and are not pre-trained.
Second, it prevents adaptation of the recognition models to new environments or new observation conditions, since that would require extensive manual labeling of the data and expensive supervised re-training. This problem, broadly termed "domain adaptation", is crucial in the robotic setting: Indeed, it is impossible to guarantee the extent to which what the robot sees during operation matches what it had seen at training time.
It is important to note that this is a departure from the standard learning approaches which rely on increasingly large amounts of supervised data given the advent of deep learning tools and this is a critical challenge that must be addressed to be able to use the best state of the art tools on the mobile robot applications.

Scientific approach

Unbounded amounts of data can be freely acquired from any mobile platform, yet effective ways to use this wealth of data remain elusive. The proposed techniques here are based on the general idea of "learning to learn" from experience, in addition to learning from data directly. One class of approaches relies on the insight that there exists a generic transformation from a model learned from few samples to a model learned from large enough samples, and that such a transformation could be effectively identified by high-capacity regressors. Intuitively, this class of approaches attempts to model the structure of the space of useful visual models; the dual class of approaches attempts to model the structure of the visual space directly. In brief, these approaches introduce an additional unsupervised pre-training stage that exposes multiple higher layer units to large amounts of unlabeled real-world images. By encouraging these units to learn diverse sets of separators across the unlabeled data, they capture a more generic, richer description of the visual world, which decouples them from biases to a specific set of categories. Our experimental results show considerable improvement over state of the art techniques in the small sample regime as well as in standard training protocol over a broad range of vision tasks, confirming the value of using experience from unsupervised data.
In 2016 we have developed and evaluated these fundamental approaches on benchmark datasets and validated them. In this subtask, we propose to build on these approaches in three ways. First, unlike in other Computer Vision scenarios in which the data consists of static, individual images, we have access to data acquired continuously over time. We propose to exploit this temporal aspect, unique to robotics applications such as CTA. First, we propose to use sequences instead of the individual images as training data in our data- and model-space approaches described above. This implies an increase in data volume by possibly orders of magnitude, thus requiring complete reformulation of the approaches to handle that volume of data efficiently. We will develop an approach to incorporate temporal consistency in our small sample learning approaches. In brief, this involves adding a regularization term in the optimization models, which penalizes models that lead to inconsistent detection

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

in consecutive observations. Essentially, this technical approach takes advantage of the fact that, even though the data is unlabeled, temporal consistency can be measured in an unsupervised manner and used for controlling the quality of the trained models. This work will build loosely on our earlier work in the CTA on weakly-supervised learning in which addition training samples are acquired through temporal consistency (tracking) in sequences.

Second, our current approaches are limited to (sequences of) images and do not use any 3D data. In practice however 3D cues are critical and need to be incorporated in recognition models. Here again, we can have access to a wealth of 3D data from LADAR and stereo as the robot explores the environment. Third, a major goal of this subtask is to move our approaches to an online robot setting.

Collaboration: This requires 1) fast computation and integration with the rest of the system, some of the anytime techniques developed in T1C1S3A will be used here; 2) Integration in a use scenario in which the user is able to designate new objects or concepts of interest, this will require close collaboration with the HRI tasks and with subtask T1C1S1D to be able to use the small sample learning concepts in the context of semantic labeling.

Linkage to previous sub-tasks:

This subtask builds on the results already obtained in the area of small sample learning in P3-6.

Expected products:

On the scientific front, the products of this subtask include new approaches for small sample learning from images, videos, and 3D data. Small sample learning is one of the most important areas of research because of the central limitations of current, data-hungry deep learning approaches. These contributions will translate to implementations suitable for use in the system, using the CMU Husky as our primary testing platform. We target integration of these implemented results in two broad classes of scenarios, which cover a large fraction of the scenarios in which the system needs to be retrained, operationally:

- User-defined concepts: Scenarios in which the user wishes to direct the system to learn from few (or only one) observations designated by them.
- Domain adaptation: Scenarios in which the domain (environment, observation conditions, etc.) has changed sufficiently that the system needs to be retrained.

End of CY2017	• Initial implementation and evaluation of temporal constraints to take advantage of "on the move" data • Integration in a user-designation scenario for rapid designation of new objects
End of CY2018	• Implementation of temporal constraints • Integration of 3D data • Integration in an adaptation scenario in which the system updates models by using unsupervised data in the new domain

References:

1. Y. Wang and M. Hebert. Learning from small sample sets using unsupervised binary codes. European conference on Computer Vision (ECCV). 2016.
2. Y. Wang, M. Hebert. Learning by Transferring from Unsupervised Universal Sources. AAAI Conference on Artificial Intelligence. 2016.
3. Y.Wang, M. Hebert. Combining Low-Density Separators with CNNs. Proc. Neural Information Processing Systems (NIPS). 2016.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.3.2 T1C1S1B - Understanding of environment at contextual and more generalized levels

PI: Fiore, Stephen M. (UCF), sfiore@ist.ucf.edu, (407)882-0298

Objective and Benefits:

Concurrent with recognizing humans, both teammates and non-teammates, and perceiving social cues, is the ability to understand the surrounding environment a robot is embedded within. The environment is capable of providing context, even in a social scenario; thus, particularities about the environment (e.g., crowded versus deserted, urban versus rural) can be utilized as a social cue and factored into social-cognitive models during decision making. Because these computational models are developed to fit a given social signal and not a specific scenario, we avoid designing a robot capable of only functioning under narrowly defined conditions. That is, the robotic platform will be capable of addressing a wide range of situations wherein it must make judgments within a tactical context. The outcome of this, then, is the identification and description of similarities and differences between social and non-social cues associated with varying operational environments (e.g., crowded, urban, hostile).

Capability Supported: Capability I.1. This research will primarily support the RCTA's capability "Optempo Mobility in Dynamic Scenes: Fast Motion through Crowded Areas". Accordingly, our scientific approach will focus specifically upon social and non-social cues as they relate to movement within a dynamic environment.

Scientific Approach:

Efforts completed in the previous years of the RCTA subtasking serve to support the work proposed herein. Subtask H9.2 enabled exploration of the judgment process concerning the appropriate behavioral response for a specific social signal given input of particular social cues (Fiore et al., 2013; Lobato et al., 2015). While this work concentrated on understanding social agents within the environment, our empirical evidence also has the potential to provide insight into judgment processes utilized to gain an understanding of the environment itself. Following our prior research (Best et al., 2016a; Best et al. 2016b), in subtask H9.3, we are currently investigating social cognitive models using simulated human-robot interaction scenarios to benefit capabilities pertaining to perceptual understanding, tactical movements, socially conscious behaviors, and environmental considerations (in progress). The results of this empirical study will directly support the current tasking, affording an extended awareness of how individuals function from a social-cognitive standpoint from environment to environment when faced with social agents presenting novel and varied scenarios. These findings can be used within a model of cognition to enable a system programmed to recognize social cues and quantify them accordingly, weighting social cues differentially depending upon the constellation of cues expressed. This would allow the system to probabilistically estimate the likelihood of a given social signal being conveyed.

In service of the current tasking, we will continue with the two-phased approach to our research. This has allowed us to contribute numerous publications to the basic science of human-robot social cognition (Kapalo et al. 2016; Warta et al., 2016; Wiltshire & Fiore, 2014; Wiltshire et al. in press). Specifically, this research subtask will identify the social and non-social cues associated with varying operational environments. First, complementing our earlier work (Wiltshire et al., 2014), we will examine the extant literature through extensive reviews of relevant research to produce a taxonomy capable of differentiating between subtypes of social and environmental cues. Second, we will conduct empirical studies to investigate the salience of environmental and social cues and their role in creating an understanding of the surrounding environment. Literature reviews will facilitate decisions concerning *what* cues to investigate and *how* best to examine them based on past empirical research. These efforts will provide the foundation needed to move forward and begin development of computational models to fit a given social signal. These models will be developed from human-generated data obtained from participants in our empirical experimentation to provide a realistic and accurate simulation of the cognitive processes necessary for supporting capabilities addressing a wide range of situations wherein a robot must make judgments within a tactical context. This human-generated data, based upon their assessments of social and environmental cues in

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

the simulation/experimentation settings, will allow us to determine the presence and numerically-weighted salience of different social cues in forming a communicated social signal. In essence, this would allow us to create a probabilistic model of social cognition. Concurrent with this, we will utilize empirical experiments to verify whether the constellation of social cues established for a social signal in a given scenario will generalize to an alternative scenario (e.g., crowded versus deserted, urban versus rural).

List of dependencies: T2C1S3A (Oh, CMU) proposes to use deep learning to associate visual features with language models. Similarly, T2C1S3B (Walter, TTI) proposes to generate natural language descriptions of the robot's surroundings. In relation to both of these subtasks, our work could provide the foundation for the labeling of social cues perceived in the environment by utilizing the participant-generated data obtained from our empirical experimentation. Additionally, our quantitative models demonstrating the salience of social cues could serve to inform CMU's proposed research of attention-based neural models since the salience of a social cue indicates the focus of attention as well as support TTI's determination of the salience of objects in the environment. With regards to T2C1S4B (Barber, UCF), our social-cognitive model could be used to support the proposed interface in determining the social signals communicated by soldiers. Given that our proposed model would take into account both social and non-social cues within the surrounding environment, this could provide important information to help determine when and how the system should adapt.

Collaboration: Previous collaboration has taken place between Fiore and Bölöni to investigate social cues and signals, and their connection to social-cultural judgments. Given that the shared goal is prediction of human behavior, future collaborations are probable and would be better supported through the use of an easily accessible central repository for adapting, testing, and storage of computational models.

Linkage to previous RCTA BPP sub-tasks:

H9.2: Applications of Social Signal Processing to HRI

- Explore the judgment process concerning the appropriate behavioral response for a specific social signal given input of particular social cues (Best et al., 2016; Fiore et al., 2013; Lobato et al., 2015).

H9.3: Robotic Social Cues and Signals Experimentation

- Investigate social cognitive models using simulated human-robot interaction scenarios to support: perceptual understanding, tactical movements, socially conscious behaviors, and environmental considerations (in progress).

Expected products:

This subtask will generate the following products:

Deliv. by end of	Products
2017	• Develop a taxonomy that specifies the numerous social and non-social cues associated with varying dynamic operational environments. • Utilize empirically-based experimentation in conjunction with simulation to investigate the salience of environmental cues and their role in creating an understanding of the surrounding environment.
2018	• Enable development of computational models to fit a given social signal and simulate the cognitive processes necessary for supporting

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

	capabilities addressing a wide range of situations wherein a robot must make judgments within a tactical context. • Verify through empirical experimentation whether the constellation of social cues established for a social signal in a given scenario will generalize to an alternative scenario.

References:

Best, A., Kapalo, K., Warta, S. F., & Fiore, S. M. (2016a). Clustering social cues to determine social signals: Developing learning algorithms using the "N-Most Likely States" approach. In SPIE Defense + Security Unmanned Systems Technology XVIII. 9837, 98370L-98370L-15. doi:10.1117/12.2223900.

Best, A., Warta, S. F., Kapalo, K., & Fiore, S. M. (2016b). Of Mental States and Machine Learning: How Social Cues and Signals Can Help Develop Artificial Social Intelligence. *Proceedings of the 60th Annual Meeting of the Human Factors and Ergonomics Society* (pp. 1361-1365). Santa Monica, CA: Human Factors and Ergonomics Society.

Fiore, S.M., Wiltshire, T.J., Lobato, E.J.C., Jentsch, F., Huang, W.H., & Axelrod, B. (2013). Towards Understanding Social Cues and Signals in Human-robot Interaction: Effects of Robot Gaze and Proxemic Behavior. *Frontiers in Psychology: Cognitive Science*. 4:859. doi: 10.3389/fpsyg.2013.00859.

Kapalo, K.A., Phillips, E., & Fiore, S.M. (2016). From Animals to Androids: Leveraging the Human- Animal Team Model to Inform Human-Robot Interaction. *Proceedings of the 60th Annual Meeting of the Human Factors and Ergonomics Society* (pp. 1225-1229). Santa Monica, CA: Human Factors and Ergonomics Society. (1st Place, Best Student Led Paper Award, Cognitive Engineering and Decision Making Technical Group).

Lobato, E. J. C., Warta, S. F., Wiltshire, T. J., & Fiore, S. M. (2015). Varying social cue constellations results in different attributed social signals in a simulated surveillance task. *Proceedings of FLAIRS-28, Twenty-Eighth International Florida Artificial Intelligence Research Society Annual Conference* (pp. 61-66). Hollywood, FL. AAAI Press.

Warta, S. F., Kapalo, K., Best, A., & Fiore, S. M. (2016). Similarity, Complementarity, and Agency in HRI: Theoretical Issues in Shifting the Perception of Robots from Tools to Teammates. *Proceedings of the 60th Annual Meeting of the Human Factors and Ergonomics Society* (pp. 1230-1234). Santa Monica, CA: Human Factors and Ergonomics Society. (3rd Place, Best Student Led Paper Award, Cognitive Engineering and Decision Making Technical Group).

Wiltshire, T. & Fiore, S.M. (2014). Social Cognitive and Affective Neuroscience in Human–Machine Systems: A Roadmap for Improving Training, Human–Robot Interaction, and Team Performance. *IEEE Transactions on Human-Machine Systems, 44(6),* 779-787.

Wiltshire, T. J., Lobato, E. J. C., Velez, J., Jentsch, F. G., & Fiore, S. (2014). An interdisciplinary taxonomy of social cues and signals in the service of engineering robotic social intelligence. In Proceedings SPIE 9084, Unmanned Systems Technology XVI, 90840F (June 3, 2014); doi:10.1117/12.2049933.

Wiltshire, T.J., Warta, S.F., Barber, D., & Fiore, S.M. (in press). Enabling Robotic Social Intelligence by Engineering Human Social-Cognitive Mechanisms. *Cognitive Systems Research*.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.3.3 T1C1S1C - Scene Understanding at the Tactical Edge

PI: Rao, Raghuveer (ARL), raghuveer.m.rao.civ@mail.mil, (301)394-0860
Additional PIs: Gurram, Prudhvi (ARL), gurram_prudhvi@bah.com, (301)394-0464
Narayanan, Priya (ARL), (301)394-2376

Objective and Benefits:

There is need for gathering real time intelligence in various tactical scenarios, e.g. Cordon and Search, to help Soldiers gain situational understanding and make quick mission-oriented decisions. In future situations, Soldiers may deploy autonomous teammates with imaging and video sensors onboard for purposes such as detection and location of adversary warfighters, vehicles, weapons, and related activities. The goal of this task is to automatically detect and localize relevant objects in the scene from the perspective of an autonomous system.

Scientific approach:

Deep convolutional neural networks (CNNs) based algorithms are the current state-of-the-art object detectors [1, 2]. They are able to provide high detection rates because of their ability to exploit non-linear features in the imagery and to build versatile non-linear classification functions. But this non-linearity comes at the cost of millions of parameters that need training, which in turn needs millions of labeled samples corresponding to different objects. There are publicly available datasets with millions of traditional images (e.g. ImageNet) on which CNNs can be trained effectively [3]. However, the class distributions of objects observed in traditional imagery and Army operational imagery are significantly different. Hence, the CNN-based object detectors trained to work well on traditional imagery (source domain) do not work at the same level of accuracy on the data in an operational context (target domain). Also, labeled training datasets are not available in the target domain to re-train the algorithms or perform traditional transfer learning such as "fine-tuning" [4]. Hence, the goal of this subtask is to investigate and develop methods to perform unsupervised transfer learning across source and target domains. The input data for the source domain are labeled images obtained from widely available traditional datasets like ImageNet [5], Microsoft COCO [6], etc. The input image data for the target domain (tactical scenario) will be collected internally at ARL. If successful, the superior performance of supervised object detection and recognition algorithms using deep convolutional neural networks (CNNs) that are trained on source domain can be leveraged in the target domain of an operational scenario without the need for labeled training data from the target domain.

List of dependencies:

The proposed subtask is closely related to the following subtasks. The collaboration can happen at both the research level (exchange of ideas) and the peripheral level (output of one subtask goes in as the principal or as an auxiliary input of the other subtask).

1. T1C1S1B – Understanding of environment at contextual and more generalized Levels.
2. T1C1S1D – Multimodal, multiscale reasoning for scene understanding.
3. T1C2S2A – Perception for maneuver in complex terrain.
4. T1C2S2D – Perception for optempo mobility in rough terrain.
5. T2C1S3B – Generating natural language scene descriptions for shared situational awareness.

Collaboration:

This effort will also include interaction with Dr. Larry Mathies' group at JPL/Caltech to find complementary aspects and integrate our work with their work for object detection and activity recognition tasks. We will also

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

explore opportunities to collaborate with Dr. Stuart Young's group at ARL to implement the developed algorithms on robotic platforms available at ARL.

Linkage to previous RCTA BPP sub-tasks:
This is newly proposed research, and does not have linkage to any subtasks in the past two years.

Expected products:

1. At the end of the three-year period, we expect the algorithm to run on a robot while generating detections (bounding boxes) of different objects – people, vehicles, and weapons in a scene. The locations and sizes of the bounding boxes will be added to the world model. The subtask will be executed with 4 different milestones: 1) Algorithm development during 2017 and 2018; 2) Data collection using a ground based sensor (preferably on a robot) from a tactical scenario during 2017; 3) Evaluation of the developed algorithm on the data during 2018; 4) Evaluation of the developed algorithm on a RCTA robotic platform (Husky) during 2019.
2. The other RCTA members will be updated on the progress of the subtask at the annual RCTA PI meetings. The following conferences and workshops are possible venues for publishing results of the proposed research work (at least 2 conference and workshop papers): Computer Vision and Pattern Recognition (CVPR), International Conference on Computer Vision (ICCV), European Conference on Computer Vision (ECCV), IEEE Winter Conference on Applications of Computer Vision (WACV), International Conference on Machine Learning (ICML), International Conference on Learning Representations (ICLR), Neural Information Processing Systems (NIPS), IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).

References:

[1] Girshick, R., Donahue, J., Darrell, T., & Malik, J. "Rich feature hierarchies for accurate object detection and semantic segmentation," Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (2014).
[2] Girshick, R. "Fast R-CNN," Proc. of the IEEE Inter. Conf. on Computer Vision (2015).
[3] Krizhevsky, A., Sutskever, I., and Hinton, G. E., "Imagenet classification with deep convolutional neural networks," Proc. of the Advances in neural information processing systems (2012).
[4] Bengio, Y., Lamblin, P., Popovici, D., & Larochelle, H., "Greedy layer-wise training of deep networks," Proc. of the Advances in neural information processing systems (2007).
[5] http://www.image-net.org/
[6] http://mscoco.org/home

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.3.4 T1C1S1D - Multimodal, multiscale reasoning for scene understanding

PI: Pires, Bernardo (CMU), bpires@andrew.cmu.edu, (412) 268-3818.

Description

The objective of this task is to develop algorithms that generate *more complete and more coherent representations of the environment*. These algorithms are relevant to this Capability since they provide more powerful perception capabilities to support robot operation under *uncontrolled observation situations*. Additionally, these algorithms will reduce the workload of the World Model by taking care of inferring complete environment representations at the perception level. This task supports primarily Thrust 1 Capability 1, and also contributes to other capabilities that require *understanding the environment statically*.

Scientific approach

In previous work we have developed approaches and algorithms that produce scene representations at different degrees of description: ranging from a coarser region categorization, down to the object level, and finally leading to a finer grained description of object details. This work addressed this integrated representation on three different fronts: the fusion of pixel-level labeling with object- and feature-level labeling; the extraction of detailed features by parsing the scene at the object-level; and the use of 3D data to resolve spatial ambiguity from images. State of the art approaches, though capable of generating multiple interpretations of the environment with varying degrees of granularity, still suffer from issues derived from ambiguity when merging dissimilar models—generated by different predictors or data types, and at several scales and resolutions—and from the use of insufficient or noisy data.

In this task, we seek to build on the approaches developed in previous years to address these shortcomings in three ways. First, we intend to further increase the ability to deduce knowledge from human demonstration and use it to correctly find descriptions at other levels to match the actual scene. In the past we have applied an inverse reinforcement learning approach—developed in previous years of the RCTA [1]—to derive knowledge from human demonstrations which in turn is used to correctly and efficiently fit grammars to describe many different architectural styles for façade parsing [2]. We plan to extend this approach to cover more instances of descriptions from different modes of data and at different levels of granularity.

Second, we propose to expand the use of geometric reasoning from 3D data—previously used mainly to correct geometric errors from image-only processing —to support the semantic labeling from images by providing additional hypothesis at the object level. To this end, we plan to leverage previous work on semantic labeling of 3D point clouds, which allows parsing a point cloud directly using geometric features. We will investigate approaches to efficiently combine hypothesis from scene parsing algorithm using images and point clouds to increase the accuracy and granularity of the scene descriptions, and for transitioning across varying levels of depth information, depending on the range at which the scene is observed.

Third, we will continue the develop of semantic labeling tools based on deep learning with an emphasis on limited computational and data resources both at training and testing time. Specifically, at training time we need to rapidly train semantic labeling tools from limited data in order to quickly adapt to changing environments and observation conditions, in contrast to traditional deep learning approaches which rely and very large volumes of training data. For this, we will leverage the results of our T1C1S1A task on learning from small samples. At run time, we must execute semantic labeling with limited computation on the fly as the robot is on the move. For this, we will leverage our work on general techniques for anytime approaches to inference in subtask T1C1S3A.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

We will explore collaboration opportunities with T1C1S1C (Gurram and Rao, ARL). This collaboration could include exchange of ideas, as well as exchange of outputs and results from one task as a principal or auxiliary input to the other.

Linkage to previous sub-tasks

This task continues the work done in subtask *P3-3 Integrated Reasoning for Scene Understanding and Object Recognition* and in subtask *P4-2 Integrated Reasoning with Non-visual Information*. It extends it further into more precise structured description and it incorporates the key constraints of the CTA scenarios: limited data and limited computation.

Expected products

The algorithms developed in this task will be integrated into a software library to simplify its use by the RCTA demonstration platforms. We will perform integration using our own robot platforms (Clearpath *Husky*). The end result will be an implementation of an integrated scene understanding system capable of generating descriptions of the scene at different levels, using images and 3D data. Additionally, we will conduct an experimental characterization of performance. The library will be developed and released progressively according to the following milestones (the "implementation" items are the products intended to be used in the integrated assessment):

End of CY2017	• Implementation of framework for descriptors learned from human demonstration • Basic implementation of multi-scale scene understanding • Basic implementation of online training and adaptation
End of CY2018	• Implementation of extended multi-scale version, including multimode descriptions from image and 3D data • Implementation of online training and anytime computation • Experimental characterization of performance

References

1. A.Laddha and M. Hebert. Improving semantic scene understanding using prior information. Proc. SPIE 9837, Unmanned Systems Technology XVIII. 2016.
2. A.Suppe, L. Ernesto Navarro-Serment, D. Munoz, A. Suppe, L. Ernesto Navarro-Serment, D. Munoz, J. Andrew (Drew) Bagnell, and Martial Hebert. An Architecture for Online Semantic Labeling on UGVs. Proc. SPIE 8741, Unmanned Systems Technology XV, April, 2014.
3. J. Oh, L. Navarro-Serment, A. Suppe, A. Stentz , and M. Hebert. Inferring door locations from a teammate's trajectory in stealth human-robot team operations. Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), October, 2015.
4. J. Oh, A. Suppe, F. Duvallet, A. Boularias, J. Vinokurov, L. Navarro-Serment, O. Romero, R. Dean, C. Lebiere, M. Hebert , and A. Stentz. Toward Mobile Robots Reasoning Like Humans. Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI), January, 2015.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.3.5 T1C1S2A - Modeling interactions and action

PI: Navarro-Serment, Luis (CMU), lenscmu@ri.cmu.edu, (412)268-6034.

Objective and Benefits:

The objective of this task is to develop algorithms to *predict human behavior in dynamic environments considering heterogeneous movers*. These algorithms are relevant to this Capability since the online planning of robot actions through multiple agents requires the timely determination of the expected motion of all the participants in the dynamic scene. This task supports primarily Thrust 1 Capability 1, but is a basic building block of all the capabilities that require operation in dynamic environments.

Scientific approach:

In previous research, our work focused on developing predictive models of pedestrian dynamics assuming interactions between movers of the same type—in our particular study, all movers were humans with the same physical characteristics. However, the more complex, uncontrolled scenarios considered in the RCTA—and in this Capability in particular—may involve dynamic scenes with heterogeneous types of movers. Different types of movers typically move purposefully towards different types of goals, while reacting in dissimilar ways to features in the environment. These differences rapidly increase the computational complexity of the interaction dynamics and make it difficult to compute a solution in real time.

State of the art approaches either ignore [1] or restrict [2] the interactions between agents. Some researchers have begun to explore game theoretic approaches to account for multiple interactions, but limited to homogeneous sets of movers [3]. In this task, we propose to build on our previous decision-theoretic approach and extend it in two principal ways. The first one is by learning models for other types of movers. To this end, we consider a progression, starting with movers whose motion is constrained (e.g. bicycles and cars), followed by additional models of less constrained movers, like pedestrians of different ages. The second way is by investigating online adaptation techniques to identify and select the sets of features that should be used for each mover type. We will leverage work done in the RCTA in static scene understanding to detect environment features and to provide a context for the observed behavior. This research will advance the state of the art by developing an approach to obtain predictive models for heterogeneous dynamic scenes. The expected outcome is a system that can generate predictions of future trajectories for a heterogeneous group of moving objects—represented as sequences of waypoints for each mover—which are made available to other modules or subtasks.

List of dependencies:

The algorithms considered in this subtask are indifferent to the type of sensor or system used to detect and track movers: they only require updates of the location of each mover. This subtask can take advantage of the tracking information produced by subtask T4C3S1D – Human Location and Platform Interfaces. Similarly, these algorithms require grid maps of environment features (e.g. types of floor, obstacles) to establish a context for predictions. To this end, we plan to use the World Model to extract the information we need to populate these grids. In particular we plan to apply the output from the semantic labeling modules developed in previous years of the RCTA [4], which currently feed the World Model, to obtain feature maps at the region level. The training data sets used in this subtask will consist of tracking information and the grid maps. Finally, the predicted trajectories can be used by subtasks that reason about the dynamic scene, such as T1C1S4B—Navigation among dynamic, reactive agents.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Collaboration:
We will continue ongoing collaborations with T1C1S4B (Jean Oh, CMU) at the research level.

Linkage to previous RCTA sub-tasks:
This task continues the work done in subtask *P5-1 Modeling interactions*.

Expected products:
The algorithms developed in this task will be consolidated into a software library or API for easy integration into the RCTA demonstration platforms. We will perform integration using our own robot platforms (Clearpath *Husky*). The result will be an implementation of a perception system capable of taking tracking data of all moving objects in the scene, and then generating expected trajectories to be traversed by all of them. This system will be integrated and deployed in an incremental fashion over the course of two years. Additionally, we will conduct an experimental characterization of performance.

The library will be developed and released progressively according to the following milestones:

End of CY2017	• Initial release, implementing preliminary version of prediction for multiple heterogeneous movers, including at least one of each constrained- and unconstrained-motion mover types. • Publication in peer-reviewed conference. • Initial experimental characterization.
End of CY2018	• Refined version, including enhanced implementation suitable for qualification by IRA process and full integration on platform. Enhancements include algorithmic improvements and additional mover types. • Evaluation of performance of enhanced algorithms. • Publication in peer-reviewed conference.

References:

1. Karazev, V., Ayvaci, A., Heisele, B., and Soatto, S. Intent-aware long-term prediction of pedestrian motion. In *Proc. of the International Conference on Robotics and Automation ICRA*, May, 2016.
2. Lee, N., and Kitani, Kris M.. Predicting Wide Receiver Trajectories in American Football. *Winter Conference on Applications of Computer Vision (WACV)* 2016.
3. Ma, W.-C., Huang, D.-A., Lee, N., and Kitani, K. M.. A game-theoretic approach to multi-pedestrian activity forecasting. *arXiv:1604.01431v2*.
4. Suppe, A., Navarro-Serment, L., Munoz, D., Bagnell, J. A. and Hebert, M. An Architecture for Online Semantic Labeling on UGVs. *Proc. SPIE 8741, Unmanned Systems Technology XV*, April, 2014.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.3.6 T1C1S2C - Human-robot teaming in a social context

PI: Bölöni, Ladislau (UCF), lboloni@cs.ucf.edu, (407)823-2320

Objective and Benefits:

This research seeks to develop techniques for mixed-robot teaming enacted in a social context. In particular, we assume that the robot needs to operate in the presence of multiple human and robotic teammates, civilian bystanders and potentially enemies. We also assume that the robot is able to take different roles in the team (scout, lookout, protecting left or right flank, guide and so on). Due to the unpredictable nature of the unstructured environments, the team often needs to adjust its mission, which might also require changes in the roles assigned to humans. Our goal is to develop algorithms that allow the robot to make communication decisions, plan and role change decisions.

Scientific approach:

In order to describe a specific human robot team structure, we need to use models that are (a) understandable by humans and (b) formally rigorous enough to be operationalizable for the use by the robot. We will investigate techniques including temporal logic based descriptions, narrative descriptions, and various models of apprenticeship learning. Some of these models can also represent explicit commands and their expected effects. Of particular interest is to investigate whether formal specification approaches or learning-based approaches are a better model for acquiring this data. We conjecture that for implicit communication, which is essentially a pattern recognition problem, learning based approaches are the only viable technique that can work in an unstructured environment. We will investigate deep neural network based approaches, in particular LTSM neural nets which our team had successfully used to learn robot trajectories in simulation and transfer them to physical robots [1]. Nevertheless, deep learning has its own limitations. It requires large training sets, especially when used in reinforcement learning mode. It also makes it relatively difficult to describe the learned behavior in a human-understandable way. Thus we will also investigate other modeling techniques, such as hidden Markov models and dynamic Bayes nets.
The typical scenario we consider, will include a human-robot team, such as a guide, a VIP, and 2-3 bodyguards executing a mission in a public space. Some of these roles are specific to humans, others can be fulfilled by robots as well. The roles can change during the mission. All team members are sensing the environment, interpret explicit and implicit communication from team members and other persons. Specific situations can lead to change of plans and/or role reassignments (abort mission, take cover etc).

We will perform integration on our own small robot platforms (Clearpath Jackal with a Kinova Mico arm, Lynxmotion AL5D+A4WD1), and transition the code to the shared repositories for integration on the RoMAN platform.

Stretch goals include focus on interaction / communication with civilians. In its current form, the proposal deals with the interactions of human and robotic teammates. The civilians present in the public space are evaluated for threat level and obstacles, their movement is predicted but not influenced by the robot. A possible extension of this research involve behaviors where the robot initiates explicit or implicit communication with the bystanders in order to better achieve its goals.

Collaborations:

- T1C1S1B (Fiore) – using the same family of computational models for the representation of social signals and metrics. We will use a shared repository for these.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

- T2C1S4B (Barber, UCF) – the research framework developed to identify the effectiveness of communication and its cognitive burden can be used to study the effectiveness of the human / robot team member synchronization tasks in this task.
- T2C1S3A (Oh, CMU) – use deep learning models based on CNN and LSTM to connect visual features and language models. While there is no direct dependency, the common use of the CNN/LSTM models open the possibility of sharing network structures, training data or learned features.
- T2C1S3B (Walter, TTI) - the content selection model proposed as part of this task can inform the implicit non-verbal communication model.

Linkage to previous RCTA BPP sub-tasks:
This work follows up the work done in H8.4 from 2015-2016 [2,3].

Expected products:

- End of BPP Year 2017: Implement the capability for a robot to perform a role reassignment as a result of explicit and/or implicit communication from human team members. Demonstrate the capability in simulation and on a small-scale mobile robot in lab environment.

- End of BPP Year 2018: Implement the capability for a robot to made role reassignment decisions in the presence of bystanders and changing mission. Demonstrate the capabilities on a small scale robot in lab and partially controlled environment (bystanders aware of the robot).

- We will target the following conferences:

    - IEEE International Conference on Robotics and Automation (ICRA)
    - IEEE/RSJ International Conference on Robotics and Systems (IROS)
    - IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)
    - International Conference on Artificial Intelligence (AAAI) and International Conference on Autonomous Agents and Multiagent Sytems (AAMAS).

References:
[1] Rouhollah Rahmatizadeh, Pooya Abolghasemi, Aman Behal, Ladislau Bölöni "Learning real manipulation tasks from virtual demonstrations using LSTM" arXiv:1603.03833
[2] T.S. Bhatia, G. Solmaz, D. Turgut, and L. Bölöni. Two algorithms for the movements of robotic bodyguard teams. In Proc. of Workshop on Knowledge, Skill, and Behavior Transfer in Autonomous Robots, pp. 2–8, January 2015.
[3] S. Arif, S.A. Khan, and L. Bölöni. Balancing predicted mission cost and social costs by mobile robots navigating a crowd. In Proc. of Autonomous Robots and Multirobot Systems (ARMS) workshop at AAMAS-2014, May 2014.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.3.7 T1C1S3A - Anytime neural networks for semantic understanding on the move

PI: Muelling, Katharina (CMU), kmuelling@nrec.ri.cmu.edu, (412)692-1097

Objective and Benefits

This subtask is devoted to enable computer vision and machine learning approaches to scale to real world tasks by trading off the computational effort (cost) and the quality of the prediction. The focus here is on anytime predictors: algorithms that are able to automatically adjust to and efficiently use any computational budget. The advantage of anytime prediction is that it can be interrupted at any test-time and still provides valid results. The more time available to the predictor, the better the result. Anytime prediction is particularly useful for this Capability, because anytime models dynamically trade off computation and accuracy to address the various needs of a robot in a dynamic, changing, and unknown environment. For example, a fast-moving robot on open plain needs fast but simple detection of obstacles, but a stationary robot making a mapping needs accurate and complex scene understanding. Unfortunately, the state-of-the-art perception and planning algorithms typically do not have anytime prediction capability.

Scientific approach

Our previous work, which was funded under Task I5-1: 'Learning Algorithms for deep understanding and real-time performance', focused on anytime linear prediction, and proved that our fast and simple greedy algorithms can make near-optimal linear predictions when interrupted [1,2]. However, this work relies on boosting, which suffers when scaling to large datasets. Furthermore, the predictions were restricted to generalized linear models. This year, we developed streaming boosting algorithms with theoretical performance guarantees, which allowed us to scale anytime algorithms to large datasets.

In this task, we develop neuronal networks with anytime properties, which allow neural networks to generate competitive anytime predictions before the full network is computed. Anytime neural network can be applied to perception task so that the robot can receive the state-of-the-art scene understanding of various computational cost depending on situations. Furthermore, we can design the training of anytime neural networks to produce usable models in an anytime fashion: simpler models converge faster and are available sooner than complexer ones.

The suggested anytime algorithms can be applied to both outdoor and indoor images provided by a camera. For initial training, we will use public available data-sets such as Pascal VOC, Imagenet, and Cityscape. The algorithms require labelled images in which at least the object of interest is labeled. The anytime networks generate labels for image pixels, x/y object positions, or bounding boxes which can be used to update the world state of the system.

List of dependencies

The project depends on the perception task T1C1S1D.

Collaboration

We expect collaborations with Martial Hebert (T1C1S1A) to integrate our models and work towards anytime learners able to learn from few samples. Other possible collaborations include Raghu's task T1C1S1 to explore the possibility to apply anytime augmentation to their learners and Oh (T2C1 S3A) to evaluate our approaches on the Husky platform.

Linkage to previous RCTA BPP sub-tasks

Task I5-1: 'Learning Algorithms for deep understanding and real-time performance'.

Expected Product

From a general scientific standpoint, the subtask will produce the first anytime approach to neural network. For a robotics standpoint, the need for anytime prediction is pervasive in CTA and the algorithms developed in this task will be integrated in the object recognition and scene understanding components of the system using the Husky platform.
The subtask will generate the following products:

Deliv. by end of	Products
2017	- Develop an anytime neural network whose anytime classifications are competitive against the state-of-the-art networks trained specifically for the same computational cost. - Evaluation on public classification data-sets and submission for publication. - Implementation for object detection tasks
2018	- Integrate anytime neural network with perception tasks such as anomaly detection scene understanding. - Implementation of anytime prediction for other relevant perception tasks, such as semantic labeling

Target publication venues for the suggested work includes NIPS, CVPR, AAAI, ICML and UAI.

References:

[1] H. Hu, A. Grubb, J.A. Bagnell, M. Hebert: Efficient Feature Group Sequencing for Anytime Linear Prediction. UAI (accepted)
[2] A. Grubb; J.A. Bagnell*: SpeedBoost: Anytime Prediction with Uniform Near-Optimality.* Fifteenth International Conference on Artificial Intelligence and Statistics, 2012

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.3.8 T1C1S4B - Navigation among dynamic, reactive agents

PI: Oh, Jean (CMU), jeanoh@nrec.ri.cmu.edu, (412) 651-6052

Objective and Benefits:

In this task, we aim to tackle the following research questions: How can a robot navigate through a dense, moving crowd whose route choices will be affected by the robot's path decisions? How can we model the interactions among the agents in the crowd and between the robot and the crowd so that the robot's navigation path is safe and socially compliant?

Scientific approach:

In our previous task I1-1 Framework for Collaborative Intelligence, we demonstrated semantic navigation in static environments [Oh2015;Oh2016]. In dynamic environments, we developed an algorithm that is complete and efficient with a quality guarantee with a bound, relying on the assumption that the trajectories of other agents are known a priori [Vemula2016]. In proposed work, we relax this assumption and focus on planning based on predicted trajectories of other agents. Toward this goal, we will investigate algorithms for learning cooperative joint collision avoidance model. We will verify our hypothesis that a joint prediction model (as opposed to a conditional prediction model) can capture collision avoidance behaviors and lead to more efficient, safe and socially compliant navigation. We will focus on developing algorithms that can perform robustly on a long-term horizon and that can provide a way to measure uncertainty. In a soldier-robot teaming scenario, it is important for the robots to be able to reason about the uncertainty in their predictions because over-confident predictions can result in critical damages to the robots, their teammates as well as their environment. Based on the estimated level of uncertainty, we propose to develop suitable high-level navigation strategies. For this, we will conduct research on Bayesian approaches to prediction that can provide not only future predictions of trajectories but also how uncertain the predictions are. Whereas T1C1S2A (Navarro-Serment, CMU) focuses on modeling various types of movers with motion constraints, this task will focus on learning the joint interaction model to plan to navigate safely in dynamic environments.

We generally follow human-like reasoning in this task, meaning that, as opposed to relying on accurate metric mapping information, our approach aims to establish a topological map that uses relative spatial relationships to known landmarks and other reference objects. We note that it is, however, still important to perform accurate and timely obstacle detection in order for a local planner to avoid immediate collisions. In the current system, both objects (detected by the perception systems) and obstacles are stored/accessed from the World Model as a central storage. We will collaborate with the perception group to define a set of RCTA standard communication messages (e.g., ROS messages) for sharing perception related input and output information; this effort will extend the existing set of ROS object messages used in the previous program.

List of dependencies:

In order to run experiments on the Husky platform, task T1C1S1A—maintaining the Husky platform in working order—will be crucial. The output from the navigation planner will be sent to a local planner (or waypoint follower) on the Husky platform. Additionally, the learned mover models from T1C1S2A could be used an input to this task.

Collaboration:

This research will require collaboration with task T1C1S2A for modeling interactions (Navarro-Serment, CMU) and T2C1S4A for trust building in human-robot teams (Jentsch, UCF). We will also collaborate with the perception tasks T1C1S1D (Hebert, CMU) and T1C1S3A (Muelling, CMU) for the perception communication protocol design.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Linkage to previous RCTA BPP sub-tasks:
I1-1 Framework for Collaborative Intelligence

Expected products:
We will evaluate the performance of our method against state-of-the-art approaches such as IGP or social LSTM on public datasets in the first year, followed by robot experiments on the IRA Husky platform in the second year. Target publication venues include AAAI, IJCAI, R:SS, IROS, ICRA, AAMAS, ICAPS, and ISER.

2017	• Continue developing GP based prediction algorithms that take social interaction into account and that performs robustly on long-term horizon. • Develop a neural net (LSTM) based algorithm and extend it to model uncertainty in prediction. • Evaluate on video + trajectory data sets
2018	• Extend the model to include visual (or environmental) features in prediction and planning. • Develop high level navigation planning algorithm to be able to choose proper strategies according to the level of estimated uncertainty. • Evaluate on RCTA IRA platforms (ClearPath Husky).

References:
[Oh2015] J. Oh, A. Suppe, F. Duvallet, A. Boularias, J. Vinokurov, L. Navarro-Serment, O. Romero, R. Dean, C. Lebiere, M. Hebert, and A. Stentz. Toward mobile robots reasoning like humans. In Proc. of AAAI Conference on Artificial Intelligence (AAAI), 2015.

[Oh2016] J. Oh, M. Zhu, S. Park, T.M. Howard, M.R. Walter, D. Barber, O. Romero, A. Suppe, L. Navarro-Serment, F. Duvallet, A. Boularias, J. Vinokurov, T. Keegan, R. Dean, C. Lennon, B. Bodt, M. Childers, J. Shi, K. Daniilidis, N. Roy, C. Lebiere, M. Hebert, and A. Stentz. Integrated intelligence for human-robot teams. In Proc. of International Symposium on Experimental Robotics (ISER) 2016.

[Vemula2016] A. Vemula, K. Muelling, J. Oh. Path Planning in Dynamic Environments with Adaptive Dimensionality. In Proc. of International Symposium on Combinatorial Search (SoCS) 2016.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.3.9 T1C1S4D - Modeling trajectories for robust perception

PI: Navarro-Serment, Luis (CMU), lenscmu@ri.cmu.edu, (412) 268-6034.
Collaborators: Wigness, Maggie (ARL), maggie.b.wigness.civ@mail.mil, (301) 394-3927.
Rogers, John (ARL), john.g.rogers59.civ@mail.mil, (301) 394-3927.

Objective and Benefits:
The objective of this task is to develop algorithms for *modeling and transferring sensor maneuvering patterns for enhanced perception*. These algorithms are relevant to this Capability since they support robot operation under *uncontrolled observation situations* by providing a mechanism to transfer predictive behavior models from humans to robots. This task supports primarily Thrust 1 Capability 1, and also contributes to other capabilities that require operation in dynamic environments.

Scientific approach:
In previous work we have focused on the application of a maximum entropy inverse reinforcement learning approach to predict pedestrian behavior. This technique allows the creation of predictive models that best describe the motion patterns of a pedestrian who moves purposefully to reach a goal and reacts to the features present in the environment. In collaborative work between researchers from ARL and CMU, we have explored the use of these models to generate paths that can be used by a robot to traverse the environment while reacting to environmental features in the same way a person would. A key aspect of this work is that motion behaviors demonstrated by humans are effectively transferred to the robot. This is an important ability that can facilitate the adaptation of robot behavior to the uncontrolled scenarios considered in the RCTA, since models do not need to be handcrafted for robots: instead, models are learned by human demonstration in the presence of different features and then transferred to the robot.
In this task, we seek to extend the ability of predictive models to include targeting maneuvers for perception tasks. Specifically, we consider situations where not only it is necessary to plan a path between two locations, but it is also important to maintain a specific vantage point with respect to a target object or area. State of the art approaches have successfully trained agents using apprenticeship learning techniques. However, these methods have focused on problems involving the agent, the environment and other agents, but without encoding any relationships that are relevant for perception tasks [1,2]. For example, a vision sensor operating outdoors not only must be aimed towards the area of interest: it also must avoid being dazzled by the sun, and placed at the proper distance. In this task, we propose to build on our previous approach and extend by creating features that encode the spatial relationship between a mobile sensor and its target. We propose to start our study by focusing on mobile vision sensors aimed at specific static targets for perception tasks (e.g. semantic scene labeling). Next, we will study encodings of spatiotemporal features which are suitable for dynamic targets. We will leverage work done in the RCTA in static scene understanding to detect environment features and to provide a context for the observed behavior. This research will advance the state of the art by applying a Maximum Entropy Inverse Reinforcement Learning approach [3] to create predictive models that represent the trajectory and orientation that a human would follow to collect information for a perception task, and to transfer this ability to a robot.

List of dependencies:
This subtasks depends on information generated by T1C1S4E - Learning traversability cost from unsupervised semantics. For this study we will rely on one of ARL's Husky robot (property of the Asset Control and Behaviors branch, and operated by the PIs of T1C1S4E) which is equipped with a Velodyne 32 Lidar, a video camera, and several proprioceptive sensors for localization. The pose estimation system in this vehicle works reasonably well, and this robot is already capable of producing grid maps of environment features (e.g. types of floor, obstacles) used to establish a context for demonstrated behavior. Similarly, we will produce sequences of

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

sensor maneuvers (i.e. positions and orientations) that can be used by T1C1S4E to learn traversability costs. We plan to use this platform for our Task Based Assessments.

Collaboration:
We will collaborate with T1C1S4E - Learning traversability cost from unsupervised semantics. We will follow our previous collaboration model at the research level, and at the data collection, testing and processing levels.

Linkage to previous RCTA sub-tasks:
This effort continues the collaborative work with ARL which started as part of subtask *P5-1 Modeling interactions*, establishing it now as a new independent task.

Expected products:
The testing and implementation work will be supported by the *Husky* robot available through collaboration with ARL. Similarly, we plan to take advantage of the ARL's facilities for testing and data collection at Adelphi, MD. The algorithms developed in this task will be consolidated into a software library or API for easy integration into the RCTA demonstration platforms. The result will be an implementation of a planning-for-perception system capable of generating a sequence of poses to be executed by the robot to accomplish a perception task using vision sensors. Additionally, we will conduct an experimental characterization of performance. The library will be developed and released progressively according to two major milestones:

End of CY2017	• Initial release, implementing preliminary set of spatial features for perception maneuvers, with particular focus on vision sensors and static targets. • Experimental determination of baseline performance.
End of CY2018	• Extended version, including spatiotemporal features suitable for dynamic targets. Enhanced implementation suitable for qualification by IRA process and full integration on platform. • Experimental characterization of performance.

References:

1. Kuderer, M., Gulati, S. and Burgard, W. Learning Driving Styles for Autonomous Vehicles from Demonstration. In *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, USA, 2015.
2. Silver, D., Bagnell, J. A., and Stentz, A. Learning autonomous driving styles and maneuvers from expert demonstration. *International Symposium on Experimental Robotics*, June, 2012.
3. Ziebart, B., Maas, A., Bagnell, J. A., and Dey. A.. Maximum entropy inverse reinforcement learning. In *Proceeding of AAAI 2008*, July 2008.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.3.10 T1C1S4E – Learning traversability cost from unsupervised semantics

PI: Maggie Wigness and John Rogers (ARL), maggie.b.wigness.civ@mail.mil, john.g.rogers59.civ@mail.mil, (301) 394-3927.
Collaborators: **Luis Navarro-Serment (CMU), lenscmu@ri.cmu.edu, (412) 268-6034.**

Description:
Operation of autonomous robotic assets in Army relevant environments requires reliable and efficient outdoor navigation. Visual perception provides environmental context, which enables a robot to move and select actions with greater awareness, e.g., traversability of terrain. However, collection of labeled training data for supervised learners is time consuming. In this task we will *develop a terrain cost estimator that learns from inverse optimal control and unsupervised environment semantics,* which provides enhanced and adaptive perception with low human cognitive workload.

Scientific approach:
Previously, we demonstrated the ability to quickly collect large sets of labeled training data with minimal human effort to train a high performing Hierarchical Inference Machine (HIM) [1] that aided in successful autonomous road navigation. In this task we introduce algorithms that eliminate human labeling intervention and learn traversal costs associated with a set of discovered abstract (i.e., no semantic label) visual concepts in the environment. These learners will provide fine-grained traversal cost estimation for path planning.
Specifically, our previous semi-supervised visual concept learning [2] will transition to a completely unsupervised approach, where semantic labels are not assigned. Instead, each learned pattern is treated as a unique abstract class. An ensemble-like approach, where multiple models are learned for overlapping windows of the training data stream, is used to help combat noise introduced by unsupervised learners. We will use inverse optimal control to gather expected traversals paths and map these trajectories to the unsupervised environment concepts learned by a HIM. Coupling the trajectories and the classifier output, the system can learn that highly traveled terrain types can be assigned low traversal costs whereas avoided terrains should be assigned higher costs. Learning these costs will be done by extending the reward function used in work that learns by example for pedestrian path planning and obstacle avoidance [3].

Linkage to previous sub-tasks: This effort continues the collaborative work with CMU which linked subtasks *I1-5 Visual Navigation in Semi-Structured Environments* and *P5-1 Modeling Interactions*. We will follow the same collaboration model at the research level, and at the data collection and processing level.

Collaboration: with T1C1S4D

Expected products: The testing and implementation work will be supported by the Husky robot. The algorithms developed in this task will be consolidated into a software library or API for easy integration into the RCTA demonstration platforms. Additionally, we will conduct an experimental characterization of performance.

References:
1. Munoz, D., Inference Machines: Parsing Scenes via Iterated Predictions. PhD Dissertation, The Robotics Institute, Carnegie Mellon University. June 2013.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2. Wigness, M., Rogers III, J.G., Navarro-Serment, L.E., Suppe, A., and Draper, B. A., Reducing Adaptation Latency for Multi-Concept Visual Perception in Outdoor Environments. In *Proc. of the International Conference on Intelligent Robots and Systems. Daejeon*, Korea, 2016.
3. Ziebart, B. D., Ratliff, N., Gallagher, G., Mertz, C., Peterson, K., Bagnell, J. A., Hebert, M., Dey, A. K., and Srinivasa, S., Planning-based Prediction for Pedestrians. In *Proc. of the International Conference on Intelligent Robots and Systems.* St. Louis, USA, 2009.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.4 Thrust 1/Capability 2: Op-tempo mobility in rough terrain

2.4.1 Overview

This part of the project addresses the ability to achieve mobility for a dynamic, autonomous legged robot at a speed that enables it to accompany human running through unstructured terrain. A typical scenario would involve a robot moving autonomously through rubble, forest, etc. keeping up or preceding soldiers as they move through rough terrain. The legged robot is able to plan its trajectory across variation in terrain (slopes, obstacles, surface properties) and enlist help from soldiers when needed as it plans and executes these maneuvers. In a dismounted team setting, such as in an urban disaster like environment, the robotic team mate should demonstrate human compatible mobility over highly complex (structured or unstructured) terrain that is not traversable by tracked or wheeled robots. This will require human or animal like platforms with dynamic mobility. Unlike the current state-of-the-art, this capability will be non-teleoperated, therefore it requires using the results of perception elements and tight integration with intelligence technical elements.

This capability requires the ability to adapt its plan in order to deal with small obstacles or changing conditions. It requires a new approach to mechanical design and control to deal with difficult terrain geometry at high speeds, and appropriate perception elements to generate informative representations of the environment at the desired rate. To enable this level of operation in unstructured environments, the current legged platforms developed and employed in the RCTA will be utilized to develop the fundamental control and adaptation strategies and then, will need to be upgraded to a next-generation platform large and powerful enough to supporting the sensing and computational power for the Intelligence architecture. Importantly, we recognize that, in extreme terrain, human intervention might be required and we built into the research program an explicit component that enables intelligent human interaction. Specifically, at appropriate times, the robot can in a minimally-invasive way, interface with a human teammate for assistance in updating its navigation plan/strategy, using virtual or heads-up display to relay possible trajectories to human teammate. It can then use human input on properties of desired path to refine its motion plan.

While the state of the art does begin to address autonomous mobility in rough terrain, it is currently only applicable to slow (quasi-static) control approaches, or through teleoperation of fast, dynamic robots. The proposed research will contribute a general approach that combines robot design and control, planning, and perception to enable rough terrain mobility. In addition, we contribute a general approach to human interaction which directly addresses the performance issues of state of the art approaches in difficult terrain.

2.4.2 Relevance to the RCTA and the Army

Motion through rough terrain, such through a forest littered with brush or an urban disaster area strewn with rubble, is something that commercial robots do not have to deal with, especially at op-tempo. While some progress has been made on developing robotic strategies for legged systems to move on difficult, unstructured terrains, these advances have been restricted to teleoperated scenarios or low speed operation. For Army purposes it is, however, critical to move at speed, keeping up, or even better, preceding soldiers as they patrol or approach an objective. This capability addresses directly this gap.

2.4.3 Required research areas

Research is required at all levels, starting with the low-level control elements and working up to higher level planning and interaction with humans. As the terrain becomes rougher the necessity for understanding the trade offs between speed, efficiency and stability become more important. Fundamental improvements in gait and control policies are necessary to enable fast motion, agile jumping, and maintaining posture as obstacles are traversed. As the perception system identifies changes to the environment the path and trajectory planning for the robot needs to update to identify the safest, fastest, or more efficient way to reach the objective. The proper planning of these routes requires both an understanding of mission objectives as well as a deep understanding of how the robot dynamics and terrain interacts to affects the resulting motion. Ultimately, a human teammate may need to intervene in order to assist the robot to deal with novel or unexpected complexities in the terrain. Eliciting this input without burdening the human teammate requiring novel approaches to human-assisted planning. While each of these technical advances can be developed and tested using smaller existing legged platforms, the integration of these mobility advancements onto a computationally autonomous platform with adequate sensing capabilities requires the design and development of new, larger robot configuration. This will be enabled by combining advances in dynamic scaling principles and efficient motor control strategies.

2.4.3.1 High-speed motion in changing terrain

This part of the work addresses the fundamental issues in locomotion control for high-speed legged motion in rough terrain. In particular, the challenge of efficient application of actuator power into forward speed has historically run counter to the strategies for careful and stabilizing motions in rough terrain. In this work we will investigate how to preserve fast, efficient motion by utilizing the natural whole-body dynamics of running (captured in reduced order models of running such as the Spring-Loaded Inverted Pendulum (SLIP) model) with principles for stabilization, particularly as the terrain changes. Both mechanically embedded feed-forward techniques such as compliance modulation as well as more reactive strategies such as Adaptive-Energy Removal (AER) and foot placement strategies will be incorporated. In particular, we will focus on how these strategies need to be adapted as the terrain changes.

In terms of low-level control aspects, this research area will perform basic research on the development of a control methodology and design analysis for legs with different morphologies, and how to modulate leg coordination for a quadrupedal runner. This will include investigating optimal properties for producing stable running gaits, analysis of the effect of leg morphology on speed, efficiency, and stability of running performance, optimization and analysis of the effect of varying control parameters on behavior as the robot runs through varying environments. These findings will support the design of future platforms as well as gait development for platforms using different leg morphologies.

Work in mobility in rough terrain is supported by input from perception to provide the basic environment representation for mobility in rough terrain has progressed through stages that first modeled terrain as an elevation map, then added terrain type classification using a variety of sensors. For vegetated terrain, analysis of point clouds was used to infer material density with more general 3-D data structures. To enable higher speed mobility, efforts have been made to generalize trafficability inferences from short-range, geometry-based analyses to long-range, appearance-based analyses. Rudimentary analysis of proprioceptive sensor data has been used to make inferences about terrain roughness, associate terrain roughness classes with terrain appearance classes, and use appearance to predict roughness ahead of the vehicle. Achieving optempo mobility in rough terrain, ultimately with limbed vehicles, requires a foundation of robust local state estimation and short-range 3-D mapping, extending proprioceptive trafficability perception to limbed vehicles, and doing near-to-far generalization of trafficability inferences based on appearance. Proposed perception research will follow that progression, starting with a small

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

number of discrete terrain classes. This will aim for a capstone demonstration in 2019 where a limbed vehicle learns to select terrain with better trafficability.

2.4.3.2 Environment and terrain-adaptive planning

This second area addresses the high-level planning functions for rough terrain mobility. This work builds upon our knowledge of the terrain and environment provided by our perception system, and the policies for efficient and stable running described in Section 2.3.1. Our basic approach is to learn the dynamic and power models of unique mobility robots and use these to develop terrain-dependent energy efficient motion planning algorithms. Our current work showed the feasibility of learning such models for different platforms, e.g., legged and skid-steered. We have also tackled directly the issue of limited computational resources by developing an anytime version of this approach. Our goal is to extend this work to generalize to generalize to a wide variety of terrain and platform configuration. We will achieve this through new techniques for train neural networks to learn motion and power models and for incorporating them into a sampling-based model predictive optimization framework. The latter will enable us to generated efficient trajectories under different conditions (payload, slope, etc.). Our key contribution to the state of the art is a unique approach to learning how to plan for rough terrain, in contrast to fixed-model approaches. In addition, our approach takes explicitly into account power models, which is an important consideration for mobility in rough terrain with a limited power budget.

This level of motion planning must be connected to directives derived from higher level mission descriptions communicated by teammates. We investigate explicitly the challenge of decomposing these descriptions into behaviors suitable for motion planning. This part of the work integrates the cognitive models of shared mental models intended to formalize human task procedures in a computational form suitable for robots with the mission planning work used to decompose mission-level tasks received from the soldier into tactical behaviors. The former assumes a procedure communicated all at once, then gradually augmented with learned experiences to improve individual decisions. The latter assumes a self-contained generative process from a one-time user input, also improved over time by learning from the outcomes of generated plans. This work would integrate the two processes while transforming mission planning from a batch process to a highly interactive one in which the robot learns to decompose missions by combining its own reasoning and previous experience with explanatory input received as needed from human teammates, who would gradually teach and train the robot.

2.4.3.3 Human assisted path planning

This research area addresses the need for a systematic mechanism for the robot to ask for help when confronted with a situation that the planner cannot deal with. It is critical that such a mechanism exists; otherwise we would need to assume that perception, mechanism, and control/planning can handle any terrain configuration that can ever be encountered. The central research questions involved in taking advantage of the teaming aspects of the RCTA scenarios involving the human teammate in the planning are: What information should be sent to the teammate? In what form should it be presented? What form should the teammate's input take and how should the decisions be shared between robot and teammate? These questions are further complicated by the fact that multiple alternatives may be equivalently considered and that the load on the teammate must be minimized.

This research area addresses these questions. For the first question, we propose a human-robot interface to communicate to the human information about the robot's planned behavior in an understandable form. The interface, which will use a heads-up or mixed-reality display, enables two-way communication between the robot and its operator. A subset of the possible robot trajectories developed in 2.3.2 will be presented to the soldier via this interface. To address the second question, we tackle explicitly the issue of multiple

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

planning alternatives and of minimizing human querying through an approach based on homotopy-aware, multi-objective path planning with a human in the loop. This focus has three parts. First, we will work with ARL to support a planned study that compares paths generated by a human to paths generated by an algorithm; we will provide our homotopy-classifier as a metric for his study, compare the results for this metric with other metrics, and use these results to inform the next studies.

Second, we will perform human studies that examine when and how often the human should be queried during execution of a navigation plan when replanning might be merited. In particular, the conventional paradigm, waiting until near system failure to query the human and replan, is known to induce high operator workload and can contribute to operator-out-of-the-loop problems. We propose that a better approach is to use a numerical representation of intent, which is necessary for selecting among tradeoffs identified by the planning algorithm, and then only replan if the current path differs significantly from human intent. Third, we will integrate these with the planning algorithms from Sec. 2.3.2, using the interface described above. An example demonstration for this could be a legged robot that can either go a long way around an obstacle to achieve a desired goal (safe but slow) or can reach the goal sooner but would have to jump over a log, or traverse muddy terrain (risky but quick).

Finally, we will develop novel ways to formalize robot's behaviors and learning abilities by (1) examining human mental models pertaining to critical thinking and logical argumentation for enhanced inductive, deductive, and abductive reasoning via both literature and field/simulation based experimentation, (2) using unsupervised machine learning principles of self-sustaining internally based reinforcement learning to build motivation models for robots, (3) focusing on modeling new state-of-the-art curiosity-driven/uncertainty-reduction learning paradigms for varying cognitive tasks, (4) incorporating Bayesian mathematics into psychological concepts, and (5) researching higher level cognitive processes in humans to build a sensory classification computational model for robots.

2.4.3.4 Design and control of a dynamic quadruped

In this area, we will perform basic research and support ongoing RCTA research for understanding aspects of quadrupedal dynamic motion and morphology at multiple scales. We will undertake this through a combination of analytical and experimental studies, including the development and verification of modeling techniques for physics-based dynamics simulations. This research will be central to the design development of the next-generation legged platform through modeling of legged robotic design concepts and larger versions of these platforms created through application of scaling methodology currently being studied in RCTA DMUM on smaller prototypes. Key to enabling a fast, agile motion on a larger platform will be the efficient design and exploitation of electromechanical actuators. Unlike most electromagnetic machines, limbed systems involve highly dynamic behaviors, and the mass and inertia of the actuators contribute to the load of the system. These additional complexities added by the intricate dynamics obscure the design processes. For example, the effect of inertia distribution along the robotic limb from the rotor of the motor to the foot during legged locomotion is not well studied. The design of the combination of motors and transmission become quite intricate in dynamic tasks. Designing electromagnetic actuators and transmission for dynamic and energy efficient robot requires thorough understanding fundamental energy loss mechanisms in every step of power flow from batteries to ground interaction. A combination of power analysis, dynamic scaling and modeling will inform the development of a dynamic runner with sufficient speed for keeping up with soldiers, agility to move through difficult terrain, and payload for autonomous operation.

2.4.4 Demonstration plan

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

The technical elements that comprise this capability will be continuously integrated as part of the hardware and software development process during the course of the program. Over the course of the next three years the component capabilities described above will be incrementally integrated. Building upon previous task-based assessments (TBAs), a series of Multi-Component Assessments (MCAs) will lead to a final Integrated Research Assesment (IRA), to demonstrate an integrated system for op-tempo mobility through difficult environments at the end of the program. We envision three milestones for integrating the technical elements for Capability 2:

- *Planning over various terrains (CY18a*): This milestone will demonstrate the ability to generate fast and stable gaits on varying terrains and efficient motion plans based on representations of changing terrain from perception.
- *Autonomous operation in complex setting (CY18):* The first milestone does not incorporate high-level directives and plans directly from waypoints. The second milestone incorporates these into behaviors dynamically and to execute them on the robot, and utilize improved low-level behaviors and perception for long range terrain modeling to enabling faster motion in complex terrain.
- *Human-assisted planning (CY19):* The third milestone further enhances the robustness of the system by injecting human interaction, optimizing the sharing of information and decisions between human teammates and robot.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.5 Thrust 1/Capability 2 Subtasks

2.5.1 T1C2S1A – Motion control policies for planning and high-speed running in varying terrain

PI: Clark, Jonathan (FSU), jeclark@fsu.edu, 850-645-0132

Objective and Benefits:

The objective of this task is to construct algorithms for control of a dynamic, autonomous quadruped LLAMA. This work will entail three primary goals; design and development of control algorithms producing both walking and fast, stable and efficient periodic dynamic gaits, policies for adapting gaits to changes in terrain conditions, and construction of appropriate motion primitives and cost-functions for high-level motion planning. The products of this research will overcome the historic challenge of bridging the gap between efficient application of actuator power into forward speed and strategies of carefully stabilizing motions over rough terrain, enabling the LLAMA platform to move effectively and at soldier's optempo. Furthermore, these advances will allow the high-level intelligence architecture to take full advantage of the dynamic behaviors of LLAMA as it moves through difficult, changing terrain.

Scientific Approach:

Previous research utilizing the natural whole-body dynamics of running (captured in reduced order models of running such as the Spring Loaded Inverted Pendulum (SLIP) model) has shown that fast, efficient legged locomotion is obtainable. Based on those results, we will pursue an investigation of vector-field based leg coordination strategies coupled with Adaptive Energy Removal stabilization techniques [Miller, 2015] to improve the speed, efficiency and stability of running performance of LLAMA. We will look at combining careful, kinematic-based footfall planning algorithms with these momentum-based high-speed running control strategies. Techniques akin to hybrid-zero dynamics will help map the high DOF workspace of LLAMA's legs with the lower-dimensional manifolds associated with stable dynamic gaits. The second aspect of this work will look at how to adapt the high-speed locomotion gaits to unknown rough terrain. Our focus will lie on how template-based gait strategies and parameters need to be modified as the physics of the ground-foot interaction change as the robot transitions between different terrains. A control methodology for these techniques is to be developed, building upon our previous RCTA work identifying terrains through foot contact [Shill, 2015] and adapting to different terrains by modulating the effective limb impedance [Miller, 2013].

Third, we will we build on our initial efforts to interface running behaviors to the RCTA intelligence architecture. We will develop the motion planning primitives and the groundwork to establish heuristic cost functions that the planning algorithms need to plan high-level trajectories for the robot as it moves at different speeds, negotiates obstacles, goes up and down hills, and over different terrain surfaces.

List of dependencies:

This work will depends on, and feeds, the results of Simulation of Legged Behaviors, Jason Pusey - ARL (T1C2S5C) Integration and Design of LLAMA, Brett Kennedy – JPL (T1C2S5A)

Collaborations:

Primary collaboration will be with Brett Kennedy at JPL who is working on the design and fabrication of LLAMA. The modeling and simulation work will build on our ongoing personal exchanges and collaboration with ARL, particularly Jason Pusey. Graduate student exchanges will continue with ARL, and be expanded to include personal traveling to MIT and JPL to assist with development and testing on

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

LLAMA. Collaboration with Emmanuel Collins at FSU will connect gait primitives and cost functions to higher-level path planning algorithms.

Linkage to previous RCTA BPP sub-tasks:
This work builds on M4.1 (Dynamic Mobile Manipulation Mechanisms) and M8.1 (Dynamic Multi-Modal Locomotion) from the 2015-2016 BPP.

Expected Products:
Year 1 (2017) -

Quadrupedal gaits will be developed for stable working, fast and efficient motion, and adaptation on multiple terrains (concrete, grass, etc.). Performance will be evaluated by quantitative measures of efficiency, stability, and turn radius as a function of speed and the terrain that the robot is running on. These results will span from simplified models and control strategies to be optimized behaviors on legged platforms. Expected performance includes running at 2.0+ m/s over a variety of surfaces. These gaits and the resulting motion planning primitives will also be fed to Thrust 1, Capability 2, Research Area 2. These primitives and cost functions will provide tools necessary for the intelligence architecture to plan efficient and safe paths that make use of the unique locomotive capabilities of the quadruped. Scientific advances will be published in collaboration with ARL researchers in technical conferences such as ICRA and IROS.

Year 2 (2018) -

Obstacle traversal strategies and transitions between these and standard gaits. Examples include overcoming obstructions or transitioning from a walking gait on one surface to a running gait on a second. This will require integration of motion strategies for careful foot placement and for dynamic running gaits. Motion primitives and cost functions for varying conditions with these gaits will be fed to Thrust 1, Capability 2, Research Area 2. Technical findings will be published in scientific conferences (e.g. ICRA or IROS) and journals.

Reference Research Papers:
Miller, B. and Clark, J.E., "Towards highly tuned mobility in multiple domains with a dynamical legged platform," Bioinspiration & Biomimetics, vol. 10, no. 4, pp. 046001, May 2015.
Shill, J., Collins, E., Coyle, E. and Clark, J.E., "Tactile Surface classification for limbed robots using a pressure sensitive robot skin," Bioinspiration & Biomimetics, vol. 10, no. 1, pp. 016012, Feb 2015.
Miller, B.D., Cartes, D. and Clark J.E., "Leg stiffness adaptation for running on unknown terrains," Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Tokyo, Japan, Nov 3-7, 2013.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.5.2 T1C2S2A – Perception for maneuver in complex terrain

PI: Matthies, Larry (JPL/Caltech), lhm@jpl.nasa.gov, (818)354-3722
Additional PIs: Karumanchi, Sisir (JPL/Caltech) sisir.b.karumanchi@jpl.nasa.gov, (818)356-3867

Objective and Benefits:

Among technology advancement priorities in the Army's Robotic and Autonomous Systems Strategy [1], "Foremost is autonomy for ground vehicles because, as a land force, the Army relies on ground combat vehicles and off-road mobility. It is the most challenging, military-specific requirement ... Autonomy is a gateway technology that, once obtained, will be integrated into all ground vehicles, combat or otherwise. Its impact will influence all warfighting functions." (p. 12). Autonomous mobility will have to cope with a wide variety of terrain, including desert, mountainous, grassland, forest, jungle, and a spectrum of urban terrain from intact to heavily damaged. The perceptual front-end for such systems must be able to represent trafficability of these terrains. Since preprogramming knowledge of such diverse conditions is infeasible, the system must be able to learn trafficability models from training data, including on-the-fly from its own experience. Elements of the environment will be dynamic, from effects of weather, to vegetation blowing in the wind, to terrain that is dislodged by the vehicle, to people, other vehicles, and animals; therefore, the scene representation must be able to deal with such dynamics.

A key element of this problem is how to model, learn, and plan for the resistance to motion presented by the environment, such as slopes, vegetation, and heterogeneous soil with varying moisture content. In this subtask we focus on vegetation, because of its ubiquity. Past research has addressed navigation in vegetation by seeking to estimate where the load-bearing surface is underneath vegetation, by classifying terrain into a small number of discrete types with known trafficability properties, or both [2]. However, to cope with the variety of vegetation types and densities in real terrain, continuous models of vegetation trafficability must be developed, with methods to estimate (learn) parameters of such models on the fly.

Figure 1.

We will address this gap, building on an efficient, voxel-based world model and perception system from past work in the RCTA and LS3 programs that supports reasoning about vegetation and detection and tracking of nearby personnel [3, 4]. In collaboration with other subtasks, this will enable demonstration of off-road movement with dismounts, potentially including taking directions from dismounts and learning about trafficability by observing where the dismounts are able to go. In 2017, this work will be done with stereo cameras, lidar, inertial, track encoder, and motor current sensors on the RoMan platform. A novel, unique element of the approach will be use of manipulator(s) with force-torque sensors on RoMan as part of the sensing strategy, to probe vegetation or other terrain in front of the vehicle. This will enable the vehicle to learn on its own where it can push through brush and where it cannot. In 2018, most of these sensors and software will be integrated into a legged vehicle prototype to enable autonomous mobility experiments with a legged platform.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Scientific approach:

The key barrier to be overcome is the lack of continuous models of vegetation trafficability or methods to learn such models in past research. We will use the efficient voxel tile data structure from [3] and the dual return mode of a Velodyne lidar, together with color stereo cameras, to estimate density and other features of voxels, as well as to find the ground surface at short range. We will design a motion resistance model for vegetation based on simple physical models. We will estimate parameters of this model in two ways. First, we will use track encoder, motor current, and inertial sensors in the RoMan vehicle, as an attempt to generalize past work on estimating traction coefficients [5] and slip on slopes [6]. Second, we will use one or both manipulators on RoMan to grasp a rod (or similar tool) and hold that in front of the vehicle as it attempts to push into vegetation; we expect that the manipulator force-torque sensors may give a more sensitive measurement than the sensors presently in the vehicle base. This is will give one of the first examples of using a manipulator to aid mobility in complex terrain.

We will test this work in simple scenarios on RoMan (and/or the JPL Surrogate, which is very similar) in 2017. In 2018 we will address more complex scenarios; assuming completion of a new prototype legged mobility platform, we will also port the perception system to the legged platform to support autonomous mobility experiments with it. We expect that this will enable capstone experiments/demonstrations in 2019 with both tracked and legged platforms, negotiating vegetated terrain with very similar sensor suites and perception software.

List of dependencies: none per se, but important collaborations, as noted below.

Collaboration:

- T1C2S2D, GDLS, M. Henke: voxel-based world model development, non-GPS localization, evaluation of proprioception
- T1C2S5A, JPL, B. Kennedy: adaptation to a legged platform
- T1C2S2G, UPenn, C. J. Taylor: reliable non-GPS localization in complex terrain
- T3C1S3C JPL, S. Karumanchi: use of manipulators for probing trafficability
- T1C2S3A FSU, E. Collins: motion planning based on models of power consumption in various terrains

Linkage to previous RCTA BPP sub-tasks:
This subtask builds on perception capabilities begun in tasks P4-1 and P6-1 early in the program, which addressed self-supervised learning of trafficability of different soil types (P4) and distinguishing static and dynamic scene elements in voxel maps (P6). Later, this work was substantially extended under the DARPA LS3 program. We leverage this combined heritage.

Expected products:

- Algorithms and software that enables learning trafficability of vegetation in off-road navigation.
- Integration and testing on RoMan, with transition to a legged platform in follow-on work for the Capstone.
- Conference publication of results in 2018.

References:
[1] *The US Army Robotic and Autonomous Systems Strategy*, draft of 30 September 2016.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

[2] D. M. Bradley, et al. Scene understanding for a high-mobility walking robot. In IEEE/RSJ International Conference on Intelligent Robots and Systems, 2015.
[3] M. Bajracharya, J. Ma, A. Howard, L. Matthies. Real-time stereo mapping in complex dynamic environments. In IEEE International Conference on Robotics and Automation, 2012.
[4] M. Bajracharya, J. Ma, M. Malchano, A. Perkins, A. A. Rizzi, and L. Matthies. High fidelity day/night stereo mapping with vegetation and negative obstacle detection for vision-in-the-loop walking. In IEEE/RSJ International Conference on Intelligent Robots and Systems, 2013.
[5] S. B. Karumanchi. *Off-road mobility analysis from proprioceptive feedback*. PhD thesis, University of Sydney, 2010.
[6] A. Angelova, L. Matthies, D. Helmick, and P. Perona. Slip prediction using visual information. In Robotics: Science and Systems Conference, 2006.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.5.3 T1C2S2D – Perception for optempo mobility in rough terrain

PI: Michael Henke (GDLS), henkem@gdls.com, (240) 446-9745.

Level: 6.2 Applied Research

Research description

Objective and benefits

The objective of this task is to develop, implement, test and evaluate an *integrated sensing/processing perception unit* to enable optempo mobility on the RCTA platforms over rough terrain (e.g. urban rubble, wooded terrain, ditches and fallen trees). These abilities are relevant since they will support robot operations in the types of unstructured environments considered by the RCTA program by providing an independent module that can be mounted on the robot and provide it with high-level perception capabilities without imposing an additional burden to its computing budget. Additionally, this effort satisfies the current need for a concerted effort to produce advanced sensors specifically tailored to operations in complex terrains—there is no commercial market or government program currently working on sensors especially suited to this application.

Sub-task description

In this task we will develop a state-of-the-art sensor suite for mobility in rough terrain. The suite will integrate 3-D and vision sensors and their corresponding processing pipelines into a single device capable of performing visual terrain classification, proprioceptive terrain sensing, and near-to-far learning to model terrain trafficability. We plan to maximize the use of COTS-based hardware and software to expedite the development process and to avoid unnecessary debugging and testing. In particular, we propose the use of a RGBD-NIR 3-D sensor capable of indoor/outdoor operation at 60Hz frame rate, in combination with JPL's stereo implementation. The corresponding processing pipelines will be implemented on a state-of-the-art GPU-based computing engine with a computational capability of over 20 TFLOPS.

During the first year of this effort, we plan to integrate the sensors described before with the GPU computers. In terms of processing, we will include the 3-D Voxel world model; the JPL stereo implementation; and a high mobility planner. We will test the synergistic combination of these elements using RCTA *Husky* and *RoMan*. Finally we will demonstrate and evaluate the performance of optempo mobility on cross country and trails.

In the second year, we will focus on the integration onto the RCTA legged platform *Robosimian*. Testing will be conducted under increased terrain complexity, using real-time fusion of visual and proprioceptive terrain perception to improve trafficability assessment, and improved reasoning about limb placement for limbed platforms, to culminate in capstone. For the tracked and wheeled platforms that have arms we will experiment with balance and push point algorithms to increase the platform stability while traversing challenging terrain.
This effort will continue through to the 2019 capstone assessment, with annual additions in complexity to meet relevant attributes for upcoming IRAs, while headed towards a final goal optempo mobility integrated onto wheeled, tracked and legged RCTA standard platforms.

Metrics & Progress Assessment

For evaluation purposes, we will measure peak and average computing speeds versus terrain complexity. To assess terrain complexity, we will compute a metric based on terrain plane inclination, mean absolute distances of terrain points to plane, and the standard deviation of the distances from point to plane.

Linkage to previous sub-tasks

This work builds on prior works done at ARL, JPL, UPENN and CMU. Leverage efforts started early on RCTA and that have been matured on other programs. In particular, this task builds on work done in *P6-2 Multi-Sensor Perception and Understanding of Complex and Dynamic Environments*. By pulling this capability back into RCTA it will facilitate higher level RCTA research to progress by utilizing this optempo mobility capability.

Collaborations

- Larry Matthies, JPL – T1C2S2A - *Perception for Maneuver in Complex Terrains*.
- Jaymit Patel, GD – T4C3S1A - *World Model Maintenance and Enhancements*.
- CJ Taylor – T1C2S2G – *Perception for Optempo mobility in rough terrain*

Expected products

2017	• Integrate upgraded sensors and GPU-based computers on tracked and wheeled vehicles. • Integrate existing 3D voxel world model and JPL stereo algorithms. • Evaluate performance of existing 3D voxel world model and JPL stereo algorithms. • Evaluation of performance against proposed metrics and SOA.
2018	• Assess lessons learned from 2017, and plan out. • Develop a model of trafficability that is estimated from proprioception. • Develop methods to predict trafficability model from lookahead terrain perception features. • Integrate upgraded sensors and GPU based computers on RCTA legged platform. • Integrate legged platform foot placement planner. • Evaluation of performance against proposed metrics and SOA.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.5.4 T1C2S2G - Perception for OpTempo Mobility

Camillo J. Taylor, University of Pennsylvania, cjtaylor@cis.upenn.edu, 215-898-0376

Objective and Benefits:

The kinds of autonomous systems envisioned in this RCTA will need to be able to autonomously navigate in a variety of environments including densely forested scenes with vegetation that is deformable and/or dynamic. This can be contrasted with the situations found in indoor environments where one can typically assume that most, if not all of the scene is static. Many of the existing approaches to visual inertial localization proceed from the latter assumption. In contrast this task will address the kinds of complex vegetated environments envisioned in RCTA task T1C2S2A.

Scientific approach:

In part this effort will draw inspiration from recent work in neuroscience that seeks to determine how animate systems maintain a sense of place as they move through the world. Experiments in this area suggest that a range of species, including rats, encode their orientation independent of position in a set of head orientation cells. These cells encode information gleaned from both the visual cortex and the animals inertial measurement system and allow the animal to maintain a consistent sense of heading even as landmarks and lighting are manipulated.

One of our first goals will be to try to develop software systems that attempt to replicate the robustness of animate systems by integrating information from the stereo vision system and Inertial Measurement Unit (IMU) that are being provided by RCTA task T1C2S2D. The approach will be to use the stereo system to identify visual features that appear to be distant landmarks where available and to integrate this sporadic orientation evidence with the measurements from the onboard accelerometer and yaw rate sensors to provide a continuous estimate of the orientation of the robot.

A key problem that we will need to address in this effort is reliably discriminating landmarks in the image that may be appropriate for navigation tasks from a sea of distracting features provided by leaves, brush and other dynamic visual elements. Initial approaches to this problem will attempt to leverage the appearance of the landmarks as well as the geometric properties derived from multiple views.

Another component of this effort will be an investigation into how best to use the measurements that can be gleaned from the onboard magnetometer. These magnetic measurements can, in principle, be used to gauge the orientation of the earth's magnetic field however in practice these measurements are often confounded by other magnetic fields such as those produced by motors or other electrical systems onboard the robot as well as other sources of disturbance such as ferrous rocks or underground cables. One aspect of this task will be to collect measurements from the magnetometer in an attempt to characterize the utility of the magnetic cues.

The next phase of this project will concentrate on the problem of estimating the translational motion of the platform over the terrain. For this component of the problem we intend to exploit information returned by proprioceptive sensors such as wheel encoders and step counters as well as visual measurements obtained from the stereo system. The proprioceptive measurements would be used to inform and constrain the analysis of the motion of the visual features that are being observed. Here again a crucial task will be to distinguish between visual features and landmarks that provide reliable information about the motion of the platform and the features that are moving independently. The goal will be to strongly leverage the

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

information from the orientation system which should powerfully inform the algorithm that analyzes visual feature motion to reject outliers and to develop an accurate estimate for the motion of the platform over the terrain

Platforms: RCTA Husky and RoMan

Connection to past RCTA work: This work will build upon our work on previous RCTA efforts (P4-3 and P6-3) which enabled wheeled and legged vehicles to localize themselves in indoor environments and to autonomously climb stairs. This work will also leverage ongoing work on visual odometry for lightweight autonomous air vehicles that is being carried out in parallel.

Collaboration: T1C2S2A - Larry Matthies, JPL/CalTech, and T1C2S2D - Michael Henke,GDLS

Deliverables:
2017:

1. Work with GDLS and JPL to acquire datasets including visual, inertial and proprioceptive data as the Husky and RoMan platforms traverse relevant terrain.
2. Develop algorithms to select features that could serve as landmarks to gauge vehicle orientation
3. Develop algorithms to fuse sporadic orientation landmark information with the measurements from the onboard IMU to gauge robot orientation. Develop methodologies to evaluate the efficacy of this approach quantitatively
4. Conduct experiments to characterize the reliability of the measurements obtained from the magnetometer under various conditions. Determine how best to incorporate these measurements into an orientation estimate

2018:

1. Collect datasets to gather odometric information from the Husky and RoMan platform as they traverse relevant terrains
2. Develop techniques to estimate the velocity of the platform in it's own body frame based on proprioceptive feedback
3. Develop algorithms to gauge the translational motion of the platform by combining information from the orientation estimation system, the proprioceptive system and the stereo camera.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.5.5 T1C2S3A - Terrain-adaptive planning for autonomous movement of legged and skid-steered robots

PI: Collins, Emmanuel (FSU), ecollins@eng.fsu.edu, 850-410-6373
Additional PIs: Ordonez, Camilo (FSU), cordonez@fsu.edu, 850-645-0104

Objective and Benefits:

Legged robots have the ability to traverse rough terrain not traversable by wheeled or tracked robots; this is of great importance in military applications of mobile robots. However, the requisite motion planning technology is more difficult to develop for legged robots. This is due to the fact that kinematic-type models as well as dynamic and power models are difficult to analytically develop for these robots (a feature shared with skid-steered vehicles) and in part to the fact that legged robots have multiple gaits to be considered in motion planning. The objective of this research is to further develop a planning methodology applicable to legged and skid-steered robots that addresses these challenges and enables the robot to learn terrain-dependent models and motion plan in multiple terrain environments involving various elevations and traversable obstacles.

Scientific approach:

The general approach to motion planning has been developed and demonstrated (in part) using skid-steered vehicles [2-3], which like legged robots have motion that is inherently dependent upon the terrain being traversed. There has also been initial demonstration on the XRL robot []. Motion planning is based upon Sampling Based Model Predictive Optimization (SBMPO), which integrates a kinematic-like or dynamic model to achieve trajectories that are compatible with that model. (This is often called "kinodynamic planning.") To avoid on-line complexity, the methodology has converged to the use of a pseudo-kinematic model as the integration model. The dynamic model is used to determine turn radius constraints, and a power model is used to determine the energy cost associated with candidate trajectories.

The current planning methodology has three key features that will be integrated in future research. First, an anytime version of the algorithm has been developed. Second, is the development of an incremental version. Third, a complementary sampling-based algorithm enables robust trajectory tracking, so that the robot is able to follow the original trajectory even when it significantly deviates from it at particular trajectory points.

While kinematic-like, dynamic, and power models of skid-steered vehicles can be obtained analytically, this is a much more difficult task for legged robots. Hence, in ongoing and future research significant focus is given to learning these models (especially, kinematic-like and power models) using neural networks. This research will focus on determining the types of neural networks that are most effective for this type of learning. Candidates are convolutional networks, recurrent networks and deep learning structures. In parallel with this research is the determination of the types of diagnostic trajectories suitable for learning the desired models.

List of dependencies: The research uses ROMAN, LLAMA and a mini-LLAMA-like platform. The research on LLAMA and "mini-LLAMA" requires that we know the gaits, which will be produced by task T1C2S1A entitled *High-Speed motion in changing terrain*. Pose estimates, a terrain and elevation map, and SLAM (of some fidelity) will be needed for hardware implementation. For the "mini-LLAMA" we can simply use an outdoor motion-capture system and a hard-wired terrain and elevation map. It is unclear who will produce pose estimates, a terrain and elevation map, and SLAM for LLAMA. We have not explicitly considered noisy state estimates (that "jump around a lot"). This phenomenon may be particularly evident in state estimates for LLAMA (or any legged vehicle).

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

The research on ROMAN will also need pose estimates, a terrain and elevation map, and SLAM. It is expected that this will be provided by task T1C2SA entitled Perception for maneuver in complex terrain. Again, we have not explicitly considered noisy state estimates (that "jump around a lot"). This phenomenon may not manifest itself as much for the movement of ROMAN, but other RCTA research may inform us of its importance.

The research interacts substantially with task T1C2SA by providing that task with knowledge of what perception is needed for high vegetation situations. In particular, it will interact with that task to ensure that the perception information can be used to develop a resistance model for a mobile robot (in this case ROMAN).

The research early-on provides task T1C2S4C entitled *Practical human-robot interaction for navigation and manipulation* with software that this task needs to produce multiple trajectories from which a human can choose.

Collaboration: Using the ROMAN robot, a skid-steered vehicle, this subtask will interact with the JPL (Larry Mathies) task T1C2SA entitled *Perception for maneuver in complex terrain*. In particular, while JPL will focus on the fundamental perception problem for thick vegetation with a unique feature being the use of a manipulator for sensing resistance, FSU will focus on converting the resistance sensing to a model of how a given vehicle will experience resistance due to that vegetation. This will enable prediction of whether momentum-based motion planning is needed [1] and the type of power it will take to traverse the vegetation, which can be used to determine the traversal cost in energy efficient motion planning. The experiments for this task will be performed primarily at JPL.

This task will interact with the BYU (Michael Goodrich) task T1C2S4C entitled *Practical human-robot interaction for navigation and manipulation*. In particular, we will work with BYU to help them use the FSU planning algorithm to develop Pareto-optimal trajectories that can be provided to a human operator as alternative trajectories. For example, these trajectories can provide tradeoffs between energy, distance, and risk.

The task interacts with FSU (Jonathan Clark) task T1C2S1A entitled *High-Speed motion in changing terrain*. It informs and receives from this task the information (e.g., gaits) needed for motion planning.

Linkage to previous RCTA BPP sub-tasks:

This research builds upon the subtask, M8.2 – Efficient Motion with Dynamic and Power Models, funded in 2015 and 2016. The basic approach to motion planning was developed under this subtask.

Expected products:

- 2017:
  - Development and implementation of a neural network algorithm for learning the kinematic-like and power models of the XRL robot and implantation of a motion planning algorithm for the XRL robot.
  - Development and demonstration of a methodology for converting the vegetation resistance model developed by JPL to a resistance model for an individual robot (with ROMAN being the platform of interest).
- 2018:
  - Application of the neural network algorithm developed in 2017 to learning the kinematic-like and power models of a small legged robot (one with non-rotating legs and some resemblance to LLAMA) along with implementation of a motion planning algorithm for this robot.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

- In conjunction with JPL, demonstrate motion planning for ROMAN through vegetation for which the resistance has been determined.

References:

[1] C. Ordonez, N. Gupta, O. Chuy, Jr., and E. G. Collins, Jr., "Momentum Based Traversal of Mobility Challenges for Autonomous Ground Vehicles," *IEEE Conference on Robotics and Automation*, May 6-10, 2013, Munich, Germany, pp. 744-751.

[2] N. Gupta, C. Ordonez, E. G. Collins, Jr., "Dynamically Feasible, Energy Efficient Motion Planning for Skid-Steered Vehicles," *Autonomous Robots*, 2016, to appear. (Available at http://link.springer.com/article/10.1007/s10514-016-9550-8.)

[3] C. Ordonez, N. Gupta, B. Reese, N. Seegmiller, E. Collins, and A Kelly, "Online Adaptation of Skid-Steered Wheeled Kinodynamic Models," *Robotics and Autonomous Systems*, submitted.

[4] C. Ordonez, N. Gupta, E. G. Collins, Jr., J. E. Clark, and A. M. Johnson, "Power modeling of the XRL hexapedal robot and its application to energy efficient motion planning," in *Adaptive Mobile Robots*, A. K. M. Azad, N. J. Cowan, M. O. Tokhi, G. S. Virk and R. D. Eastman, Eds., World Scientific, New Jersey, 2012, pp. 689-696.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.5.6 T1C2S4C - Practical human-robot interaction for navigation

PI: Goodrich, Michael (Brigham Young University), mike@cs.byu.edu, 801-422-6468

Objective and Benefits:

This research area addresses the need for a systematic mechanism for the robot to ask for help when confronted with a situation that the planner cannot deal with. Interactive human-robot planning requires algorithms and interfaces that allow humans and robots to find mutually acceptable tradeoffs in complex missions. The research seeks to answer the following questions, which take advantage of the teaming aspects of the RCTA scenarios involving the human teammate in the planning are: What information should be sent to the teammate? In what form should it be presented? What form should the teammate's input take and how should the decisions be shared between robot and teammate?

Scientific approach:

Existing algorithms will be extended and adapted so that they work together and include energy-based and traversability-based cost functions. The utility of the algorithms under human-interaction will be demonstrated using the Palette Interface and evaluated under conditions when replanning is required. The algorithms and interaction ideas from the Palette will be extended so that they apply to AR-based perspective displays.

The first part of the research includes extending and adapting Multi-Objective Rapidly-Exploring Random Forest (MORRF) to approximate the set of Pareto optimal solutions to allow inclusion of critical mission properties of minimum-energy plans and plans through traversable regions. Collaborative work with Emmanuel Collins will also consider using Model-Predictive Sampling-Based Goal-Directed Optimization to identify Pareto optimal solutions. The working hypothesis is that human intent, within the context of interactive path-planning, can be expressed as a tradeoff between competing objectives, so an algorithm is required to identify these competing objectives. More precisely, an acceptable path is one that satisfies a numerical representation of intent, one which balances tradeoffs. The algorithm will need to be able to work in real-time, enabling replanning, but only if the current path differs significantly from human intent. We believe that MORRF-based planning will work with noisy state estimates, but need to study this explicitly.

The second part of the research includes developing a human-robot interface to communicate to the human information about the robot's planned behavior in an understandable form. The interface, which will use a heads-up or mixed-reality display (working with Emmanuel Collins), enables two-way communication between the robot and its operator. The Palette Interface previously developed for the RCTA will be used to evaluate whether a human and a robot can interactively and efficiently find an acceptable tradeoff given a commander's intent. An Augmented Reality interface will also be developed and used to enable interactive planning in dismounted, dynamic situations. The interfaces will enable a user to interactively select which path best satisfies intent. User evaluations of the interfaces will be conducted to determine usability and to apply the work to intent-based replanning. The interfaces and algorithms will be demonstrated on physical robots in different terrains.

List of dependencies: This research depends on input from T1C2S3A to allow us to generate energy-aware paths. The research will then generate paths that can be used in task T1C2S1A to generate motions/gaits. The dependencies create a link between human-robot interaction and legged locomotion.

This research requires a metric-based map of the world. Anticipating that there may be some inconsistencies in the map, we may need to integrate homotopy-aware planning to make planner more robust to these inconsistencies.

Collaboration: T1C2S3A -- We propose to collaborate with Emmanel Collins (T1C2S3A) to identify Pareto-optimal paths when the objective includes low-energy considerations or navigability over rough terrains.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

We also anticipate continuing to collaborate with Brandon Perelman, ARL, in his work in understanding differences between human-planned and algorithm-planned paths. I couldn't find a task number for this, but thought the lnk to ARL would be useful to identify.

Linkage to previous RCTA BPP sub-tasks:
H8, Social Dynamics Modeling and Simulation. More specifically, it builds from H8.1 – Information-Based Soldier-Robot Teamwork. Note that the proposed work does not include one element of the previous work under H8.1, namely human-robot manipulation.

Expected products:

- End of BPP year 2017
    - Integration of Palette-based planner with energy-aware path planners.
    - TBA: Demonstration of capability of finding Pareto optimal solutions using speed, energy, and stealth. Demonstration of usability of Palette Interface for finding tradeoffs among these three objectives for a ground robot on flat terrain. Demonstration on Turtlebot
    - Evaluate use of Palette for intent-based replanning.
    - Submission of paper to ACM/IEEE International Conference on Human-Robot Interaction on Replanning Using a Palette Interface. Study will be for a fixed planning interval (10-20 seconds).
    - Submission of paper to IROS on Combining Energy-Aware Motion Planning with MORRF.
    - Submission of paper on MORFF and Homotopy-Aware Planning to IEEE Transactions on Robotics.
- End of BPP year 2018: Demonstrate capability of finding Pareto optimal solutions using speed, energy, stealth, and traversability. Demonstrate usability of Palette Interface for a four-tradeoff problem. Evaluate display for mobile robot on difficult terrain.
    - TBA: Integration of Palette interface and MORRF planner on Roman platform. Demonstration interactive planning of finding Pareto optimal solutions using speed, energy, and stealth on FSU's Roman platform.
    - Submission of paper to Journal of Human-Robot Interaction on Interactive, Tradeoff-Based Planning and Replanning. Study will be for variable planning interval, an interval that depends on how quickly things in the world change including when errors are corrected in maps.

References:

[1] M. T. Shaikh and M. A. Goodrich. Design and Evaluation of Adverb Palette: A GUI for Selecting Tradeoffs in Multi-objective Optimization Problems. To appear in Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction, Vienna, Austria, March 2017.
[2] D. Yi, M. A. Goodrich, T. M. Howard, and K. D. Seppi. Topology-Aware RRT* for Parallel Optimal Sampling in Topologies. Proceedings of the International Conference on Systems, Man, and Cybernetics, Budapest, Hungary, October 2016.
[3] D. Yi, M. A. Goodrich, and K. D. Seppi, Homotopy-Aware RRT* : Toward Human-Robot Topological Path-Planning. Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction. New Zealand, March 2016.
2015
[4] D. Yi, M. A. Goodrich, and K. D. Seppi, MORRF*: Sampling-Based Multi-Objective Motion Planning. In Proceedings of the 24th International Joint Conference on Artificial Intelligence, Buenos Aires, Argentina, July 2015.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.5.7 T1C2S5A - Next-gen legged platforms

PI: Kennedy, Brett (JPL/Caltech), bkennedy@jpl.nasa.gov, (818) 223-1792

Objective and Benefits:

This effort will develop hardware design, software development and integration support for future RCTA legged platforms. The focus is getting to a point where end-to-end off-road rough terrain testing can be performed in a short timeframe. The goal is to incorporate onboard computing and autonomy from the get go in a cost-effective manner. The prototypes will be termed **LLAMA** (**L**egged **L**ocomotion **A**nd **M**obility **A**daptation). This task with work closely with T1C2S2A (Perception for Maneuver in Complex Terrain, Matthies) to develop **terrain adaptive gaits and precise foothold placements from perceptual input**.

The efforts in CY17 focus on a short design/development/integration path to the **most-feasible** legged platform for CY17. In CY18 the efforts will shift towards development of the most-functional platform based on lessons learned from CY18.

Scientific approach:

Will leverage the following resources to achieve the above

1. JPL did the LS3 terrain sensing and localization for Boston Dynamics – We will leverage this past work and couple this with precision foothold planning from RoboSimian that was developed for the DRC.
2. Aiming for a first-gen prototype for low durability and short development to outdoor testing.
3. Not designing from scratch. Lessons learned from RoboSimian build apply. Mainly choosing OTS motor components and gearboxes for locomotion as opposed to mobile manipulation.
4. Staged strategy:
   - 1 year to first ***most-feasible*** prototype and outdoor demo.
   - 2nd year – two ***most-functional*** prototypes for east and west coasts.

List of dependencies:

- Matthies - T1C2S2A, Karumanchi & Bergh - T4C1S1B, Gonzalez- T4C1S1A, Taylor - T4C1S1C**,** Westrich - T4C1S1D

Collaborations:

- *Ongoing year-long JPL/Caltech collaborations*
  - Matthies - T1C2S2A - Perception for Maneuver in Complex Terrains
  - Karumanchi & Bergh - T4C1S1B – Maintain current platforms – RoboSimian (will leverage low level motor/EtherCAT drivers from this task).

- *Other collaborations*
  - Kim, MIT - T1C2S5B- Power analysis and design of dynamic mobile manipulators.
  - Pusey, ARL - T1C2S5C - Simulation for Dynamic Scaling
  - Collins, FSU - T1C2S3A - Autonomous movement of legged robots over varied terrain

- *Tech transfer with GDLS*
  - Gonzalez, GDLS - T4C1S1A- Maintain current platforms - Husky, RoMan, RoboSimian.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Linkage to previous RCTA BPP sub-tasks:

- This work is related to work done under subtask IR3-7, M2-2, M3-1 for the Roman Platform design and development under past RCTA work.

Expected products:

	2017 Quarterly Milestones and Deliverables
Q1	Acquire components and setup computers.
Q2	Scaled build of gen-1 LLAMA with onboard sensing.
Q3	First cut adaptive gait control & port-over RoboSimian foothold planning from DRC
Q4	Outdoor testing to develop end-to-end autonomy

	2018 Quarterly Milestones and Deliverables
Q1	Assess lessons learned from 2017, and spec out most-functional platform.
Q2	Full scale build of two gen-2 LLAMA prototypes (one for east coast and one for west coast)
Q3	Outdoor testing to develop end-to-end autonomy
Q4	Deliver platform for west coast and design documents to GDLS.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.5.8 T1C2S5B - Data-driven Viability state network (VSN) framework for dynamic locomotion control

PI: Kim, Sangbae (MIT), sangbae@mit.edu, (617)452-2711

Objective and Benefits:

This subtask addresses the main challenges in stable operation of legged systems. Although legged platforms promise better coverage of the types of terrains and obstacle using articulated limbs providing more versatile ground interaction, the stable gait control of such complex and naturally unstable systems remains as a big challenge. Most control methods are limited to a few specific behaviors at specific speed and not capable of smoothly and stably transitioning among diverse behaviors. For example, in order to stop and turn at the corner of the corridor, the robot should be able to slow down and stop followed by turning and acceleration. Even seeming simple behaviors requires frequent gait mode changes and speed variations. Implementation of most conventional control approaches in legged machines has not been successful due to 1) lack of prescriptive or analytic stability criteria widely applicable to transient behaviors in a real-time controller and 2) complexity of high dimensional hybrid the system that entails large computation power even with powerful modern computers so that we cannot execute the dynamic control algorithm in real-time machines. The result of this task will be integrated with the main gait controller of LLAMA platform to expand the capability in real missions.

Scientific approach:

This subtask attempts to establish a new control paradigm to enable truly dynamic and versatile but robust and stable enough to dispatch the robot in a highly unstructured environment. We aim to develop a non-

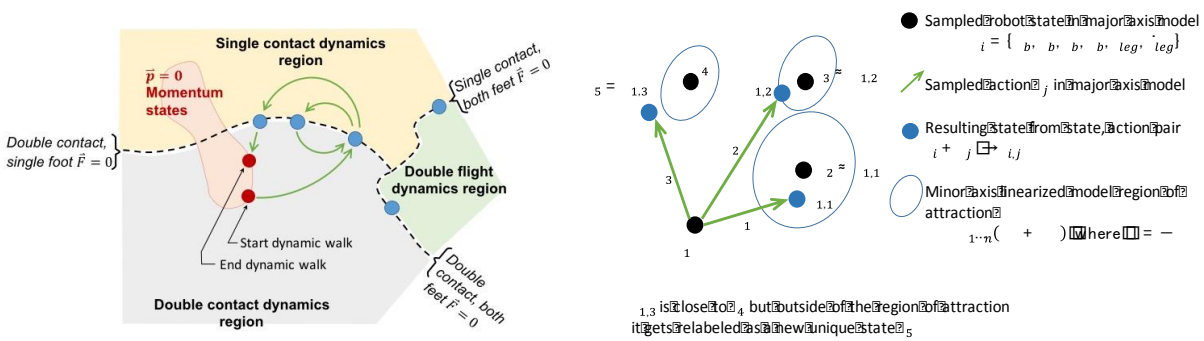

Figure 1. Viability State Network (VSN) growing scheme. The momentum-based sampling state will identify the reachable states and create a growing state network which eventually will be a database for the real-time machine.

conventional control framework that allows the robot to execute combinations of diverse behaviors without designing or optimizing each gait. First, we will explore viability state space and identify failure paths using dynamic simulation. Instead of quantifying stability, we will create stably reachable state space through sampling. Figure 1 shows the exploration method. This exploration allows us to discover reachable state spaces to start creating a VSN. It will start from a known stable state (e.g. standing upright) and start exploring states it can reach. As the exploration region grows, the state-network and the capability of the robot expand. We aim to create various cycle motions such as walking and running gait by closing

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

the loop in VSN instead of manually designing each gait. The hybrid transition of the robot will be one of the main focus of the approach. For example, the hybrid transitions between single contact (two legs on the ground in the case of trot gate) and double stance (all four legs are in the ground in the case of trot gait) provide the opportunities to create transient behaviors.
Once the network grows and identifies failure mode (e.g. fall on the ground), the viability network will be able to define possible region in state space where robot can robustly navigate with significant disturbance margins. This network will be able to provide invaluable stability information accumulated through off-line experiment in simulation environment. Ultimately, when the exploration methods is developed and matures, similar approach can be implemented to create more realistic network in real hardware with safety harnesses. We aim to explore exploring methods that that enable viable network growing starting from a small subset of the entire state space.

There are two major challenges in this approach. First of all, the state space of such a complex robot is too big even in a small subset. In order to mitigate the vast size of the state space, the controller framework will be designed based on the simplified model. For example, we will ignore the inertia of the limbs as we designed the limbs of physical machine to have a minimal mass. Second principle is to utilized momentum-based control planner. Instead of exploring torque profile as a function of time, we will explore the state through momentum scaling method we used in the MIT Cheetah. Another strategy we plan to employ is the sequential exploration. We will grow the VSN according to the momentum scale. The VSN will start from zero-momentum states to explore smaller momentum state and reach larger momentum states. The angular momentum will be strictly regulated unless the behaviors critically require angular momentum (e.g. turning).

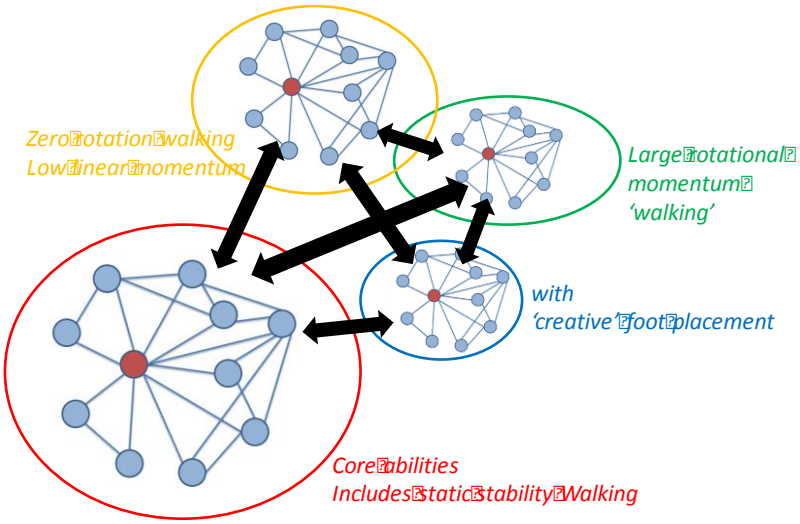

Figure 2. Viability State Network (VSN) growing scheme in different difficulty level. Typically, angular momentum is discouraged and heavily regulated until the behaviors become complex. The VSN will grow from standing to one step motion, multi-step motion, and eventually motion with large rotational momentum.

The second challenge is the model accuracy. There are variety sources that can cause model inaccuracy such as bearing friction model, gear friction modeling, contact dynamics, body frame flex and the motor dynamics. Instead of trying to refine the model to be close to the real-machine, we will use reduced model for VSN exploration and use the network as a planner, but not complete solution for controlling the robot. Our control algorithm will be composed of two part - feed forward planning and feedback regulation. The VSN will only provide feed forward planning controller but not provide the regulatory feedback. For

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

example, when human walks, major muscles that provide propulsion and global momentum change will be similar at every walk cycle. However, the regulatory muscles, which typically small and fast, that stabilize the body change its action at almost every step. Disturbance rejection should be done via real-time feedback controller while global momentum shifting should be done through the pre-calculated viability state network. The regulatory controller will be designed separately and implemented in the region where the VSN is already formed. The controller validation model will be more complex and realistic than the model we use in VSN and combine the regulatory controller to test the approach.

List of dependencies: T1C2S5A, T1C2S5C

Collaboration:

The main collaboration will be with Brett Kennedy at JPL, in charge of the development of LLAMA platform. The network development work along with modeling and simulation will build on our ongoing collaboration with Jason Pusey and Harris Edge at ARL. Graduate students in our laboratory will continue to exchange the progress and discussions with ARL and JPL. This collaboration includes personal traveling to ARL and JPL to assist with development and testing on LLAMA.

Linkage to previous RCTA BPP sub-tasks:

Expected products:

Deliv. by end of	Products
2017	- A method and a software package for VSN formation. This package will allow the users to grow a VSN for operating a legged machine. A successful example of implementation of the method in a simulation quadruped model. - A measurement tool of stability based on the actuator effort to reach the viability boundary. - Demonstration of diverse behaviors within viable states in dynamic simulation. - One conference publication on the concept and the methodology
2018	- Implementation of the stability measurement in a real hardware to improve the locomotion stability - Exploration of new gaits and new behaviors created from the VSN - Implementation of the new behaviors verified from VSN - One journal publication on the concept and the methodology

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.5.9 T1C2S5C – Analysis of Legged Morphological Design and Control for Traversal of Complex Terrain at OPTempo for Different Scales

PI: Pusey, Jason (ARL), jason.l.pusey.civ@mail.mil, (410)278-9098

Objective and Benefits:
The objective of this task is to analyze the effect of variations of morphological, power distribution and other design features on the stability and performance of LLAMA. As our previous RCTA-related efforts (Blackman, 2017) have indicated there is a close coupling between the kinematic structure of a robot's leg, it's control policy and the resulting dynamic performance. Small changes in posture can result in dramatic changes in efficiency or performance. In this task we will look at how leg posture, motor specialization, terrain, and leg compliance affect the behavior of the LLAMA robot.

Scientific Approach:
Utilizing dynamic Simulation of LLAMA in LMS, and building on work from previous research leg morphology in a conservative Matlab model, we will examine the coupling between control morphological parameters on the speed, stability, and efficiency of the resulting motion.
Simulations of the robot at different sizes will be utilized to verify the dynamic scaling predictions. In particular the power requirements for running will be considered. Motor models will be verified and used to suggest refinements to both mechanical design and control and how to optimally distribute motor power. These simulations will also be used to investigate dynamic maneuverability, such as turning. For each specific dynamical gait, such as stair climbing, up hills, jumping, different surfaces, evaluate the energetic cost and dynamic limits. These metrics are essential for connection to planning and for determining when each will be used, and for determining how actuator demands vary for each case.

Dependencies and Collaborations:
Primary collaboration will be with FSU (Jonathan Clark, Daniel Blackman, Jason Brown), JPL (Brett Kennedy, Jay Jasper, Sisir Karumanchi) and UPenn (Dan Koditschek, Wei-Hsi Chen).

Note: UPenn (Dan Koditschek, Wei-Hsi Chen) supported through ONR grant #N00014-16-1-2817, a Vannevar Bush Fellowship held by the last author, sponsored by the Basic Research Office of the Assistant Secretary of Defense for Research and Engineering.

Expected Products:
Year 1 (2017) - Develop Matlab and LMS models of LLAMA. Use Matlab model to compare stability regions of different leg designs in a conservative model. Use non-conservative LMS models and physical experiments to analyze the realization of these designs (and their conservative predicted stability regions) within a natural real-world environment. Develop 3D multi-body dynamic models of the Llama and MiniTaur quadruped models within LMS Virtual Lab. Use Llama LMS model to assist in design of RCTA Llama. Use Matlab and LMS models to support scaling study.

Year 2 (2018) - Use LMS models for to test different controller concepts developed in Matlab Simulink. Controllers developed in Matlab-Simulink can be directly linked to and control the LMS multi-body dynamics models in a co-simulation.

References:
Blackman, D., Pusey, J., Nicholson, J., Austin, M., Brown, J., and Clark, J. "Running and Jumping Dynamics of a Symmetric 5-Bar Leg Design," IEEE ICRA 2017 (submitted).

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2.6 Thrust 1/Capability 3: Optempo mobility in confined spaces

2.6.1 Overview

Imagine a squad consisting of a mix of human soldiers and robots searching for an object of interest in an urban area. The robot is told to look in a building down an alleyway, which will require it to navigate over an uneven road, climb a few steps up to the door, enter the building, look around a small room, and return. There are many obstacles on its path, including both unrelated clutter and human civilians. The robot must, at a speed relevant to the operation, detect and track these obstacles, reason about which are static and which may be moving, factor in social cues (both coordinating with the squad and being respectful of the civilians in the area), generate and update a motion plan, and execute the mission.

In this scenario the environment makes all aspects of the mission more challenging. Moving in these confined spaces with clutter and uneven terrain (both structured, e.g. stairs, and unstructured, e.g. cluttered room) requires careful planning and continual online replanning. The squad does not have a detailed map of the area and all of the obstacles ahead of time, and so the robot must interpret the commands in the context of the environment it finds. As it is moving through the space at the pace required, its perspective on the world may be jostled but must be maintained. The confined space means that the robot cannot simply avoid any civilians in the area – instead it must respond to social cues and interact in a respectful and predictable manor. These challenges will hold for any robotic platform, whether it be legged, wheeled, or tracked.

2.6.2 Relevance to the RCTA and the Army

This capability addresses the challenges involved in operating in confined spaces. In particular, the focus is on urban spaces which are of high importance to the Army. These environments are very hard to handle with state of the art techniques. Particularly in urban spaces, interaction with humans is unavoidable and doing so while responding to social cues will help the overall operation. This capability also focuses on moving at optempo, which requires fast perception that can accommodate a rapidly changing perspective as the robot rumbles over the terrain. For army-relevant missions, we cannot choose the terrain or engineer it to be smooth and perfect. Therefore the perception, planning, and control must all seamlessly integrate to be successful in these environments.

2.6.3 Required research areas

The state of the art for this capability involves motion planning in static environments, limited understanding of interactions and behaviors, slow operation to help perception and allow for complete re-planning, and is greatly aided by a priori knowledge. To move beyond this, this capability requires:

- Perception primitives for understanding static environments, especially with interior constraints.
- Perception primitives for understanding dynamic environments with occlusions and complex visibility issues in confined spaces.
- Dynamic motion primitives for high-energy self-manipulation.
- Anticipating interactions between agents and environments.
- Control and planning in dynamic, obstructed environments.
- Generating planning directives from human teammates and reporting observations to human teammates.
- Robust perception in the presence of chaotic/rough motions.

To achieve this, the capability is broken into three sub-areas:

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

1. **Robust Perception** in confined, unstructured environments
2. **Planning in confined, dynamic environments**, including interactions
3. **Human interaction,** both with teammates and other agents

2.6.3.1 Robust Perception

A mobile robot running at optempo in a confined environment cannot have a fixed or even perfectly smooth perspective on the world. As the pace of the mission and the roughness of the terrain increase, this perspective will get even more uneven. State of the art perception algorithms are mostly designed for fixed or smoothly varying motions, and can easily loose tracking in the presence of large cyclic oscillations or impulsive motions. These motions are worse in legged platforms but are present in any mobile platform, including wheeled or tracked vehicles where compliance in the tread, suspension, and mechanism can magnify any roughness in the terrain. The collaborative and integrated research focus of the CTA is uniquely positioned to explore these fundamental issues that arise when trying to field a complete system. This part of Capability 3 addresses these issues by making novel changes to the perception algorithms, sensors, and the platform mechanics.

Algorithms: At the algorithmic level, this capability will investigate how to incorporate an understanding of the dynamics of the system into the vision algorithm in order to work better on systems with non-trivial dynamics. This will take the form both as conventional, first principles models as well as learned CNNs. These dynamic motion models will be informed by the locomotor control and use novel proprioceptive feedback. They can then be used both to reweight motion estimations and also to learn when the system as a whole can and cannot trust the output of the algorithm.

Sensors: Perception in these challenging and dynamic environments must be fast enough to keep up with the motion of the robot and features of interest. Therefore this capability will also look into novel ways of using embedded hardware (FPGAs) to fuse visual, inertial, and other sensors faster than could be done in software. This provides a unique capability to keep up with the pace of the mission, but there are many questions that remain about how to process the data at this lower level.

Mechanics: In a complete system, the perceptual and locomotion control objectives should not be considered separately. Instead, a tight coupling between perception and mechanical control will lead to the most accurate perception and the most robust closed loop control. In order to improve the visibility on the scene, this can be achieved both by modulating the motion of the legs/wheels/tracks and also by modifying the mechanics itself to use novel appendages. Unlike with conventional off-the-shelf gimbals, adding appendages for this purpose must be done in an integrated manor with the locomotion system in order to have sufficiently precise state estimation and closed loop control.

2.6.3.2 Planning in confined, dynamic environments

Generating motion plans for complex robots in confined environments is very challenging. State of the art algorithms for learning and planning are often designed for simple bodies (assuming the robot is a point particle or skid-steer) and/or simple environments (clear open hallways, 2D, etc). To move beyond this we need new algorithms that can learn from experience and from humans, but are not limited by these inputs. The only way that a robot can succeed in the most challenging environments is to have a more complex body morphology, and even then only if the planning and control software are advanced enough to take advantage of these affordances. The RCTA, with its diverse group of researchers, is uniquely positioned to explore this fusion of novel mechanisms and algorithms. The research in this part of Capability 3 falls into three categories: multi-modal & symbolic reasoning, semantic perception, and learning whole body behaviors.

Multimodal and symbolic reasoning: Intelligent systems in challenging environments need to take advantage of any and all sources of knowledge: human demonstration, geometric reasoning, symbolic labels, past experience, etc. The research tasks in this capability address all of these by fusing different learning modalities at different scales, global reasoning over multiple hypotheses, taking advantage of human demonstration, and moving beyond the semantic symbols to learn more abstract planning methods over the lifetime of the system.

Semantic perception: The level of planning described above depends critically on accurate, semantic description of the scene from perception. State of the art techniques for indoor scene understanding, including from earlier stages of the CTA, include inference machines and CNN techniques for semantic segmentation that incorporate geometric constraints of indoor scenes, e.g., frequent orthogonality of surfaces. The SoA is limited to off-line operation, does not suitably reason about occlusion and physical interactions, and does not reason about ambiguity. All three limitations have impact on the feasibility of planning in complex environment. We propose to extend the anytime aspects developed in Capability 1 to the type of constrained semantic segmentation used here, enabling operation with limited computation. Reasoning about geometry is not sufficient in confined spaces, we extend to reasoning about physical interaction, building on recent work on correlating visual observations with physical interactions. Finally, forcing a single interpretation of a scene is guaranteed to lead to errors, which translate to catastrophic planning failures. Instead, we modify our perception algorithms to directly generate multiple hypotheses, using techniques based on contextual submodular list optimization developed in part in earlier phases of the program. Overall, this part of the work will produce a unique approach combining semantic segmentation constrained through geometry and physics, anytime design for limited computational budget, and direct representation of multiple hypotheses.

Learning whole body behaviors: Confined environments that are changing dynamically require the robot to leverage all of the available affordances. The research in this capability will develop new methods of taking advantage of the whole require by generating new compositional behaviors, learning from past experience with the whole body, and closing the loop back to the design of the body itself.

2.6.3.3 Human Interaction

A robot on a mission in an urban area will naturally interact with at least two groups of humans around it: the robot's teammates, and the unrelated humans who it happens upon. Interacting with these two groups is not avoidable, and therefore it must be aware of its role in the social interaction. State of the art methods either assume a controlled interaction, such as a collaborative assembly task where all agents are instructed ahead of time or direct control through e.g. a joystick, or use rule based reasoning.

This work goes beyond that in a number of different ways. First we propose to use fine-grained sensing of human intention to better predict the dynamics and social interactions in the scene. This will enable more naturalistic interfaces between the robot and its team, as well as inform social cues from all humans in the environment. Then we need to develop a better language for shared situational awareness among the team, and generate succinct, mission-relevant descriptions. Finally, we will develop new methods of planning to leverage these social cues and naturalistic descriptions.

2.6.4 Demonstration plan

If funded, the research tasks for this capability will continually work to combine into task and pair-wise demonstrations, culminating in the following:

- Robust Perception (CY17): The robust perception tasks are integrated onto a legged platform and demonstrate accurate perception while running over challenging terrain. This milestone demonstrates the capability in static environments.
- Robust Perception on Next Gen Platform (CY19): This milestone extends to dynamic environments and general platforms. Further developments in the robust perception are tested on the next gen integration platform. The perception is also used with the human intention sensors to inform the multi-modal reasoning methods.
- Pushing and Complex Terrain Motions (CY19): As commanded through the naturalistic interface, this milestone demonstrates the learned whole-body behavioral capabilities to interact with the environment.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.0 Thrust 2: Human-Robot Execution of Complex Missions

3.1 Overview

In addition to maneuvering at optempo speeds through unstructured terrain (i.e., the focus of the capabilities in Thrust 1), a crucial capability for RCTA systems is the ability to carry out complex missions as part of a human-robot team. A key goal for the RCTA is to enable robots that can do more than simply travel from point to point, but instead can execute multi-stage missions that do not have a specific location as an endpoint. Further, the robots should be designed to increase Soldier situational awareness, to maintain security, and to react to changing environmental conditions. These different mission goals require a range of capabilities, from simply understanding the environment, to planning missions with a high level of autonomy, to sharing information across a team, while doing all of these with a high level of reliability that engenders Soldier trust.

Various extant, state-of-the-art systems have focused on narrow slices of these capabilities; for example, self-driving cars clearly require a high level of reliability, but execute missions that have limited scope in terms of unstructured environments and limited levels of autonomy, resorting to a human driver as soon as the world changes in unexpected ways. Manufacturing and service robots require a high degree of reliability and trust, but do so with very limited interaction modalities, requiring an enormous degree of effort on the human partner to model the robot partner and manage the overall performance.

There are substantial technical gaps that must be addressed to deliver systems that are capable of building human-understandable, semantic models of the environment, execute a mission as part of a human-robot team, share information across that team, and be trusted to carry out missions reliably. As shown above, today's state of the art is that pockets of the problem space are addressed within relatively narrowly defined applications which often are brittle and not resilient to even small perturbations. Consequently, the overarching goals of this thrust are to develop robotic capabilities that are applicable to a range of mission contexts, robust to changes in the environment, and verifiably worthy of trust by their human teammates.

In pursuit of these goals, we are conducting the necessary research for three capabilities under this Thrust. These are (a) that the robot can perceive and make sense in an unstructured world and communicate the critical situation elements to the Soldier, in order to build Soldier situation awareness; (b) planning and execution of militarily relevant tasks in a collaborative and interactive team, something we have termed "distributed mission execution" and which requires the ability to understand mission intent, ability to decompose the mission into tasks, and the iterative process of sequencing and executing tasks to take advantage of the strengths inherent in both Soldier and robot; and (c) the underlying scientific and technological advances necessary for the verification of robot behavior and for the development of Soldier trust in their robotic teammates.

3.2 Thrust 2/Capability 1: Situation awareness in unstructured environments

3.2.1 Overview

In order for a Soldier-robot team to perform a team mission at optempo in unstructured, unfamiliar environments, it is crucial for the team to develop accurate and cohesive situation awareness. Consider a scenario where a team of Soldiers and robots operate in unknown, unstructured environments, screening a large area in a distributed manner. Given a command **"Robot – go tell me what's out there,"** the robot is expected to explore an environment and report back to the Soldiers an intelligible description of the environment. For instance, a robot enters and moves through indoor and outdoor environments that include rubble, stairs, doors and people moving about. The robot populates its world model with semantic concepts in terms of scene types, objects, people, and behaviors and can report back in a natural way a concise

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

description about objects of interest and other physical constraints in the environment, as well as various human activities and their intentions.

In order to support such a capability, the robot needs to be able to *perceive/understand* the semantic meanings of both static and dynamic environmental features, e.g., challenging types of terrains, ill-structured or naturally-shaped objects, and dynamically moving people as well as other robots. The robot needs to be able to autonomously *navigate* in challenging indoor and outdoor environments. The robot needs to be able to *communicate with soldiers* effectively to understand the description of a scene, task, and plans expressed in natural language. At the same time, the robot also needs to be able to generate descriptions and explanations for what it has observed, what progress has been made in the team's task, and what their future plans are. In communicating such information with soldiers, the robots need to be able to choose and adjust the type of modality and the level of details to meet the soldiers' needs and incorporate a description of complex scenes and plans into a *multimodal display*.

3.2.2 Relevance to the RCTA and the Army

To develop a competent robot teammate that can work in a team with Soldiers in support of the RCTA vision, it is essential to address research questions needed for establishing solid technical basis for the situation awareness capability. For example, in unstructured environments and unfamiliar environments, robots need to be able to go explore and report back an intelligible description of the environment.

3.2.3 Required research areas

This research addresses the following research questions:

- How can we design the perception system to deal with challenges in dynamic environments? To address this issue, we propose to investigate approaches in understanding human activities and dynamic events.
- How can we generate a concise description of an environment in natural language? To address this challenge, we propose to extend our prior work on language understanding and also explore neural-net based approaches that has been used for image caption generation.
- How can we design the interface between soldiers and robots that will maximize the quality of information shared under limited resources? We propose to approach this problem through user-centric models, and multimodal study on human-soldier exchanges.

3.2.3.1 Dynamic scene understanding in unstructured environments

Dynamic scene understanding across length scales (big and small) and integrated with gesture moves well beyond the state-of-the-art in person perception. The purpose of this research area is to investigate perception and tracking of human pose, articulated motion, detection of fine-grained contextual cues, such as holding or interacting with small objects, and motion reasoning to support gestural and activity understanding. Current research has investigated the use of deep networks to generate high quality articulated human pose estimates on benchmark datasets using monocular imagery. In addition, research has explored pose tracking in video through multi-frame and temporal joint trajectory modeling at minimal computational cost.

This research area has two objectives: (a) move from academic datasets to real-world environments, and (b) develop new techniques to:

- Recognize pose context such as items held, worn, interacted with, etc.
- Basic recovery of 3d pose such as localizing gaze or pointing directions in the world frame
- Improving body motion understanding

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.2.3.2 Synthesizing semantic content into human-understandable descriptions

The goal of this research area is to use language as an efficient means of establishing shared situational awareness that focuses on task-relevant descriptions between robots and Soldiers. Current research has demonstrated a very basic ability to understand a task command given by a human commander and perceiving an environment within that task context. This, however, is currently limited to operating in an imitation-learning framework that is mainly based on hand-picked spatial features and, as a result, provides only a limited understanding of task-specific information. In addition to understanding language, robots need to be able to generate descriptions and explanations in natural language to communicate back to Soldiers.

To address these needs, this research area proposes to develop new techniques that combine probabilistic graphical models (inference) with deep learning (feature acquisition) to: (a) learn new features correlating language with image and 3D data, (b) train language models together with vision-based recognition such as CNN-RNN models to better associate language features with visual features, and (c) generate descriptions and explanations that are natural to human users (when compared to template-based approaches). In addition, this research aims to provide an analysis of graphic types of communication (e.g., visual: augmentation on the sensor feed, or as still images) and how they can best be used for communication within the Soldier-robot team. We also want to develop new evaluation methods to measure impact on mission and team performance and SA. These objectives would allow to extend across time and length scales while incorporating environmental and temporal context.

3.2.3.3 System design for effective teaming

Although system design research has investigated interaction models that focus on mobile device (e.g., smart phone) and desktop type interactions, very little work has been done with respect to the development of new design models for devices and procedures in human-robot interaction that take into account task execution and multimodality. The purpose of this research area is, therefore, to help inform user-centered design (UCD) of naturalistic Soldier-robot communication interfaces. These UCD interfaces need to benefit from state-of-the-art engineering capabilities but should have a special focus on user experience to minimize risk of disuse (abandonment) or their misuse. This approach will maximize the team's ability to communicate effectively when executing missions. Current research has explored multimodal communication interfaces (e.g., MMI), however, these interfaces have yet to be tested with end-users (i.e., Soldiers) to evaluate overall usability.

This research area proposes to inform communication interface design by conducting interviews with projected end-users (e.g., Soldiers) to identify their expectations and preferences. In addition, user testing will be conducted to evaluate usability of developed interfaces (e.g., MMI).

3.2.4 Demonstration Plan

The effectiveness of a Soldier-robot team's situation awareness will be demonstrated within a larger mission context, such as *search and rescue* or *cordon and search*. An example of envisioned demonstration is the following: A Soldier commands a robot, "Go tell me what's out there," to send it to several different sites that include rubble, stairs, pedestrians, etc. The robot autonomously navigates through the challenging environments, asks for disambiguation when needed, and reports back to the soldier via a multimodal display about the environments and anomalous events, e.g., "There are people in uniform standing in rubbles in the back of the tall red building," accompanied with key image frames if there is enough network bandwidth. This example illustrates the challenges that we propose to address in

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

the areas of perception, language synthesis, and human-robot communication modalities. We expect these milestones to be executed on the existing Husky platform.

We envision four milestones for integrating the technical elements for capability 1:

- Building A Semantic World Model in Unstructured Environments (CY17): This milestone will demonstrate the ability to identify objects and concepts and add them to the world model, focusing on properties of unstructured environments. The evaluation will be in terms of the concepts added to the world model as the robot explores.
- Synthesizing Semantic Content into Human Understandable Descriptions (CY17/C18): This second milestone adds the ability to explain the world model in human-understandable terms, and interact with the human team-mates for disambiguation.
- Dynamic Scene Understanding in Unstructured Environments (CY18): The third milestone addresses the need to characterize dynamic scenes. Changes to the world model to represent dynamics, as well as the ability to infer dynamic entities from sensor data will be required.
- Designed System Demonstrating Effective Teaming (CY19): The first three milestones address the underlying algorithmic capabilities. This milestone evaluates the performance of the system in terms of how human team-mates acquire improved situational awareness, and how the ability to build the world model and share it across the team is dependent upon a user-centered design.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.3 Thrust 2/Capability 1 Subtasks

3.3.1 T2C1S2A - Long term learning and prediction of 'action objects' interaction

PI: Shi, Jianbo (UPenn), jshi@cis.upenn.edu, (215)746-2851

Objective and Benefits:

Human-robot teaming in unstructured environments has been one of the most important robotics problems. For this effort to be successful, robot needs to understand human intention and anticipate their movements, and human needs to be able to teach robot how they prefer to carry out skilled tasks. For example, as a soldier runs to evacuate from a building, the robot nearby needs to know which exists he is taking, and provides support and covers for him. At the meantime, as a robot observes how a soldier evacuate from an unfamiliar building, it should learn from his experience and action the navigation and interaction skills so the robot can acquire more 'human-like' behaviors.

Our research supports this thrusts through two efforts: 1) enable robots to understand human intention and attention by "seeing through the person's eyes"; and 2) enable a human to teach robot human-like skills using first-person perspectives.

Scientific approach:

The use of first-person camera is a key aspect of our approach. A first person camera placed at the person's head captures candid moments in our interactions, providing detailed visual data of how we interact with people, space, and objects. It reveals our attention, intention and momentary visual sensorimotor behaviors. With the first person vision, our goal is to build a computational model for personalized intelligence that predicts what we see and act by "putting yourself in her/his shoes".

There are multiple challenges to recognition human intention and prediction of movement in complex environments. First, individuals can carry out unexpected actions intentionally or by mistake. Second, first-person videos can be highly jittery, resulting in fast and unpredictable object motions in the field of view. Third, any given individual might carry out the same activity in a variety of ways, due to variances in individuals' preferences and abilities.

We build on our recent work establishing computational foundations for modeling social (people-people) and personal (people-scene, and people-object) interactions using first person cameras.

We will focus on two most important aspects of human-robot teaming: motion planning and object interaction. Motion planning in unstructured environments has been one of the most important robotics problems. Integrating such a motion planning mechanism into a machine is a challenging task, because it requires a good understanding of different physical environments, in which a robot needs to operate. Humans have lots of prior experience navigating in many different environments, which makes them pretty good at this task even when placed in the completely novel environments. Thus, if machines could learn navigation skills directly from humans, they could presumably attain equally good capability for motion planning in unstructured environments. In order to achieve this goal we propose to use first person cameras, which allow us to see exactly what that person sees. Such an ability to see ``through the person's

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

eyes" allows us to better understand (1) how that person navigates in his environment, and also (2) how he interacts with various objects in the environment around him.

To learn common navigations patterns of a person from a first-person video we could first use the standard structure from motion pipeline to automatically extract that person's movement trajectory over time. Using this trajectory data as the ground truth annotations, we could then build a predictive navigation model that learns common person's navigation patterns in specific environments directly from the data. Furthermore, in addition to the navigation capability, we also want to infer the locations of exits or other areas in the environment (e.g. stairway passage) that may not be directly visible from the visual signal. To achieve this goal, we aim to employ a temporal structure of the first person video data. More specifically, we can first pair up the frames from the first person video data that are T frames apart from each other. Then the location of a person at a frame t+T will likely not be visible at a frame t. Thus, given a frame from time t we can use the location from the frame t+T as the ground truth optimization criteria for our predictive ``hidden area discovery'' model. Such a learning scheme would allow our model to discover ``hidden'' or not clearly visible locations, which are beneficial for general navigation capability too.

Finally, in order to navigate in the unstructured environments effectively, one typically needs to interact with various objects (e.g. open a door, push a box that blocks a pathway, etc). Thus, it is important to know, (1) which objects to pay attention to and (2) how to interact with those objects. Learning to recognize such objects from the first-person view has significant benefits because they are spatially normalized to the first person camera: they typically appear at a constant spatial location in the first person view (e.g. around the center). Additionally, due to the fixed length of an arm, the distance from a person and an object remains approximately fixed (0.3-0.5m). Thus, by employing first-person stereo cameras, which capture visual, and depth information of the scene, we can leverage spatial 3D and visual appearance cues to detect such objects in a data-driven manner.

We reason our trajectories in both first person ego-centric 3D coordinate system, and global 3D coordinate system. We use visual odometry to estimate the person's 3D position and orientation from a wearable camera + IMU. For ego-centric 3D coordinate system, we use a common ground plane to link different first person's 3D coordinate into a consistent 3D coordinate. We will use our visual odometry algorithm combining visual sensor and IMU. It will reason over time to allow noisy measurement. The inferred interaction will be trajectory of body motion, or hand motion, or gaze motion in the 2.5 ego-centric first person view. Since we compute body motion in 3D using visual odometry, we can translate those motion trajectories in to a global 3D coordinate frame.

We use both image semantics, ego-motion of the solider, context movement of the nearby human and objects, as well as the global map to predict a person's future trajectory and objects of attention.

Prediction of human attention can be used directly in the visual perception, to focus object recognition on only relevant area of images/video. Future trajectory of body motion, object actions and attention would be also used in both 1) action recognition in **T2C1S2D**, since action recognition requires understand what objects a solider is paying attention to; and 2) two way translation of language and vision in **T2C1S3A**, since attention would allow the system to disambiguate which objects a solider is currently talking about.

List of dependencies: We have our own object detection and people tracking routines. If the tasks in T1C1S1A and T1C1S2A results in better system, we can integrate those.

We only expect images/videos and IMU as input. We can integrate CMU's if those are better.

Collaboration: Possible joint research through student exchange with ARL.

Linkage to previous RCTA BPP sub-tasks: P3-1

Expected products:
Publications, presentations and workshops at CVPR, ICCV, ECCV or NIPS. Algorithms and codes for

1) Software for First-Person prediction of future location of a moving person in complex unknown environment, at 5, 10, 15 seconds ahead
2) Software for Frist-Person object action prediction of a moving person in complex unknown environment, hand motion and gaze motion at 5, 10, 15 seconds ahead
3) Software for First-Person prediction of human attention without gaze tracking. Prediction of what people are looking, and where they will interact with at 5, 10, 15 seconds ahead.

References:
[1] Gedas Bertasius, Hyun Soo Park, Stella X. Yu, and Jianbo Shi. First person action-object detection with egonet. CoRR, abs/1603.04908, 2016.
[2] Hyun Soo Park, Jyh-Jing Hwang, Yedong Niu, and Jianbo Shi. Egocentric future localization. In CVPR, 2016.
[3] Hyun Soo Park, Kris Kitani, Jianbo Shi, CVPR 2016 Tutorial on First Person Vision

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.3.2 T2C1S2B – Dynamic fine-grained understanding of human scenes

PI: Rothrock, Brandon (JPL/Caltech), rothrock@jpl.nasa.gov, (818)354-4689

Objective and Benefits:

The objective of this subtask is to develop and incorporate robust and detailed perception of human agents within dynamic environments to enhance situational awareness of the robot and enable capabilities that utilize human-robot teaming and collaborative execution. This subtask will develop capabilities to visually localize humans, understand their body motion and dynamics, and characterize basic activities or interactions. These capabilities can support navigation in complex human environments, enable gestural communication with human team members, and assess threat levels of unknown parties. Our proposed approach consists of developing three complementary capabilities: perception of 3D human attributes, articulated pose tracking and gesture recognition, and fine-grained understanding of human interactions with the environment. 3D human perception consists of estimating the true location of humans in the environment, as well as spatial properties of the pose such as gaze or pointing direction. By incorporating motion reasoning to the articulated pose, a set of standardized gestures can be recognized to convey commands, identify locations in the environment through pointing, and provide a general visual communication pathway between a human operator and the robot. Lastly, many important human activities are characterized by manipulation or interaction with the environment such as the use of tools or weapons. By utilizing interaction context of the human, the local environment, and body motion, small items being held or manipulated can be detected and classified more reliably. This can help provide important intelligence for assessing threats in the environment, and establishing situation awareness to team members.

Scientific approach:

Dynamic scene understanding containing human gestures and trajectories across both big and small time scales moves well beyond the state-of-the-art in person perception. The technical approach for the proposed articulated pose estimation and pose tracking leverages previous work using recurrent fully-convolutional neural networks for direct prediction of spatial positions of the body joints, as well as temporal tracking of those joints. We propose to extend this framework to incorporate 3D reasoning to pose tracking by utilizing kinematic priors from motion capture datasets, as well as exploiting scene structure such as planar regions and vanishing points to inform the scale and position of humans. Motion analysis for gesture recognition and prediction of object interactions will utilize novel models to represent motion trajectories with hierarchical and compositional motion primitives combined with learned image appearance models. The state-of-art for gesture recognition is generally dominated by Kinect-based methods that rely on accurate articulated pose as input, or through the use of instrumentation on the human to sense arm and body motion. Kinect-based methods are limited to indoor environments and the short range of the sensor, however, and are not applicable to many practical field applications. Our approach advances the state-of-art by employing end-to-end learning of human gestures through pose trajectories and interaction context from monocular video that is much less susceptible to distance from the subject, lighting, or environment.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

List of Dependencies: The proposed method takes input from monocular or stereo cameras, and optionally semantic scene context from subtasks such as T1C1S1A. Output consists of human kinematic skeletons, part bounding volumes, and classification labels in the world or robot-centric frame.

Collaboration:
The most direct integration path to the system architecture is the deployment of human gestural control of the robot without explicit representation into the world model. More general perception of human dynamics in the environment will require additional development of the world model to incorporate dynamic entities, and the corresponding planners to utilize this information. Additional collaborations within the consortium include T2C1S2D at Indiana University to utilize the low-level human motion perception developed from this task to infer high-level activity semantics. Collaboration with ARL will include data collection and experimental evaluation with Craig Lennon. TBAs and MCAs will be conducted at JPL on the Surrogate platform, and will ultimately transition to RoMan.

Linkage to previous RCTA BPP sub-tasks:
This work is a continuation of the P6-1 subtask under past RCTA work entitled "Fine-grained Video Understanding of Humans and their Attributes in Complex Environments."

Expected products:

Year	Deliverables
2017	**Integration deliverable**: modular software package for human detection, pose tracking, and 3D pose reasoning. Package will provide structured motion trajectories of humans and their corresponding body pose in the environment, and support perception for semantic human activity models from T2C1S2D. **TBA at JPL** to evaluate gesture recognition performance against relevant video datasets. Tech transfer to RoMan/Robosimian platform to incorporate human perception capabilities on robot hardware. **Dataset collection** and quantitative evaluation to assess pose tracking accuracy, and for learning and evaluating classification accuracy of human gestures relevant to RCTA scenarios. **Publication**: academic paper submission targeting ICCV, CVPR, or ICRA
2018	**MCA at JPL** to demonstrate basic command and control of robot from human gesture input. **TBA at JPL** to provide quantitative assessments of human perception capabilities in scenarios relevant to the capstone. **Integration deliverable**: software module extensions to include human gesture recognition at optempo framerates to support gestural control and communication with the RoMan or Husky platforms. **Publication**: academic paper submission targeting CVPR or ICRA

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.3.3 T2C1S2D – Predicting human intent and activity possibilities

PI: Ryoo, Michael S. (Indiana University), mryoo@indiana.edu, 812-855-9190

Objective and Benefits: This task is about human activity recognition from videos, particularly focusing on robot learning to predict future human activities. This not only includes video feature-level representation and recognition of activities, but also covers higher-level modeling of onsets and sub-events composing human-human and human-robot activities for their prediction. The outcome of this basic research (6.1) task will provide robots an ability to automatically understand human activities and predict human intent, which directly supports "human-robot execution of complex missions". Activity-level situation awareness of a robot will be enabled, which is necessary for robots operating with friendly forces, moving near other people, or conducting surveillance. With the development of the proposed technologies, the robots will not only recognize human actions, human-human interactions, and human activities targeting the robot with better accuracies, but also predict future activities and infer possibilities, ultimately taking advantage of scene, object, and people context.

Scientific approach: Scientific challenges to be addressed in our 2017-18 effort are (1) end-to-end learning of latent onsets (i.e., subtle actions observed prior to activity occurrences) optimized for the forecasting of future activities and (2) robot recognition of explicit activity possibilities by visually predicting future video frames.

In our previous RCTA subtasks (P5-2 "Semantic understanding of human activities" in 2015 and P5-5 "Human activity recognition with context learning" in 2016), we developed a general framework to recognize different types of human activities using human pose features [Gori et al, 2015]. Furthermore, we studied the use of pre-defined onset activities [Ryoo et al., 2015] for the prediction of ongoing activities. In our 2017-18 effort, we will extend such onset-based activity prediction by automatically mining important onsets and taking advantage of them for the better activity prediction in an end-to-end fashion. We propose to represent each onset and sub-event using a fully-differentiable temporal attention filter, model relations between the filters, and learn iterative inference behaviors taking advantage of such relations with recurrent networks. Learning of latent onsets and their relations will be done together with features/representations for the video-based convolutional network such as [Piergiovanni et al., 2017], jointly optimizing them for the human intent prediction.

In addition, we will newly investigate approaches for robot recognition of explicit activity possibilities by visually predicting future video frames. New network models combining fully convolutional regression and deconvolutional layers will be designed for frame-level prediction under each possible activity. This will allow synthesizing expected future frames based on the current frame, and describe how the humans in the scene would move to perform the activity. Our method will take advantage of such frame-level predictions to evaluate possibilities of each activity, measuring what activities are possible (and likely) and what are not. Although CNN feature-level prediction of future frames have been studied previously [Vondrick et al, 2016], it was learned without considering explicit activities or having an ability to generate activity possibility videos.

List of dependencies: Our component will take pure video frames as inputs, without depending on outputs from other tasks. In addition, it will benefit from the fine-grained estimation of human skeletons and other objects from T2C1S2B (Rothrock, JPL) when its results are available. Use of objects/scene recognized from T1C1S1A (Hebert, CMU) is also a possibility. Our component will not depend on GPS-based or other localization of the robot or a map from the world model. The outputs of our component will be provided as time intervals (i.e., starting and ending time) of detected human activities. The cognitive architecture and the world model will need to be extended to capture such outputs. In terms of the platform, the proposed approach is able to cope with any of RCTA platforms as long as videos are taken with the robot camera. Use of RGB-D sensors will benefit more reliable estimation of human skeletons and thus benefit the proposed component.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Collaboration: We plan to collaborate with JPL researchers working on T2C1S2B. Our planned effort will benefit from the fine-grained estimation of humans/scene/objects, which will become possible with the outcome of T2C1S2B. In addition, we will collaborate with Craig Lennon at ARL for the video data collection appropriate for the evaluation of the human activity recognition methods, and with Prudhvi Gurram at ARL for the data collection and research discussions.

Linkage to previous RCTA BPP sub-tasks: This work is a continuation of the past P5-2 (2015) and P5-5 (2016) perception tasks.

Expected products: The algorithms developed in this task will be packaged as a modular software library, to be used by the RCTA platforms. They will be tested and evaluated through experiments using publically available datasets as well as using the video dataset jointly collected with ARL and JPL. We will perform the integration in coordination with ARL researchers, for the desired platform.

End of CY2017	• Implementation of human activity recognition and intent prediction using onset learning. A software package for early recognition of ongoing activities in the environment. • TBA at Indiana University; evaluation of the developed software using public and new video datasets.
End of CY2018	• Implementation of a new approach for predicting future video frames and their visualization. A modular software package will be delivered. • Implementation of an approach for inferring activity possibilities by visually predicting future video frames • Performance evaluation of this capability using the data jointly collected with ARL. MCA.

We will target publications at conferences including CVPR, ICCV, ICRA, IROS, AAAI, IJCAI, and NIPS. A publication at AAAI 2017 has already been confirmed [Piergiovanni et al., 2017]. There also will be an invited paper at IJCAI 2017. The PI will provide a keynote talk at the workshop of WACV 2017, and also plans to organize workshops/tutorials.

References:

[Gori et al, 2015] I. Gori, J. K. Aggarwal, L. Matthies, and M. S. Ryoo, "Multi-Type Activity Recognition in Robot-Centric Scenarios", IEEE Robotics and Automation Letters (RA-L), 1(1):593-600, February 2016.

[Ryoo et al., 2015] M. S. Ryoo, T. J. Fuchs, L. Xia, J. K. Aggarwal, and L. Matthies, "Robot-Centric Activity Prediction from First-Person Videos: What Will They Do to Me?", ACM/IEEE International Conference on Human-Robot Interaction (HRI) 2015.

[Piergiovanni et al., 2017] A. Piergiovanni, C. Fan, and M. S. Ryoo, "Learning Latent Sub-events in Activity Videos Using Temporal Attention Filters", AAAI 2017.

[Vondrick et al, 2016] C. Vondrick, H. Pirsiavash, and A.Torralba, "Anticipating Visual Representations from Unlabeled Video", CVPR 2016.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.3.4 T2C1S3A - Two-way translation between vision and language

PI: Jean Oh (jeanoh@nrec.ri.cmu.edu) (412) 651-6052, CMU.

Objective and Benefits:

In this task, we address the following research question: How can a robot combine and communicate back the information that it perceived through its sensors and that has been verbally described by human teammates to effectively build up shared situation awareness? In order for a human-robot team to build accurate and cohesive situation awareness in unstructured environments, robots need to have a two-way communication capability to interpret human teammates' descriptions of an environment and convey similar types of information to human teammates by describing what they have perceived, have done or plan to do.

Scientific approach:

Building on our prior work on understanding verbal commands given by a human commander and perceiving an environment within the task context [Boularias2015, Boularias2016, Oh2015, Oh2016], we will continue developing the intelligence framework to support deeper understanding of rich and complex information, and support the reverse direction of information flow, that is, generating multimodal descriptions and explanations for the robot's perception and plans.

We view human-robot communication as a team perception problem where each agent may have one's own perception of a shared environment. Here, the goal of the team is to gradually build a more accurate and cohesive situation awareness as the teammates share their views through communication over time. In this context, we focus on two-way translation between vision and language such that what has been perceived (through vision) can be described in language to communicate with other teammates, while what has been described (in language) by other teammates can also be interpreted and fused with each agent's perception of the shared environment.

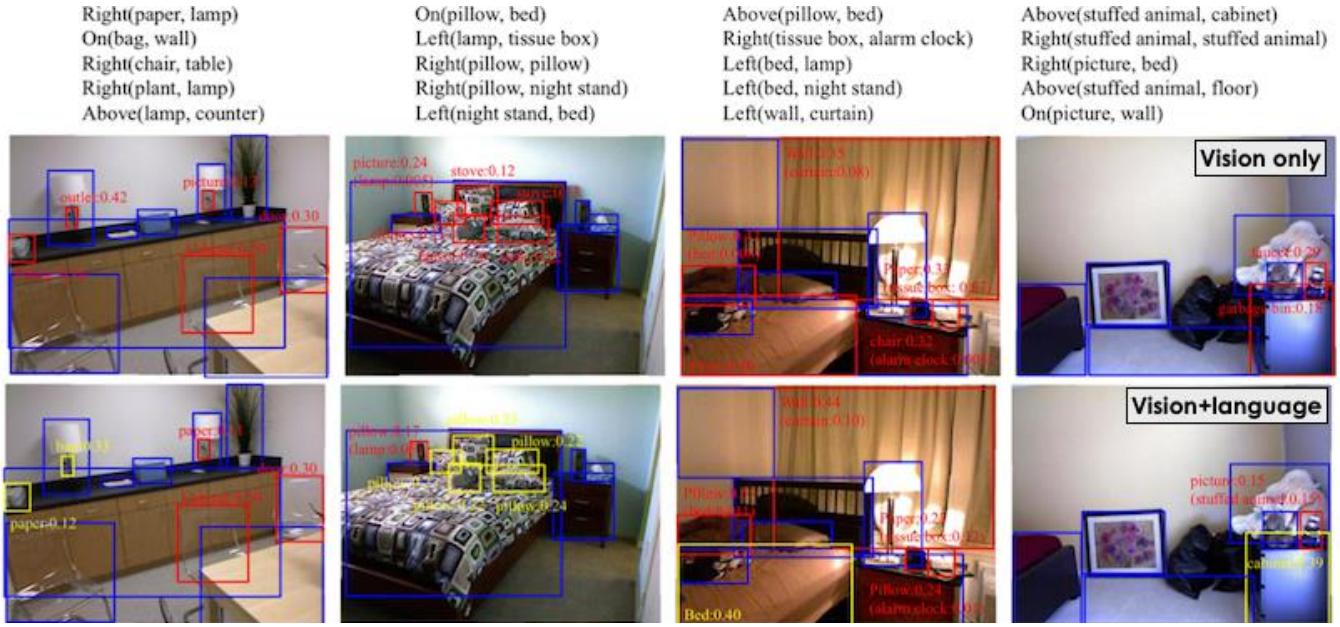

Fig. 1 Vision-language fusion to improve object recognition (blue: correct, red: incorrect, yellow: corrected after adding language features).

In the previous program, to support the fusion between vision and language, we have developed a random walk based graph model [Shiang2017]; sample results are shown in Fig. 1. Although the fusion approach

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

shows significant improvements over a vision-only approach, the current algorithm is difficult to generalize to support n-ary relations or aggregate constraints, and it also does not model location information as needed in semantic maps. To overcome the limitation in expressiveness, we will investigate other graphical models including Conditional Random Fields (CRF). For instance, in the CRF approach, a node can represent a bounding box in an input image that is associated with a random variable representing the probability distribution over candidate labels, and each link (or edge) between nodes can have a potential function representing the dependence among multiple objects enclosed in bounding boxes.

We will also conduct research on attention-based neural models that will give more flexibility in the natural language front end and that can be easily integrated with the synthesizing language task. Specifically, to generate descriptions from visual features, we will investigate CNN/LSTM models to connect visual features and language models. We will collaborate with task T1C1S2C (Bölöni, UCF) on accounting for social context here. Existing works on this end leave rooms for improvement especially in our target problem of describing the relationships among objects of interest in noisy unstructured environments.

List of dependencies: This task will depend on the perception tasks T1C1S1D (Hebert, CMU) and T1C1S3A (Muelling, CMU) since the objects detected from the perception systems will be used as inputs to this task.

Collaboration: We will collaborate with the perception tasks T1C1S1D (Hebert, CMU) for reasoning about 3D geometric object constraints and T1C1S3A (Muelling, CMU) for resource-bounded anytime perception. We also plan to collaborate with task T2C1S1D (Summers-Stay, ARL) for semantic analysis and T4C5S1E (Barber, UCF) for multimodal communication.

Linkage to previous RCTA BPP sub-tasks: I1-1 Framework for Collaborative Intelligence; I6-1 Self-Directed, On-Line Learning for Understanding the Environment and Improving Robot Performance

Expected products: The resulting software will be integrated on the IRA Husky platform to support indoor/outdoor navigation tasks. Since this research is interdisciplinary, target publication venues include AI conferences, e.g., AAAI, IJCAI, AAMAS, ICAPS; robotics conferences, e.g., R:SS, IROS, ICRA, HRI, or vision conferences, e.g., CVPR, ICCV, ECCV.

2017	• Develop a CRF model for vision-language fusion • Develop a neural net based approach with attention • Evaluate the fusion approach on public data sets, e.g., NYU Depth datasets • Develop CNN/LSTM description generation • Evaluate the description generation system on public data sets, e.g., VisualGenome
2018	• Extend the attention based model to support composite descriptions. • Develop algorithm for prioritizing descriptions to support concise & informative communication. • Evaluate on RCTA IRA platforms (ClearPath Husky). • Create RCTA image-text dataset for dissemination.

References:

[Shiang2017] Shiang et al. Vision-Language Fusion for Object Recognition. In Proc. of AAAI Conference on Artificial Intelligence (AAAI), 2017 (to appear).

[Oh2016] Oh et al. Integrated intelligence for human-robot teams. In Proc. of ISER, 2016.

[Boularias2016] Boularias et al. Learning Qualitative Spatial Relations for Robotic Navigation. in Proc. of International Joint Conference on Artificial Intelligence (IJCAI), 2016.

[Oh2015] Oh et al. Toward mobile robots reasoning like humans. In Proc. of AAAI, 2015.

[Boularias2015] Boularias et al. Learning to ground spatial relations for outdoor robot navigation. In Proc. of IEEE Conference on Robotics and Automation (ICRA), 2015.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.3.5 T2C1S3B - Generating Natural Language Scene Descriptions for Shared Situational Awareness

PI: Walter, Matthew R. (TTI-Chicago), mwalter@ttic.edu, (773)834-3637

Objective and Benefits:

We propose to develop models and algorithms that allow robots to synthesize natural language descriptions of their environment as a means of establishing shared situational awareness among human-robot teams. Shared situational awareness is essential in order for soldier-robot teams to accurately carry out tactical operations at optempo. However, establishing a common understanding is challenging when the robot is not nearby the soldier, particularly when the robot is operating in complex, cluttered environments. Current solutions to establishing situational awareness involve transmitting images from one or more cameras mounted to the robot or synthesizing maps of the robot's surround that are then visualized on a remote interface. These approaches have several drawbacks as a large amount of bandwidth is required to transmit images at sufficiently high resolutions and frame rates; images only convey information about the robot's immediate surroundings; and the soldier is required to devote their full attention to interpreting these image streams or maps.

Scientific approach:

Language provides a flexible, bandwidth-efficient medium for conveying environment information that complements existing methods for establishing shared situational awareness. We propose the development and implementation of models and algorithms that enable robots to generate natural language descriptions that express spatial, semantic, and topological properties of both their immediate and distant surroundings, in the form of potentially noisy estimates that result from a variety of inference methods. These may include the presence of people nearby and their perceived activities (e.g., sitting, running, etc.), the existence of dangerous or otherwise salient objects (e.g., IEDs), and the layout of buildings (e.g., number and location of rooms). There are two primary challenges to producing free-form descriptions that efficiently and effectively establish shared situational awareness. The first is to decide what to talk about, i.e., to reduce a potentially large, exhaustive model of the environment to a small set of properties most relevant to soldier according to the mission. The second is to synthesize succinct natural language utterances that accurately express these properties in a manner that the soldier can easily consume. We will address these challenges by framing the task as a *selective generation* problem [4], whereby we jointly reason over what information to convey (i.e., content selection) and how to best convey the information (i.e., surface realization). Specifically, will develop probabilistic models that express the mission-relevance of information expressed in the world model and algorithms for inferring which attributes to share with the soldier. To address the challenge of surface realization, we will develop sequence-to-sequence language models that convert the subset of relevant environment information to free-form descriptions. We will train the content selection and surface realization models using datasets collected as part of this effort using in-person and/or crowd-sourced user studies. Our effort will leverage existing models developed by PI Walter's group in the areas of natural language understanding [1,3] and synthesis [2].

In practice, we will assume access only to noisy estimates of the robot's pose (i.e., GPS-level accuracy) and estimates of various properties associated with the objects, people and locations in the environment (e.g., their name, size, location, and actions/behavior), which need not be maintained in a globally consistent metric map. In deciding which information contained in the world model to share with the soldier(s), the content selection mechanism will reason over various measures of estimation accuracy including the confidence

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

assigned by the inference methods (if available), the time since the properties were inferred, and the confidence in the robot's pose estimate.

List of dependencies:
Our model reasons over a world model that expresses various properties associated with the objects, people, and locations in the environment (e.g., their name, location, actions, intentions, etc.). To that end, the model will take as input the result of the inference capabilities developed under other subtasks. These include the predicted intent and motion of human's in the robot output by the module developed under T2C1S2A "Long term learning and prediction of 'action objects' interaction" (Shi, UPenn); the location and type of objects inferred as part of T3C1S2A "Geometry learning of rigid and deformable objects" (Daniilidis, UPenn) and T1C1S3A "Anytime Neural Networks for Semantic Understanding on the Move" (Muelling, CMU) and T1C1S1A "Rapid learning of models for unique mobility platforms" (Hebert, CMU); the mission-relevant objects in the robot's surround detected and localized as part of T1C1S1C "Scene Understanding at the Tactical Edge" (Gurram and Rao, ARL). We will interact with these capabilities through the world model.

The model and algorithm that we developed will interface closely with the multimodal interface (MMI) capabilities developed as part of T2C1S4B "Adaptive Multimodal Communication" (Barber, UCF). Specifically, we assume that our method will be initiated through a call from the MMI, upon which the natural language description (output) will be conveyed as text to the MMI, which will then synthesize the description in audio form.

The proposed content selection method will identify aspects of the environment relevant to the human-robot teams. In this way, this (intermediate) output will be relevant as an input to the module developed under T1C1S2C "Human-robot teaming in a social context" (Boloni, UCF).

Collaboration:
The subtask will have close ties and provide opportunities for collaboration with other researchers in the consortium. Of particular note are the 2017–2019 BPP efforts that are similarly focused on situational awareness, communication, and semantic content generation:

This subtask shares similar objectives with that of T2C1S3A "Two-way translation between vision and language" (Oh, CMU), which proposes the development of models that capture the relationship between images of a scene and their corresponding descriptions. Whereas their work focuses on descriptions of the robot's immediate surround (i.e., what is immediately visible using the robot's cameras), our proposed effort focuses on generating descriptions of spatially extended scenes, some or all of which may be outside the robot's field-of-view. We will work with Oh and her team at CMU to investigate the complementary nature of our tasks. For example, our proposed content selection method can be used to reduce the set of objects that they detect in images to a subset relevant to the mission.

The subtask will also work closely with D. Barber (UCF) and his team, with whom Walter has previously worked with, to interface our capability with the multimodal interface. This will take the form of shared code, in-person team meetings coinciding with RCTA events, and regular conference calls.

Linkage to previous RCTA BPP sub-tasks: H7-4

The proposed subtask will extend previous work by PI Walter performed under the RCTA, notably H4-2: "Human-Facilitated Learning of Shared Semantic World Models", which used environment descriptions as an observation (rather than an output) for learning environment models; and H3-3: "In-Situ Monologue Comprehension for Faster Mission Tempo", which was a joint effort between PI Walter and Thomas Howard (University of Rochester) that developed models and algorithms for understanding multi-sentence natural language instructions

Expected products:
We will publish papers describing the selective generation models and algorithms at workshops and robotics conferences (e.g., RSS, ICRA, and IROS) and in robotics journals (e.g., IJRR and TRO). Walter and Howard (University of Rochester) also plan on organizing a workshop on model learning for human-robot interaction, as they have done the previous two years at RSS. In FY2017, we will conduct a task-based assessment (TBA) or our capabilities and integrate them with the MMI and natural language understanding (NLU) models as part of an integrated research assessment (IRA) in FY2018.

References:
[1] I. Chung, O. Propp, M.R. Walter, and T.M. Howard. "On the Performance of Hierarchical Distributed Correspondence Graphs for Efficient Symbol Grounding of Robot Instructions", In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015.

[2] A.F. Daniele, M. Bansal, and M.R. Walter. "Navigational Instruction Generation as Inverse Reinforcement Learning with Neural Machine Translation," Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2017.

[3] H. Mei, M. Bansal, and M.R. Walter. "Listen, Attend, and Walk: Neural Mapping of Navigational Instructions to Action Sequences," In Proceedings of the National Conference on Artificial Intelligence (AAAI), 2016.

[4] H. Mei, M. Bansal, and M.R. Walter. "What to talk about and how? Selective Generation using LSTMs with Coarse-to-Fine Alignment," In Proceedings of the Annual Conference of the North American Chapter of the Association of Computational Linguistics (NAACL), 2016.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.3.6 T2C1S3C - Investigating Tactical Multi-modal Soldier-Robot Exchanges

PI: Oron-Gilad, Tal (Ben-Gurion University of the Negev), orontal@bgu.ac.il, +972-8-6461434
Additional PIs: Parmet, Yisrael (Ben-Gurion University of the Negev), iparmet@bgu.ac.il, 972-8-6461434; Ilit Oppenheim, (Ben-Gurion University of the Negev), ilit@bgu.ac.il, 972-8-6461434.

Objective and Benefits:

Digitization in the battle field enables the creation of new communication means among distributed elements. We aim to examine how the use of graphic imagery and means of graphic communication facilitate performance, shorten the decision cycle (OODA), and enhance Soldiers' situation awareness particularly by improving their coordination with robotic autonomous or semi-autonomous systems. This research aims to implicate on how means of bi-directional graphic communication should be designed and implemented for operations in rural built area environments.

The goals are to: (1) identify if and when bi-directional graphic communication is beneficial, (2) define how it should be implemented; e.g., type of info (e.g., augmentation on reality, video feed or still images), interface components (e.g., annotations and marking tools), and with regard to differences between ground and aerial autonomous system, (3) asses how soldiers' situation awareness and verbal communication patterns with the systems or their operators is affected, and (4) provide initial ground toward the development of dynamic algorithms that can facilitate graphic coordination automatically. Lastly, since visual load on soldiers, especially commanders at the squad and battalion level can be overwhelming, we aim to examine whether 'tactile annotations' can be integrated as an additional mean to direct soldiers' attention toward critical graphic information or toward feedback information (i.e., acknowledgments).

The outcomes of this study will be empirical evidence as well as design guidelines for initial implementation. Further along the way, based on the behavioral findings, we aim to develop dynamic algorithms that will facilitate the graphic communication means.

Scientific approach:
Four experiments are planned for the period of 2017-2018. Based on the findings of our previous experiments (Battlelab Experiments #1-2 see [1-2]) conducted in collaboration with the Israeli Defense Forces (IDF) we intend to continue the investigation of bi-directional graphic communication in two additional large scale team experiments in the Battlelab in 2017-2018. In parallel to the 2017 Battlelab Experiment #3, we intend to further develop the 'tactile annotations' in the BGU dome laboratory and incorporate the tactile system and tactile taxonomy in the second 2018 Battle lab experiment (Battlelab Experiment #4).

Bi-directional graphic communication. Our upcoming experiment (Battlelab Experiment #3; planned for early 2017, around March) is aimed to investigate the added value of using bi-directional graphic communication either by still images or temporary markings on dynamic scenes (augmentation of reality), with predefined symbols for teams of ground-ground or air-ground soldiers and robotic systems. Participant-teams will consist of ground or aerial Intel collectors (operators of unmanned systems) and ground attackers who rely on this Intel (dismounted soldiers), and must understand it and trust it (e.g., [3, 5-6]). The operational scenarios will take place in a built rural area and during day time. For this evaluation we have already prepared means to asses SA that are now implemented in the experimental design and measurement (e.g., [4]). These metrics can be generalized towards other studies of the RCTA. This upcoming experiment will investigate the value of bi-directional means of graphic communication on objective performance, team preferences, and situation awareness.

A Tactile alert taxonomy for commanders in teams of manned-unmanned. The operating units are most likely to operate in hostile environments where there is an immediate threat to their safety. There is a need to balance between attending to graphical information that is available on mobile or portable devices, and the immediate environment. The need for a tactile alert taxonomy [7] was already identified. In 2016 we conducted brainstorming sessions with Subjects Matters Experts (SME's) and based on their insight we determined four key topics for 'tactile annotations': Warning, Uncertainty, Awareness and Commands. Each topic includes one or more level of urgency or complexity and therefore requires a different response. We now aim to transfer this information to the commander/soldier via tactile alerts in order to improve the Soldier's awareness to the accumulation of new information and its criticality. Adding these types of information about the immediacy and urgency should affect Soldier's situation awareness, but we need to ascertain that in a positive way.

The goal of the upcoming laboratory studies is to examine how the four types of tactile information are perceived; whether as distributed knowledge or shared knowledge [4] and how they integrate with visual/graphic critical or urgent information. To achieve this, in parallel to the large scale studies conducted in the IDF battle lab, two smaller scale investigations in the BGU dome simulation laboratory will be geared toward building the tactile taxonomy and determining what type of information should be transferred via tactile alerts and annotations. Based on feedback from our collaborators in the RCTA, attempts will be made to approximate RCTA platforms or payloads (sensor capabilities), as well as scenarios.

Laboratory experiments in the BGU dome projection facility: Tactile #1 Experiment is planned for April 2017. Tactile #2 Experiment is planned for October 2017. The tactile system is currently being integrated in the IDF battle lab as well. Once we establish the taxonomy, our design will be implemented in the IDF battle lab to enable these capabilities for upcoming large scale experiment. Notably, this design can be transferred also to other RCTA laboratories.

Integration of multi modal communication. The second experiment (Battlelab Experiment #4) in the IDF battle lab will be geared toward more complex contexts (e.g., different sensors or operational scenarios) to broaden the scope of investigation. We aim that by the end of 2017 we will have enough understanding of the RCTA robotic platforms and sensors to incorporate or approximate some of them in our study. We will aim to use as many common RCTA scenarios as possible in our implementation. In addition, this will be the first attempt to combine graphic and tactile communication capabilities into one multi-modal experiment. This experiment will be held in 2018.

Joint US-IL experiment in Ft. Benning. A questionnaires regarding the use of tactile alerts [7] has been sent to the US and its administration is coordinated with the Jentsch team at UCF and the US Army's Research Laboratory's (ARL's) field element at Ft. Benning led by Elliot. The purpose is to develop a joint experimental plan in order to derive general design principles of interfaces including possible bi-directional (simulated or real) communications between UVs and Soldiers operators. The proposed experimentation plan will also involve close coordination with University of Central Florida's (UCF) research program on language understanding as part of ARL's Robotics Collaborative Technology Alliance (RCTA).

List of dependencies:

BGU will work closely with its collaborators to examine how, based on feedback from our collaborators in the RCTA, to attempt to approximate RCTA platforms or payloads (sensor capabilities) and/or operational scenarios in our experimental designs.

Collaboration:
This subtask will involve collaboration with Dr. Daniel Barber (T2C1S4B Adaptive multimodal communication and T4C3S1E HRI Driven Enhancements) as he is the primary developer/integrator for the RCTA MMI. This subtask will coordinate with Dr. Barber on developing interview questions, experimental design, sharing results and guidelines derived from our experiments. Upon request we will take part in evaluations of how the changes made to the MMI impact SA and usability.

This subtask will involve close coordination and collaboration with other RCTA researchers such as Linda Elliott (T2C2S3C Multimodal comms research) and Florian Jentsch (T2C1S4A A user- centered design (UCD) approach to creating usable naturalistic communication interfaces for Soldier- robot teaming) to establish connections with end-users (e.g., Soldiers), consult on experimental designs, facilitate data collection with end-users, and produce joint papers/publications. In addition we plan on potentially collaborating/coordinating with Dr. Jessie Chen, ARL (T2C2S2D Transparency of agent reasoning and operator trust) given her previous and current work on trust.

Additional collaboration:

- ARL RDECOM Michael Barnes and Eric Holder on the ability of humans to convey their intent to a robot/autonomous agent (UAS/UGV).
- AMRDEC Grant Taylor and Tom Alicia – Human Systems UAS coordination Task 31 **Rapid Interpretation of Temporal-Spatial UAS Data**. We are working on visualizations to enhance UAS operators' performance.
- Michael Goodrich will be the ABC robotics visiting professor in June 2017 at BGU. Our students will meet with Mike and we will have time to learn and discuss research approaches and collaboration.

Linkage to previous RCTA BPP sub-tasks:

This subtask primarily builds upon research performed in 2015-2016 under subtask H4-4, Investigating Tactical Multi-Model Soldier-Robot Exchanges. Research previously performed under subtask H5-2 in 2015 and 2016 was focused on understanding bi- and one way graphic communication means and how they impact a Soldier's ability to perform a mission. From these studies we can confirm that graphic communication means are important in operational scenarios that require coordination between autonomous or semi-autonomous Intel gathering system and Soldiers who need to acquire targets or prepare for a mission. In addition, an extensive literature review of SA, decision loops (OODA) in human-robot teams was also conducted during this time. This review identified the need for development of additional metrics of assessment for SA and collaboration when multiple robotic assets and multiple communication means are available, aside from verbal communication.

The research proposed in this subtask for 2017-18 seeks to continue to support the goals of the 2015-2016 research by studying how bi-directional communication facilitates coordination and performance and how the addition of 'tactile annotations' integrate into this type of communication. This research will aims to build upon research previous multimodal RCTA research such as research performed under H5.1 Tactical Display Design.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Expected products:

1. Empirical results will be published in three types of formats: Technical reports, peer-reviewed publications and conference proceedings.

2016 work will be presented in the HCII 2017 conference**:** Tal Oron-Gilad and Ilit Oppenheim (2017). Use of graphic imagery as a mean of communication between an operator and unmanned systems, proceedings of the HCII 2017 Annual meeting.

We intend to submit a proposal to HFES 2017.

Target publication venues include IEEE THMS, HF and CEDM, TIES, *Communication Design Quarterly (CDQ)*.

2. Ongoing integration with other subtasks, primarily with Linda Elliott with regard to the tactile taxonomy and integration of tactile annotations.
3. Work meeting with Mike Barnes and Eric Holder will probably take place in May 2017. So far, we have been sharing our data with Mike and Eric for their development use.
4. As we are attempting to provide initial ground toward the development of dynamic algorithms that can facilitate graphic coordination automatically, we would like to initiate a collaboration with one of RCTA parties that would be interested in moving forward in this direction.

References:

[1] Oron-Gilad T. and Oppenheim I. (August 2016). Scalable interfaces for operator control units; a common display to conduct MOUT operations with multiple video feed sources - Final Research Report - Task (6), DCS subcontract no: APX03-S010 (B.G. Negev Technologies and Applications Ltd) under Prime Contract no W911NF-10-D-0002.
[2] Oron-Gilad T. and Oppenheim I. (December 2015). Scalable interfaces for operator control units; Bi-directional Graphic Communication between Dismounted Soldiers and operator control units- Final Research Report - Task (5), DCS subcontract no: APX03-S010 (B.G. Negev Technologies and Applications Ltd) under Prime Contract no W911NF-10-D-0002.
[3] Oron-Gilad T. and Parmet Y. (2016). Close target reconnaissance: a field evaluation of dismounted soldiers utilizing video feed from an unmanned ground vehicle in patrol missions, Journal of Cognitive Ergonomics and Decision Making,
[4] Kitchin, J., & Baber, C. (2015). A comparison of shared and distributed situation awareness in teams through the use of agent-based modelling. Theoretical Issues in Ergonomics Science.
[5] Barnes, M., Chen, J. & Hill, S. Humans and Autonomy: Implications of shared decision making for military operations, ARL-TR, in press.
[6] Chen, J., Lakhmani, S., Stowers, K., Selkowitz, A., Wright J., & Barnes, M. Establishing a dynamic situation awareness-based agent transparency model. *Theoretical Issues in Ergonomics Science* (in press)
[7] Katzman, N., Oron-Gilad, T., & Salzer, Y. (2015, September). Tactile Interfaces for Dismounted Soldiers User-perceptions on Content, Context and Loci. In Proceedings of the Human Factors and Ergonomics Society Annual Meeting (Vol. 59, No. 1, pp. 421-425). SAGE Publications.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.3.7 T2C1S4A - A user-centered design (UCD) approach to creating usable naturalistic communication interfaces for Soldier-robot teaming

PI: Jentsch, Florian (University of Central Florida), Florian.Jentsch@ucf.edu, (407) 882-0304

Objectives and Benefits: To establish and maintain effective team dynamics when teaming with an autonomous robot, there is a need to develop and design user-centered interfaces. Given that information sharing is critical to the adequate execution of missions, interface designs should be developed considering factors such as team situational awareness (SA), communication, and trust. That is, an interface that facilitates a Soldier's understanding of the complex nature of an environment from the robot's perspective, helps the Soldier maintain SA and fosters trust. Multimodal communication interfaces used to interact with one (or more) robots have been developed (e.g., MMI) but have yet to be tested with the end-users (i.e., Soldiers) to evaluate the usability of these interfaces. As such, the objective of this research task is to utilize user-centered design (UCD) methods to derive guidelines for developing naturalistic communication interfaces that help Soldiers develop/maintain situation awareness (SA) and are usable. This subtask seeks to establish generalizable guidelines for designing human-robot interfaces that are both usable and help Soldier's maintain SA when interacting with an autonomous robot.

Scientific Approach: We propose to begin by conducting literature reviews on existing SA theories and usability design theories which will inform our experimental methods and designs (in 2017). Based upon our findings within the literature, we will conduct interviews with subject matter experts (e.g., individuals with prior military experience) and, if possible, projected end-users (e.g., Soldiers) of these robot systems, to better understand specific needs, goals, and problems relevant to building, and maintaining, situational awareness when teaming with an autonomous robot (in 2017). Based on the insights garnered from these interviews, laboratory human-in-the-loop simulation studies will be conducted with prospective end users (e.g., Soldiers) to empirically evaluate different interface designs within controlled settings (in 2017). Follow-on field experiments (in 2018) will be conducted to evaluate prototype interfaces, such as the MMI (both in terms of usability and SA development), in a realistic setting (e.g., using the Husky platform in a field setting). We plan to use the Husky platform (either simulated or real) for our experiments.

This research builds upon previous SA and multimodal interface research conducted under H1.2 (Jentsch; *Model and Improve Mission-Level Situation Awareness*), H4.3 (Elliott; *Naturalistic Interfaces -- Investigating Multi-Modal Soldier-Robot Bidirectional Exchanges*), H5.1 (Barber; *Tactical Display Design*), and H7.2 (Barber; *Testing User Interactions with High Mobility Platforms*). In addition, this research directly links with multimodal interface research being performed by Daniel Barber (T2C1S4B *Adaptive multimodal communication* and T4C3S1E *HRI Driven Enhancements*).

This research primarily will support the RCTA's Thrust II/Capability 1 "Situation Awareness in Unstructured Environments" in that our experiments will be situated within the context of a human-robot team performing a mission that requires building team SA through communication between a human and a robot via a human-robot interface. Given the relationship between maintaining team SA and communication, this research is also well suited to support other capabilities that involve team collaboration during mission planning and execution such as Thrust II/Capability 2 "Distributed Mission Execution." Furthermore, the research proposed under this subtask helps achieve the RCTA Capstone vision by providing guidelines for (a) enhancing the usability of robot communication interfaces which will enable human teammates to develop, and maintain, SA when teaming with one (or more) robots.

List of Dependencies:

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

T4C3S1E HRI Driven Enhancements
This research depends on input from Daniel Barber (UCF) since he is the developer of the MMI. We will need to work closely with him to utilize the MMI for our experiments and request new features/functionality (as needed) for our experiments. We will also provide testing and evaluation (output) of the MMI interface. Design guidelines and specifications will be derived from the results of our experiments and evaluations.

Collaboration: This subtask will involve extensive collaboration with Daniel Barber (T2C1S4B *Adaptive multimodal communication* and T4C3S1E *HRI Driven Enhancements*) as he is the primary developer/integrator for the RCTA MMI. Coordination will involve developing interview questions, experiment design, producing joint papers/publications, sharing results and guidelines derived from our experiments, and evaluating how the changes made to the MMI (based upon our provided guidelines) impact SA and usability at integrated assessments.

In addition, this subtask will involve close coordination and collaboration with other RCTA researchers such as Linda Elliott (T2C2S3C *Multimodal comms research*) and Tal-Oron Gilad (T2C1S3C *Investigating Tactical Multi-modal Soldier-Robot Exchanges*) to establish connections with end-users (e.g., Soldiers), consult on experimental designs, facilitate data collection with end-users, and produce joint papers/publications.

Linkage to Previous RCTA BPP Sub-tasks:

- H1.2 *Model and Improve Mission-Level Situation Awareness* (see references)

Expected Products:

- **End of BPP Year 2017:** Develop initial design guidelines and specifications (e.g., user requirements) to both robot developers and communication interface developers (e.g., Daniel Barber) to ensure the end-users consider a robot team member to be adequately usable (i.e., effective, efficient, satisfactory). These guidelines will be based on lab simulation experiments. The results of our lab simulation experiment will be presented at an appropriate conference and made available via a technical report. Candidate conferences include the *Annual Meeting of the Human Factors and Ergonomics Society (HFES)*, *IEEE CogSIMA*, or *HCI International (HCII)*.

- **End of BPP Year 2018:** Develop scientific and technical design guidelines and specifications (e.g., user requirements) to both robot developers and communication interface developers (e.g., Daniel Barber) to ensure Soldiers consider robot team member to be adequately usable (i.e., effective, efficient, satisfactory). These guidelines will be based on lab/field experimentation. The results of our lab/field research will be published in an appropriate journal. Candidate journals include *Journal of Human-Robot Interaction, Human Factors, or International Journal of Human-Computer Studies.*

References:

Previous H1.2 publications

1. Abich, IV, J., Talone, A. B., Barber, D. J., Phillips, E., Jentsch, F., Pettitt, R., & Elliott, L. (in press). *Soldier-robot team communications: An investigation of exogenous orienting visual display cues and robot reporting preferences* (Contract No. W911NF-10-2-0016). Army Research Laboratory.
2. Barber, D., Abich, J. A., Phillips, E., & Talone, A., Jentsch, F., & Hill S. (2015). Field assessment of multimodal communication for dismounted human-robot teams. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. Santa Monica, CA: Human Factors and Ergonomics Society.
3. Phillips, E. K. (2016). *Supporting situation awareness through robot-to-human information exchanges under conditions of visuospatial perspective taking* (Doctoral dissertation). University of Central Florida, Orlando. Retrieved from http://stars.library.ucf.edu/etd/4936
4. Talone, A. B., Phillips, E., & Jentsch F., (addressing revisions). *Towards effective robot-to-human information sharing in distributed human-robot teams.* Manuscript submitted to the Journal of Human-Robot Interaction.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.3.8 T2C1S4B - Adaptive Multimodal Communication

PI: Barber, Daniel (UCF), dbarber@ist.ucf.edu, (407)882-0304

Description

Proposed Work: Maintaining optimal situation awareness is critical as mission demands require mixed-initiative teams of Soldiers and robots to perform coordinated tasks both in- and out-of-line-of-sight. To ensure task performance remains effective, bidirectional multimodal interfaces must identify the environmental demands, cognitive burdens, and physiological state of the Soldier to adapt the mode(s) in which information is delivered. Through the flexible selection of single and redundant modes of communication it is expected that Soldiers will have better retention of information from robot teammates while minimizing effects of attention, stress, fatigue, and workload. This task seeks to understand and identify when communication modalities should adapt, how information should be presented to both the Soldier and robot teammates, and at what frequency information should be delivered to team members to facilitate highly-efficient dismounted Soldier-robot collaboration. To this end, this subtask will provide the scientific community with answers to research questions such as:

- How will sharing of information change between Soldiers and robot teammates with respect to adaptive strategy, environmental demands, and team distribution (in-and out-of-line-of-sight)?
- What HRI factors affect a Soldier's mental workload & usability preference when communicating multimodally with robot teammates during distributed missions?
- What are the best practices for how and when robot teammates should transmit reports without creating additional demands on human teammates?
- What are the modality switch costs for environmental demand as assessed by primary task performance and situation awareness?
- What are the physiological correlates associated with modality of robot reporting presentation, adaptive strategy for modality implementation, and primary task demand that can indicate a state of cognitive burden?

Thrust and Capability Supported

This research primarily will support the RCTA's Thrust II/Capability 1: Situation Awareness in Unstructured Environments.

Relation to Previous RCTA 2015-2016 Research

This subtask builds upon research performed in 2015-2016 under subtasks H5.1 *Tactical Display Design* and H7.2 *Testing User Interactions with High Mobility Platforms*. Research in the last few years focused on leveraging the results of the previous four years to inform the design requirements for the multimodal interface that was utilized for assessment which showcased the ability of interacting multimodaly with a robot through the multimodal interface (MMI) but the research up to this point primary addressed unidirectional communication exchanges.

The research proposed in this subtask for 2017-18 seeks to continually address the impact communication modalities have on Soldiers SA and task performance from a workload and usability perspective but expanding the modes of interaction (e.g. augmented reality) and close the dialogue loop between the Soldiers and robots through bidirectional interaction. Further investigation will seek to address how and when the interactive modalities should adapt dependent on environmental and task demands.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Collaborations

This subtask will involve collaboration with T4C3S1E *HRI Driven Enhancements* to support integration of research results. In addition, this subtask will involve close coordination and collaboration with other RCTA researchers:

- Barber (T4C3S1E: *Multimodal Display integration*)
- Florian Jentsch (T2C1S4A: *A user-centered design (UCD) approach to creating usable naturalistic communication interfaces)*
- Linda Elliott (T2C2S3C *Multimodal comms research*), Tal-Oron Gilad (T2C1S3C *Investigating Tactical Multi-modal Soldier-Robot Exchanges*)
- Michael Goodrich (T1C2S4C *Practical human-robot interaction for navigation and manipulation*)
- Jean Oh (T2C1S3A Two-way translation between vision and language), Matthew Walter (T2C1S3B Generating natural language scene descriptions for shared situation awareness)
- Hebert (T1C2S1D reasoning for scene understanding)

Data from each of these tasks will drive requirements for each other, and therefore require close coordination of activities. These collaborations will directly inform experimental design, capability development, facilitate data collection with end-users, and produce joint papers/publications.

Scientific Approach

The approach will first review the literature on adaptive modality within the context of HRI. Experimentation of a cooperative task with multi-robot teammates and varying environmental demands using an adaptive multimodal interface will be conducted with a focus on the usability assessment of adaptive communication modalities for content presentation and delivery. Follow-on joint experimentation with end-users will be conducted.

Product/Outcomes of this Research

- End of BPP Year 2017: Research deliverables will be in the form of journal articles, conference proceedings, and technical reports which will outline and describe recommendations based on empirical evidence for the appropriate modality of communication and frequency of interface adaption should occur during dismounted HRI teaming.
- End of BPP Year 2018: Research deliverables will be in the form of journal articles, conference proceedings, and technical reports which will outline and describe recommendations based on empirical evidence for the appropriate modality of communication and frequency of interface adaption should occur during dismounted HRI teaming.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.4 Thrust 2/Capability 2: Distributed Mission Execution

3.4.1 Overview

Meaningful use of robots, and autonomous systems in general, revolves around collaborative interactions between humans and robots to accomplish a task. We focus on two types of interactions inherent to general tasks, namely those involved in task planning and task execution. Soldier-robot collaboration can occur during both aspects and can take on different forms, depending on whether the team is co-located or distributed. For example, task execution can involve humans and robots collaborating both directly, such as when pushing an object together, and indirectly, such as when searching a divided area.

The purpose of interaction for collaborative planning is to create a detailed understanding of the task to be completed, including task goals, constraints, context, and roles. Given a suitably detailed task description, an autonomous system can be reasonably expected to plan its own actions and execute them appropriately, but it will often fail if the task allows so many acceptable plans that the system cannot quickly decide between them, or if the task is so complicated that searching for even a single satisfying plan is a challenge. Interaction provides a solution to both failures: clarifying the intended task or underlying assumptions about the world provides a means to reduce the number of solutions, and providing suggestions of what the solution should look like provides a means to guide the planning process. Training with humans can be considered a form of offline interaction where we can shortcut these vague or difficult planning problems by learning how to fill in the detailed task description from limited online interaction.

The purpose of distributed mission execution is to actually perform the required actions and to coordinate the status of actions when carrying out the plan that has been developed through collaborative planning. This coordination often involves the sharing of information between all players, both human and robot, and so design of the interaction focuses on questions such as what information must be shared, how it should be presented, and when should it be delivered.

Example of Capability in Action

A dismounted Soldier issues the command “secure that street corner” to his/her robot teammate while he/she continues on his/her patrol through the neighborhood. Because the soldier is in motion and the street corner is only of interest temporarily, the command to “secure” means here to search the street corner and surrounding area for activity that is currently interesting, according to the orders for this patrol, and to finish before the Soldier moves on. This critical mission intent and knowledge is encoded in the autonomy architecture of the robots through prior training using machine learning, giving the robots a library of behaviors to apply. The system selects the appropriate behavior, adapts the parameters to match the current environment, performs coarse mission-level planning to adhere to high-level constraints, and finally presents a proposed search area and provides a summary of the activities and targets that the robots will search for. This type of mission decomposition, and its associated visualizations and transparent view of the system capabilities, is familiar to the Soldier because he/she has trained extensively with the robotic system in many simulations. The Soldier then interacts, aids in the decomposition, and clarifies which street corner to secure. The Soldier modifies, where appropriate, the list of activities that require notification. Once the decomposition is satisfactory, the system begins execution and provides periodic updates on mission progress and on any unexpected changes in the environment.

In different circumstances, the dismounted Soldier issues the command “secure that street corner” to the robot teammate while engaged in a conversation inside a nearby building. Because this time, the Soldier is stationary and fully engaged in an activity, the command “secure” now means to continuously monitor the area for changes and notify when important events occur. This different interpretation is also encoded in the autonomy architecture through training and is activated through observations of the environment and the commander’s current activity/state. Under these circumstances, the system responds with proposed

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

observation points and confirms that the robots will inform the Soldier not only of incoming people, but also of other important changes in the environment. In this case, the Soldier does not want to be distracted and so tells the robot to "keep the area clear and only notify if important changes occur", updating the meaning of "secure" to be physically preventing access from this direction. The system uses this new information to select a set of corrected behaviors for each robot, given the updated circumstances. The robot system responds with intended patrol routes and location of observations points near the street corner and confirms that that the robot(s) will intercept people and ask them to stay back. In this instance, the system concludes multiple possible classes of changes that could be considered "important" and asks for clarification from the Soldier as to which are intended. Again, based on the presented decomposition, the Soldier compares their expected notions of the system's behavior against the system's proposition and quickly corrects any errors in behavior selection and application using a Multi-Modal Interface that combines speech, gesture, and gaze sensing. The mission decomposition is accepted and dispatched, and the system provides updates as normal.

During execution, a person intercepted by the robots ignores the request to stay back and continues walking through the area. The system notifies the Soldier that his/her request is in jeopardy and presents a plan to use a robot to physically block the entrance to the Soldier's location. He/she looks at the pictures of the interloper, recognizes them as a local community leader, and asks the system to escort them to his/her location. The system plans a route and executes the task since it needs no clarification about what was intended.

3.4.2 Relevance to the RCTA and the Army

Robots operating in unstructured and unfamiliar environments need to follow instructions from human teammates and support teamwork behaviors comparable to other humans to enable mixed-initiative team collaboration. Specific coordinating mechanisms identified for successful human-human teaming include: team leadership, mutual performance monitoring, adaptability, shared mental-models, team orientation, and mutual trust. The capability described enables these mechanisms in human-robot teaming with the robot adapting to environmental demands to deconstruct complex tasking into appropriate sub-tasks and goals, while appropriately informing soldiers of robot state, goals, sub-goals, and predicted outcomes with and without direct line-of-site while on the move to facilitate collaboration.

3.4.3 Required research areas

In order to achieve team effectiveness during mission execution, this capability focuses on three research areas: (a) collaborative mission planning and execution, (b) distributed situational awareness and communication, and (c) shared-multi-modal dialogue. Specifically, research performed under the "collaborative mission planning and execution" area aims to develop an ability to take mission-level tasks and decompose them into tactical behaviors across the human-robot team. Research performed under "distributed situational awareness and communication" has a special focus on determining how to represent unknown groundings due to partial sensor coverage and build more complete knowledge captured across team members. Finally, research performed under "shared-multi-modal dialogue" aims at moving beyond the state-of-the-art by taking advantage of language as an efficient means to establish shared situational awareness between soldiers and robots.

3.4.3.1 Collaborative Mission Planning and Execution

Though we have a high-level goal such as "go to the street corner", and we have properly grounded the street corner in the world model so that we know which region we are referring to, we still do not have a well-formulated problem statement suitable for handing to the lower levels of our robot planning and

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

control system. Goal elicitation is the process of turning the vague goal of "going to the street corner" into precise goal coordinates, cost functions, and problem constraints that are meaningful to an optimal path planner. Prior work in this area has focused on mission planning as a batch process, with human tasks given as enumerated procedures and robot tasks as single tactical behaviors. This approach has led to mission planning as a one-way process, with no way for development of an interactive plan that coordinates multiple activities across team members.

Advancing prior work in the RCTA, the subtasks supporting this capability will develop an ability to take mission-level tasks and decompose them into tactical behaviors across the human-robot team. This will be accomplished through creation of a robust intelligence architecture that is capable of independently recognizing and acting upon the stimuli and observed capabilities of human and robot teammates to support actions within unstructured environments where local networks may be slow or unavailable. We will meet this challenge through development of new computational models of structure prediction that allow a robot to learn to decompose a mission by combining its own reasoning, previous experiences, and inputs given by human teammates. Through interaction with human teammates (i.e., Soldiers), computational models will learn task descriptions and tactical behaviors, and the ability for decentralized plan inference from natural language interaction.

3.4.3.2 Distributed Situation Awareness and Communication

In order for the human-robot team to build accurate and cohesive situation awareness in unstructured environments, they must be able to share semantically-rich and persistent knowledge with mission-relevant facts seamlessly. Currently, the type of information stored internal to the robot and available to human counterparts is restricted to perceptual entities (i.e. objects, regions, paths). This repository of "world" information is typically centralized to a single robot and includes many a-priori assumptions that remain static during mission execution. Having such limited data representations, which are typically kept in isolation from human counterparts, results in discordant representations and models impacting soldier-robot teaming and trust, reducing overall op-tempo.

Building on, and interdependent with, the research needed for situation awareness in unstructured environments that is being pursued under Thrust 2, Capability 1 (above), the focus for this research area is to determine how to represent unknown groundings due to partial sensor coverage and build more complete knowledge captured across team members, what information in regards to the environment and mission must be shared, and how these situational and human factors influence human robot trust. The approach to accomplish this is to develop novel, unified computational models using probabilistic inference and experimental solider studies. New computational models will advance the state of the art allowing: (1) joint reasoning about imperative commands and world knowledge, (2) shared partially grounded symbolic representations, (3) learning of higher-order agent-environmental interactions and mission-relevant facts, (4) formulation of information gathering approaches to active acquisition of novel concepts through ontological reasoning and dialogue.

3.4.3.3 Shared-Multi-modal Dialogue

In pursuit of shared situational awareness within human-robot teams, mechanisms are needed to facilitate delivery and synthesis of bi-directional communication within the continuum of military operations. Previous research has been primarily focusing on one-way communication from the human and environment "into" the robot, ranging from communicating a commander's orders, to perception of sensor data. However, successful teams must have a means to not only coordinate their activities bi-directionally, but adapt to changing environmental demands. This task moves beyond the state-of-the-art by taking advantage of language as an efficient means to establish shared situational awareness between soldiers and

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

robots, allowing robots to explain the rationale behind their decision making and planned actions, and facilitating soldier-robot interactions that are robust to the effects of rapidly changing environments.
To accomplish these goals, this task incorporates multiple interdisciplinary approaches including: (1) development of new graphical model representations and algorithms for decentralized model inference of natural language instruction, (2) new computational models using recurrent neural networks (RNN) based approaches to synthesize natural language descriptions, (3) analysis of tactical teaming influenced by comprehension in addition to decision-making, problem solving, and social cognition, (4) human-in-the-loop experimentation to provide interfaces and guidelines for multi-modal information delivery and human-robot dialogs that support team comprehension and tactical teaming that is most relevant and adaptable to environmental demands on the team. Research performed under this subtask will (a) expand upon the limited understanding of performance costs and benefits of redundant versus single-modality presentation of information for an interrupting and ongoing task within complex scenarios, (b) provide mechanisms for how and when to adapt communication modalities based on environmental demands, cognitive burdens, and physiological state, (c) minimize working memory effects, and (d) inform multimodal interface design using natural language.

3.4.4 Demonstration plan

Demonstration of this capability will be performed within a reconnaissance and surveillance task incorporating at least one human (Soldier) and one robot. To fully assess the efforts of the combined subtasks, at variations of the “Robot, let’s go together and secure the street corner” command will be employed which captures the essence of the vignette previously described. Prior to fully integrated tests are conducted, individual unit tests demonstrating component functionality will be performed. Example tests include, the ability of the robot to decompose a mission command via natural language into constituent components with appropriate representations presented to a Soldier, and integration of natural language inputs within intelligence and world model architectures, and variations in information delivery based on current mission execution status and environment.
Upon completion of unit tests, an on-the-move reconnaissance and surveillance mission will be performed between at least one Soldier and one robot at different locations in a test facility and with different underlying goals. One test will include collaborative reconnaissance of an area to identify people, and will be completed upon successful coverage of the agreed upon area. A second test will extend the prior mission through the addition of requesting the robot to secure specific target areas until further input is given. The final demonstrations will include combinations of the previous two that include changes in environment which will require the robot to re-plan and adjust how it coordinates information with the human counterpart. In all cases it is expected that bi-directional multi-modal communication will be used (speech and gestures) to coordinate activities and share information about the environment between teammates as they move from being co-located to distributed within the environment. Additional surveys and performance measures will be used to assess trust, workload, fatigue, comprehension, and situation awareness of the Soldier during and after completion of a mission to verify effectiveness of multi-modal interfaces.
We envision three milestones for integrating the technical elements for capability 2:

- Demonstrating Collaborative Mission Planning and Execution Using Fixed World Model (CY17): This milestone will demonstrate the ability to take a fixed world model, and establish a plan across a human-robot team. This first milestone will crucially test the cognitive architecture in terms of symbolic reasoning from world model to an actionable plan.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

- Demonstrating Distributed Situational Awareness and Communication in Long-duration Mission (C18): This second milestone adds the ability to gather information during plan execution and update the world model.
- Shared Multi-modal Dialogue, Updating Plan (CY19): The final milestone demonstrates the ability to share information, query team-mates and allow the cognitive architecture to react by updating the plan as necessary.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5 Thrust 2/Capability 2 Subtasks

3.5.1 T2C2S1A - Interactive Multi-Agent Planning

PI: Lebiere, Christian (CMU Psychology Department), cl@cmu.edu, (412) 268-6028

Objective and Benefits:

This subtask focuses on supporting the Human-Robot Execution of Complex Missions Thrust (T2), specifically the Distributed Mission Execution capability (T2C2), most directly the Collaborative Mission Planning and Execution (S1) but also the Shared Multi-Modal Dialogue (S3), and also the SA in Unstructured Environments capability (T2C1), most specifically informing Design of System for Effective Teaming (S4). The core capability provided is a cognitive mission planner that supports (a) the context-sensitive decomposition of complex missions into tactical behaviors, (b) the interaction between the intelligence architecture and the human teammate(s) to jointly refine the mission plan, and (c) the monitoring of the plan execution and subsequent adaptive replanning to reflect a changing environment.

Scientific approach:

The scientific challenges for each aspect of the capability are numerous but related. As the conceptual scenario for the Distributed Mission Execution capability makes clear, the mission planning capability that performs the decomposition of the overall mission goal into tactical behaviors must be sensitive to the context of the instruction given by the human teammate, related both to the events in the overall environment and to the specific goal of the teammate [Oh2015;Oh2016]. Managing the interaction between the intelligence architecture and the teammate for purposes of jointly refining the mission plan requires being aware of the teammate constraints to minimize the workload imposed on the teammate by the interaction and maximize the effectiveness of his inputs. Monitoring the plan execution, reporting any significant issues, and working jointly on a modification of the plan requires an awareness of what changes in the environment are relevant, when an issue is urgent enough to be brought to the immediate attention of the human teammate, and which aspects of the plan need to be revised or can be preserved. The common aspect of those challenges is that each requires the intelligence architecture to maintain an awareness of the overall context of the human teammate: his task goal and workload, the intent of his requests, the extent of his preferences; in order to use that shared mental model to maximize the effectiveness of the interaction.

Our approach to solving these challenges taps the ACT-R cognitive architecture's ability to capture the mental processes of a human teammate in order to work together most effectively. The mission planning capability relies on a combination of instructions to structure the mission decomposition and examples to instantiate specific tactical behavior decisions [Lebiere2013]. Decomposition of mission-level objectives into tactical behaviors is accomplished using a general model of decision-making to apply the mental model of planning represented in the form of instructions, together with cognitive mechanisms such as instance-based learning to make individual decisions from decision examples provided by experts. Data input to that learning process can come from a variety of sources: ideally from human experts "teaching" the robot, but also from heuristic rules or algorithms. We will combine this approach with a model of robot competence and teammates' decision making to develop a plan that best reflects the abilities of the robot and optimizes its actions to the expected teammates behavior. The planning process will dynamically reflect expectations of the robot and teammates actions as well as their interactions during plan execution. This planning process will maximize the effectiveness of limited communication between robot and human teammates in sharing plans, resolving roles and jointly executing them. The mission planner works at a higher level that does not require precise state estimation, mapping or physics, but expects the detailed status of all human and robot entities and relevant environment attributes (buildings, major objects, etc.) to be accurately represented and promptly updated in the World Model. The mission planner can integrate readily new

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

symbols (objects, entities) that are added to the world model, although it expects a reasonable implementation of object permanence properties.

Similarly, the model will oversee communications between the adaptive planning process and the human teammate through the world model and the multi-modal interface. Using the same instance-based learning methodology, the cognitive oversight module will gradually learn from the human user input a model of the user planning preferences. When the cognitive model has developed a sufficient confidence in the human response to a planning query, it will supply the answer directly, gradually reducing the human involvement with experience. This capability can be trained for individual users for increased personalization, or generalized across users for more efficient learning. The same approach will be applied to the monitoring capability, where the model will gradually learn the relevance of various events to the human teammate, which conditions to report and which to handle on its own, and which require a replanning process.

The cognitive capabilities developed in this subtask are relatively high-level and therefore platform-agnostic. Additional capabilities, e.g. manipulation for the Roman platform, can be added to the model as needed. Metrics for the assessment of the capability include (1) accuracy measured as the probability that a planning decision is the same as the human user; (2) efficiency measured as the number of training instances required to achieve a given accuracy; and (3) autonomy measured as the probability that the cognitive module can correctly handle a situation independently instead of relying on human input.

List of dependencies: This work will require collaboration with other researchers in the Distributed Mission Execution capability, including Nick Roy (T2C2S1C), Tom Howard (T2C2S1D), Florian Jentsch (T2C2S3B), and in the SA in Unstructured Environments capability, including Douglas Summers-Stay (T2C1S1D), Florian Jentsch (T2C1S4A) and Daniel Barber (T2C1S4B). Integration will require collaboration with the RCTA SW Infrastructure and Integration capability of the Integrated Research Thrust, including Stanley Friesen (T4C3S1A), Ed Weller (T4C3S1C) and Daniel Barber (T4C3S1E).

Collaboration: In addition to the subtasks above, we also plan to collaborate closely with MaryAnne Fields (T3C1S1C) to develop and apply ACT-R techniques to other aspects of the architecture.

Linkage to previous RCTA BPP sub-tasks: I2-3 and H2-1
Our work in this subtask will leverage, integrate and extend work performed in a number of previous RCTA tasks in both the Intelligence and HRI thrusts, including: the mission planning capability developed under Framework for Collaborative Intelligence, the world model cognitive processes developed under task I2-3 (World Model for Situation Awareness), the shared mental models developed under task H2-1 (Instantiating Human-Compatible Computational Models of SMMs in Soldier-Robot Teams), and the structural matching cognitive mechanism developed under task I4.5 (Exploiting Priming Effects in Autonomous Cognitive Systems).

Expected products:
This research will result in the following products:

2017	• A cognitive mission planning software module to decompose user-supplied commands into a mission plan composed of tactical behaviors • A cognitive interaction software module that manages the interaction between architecture and human to reduce user workload and improve autonomy

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

	• Assessment of the efficiency, accuracy, and autonomy of the cognitive capabilities through integration and evaluation in the Fixed WM Scenario. • Conference paper(s) (e.g., HFES, CogSIMA) presenting the research results.
2018	• An execution monitoring software module to detect changes in the environment impacting plan execution and manage replanning with human • Refinements of mission planning and cognitive interaction modules to address issues of long-duration mission such as maintaining prolonged awareness • Assessment of the efficiency, accuracy, and autonomy of cognitive capabilities through integration and evaluation in Long-Duration Mission Scenario. • Conference paper(s) (e.g., ICCM, CogSci) presenting the research results.

References:

[Lebiere2013] C. Lebiere, F. Jentsch, and S. Ososky. Cognitive models of decision making processes for Human-Robot Interaction. *Proceedings of the HCI International Conference (HCII)*, 2013.

[Oh2015] J. Oh, A. Suppe, F. Duvallet, A. Boularias, J. Vinokurov, L. Navarro-Serment, O. Romero, R. Dean, C. Lebiere, M. Hebert, and A. Stentz. Toward mobile robots reasoning like humans. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.

[Oh2016] J. Oh, M. Zhu, S. Park, T.M. Howard, M.R. Walter, D. Barber, O. Romero, A. Suppe, L. Navarro-Serment, F. Duvallet, A. Boularias, J. Vinokurov, T. Keegan, R. Dean, C. Lennon, B. Bodt, M. Childers, J. Shi, K. Daniilidis, N. Roy, C. Lebiere, M. Hebert, and A. Stentz. Integrated intelligence for human-robot teams. In *Proc. International Symposium on Experimental Robotics (ISER)*, 2016.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.2 T2C2S1C – Learning Symbolic Planning Models

PI: Roy, Nicholas (MIT), nickroy@csail.mit.edu, +1-617-875-2506

Objective and Benefits:

Symbolic (or semantic) planning models allow complex plans to be computed efficiently by abstracting away parts of the underlying physical system, such as abstracting away some of the degrees of freedom. A world model that explicitly represents the connection between symbolic representations or behaviours and low-level continuous perception or action is a critical component of the Intelligence Architecture, However, the existing world model currently may not capture every property of the world in every situation. For example when a human team-mate gives commands to the robot, the inferred sequence of Tactical Behaviour Specifications (TBS) [2, 3] may be uncertain because the model does not include an explicit model of how the world works or the context for the command. Similarly, the effects of action templates may not be an incomplete abstraction of the true dynamics of the world. Our goal is to learn abstract models from data that ensure a complete description of the underlying physics of the world and enable computationally efficient planning of complex missions and tasks that can be used by ACT-R.

Scientific approach:

Existing approaches to creating symbolic planning models have several problems that have limited their applicability. Firstly, almost all approaches rely on hand-coded or otherwise brittle abstractions; these approaches fail when the underlying physical system cannot be easily abstracted by hand, or when the appropriate abstraction must come from an environmental model that is being constructed online by perception. Consider the example of a robot carrying out an assembly task. An appropriate symbolic abstraction would represent a grasp or an object contact relation as an abstract state in the system, but without a prior model of the objects and environment geometry, this abstraction cannot be manually encoded, and exactly where each grasp should occur and where each object should be placed during an abstract assembly action to ensure overall plan execution is not at all obvious.

Addressing this problem requires us to learn symbolic state and action representations, and learn how the symbolic actions can be composed or chained into an overall plan that is physically feasible, such as learning action preconditions. One type of model could be a signal temporal logic (STL) formula which can represent both the spatial and temporal properties of (continuous) trajectories. One of the benefits of inferring temporal logic formulas from a set of trajectories is that it does not require prior information such as relevant features or predicates. Then, we aim to find a good STL formula by searching over the structure space (determining the parametric predicates and which boolean and temporal operators will connect them) and the parameter space (identifying the parameters of the predicates). Overall, if STL formulas are learned for various actions, then we can create a large STL formula (possibly combined with conjunctions) and use it as an objective to find the trajectory of a more complex behavior via a chain of actions.

Additionally, we require an optimization procedure that can create the low-level motion plan from the high-level symbols. We will use our recent work in task-and-motion planning to develop a planner that generates optimal motion plans given a sequence of abstract states and actions (which can also be specified as STL formulae). We have shown this approach provides asymptotic guarantees on optimality [1], and we aim to further develop this approach to use learned models of symbolic action compositionally to improve the overall computational efficiency.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

This proposed world and accompanying symbolic planning actions are crucial components of the envisioned cognitive architecture of the RCTA. Without the ability to acquire a symbolic world model and connect it to the low-level geometric model, complex mission execution is difficult. We expect to continue collaborations with Tom Howard (T2C2S1D) and Daniel Barber (T2C1S4B), and continue developing a new collaboration with Jean Oh (T2C2S2A).

List of dependencies:
Requires input from T1C1S1A (Hebert, CMU), and T1C1S3A (Muelling, CMU), and provides output to T2C2S1A (Lebiere, CMU), to planning T1C1S4B (Oh, CMU), T1C1S4D (Navarro-Serment, CMU), T2C2S1G (Pappas, Penn) and to language T2C1S3A (Oh, CMU), T2C1S3B (Walter, TTI), T2C1S4B (Barber, UCF), T2C2S2A (Oh, CMU)

Collaboration:
This subtask will support a collaboration with Ethan Stump at ARL entitled "Goal Elicitation to Formulate Problems from Human Guidance".

Linkage to previous RCTA BPP sub-tasks:
This research builds on our previous work in I5-3 and H1-4.

Expected products:
1. Submission to Robotics Science and Systems and ICRA.
2. In FY2017 we will conduct a task-based assessment of our model learning to demonstrate the performance of our system in specific tasks.
3. Integration with perception and world model systems to ensure that our symbolic world models and planners can interact with the TBS behaviours and the higher-level ACT-R processes at CMU.
4. Integration with the MMI and language processes.

References:
[1] W. Vega-Brown and N. Roy. "Asymptotically optimal planning under piecewise-analytic constraints". Proceedings of the Workshop on the Algorithmic Foundations of Robotics (WAFR) 2016. To appear.

[2] D. Barber, T.M. Howard, and M.R. Walter. "A multimodal interface for real-time soldier-robot teaming". In: SPIE Defense+ Security. International Society for Optics and Photonics, Apr. 2016.

[3] J. Oh, T.M. Howard, M. Walter, D. Barber, M. Zhu, S. Park, A. Suppe, L. Navarro-Serment, F. Duvallet, A. Boularias, O. Romero, J. Vinokrov, T. Keegan, R. Dean, C. Lennon, B. Bodt, M. Childers, J. Shi, K. Daniilidis, N. Roy, C. Lebiere, M. Hebert, and A. Stentz. "Integrated Intelligence for Human-Robot Teams". In: Proceedings of the 2016 International Symposium on Experimental Robotics. Oct. 2016.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.3 T2C2S1D – Adaptive Models for Human-Robot Interaction in Diverse Decentralized Human-Robot Teams

PI: Howard, Thomas (University of Rochester), thomas.howard@rochester.edu, (585) 275-3755
Additional PIs: None

Objective and benefits:

In recent years, significant advances in probabilistic models for natural language understanding for human-robot dialog has led to real-time symbol grounding in non-trivial symbolic representations that are necessary to express intent to robot intelligence architectures. As we move towards multi-robot teams composed of different mobility systems and capabilities performing complex missions developing models of bi-directional human-robot interaction that are both efficient and effective becomes increasingly difficult. Contemporary models for human-robot interaction are based on probabilistic inference in graphical models formed using structure from linguistic inputs and environmental context [1]. Recent variants of these models, such as the Hierarchical Distributed Correspondence Graph [2], Adaptive Distributed Correspondence Graph [3], and Verifiable Distributed Correspondence Graph [4] refine the symbolic representation to make accurate probabilistic inference efficient and ground language in Linear-Temporal Logic for verifiable human-robot interaction. As exhibited in experiments on natural language symbol grounding of robot instructions described by the Tactical Behavior Specification [5,6], such models are effective for real-time human-robot interaction for interaction between a single human-robot pair. More recent efforts in the RCTA explored approaches that considered monologic inputs or between multiple robots with a single human operator. As we transition towards human-robot teams composed of multiple robots and humans, the complexity of the symbolic representation increases, which will impact the runtime performance of such algorithms. In this effort, we propose to exploit characteristics and dynamics of objects in the world model (humans, robots, etc.) to proactively adapt the distribution of symbols that are likely to be encountered during mobile robot navigation to increase the efficiency of natural language symbol grounding.

Scientific approach:

This addresses scientific challenges of how to accurately infer the meaning of multi-modal interactions between humans and robots working cooperatively to perform complex tasks and how to infer reduced-order representation of complex probabilistic graphs. This research directly addresses technical gaps in Thrust 2 (Human-Robot Execution of Complex Missions), specifically the ability for natural language models to address a wider variety of interactions and interpret their meaning in a variety of representations and to express instructions to multiple platforms to decentralized mission execution. Specifically, we seek to proactively adapt models based on observed behaviors in teams composed of multiple humans and robots to generate efficient symbolic representations that include or exclude particular behaviors in the context of their observations. Part of this effort will involve generalization of models across multiple platforms (Husky, RoMAN, etc.) by incorporating features and symbols for communicating intent to motion planning algorithms that understand how to satisfy constraints and cost functions with their own degrees of freedom and through symbols used to represent activities in the mission planner. We will experimentally validate this through simulations and physical experiments with multiple mobile robot platforms.

List of dependencies: As part of this effort, we will collaborate with those defining both the specification for atomic commands for navigation and manipulation (Subtasks T3C1S1B and T3C1S3C) in addition to those for mission planning (Subtask T2C2S1A).

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Collaboration: This research effort will involve interactions with researchers at MIT, CMU, TTIC, and UCF on the topic of multi-modal interactions in human-robot teams. T2C2S1A, T2C2S1C, T2C1S3B, T2C1S4B

Linkage to previous RCTA BPP sub-tasks:
Technologies that his approach builds upon include Model Learning for Decentralized Human-Robot Teams (2016 I5-4) and A Software Architecture for Human-Robot Communication (2016 IR5-6).

Expected products:
We intend to publish our research on novel decentralized models for human-robot teams at international robotics conference (e.g. ICRA, IROS) and participate in workshops at those conferences and discuss our progress at lectures at universities and national laboratories. In FY2017 we will conduct a task-based assessment of our developed models to evaluate the interfaces required and performance of the models in relevant scenarios. In FY2018 we will integrate these models into the NLU module in the RCTA software architecture. At the University of Rochester, we will utilize our research platforms (e.g. Clearpath Robotics Husky A200 UGVs) in experiments to evaluate our models against suitable baselines from contemporary research. Transition opportunities for this research include integration into H2SL (https://github.com/tmhoward/h2sl), an open-source toolkit for natural language symbol grounding developed by PI Howard and others researchers in the RCTA.

References:

[1] T.M. Howard, S. Tellex, and N. Roy, "A Natural Language Planner Interface for Mobile Manipulators," in Proceedings of the 2014 International Conference on Robotics and Automation, 2014.

[2] I. Chung, O. Propp, M.R. Walter, and T.M. Howard. "On the Performance of Hierarchical Distributed Correspondence Graphs for Efficient Symbol Grounding of Robot Instructions", In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015.

[3] R. Paul, J. Arkin, N. Roy, and T.M. Howard. “Efficient Grounding of Abstract Spatial Concepts for Natural Language Interaction with Robot Manipulators”. In: Proceedings of the 2016 Robotics: Science and Systems Conference. June 2016.

[4] A. Boteanu, T.M. Howard, J. Arkin, and H. Kress-Gazit. “A Model for Verifiable Grounding and Execution of Complex Language Instructions”. In: Proceedings of the 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems. Oct. 2016.

[5] D. Barber, T.M. Howard, and M.R. Walter. “A multimodal interface for real-time soldier-robot teaming”. In: SPIE Defense+ Security. International Society for Optics and Photonics, Apr. 2016.

[6] J. Oh, T.M. Howard, M. Walter, D. Barber, M. Zhu, S. Park, A. Suppe, L. Navarro-Serment, F. Duvallet, A. Boularias, O. Romero, J. Vinokrov, T. Keegan, R. Dean, C. Lennon, B. Bodt, M. Childers, J. Shi, K. Daniilidis, N. Roy, C. Lebiere, M. Hebert, and A. Stentz. “Integrated Intelligence for Human-Robot Teams”. In: Proceedings of the 2016 International Symposium on Experimental Robotics. Oct. 2016.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.4 T2C2S1E - Mixed Optimization Approaches to Route Planning in Humans and Intelligent Agents

PI: Perelman, Brandon S. (ARL), Brandon.s.perelman.ctr@mail.mil, 410-278-5968

Objective and Benefits:

Route planning for spatial navigation is an important behavior for humans and robots moving through dynamic and complex environments. Efficient, transparent algorithms exist to produce this behavior in robots, however, the mechanisms underlying route planning in humans are comparatively poorly understood. Disparities in decision-making processes between humans and robots may produce different emergent behaviors, and the effect of these differences in aspects of human-robot teaming, such as the human's ability to predict the robot's path and trust in the robot's route planning ability, is not yet fully understood. The goals of this research are to investigate parameters of route planning that are important to humans, identify differences in these parameters between humans and robots, and model the interplay between these effects on trust.

Scientific approach:

The first component of this research is to identify characteristics upon which human and algorithm-generated paths differ. These characteristics include measures of complexity, overlap, distance-discounted reward, and overall length. We expect humans to produce simpler routes that produce higher distance-discounted reward that minimize overlap, while algorithm-generated routes will prioritize overall path length. These characteristics are intended to inform the development of future robot path planning algorithms as well as system transparency manipulations for future experiments.

The second component of this research concerns path congruency and trust. When the paths produced by the robot's algorithm differs from the path a human would take in the same situation, we expect degraded trust in that system. Furthermore, we can expect that humans with strong trust in robots will be more likely to replan their paths.

To investigate these hypotheses, we will leverage cost function mapping techniques in order to quantify parameters of paths known to be important in human route planning, and novel path mapping algorithms such as the Algorithm for finding the Least Cost Areal Mapping between Paths (ALCAMP) to measure the similarity of human and algorithm-generated paths. Participants will be presented with 2D maps of road networks containing a starting and ending location, and a number of reward locations. Participants will plan routes through these networks with and without the help of two intelligent agents. Parameter values for each path, as well as similarity between the human-generated paths and those recommended by the intelligent agents, will be compared to demographic data and participant responses on a number of different scales measuring trust, willingness to accept cognitive uncertainty, and subjective ratings of route quality.

This research project is conducted in support of Thrust 1 (Human-assisted path planning) in parallel with related work by Dr. Michael Goodrich, BYU (T1C2S4C - Practical human-robot interaction for navigation). The results of this research hold implications for future system transparency manipulations in collaboration with Dr. Jessie Chen, ARL (T2C2S2D – Transparency of agent reasoning and operator trust). The project supports the 2019 Capstone demonstrations by informing algorithm development for more predictable and trustworthy route generation in robots.

Collaboration: This research is conducted in collaboration with coinvestigators Dr. Kristin E. Schaefer, Dr. Arthur W. Evans, and Dr. Susan G. Hill, of ARL. Future development of system transparency manipulations will be conducted in collaboration with Dr. Jessie Chen of ARL (Orlando). Finally, we expect to continue fostering a collaborative relationship with Dr. Michael Goodrich of Brigham Young University. T2C2S2D, T1C2S4C

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Linkage to previous RCTA BPP sub-tasks:
This subtask is new research.

Expected products:

1. Evaluating Path Planning in Human-Robot Teams: Quantifying Path Agreement and Mental Model Congruency – CogSIMA 2017
2. Route Planning in Humans and Intelligent Agents: Differences and Implications for Trust in Autonomous and Mixed Initiative Systems – Proposed submission to IEEE Transactions on Robotics.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.5 T2C2S1G - Adaptive Provable Missions over Unknown Environments

PI: Pappas, George J. (UPenn), pappasg@seas.upenn.edu, (215)898-9780

Overall objective/goal:

The goal of this subtask is to support this capability by developing algorithms and tools that receive mission-level tasks in a suitable formalism and decompose them into actionable behaviors across Soldier-robot teams. The behaviors of the Soldiers as well as the robots will be abstracted into behavioral primitives (such as "go to the corner") and then dynamically composed in order to achieve more complex missions that are formally expressed (or translated through linguistic descriptions). Prior work in this area focused mostly on off-line compositions of tactical behaviors over static environments. The overall objective of this subtask is to push the state-of-the-art by considering online behavioral compositions in dynamically changing environments. This task will be performed in the context of the reconnaissance and surveillance task incorporating at least one human (Soldier) and one robot.

Research Description:

The main challenge we want to address in this subtask is to perform such complex tasks but in unknown environments, viewed through imperfect perception models. Prior work in this area focused mostly on offline computations that are computationally very challenging as they plan for all environments as well as all possible compositions. Furthermore, they assume a static word model with some situational uncertainty (existence of certain objects in the scene) as well as perfect perception. In this subtask, we will address the following two research questions: 1) <u>Adaptive Primitive Compositions for Complex Missions:</u> We assume that different agents (robot, Soldier) are abstracted by a set of behavioral (mobility) primitives enabling us to have trusted execution of control algorithms. Given a formal mission description (as a temporal logic formula or an optimization cost) we formulate suitable optimization/feasibility problems that search through dynamic compositions of the primitives as well as human primitive in order to meet the complex mission. While such problems are computationally very hard, we will leverage our prior work that relies on satisfiability modulo theories (SMT) solvers to address the system of constraints that mix Boolean variables with real variables [1]. One of the major challenges we will address is to develop reactive plans that dynamically compose primitives in order to quickly respond to changes in the environment, mission, or teaming. 2) <u>Complex Missions in Unknown Environments:</u> The second challenge that we will pursue in this talk is how to perform higher level missions in partially known, or unknown but learned environments. Virtually all previous work in this area assumes a known environment [2,3]. In this subtask, we will pursue the challenge of balancing learning partially known (or unknown environments) and mission planning while providing rigorous guarantees. This approach considers for the first time robot motion planning under temporal logic constraints in probabilistic maps obtained by semantic (Simultaneous?) localization and mapping (SLAM). Preliminary results of the proposed approached can be found in [3]. Going forward, our approach for temporal logic mission planning in an unknown environment needs to be become reactive to dynamic changes in the following competing objectives during online execution: (i) explore the unknown environment to discover landmarks relevant to the mission and construct a map, (ii) follow sensing trajectories in the known portion of the map to reduce uncertainty about landmarks and robot states, and (iii) decide if and how the mission may be satisfied given the current information on the map.

How the proposed research supports the thrust/capability:

One of the major challenges in Thrust 2 is moving beyond point-to-point mobility tasks in Soldier-robot teaming. Collaborative mission planning and execution should be at a higher cognitive level which requires

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

situational awareness, semantic understanding of a environment as well as linguistic mission descriptions. This effort leverages such capabilities as well as mobility primitives of RCTA platforms in order to plan in the space of mobility primitives in order to meet with mission while ensuring a level of confidence in both planning and execution.

Scientific challenges:
There are two fundamental challenges that will be addressed in this subtask. Virtually all research in the area [1,2,3] of complex mission planning (expressed using temporal logics) suffers from two limitations. First, the plan is developed off-line and this makes most reactive approaches not scalable to larger numbers of teams as they need to develop plans for all possible contingencies. Second, all work in temporal logic mission planning assumes the existence of a known environmental map. We will push the state of the art in both directions as we will consider adaptive, online compositions of motion primitives for robots that operate in unknown environments.

Connections to the past work under R-CTA:
There are also clear connections to the work of Kostas Daniilidis (Penn) on semantic mapping and localization, as well as the work on grounding object semantics for natural language instructions for robots by Nick Roy (MIT).

Collaborations:
Our work directly connects to the work of Nick Roy (T2C2S1C, T2C2S2B) on Learning Symbolic Models as well as Grounding Natural language to precise semantic meanings for our complex missions. Learned symbolic representations can be used as our behavioral primitives in our approach. There are also direct connections to the subtask T3C1S1B by Max Likhachev whose algorithms for planning with uncertainty can be used as primitives in our larger construction of complex mission plans. There are also direct connections to the scene understanding subtasks (T2C1S2A) by J. Shi who extract symbolic scene models as well as with Kostas Daniilidis's work at UPenn on generating perception hypotheses about the object (subtasks P-5 and P-6).

Stretch goals:
If additional funding where available, we could consider this problem in a multi-robot multi-Solider as well as heterogeneous robot setting, where some robots may be better equipped for exploration and mapping, while other robots may be able to navigate and manipulate the environment to complete the mission. Another stretch goal is to connect complex mission planning with natural language fragments. An additional consideration to explore is how brittle our mission planning is to security attacks in sensor or communications.

Deliverable products:
The deliverables include papers at conferences such as RSS, ICRA, CDC as well as premier journals like IEEE Transactions on Robotics. The algorithms and tools will be developed as software tools for complex mission planning in known/unknown environments. Algorithms and the tool will be ported onto the Roman platform in 2018.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Year	Qtr	Subtask Quarterly Milestones and Deliverables
2017	**Q1,Q2**	Formalize the complex mission planning problem in a known environment using primitives from the Roman/RCTA platforms. Develop non-reactive algorithm for adaptive composition for on-line complex mission execution.
2017	**Q3,4**	Develop reactive algorithm for adaptive composition which explicitly models the environment as well as other agents in the mission as dynamic mission context. Develop algorithms for mission planning in unknown semantic environments.
2018	**Q1,2**	Integrate algorithms and tools for experimental testing of reactive planning algorithms. Incorporate semantic mapping algorithms in temporal logic mission planning.
2018	**Q3,4**	Perform experimental evaluation of algorithms on the Roman/RCTA platform

References:

[1] **Automated Composition of Motion Primitives for Multi-Robot Systems from Safe LTL Specications**. Indranil Saha, Rattanachai Ramaithitima, Vijay Kumar, George J. Pappas, Sanjit A. Seshia. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Chicago, IL, September 2014.
[2] **Temporal logic planning for mobile robots**, Georgios Fainekos, Hadas Kress Gazit, and George J. Pappas In Proceedings of IEEE Conference on Robotics and Automation, Barcelona, Spain 2005.
[3] **Temporal-logic-based reactive mission and motion planning**. Hadas Kress Gazit, Georgios Fainekos, and George J. Pappas IEEE Transactions on Robotics, 25(6):1370-1381, December 2009
[4] **Optimal temporal logic planning in probabilistic semantic maps**. Jie Fu, Nikolay Atanasov, Ufuk Topcu, and George J. Pappas. In Proceedings of the 2016 IEEE Conference on Robotics and Automation, Stockholm, Sweden, May 2016.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.6 T2C2S1J - Distributed Situation Awareness and Communication through Collaborative Mission Planning and Execution

PI: Kumar, Vijay (UPenn), kumar@seas.upenn.edu, (215) 898-3630
Additional PI: Hsieh, M. Ani, mhsieh1@drexel.edu, (215) 895-5870

Objective and Benefits:

The objective of this work is to develop collaborative strategies that exploit synergies within human-robot teams. Specifically, the work focuses on strategies that will enable human-robot teams to co-adapt and learn to improve their ability to decompose and execute given tasks over time. The proposed work will center on developing and validating: 1) appropriate robot- and human-centric world representations with varying spatio-temporal granularities, and 2) collaborative task execution strategies for human-robot teams. Toward this end, we will consider the collaborative exploration of office-like-environments by human-robot teams searching for an unknown number of target objects as a representative task. Given a generic office-like environment, humans are better equipped to reason about semantic and topological relationships between spaces, e.g., offices are often located along hallways or lobbies are often located at the ends of hallways. On the other hand, robots are better equipped to reason about metric properties about spaces, e.g., the hallway is 10 meters long or the room is 3 meters by 5 meters. Successful human-robot teaming requires suitable decomposition of the task(s) at hand that allows humans and robots to leverage each other's strengths and capabilities. Towards this end, the first aim of this work is to address the scientific and technological challenges needed to develop appropriate world model representations for both robots and humans with the ability to map information between one another. Additionally, robots must have the ability to interpret and leverage human inputs, identify similarities between subtasks, and track each team member's performance of subtasks over time to improve the team's collaborative decision. The goal is to develop tools that enable an online dialogue between the human and robot members in the team as they perform their tasks rather than simply having humans command the robotic members within a team.

How the proposed research supports the thrust/capability:

This research primarily will contribute distributed situation awareness capabilities and collaborative mission planning and execution capabilities in Thrust 2 situated within the context of human-robot teams performing exploration and mapping missions. Specifically, this task focuses on strategies that rely on collaborative reasoning of the environment, adaptation, and task allocation within human-robot teams. Thus, this research is well suited to support other capabilities that involve distributed sensing, mapping, reasoning, and task allocation. The proposed work could also be adapted to support capabilities focused on execution of more tightly coupled and complex tasks by human-robot teams (e.g., human(s) and robot(s) working together to clear the workspace of debris).

Scientific approach:

This research will leverage the PIs' existing work in multi-robot information theoretic exploration and mapping to provide initial world representation and exploration strategies for robots [1,2]. By incorporating sensor information into their world models, robots can compute the information landscape for the current workspace and help humans identify candidate locations for further exploration that maximizes information gain. Humans can assist robots to discriminate between candidate locations that may yield quantitatively different information content, but bias the search towards locations that yield "actionable" information. For example, consider the environment shown in Fig. (a). Fig. (b) shows the corresponding information reward surface where we note various locations exhibiting high information content. A human is capable of extrapolating the surrounding

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

known geometry and reason about the topology of the environment to determine that the region indicated by the white rectangle in Fig. XX (b) is likely to yield better actionable information than the region indicated by the white oval since it offers a view into a yet-to-be-explored hallway [3]. This enables the human to assist the robot in prioritizing frontier locations to visit. On the other hand, the robot assists the human in identifying boundaries between explored and unexplored space and quantifying the information along these regions. Furthermore, the robot can improve the team's exploration strategy over time by learning how humans discriminate between one frontier location versus another when encountering regions in the workspace that are similar to previously observed locations. By tracking each member's performance across various subtasks over time, robots can better match each member's capabilities to new subtasks as they arise, thus enabling humans and robots adapt to one another as they execute missions of varying complexities.

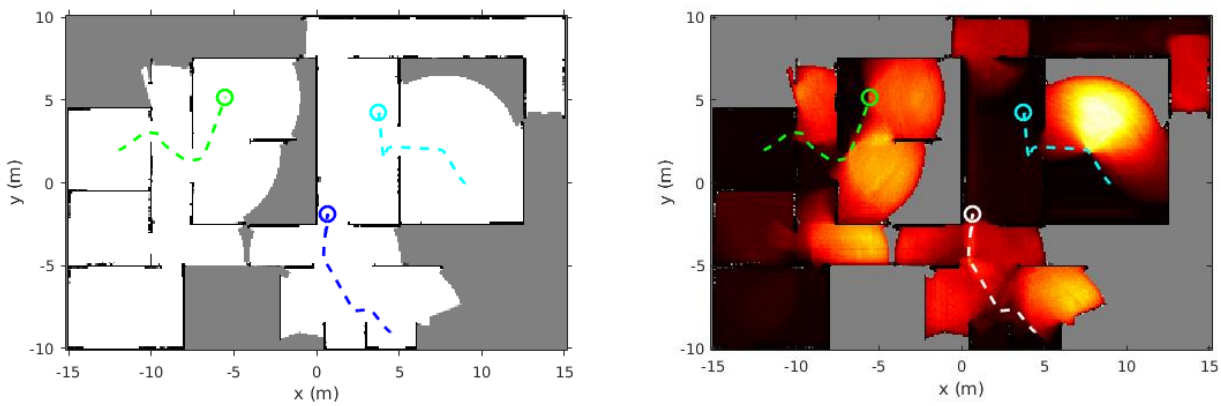

(a) (b)

Fig. XX (a) Example of three robot team mapping and exploring an office environment. White denotes explored space, gray denotes unexplored/unknown space, circles denote current robot locations, and dashed lines denote robot trajectories. (b) Corresponding information reward surface where black denotes regions of low information and white denotes regions of high information content. White rectangle indicates high information region that may yield more information in the long run compared to the region indicated by the white oval.

To achieve this, human and robot world representations must enable easy and natural assimilation of new sensor and perceptual data and model updates. These requirements necessitate representations of varying granularity since humans and robots reason about their environments at differing levels of abstraction. In this work, we assume two different maps, M_r and M_h. The robot will use M_r to reason about actions while the human will use M_h where M_r is mostly a metric map while M_h is mostly a toplogical map. Human and robot inputs would then be used to annotate the respective maps and the information would be percolated throughout maps via pattern/template matching. This work will focus on representations that can encode and blend low-level metric information with high level topological information and investigate techniques to efficiently map information from one representation to another, e.g., human to robot and vice versa. We envision the information exchange from human to robot to be some blended combination of semantic and topological information while information exchange from robot to human would be mostly semantic. Metric information can then be visualized depending on the robustness of the communications link. Furthermore, we are interested in strategies that would result in no information loss in the sense that both the robot's and human's information will increase as each provides the other with more data. Towards this end, we will evaluate the information content both before and after any mapping/translation process and identify any impact to the performance human-robot teams executing complex missions over time.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Collaboration: This work will involve collaboration with Dr. Ethan Stump and Dr. Jonathan Fink who are interested in generalizing information theoretic approaches for mapping and exploration strategies to human-robot teams and more complex tasks. In addition, this work will require collaborations with Nick Roy (T2C2S1C), Tom Howard (T2C2S1D), and Florian Jentsch (T2C2S3B) in the Distributed Mission Execution capability. Integration will require collaboration with the Integrated Research Thrust, including Ed Weller (T4C3S1C) and Daniel Barber (T4C3S1E).

Linkage to previous RCTA BPP sub-tasks: I1-2

Expected products:
The first results of this research will be submitted to ICRA, RSS, and/or ISER. In FY2017 we will develop the initial world representations for humans and robots and tools for mapping between them. These products will be provided to the integration team and program team as well as to key collaborators. Extension of information theoretic exploration and mapping strategies for robots with human input will also be completed by the end of FY2017. In FY2018 we will focus on developing task allocation strategies for human-robot teams based on each member's capabilities. Specifically, we will consider collaborative scenarios where robots and humans collective improve their performances based on the shared history of inputs, performance on subtasks, and ability to identify new versus previously encountered scenarios and take appropriate action.

References:

1. Benjamin Charrow, Sikang Liu, Vijay Kumar, and Nathan Michael. "Information-theoretic mapping using Cauchy-Schwarz Quadratic Mutual Information," in the *Proc. of the 2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp 4791–4798, 2015.

2. Daniel Mox, Anthony Cowley, M. Ani Hsieh, and C. J. Taylor. "Information Based Exploration with Panoramas and Angle Occupancy Grids," in the *Proc. of the International Symposium on Distributed Autonomous Robotic Systems (DARS 2016)*, Nov 2016, London, UK.

3. R. G. Colares & L. Chaimowicz "Planning for Simultaneous Localization and Mapping using Topological Information," in the *Proceedings of the 2012 Brazilian Robotics Symposium and Latin American Robotics Symposium (SBR-LARS)*, pp. 214-219, 2012.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.7 T2C2S2A - Learning to explain route choices

PI: Oh, Jean (CMU), jeanoh@nrec.ri.cmu.edu, (412) 651-6052

Objective and Benefits: In this task, we aim to answer the following research question: How can we make robots being able to explain what it has done and what its current and future plans are? Being able to explain its own behavior is one of the key capabilities to be a competent and trusted teammate especially when a team operates in a distributed fashion. In previous subtask I1-1 Framework for Collaborative Intelligence, we developed semantic navigation for ground vehicles. During IRAs and other experiments, we have observed various cases of robot failures, especially due to a sudden re-planning during execution, e.g., changing the direction toward another target building, moving away from its originally correct target. To find out the reasons behind those failures, researchers look through long lines of log files to follow the robot's reasoning process based on the researchers' understanding of various algorithms used inside the intelligence architecture. In this context, we propose to investigate an approach for a ground robot to be able to explain its navigation paths.

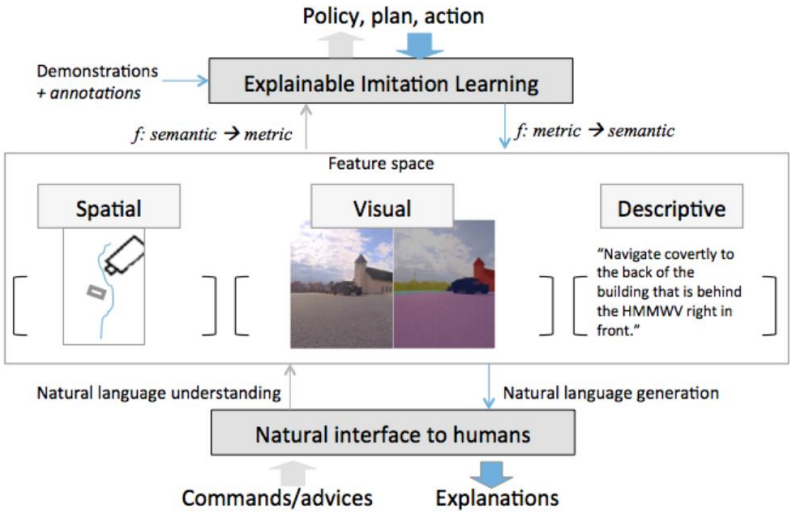

Figure 2 A conceptual model of explainable imitation learning for navigation.

Scientific approach: The proposed approach to tackle this challenge builds on our prior work on imitation learning for language understanding in outdoor navigation [Oh2016, Boularias2016, Boularias2015, Oh2015], where the focus was on learning the mapping from language to actions (the flow on the left hand side in Figure 1). In the proposed work, the focus will be on the reverse direction from actions to language (the flow on the right hand side). As illustrated in Figure 1, we will also investigate approaches to augment the feature space where a robot's decision making is depicted as a function of spatial, visual, and descriptive features. In particular, the use of semantic visual features and language features (as opposed to abstract features) will enable generating human-understandable explanation. Specifically, we will include semantic features that can be extracted via computer vision, e.g., terrain types such as grass, gravel, or concrete and object types such as buildings or cars. To this end, we will closely collaborate with perception researchers, e.g., T1C1S1D (Hebert, CMU), as we have been using their semantic perception system in the previous program.

To generate explanations, we will take a discriminative approach, i.e., we explain the rationale supporting the chosen decision by comparing that to alternatives. For instance, we want to be able to generate explanations such as "I am going to the right of the car because I would be moving too close to a wall if I take the left route." We will use sampling techniques to generate alternative routes to compare against a chosen path. In addition, we will conduct MTurk study to collect verbal descriptions of different route

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

choices; the data collected will be used for both training our imitation learning for route generation and generating explanations in natural language. We will investigate the use of recurrent neural networks (RNN) for language generation.

A route in the current system is represented only in a metric form as a set of waypoints. In the proposed work, we plan to generate semantic description similar to navigation commands or directions, using landmarks and the waypoints' relative spatial relationships to those reference objects. A route object can be encapsulated in a message similar to the JSON object that has been used for sharing updated routes with the MMI in the previous program. In the proposed work, the route object can have descriptive text, embedded images, and meta-information explaining the current route, e.g., my path has been modified because the grounding of the goal building has been changed from X to Y.

List of dependencies: This task will depend on the perception tasks T1C1S1D (Hebert, CMU) and T1C1S3A (Muelling, CMU) since we plan to utilize semantic features from the perception system. The output from this task can be communicated back to human soldiers through the MMI in task T4C5S1E (Barber, UCF).

Collaboration: We also plan to collaborate with HRI researchers to study how the capability of being able to explain its behavior affects the trust building process in a team, e.g., T2C2S2D (Chen, ARL) for agent reasoning and operator trust and T2C1S4A (Jentsch, UCF) for naturalistic human-robot trust building.

Linkage to previous RCTA BPP sub-tasks: I1-1 Framework for Collaborative Intelligence

Expected products:

- The resulting software will be integrated on the IRA Husky platform to support indoor/outdoor navigation tasks.
- Target publication venues include AAAI, IJCAI, R:SS, IROS, HRI, ICRA, AAMAS, ICAPS, and ISER.
- The data sets collected from robot experiments and MTurk study will be disseminated.

2017	• Augment feature space to include visual features. • Design and conduct MTurk study to collect language features and matching routes. • Develop a sampling method to generate a set of alternative routes given a command. • Develop a base system of explanation generation using an RNN model. • Evaluate on the data set from MTurk study.
2018	• Develop the explainable imitation learning framework using multimodal features • Continue developing natural description generation with a richer language model. • Evaluate on RCTA IRA platforms (ClearPath Husky) on both indoor and outdoor navigation tasks. • Create RCTA trajectory-text datasets for dissemination.

References:

[Oh2016] J. Oh, M. Zhu, S. Park, T.M. Howard, M.R. Walter, D. Barber, O. Romero, A. Suppe, L. Navarro-Serment, F. Duvallet, A. Boularias, J. Vinokurov, T. Keegan, R. Dean, C. Lennon, B. Bodt, M. Childers, J. Shi, K. Daniilidis, N. Roy, C. Lebiere, M. Hebert, and A. Stentz. Integrated intelligence for human-robot teams. In Proc. ISER, 2016.

[Boularias2016] A. Boularias, F. Duvallet, J. Oh, and A. Stentz. Learning Qualitative Spatial Relations for Robotic Navigation. in Proc. of International Joint Conference on Artificial Intelligence (IJCAI), 2016.

[Boularias2015] A. Boularias, F. Duvallet, J. Oh, and A. Stentz. Learning to ground spatial relations for outdoor robot navigation. In Proc. of IEEE Conference on Robotics and Automation (ICRA), 2015.

[Oh2015] J. Oh, A. Suppe, F. Duvallet, A. Boularias, J. Vinokurov, L. Navarro-Serment, O. Romero, R. Dean, C. Lebiere, M. Hebert, and A. Stentz. Toward mobile robots reasoning like humans. In Proc. of AAAI, 2015.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.8 T2C2S2B – Grounding Natural Language Instructions with Learnt Visual-linguistic Knowledge Representation

PI: Roy, Nicholas (MIT), nickroy@csail.mit.edu, +1-617-875-2506

Objective and Benefits:

Natural language interaction provides a rich, intuitive and flexible medium for humans and robots to communicate mission objectives and situational awareness, facilitating teaming and peer-to-peer tactical behaviours. Significant progress has been made in the development of probabilistic models capable of grounding rich linguistic data (possessing composition, hierarchy, monologues etc.) [1, 2, 3, 4] with complex symbolic representations (like relations, abstraction, constraints etc.) describing the varied workspaces the robot may be operating in. Contemporary models can efficiently ground instructions like "put the pallet on the the truck" or "pick up the nearest block on the table" or "go to the fire hydrant behind the traffic cone". The grounding or interpretation of a language command is intimately tied to the robot's knowledge about the world. SOA models embed the input instruction in a world representation that includes entities like object, regions and constraints that can be currently perceived by the robot.

This excludes knowledge related to behaviours/events that may have occurred in the past (interactions between objects and agents) as well as declarative facts that may be stated by the human ("Post A is guarded by a soldier", "this is my bag" etc.). Further, the symbol grounding process is treated independent of the process for acquiring the world model (typically delegated to a perception or semantic mapping system) is treated independent of the symbol grounding process. However, though the two problems are fundamentally coupled. For example, consider the instruction "pick up the item that the soldier put down" that includes both imperative and declarative components which contemporary approaches cannot model. Our goal is to develop language understanding models that can incorporate both imperative and declarative language, tightly coupled with visual perception of the scene dynamics.

Scientific approach:

We propose to address this problem by formulating a unified probabilistic model that jointly expresses the robot's knowledge about the world (from visual and language percepts) as well as inferring the groundings (objectives and constraints) for intended actions that it should execute. Firstly, we will incorporate factual information or mission relevant statements that may be communicated by human team mates via declarative language commands. Secondly, we will expand the space of symbols to include temporal events (higher-order interactions) as observed by a visual perception system. Joint inference will allow the symbol grounding model to determine intended actions by appropriately accounting for uncertainty in its belief over the world knowledge, enabling grounding of language instructions that involve simultaneous grounding of imperative and declarative statements. For example, "retrieve the canteen that the human placed on the ground", "guard the soldier who is lying down" or "the equipment on the left is mine, bring my equipment". Further, we will model the stochastic effects of the robot's actions and integrate suitable simulation models for environmental entities to update the belief over the world model. Experimental demonstration will include an integrated system incorporating lexical analysis, perception and symbol grounding capable of persistent grounding during operation. Without the ability to acquire new information from natural language and connect this information both to the existing cognitive model and to instructions from human team-mate, effective teaming will be difficult.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

We estimate a semantic representation of the workspace derived from visual observations or language descriptions provided by the human. Given the language command and the world model, we infer the sequence of actions that best satisfy the instruction. Further, we propagate the world model as using past observations and language commands. Formally, we estimate the posterior distribution over the intended sequence of actions and an updated world model given the visual-linguistic percepts and current world representation.

Learning this distribution is intractable since the joint space of possible natural language inputs, percepts, actions and possible world configurations is continuous and high- dimensional. The approach we propose to take in this work is to introduce a symbolic representation that induces a joint factorisation enabling tractable factor learning and model inference. The factorisation renders the metric actions as condition- ally independent of the input visual percepts and language given the knowledge of the grounding variables. Instead of directly reasoning about robot actions from input visual-linguistic percepts, the factorisation implies first estimating the symbolic world model and objects of motion from raw percepts and then determining the fine-grained actions from the inferred symbolic objective.

List of dependencies:
Requires input from T1C1S1A (Hebert, CMU), and T1C1S3A (Muelling, CMU), and provides output to T2C2S1A (Lebiere, CMU), to the world model T2C2S1C (Roy, MIT) and the MMI T2C1S4B (Barber, UCF).

Collaboration:
This subtask will support a collaboration with Ethan Stump at ARL entitled “Goal Elicitation to Formulate Problems from Human Guidance”.

Linkage to previous RCTA BPP sub-tasks:
This research builds on our previous work in I5-3 and H1-4.

Expected products:
1. Submission to Robotics Science and Systems and ICRA.
2. In FY2017 we move from our TBA completed in 2016 to integrate our symbol grounding into an integrated research assessment.
3. We will also participate with the integration teams to ensure that our symbol grounding can interact with higher-level ACT-R processes at CMU and the interfaces being developed at UCF.
4. Integration with the MMI and language processes.

References:
[1] T.M. Howard, S. Tellex, and N. Roy, "A Natural Language Planner Interface for Mobile Manipulators," in Proceedings of the 2014 International Conference on Robotics and Automation, 2014.
[2] R. Knepper, S. Tellex, A. Li, N. Roy and D. Rus. “Recovering from failure by asking for help”. Autonomous Robots, 39(3):347-362, October 2015.
[3] R. Paul, J. Arkin, N. Roy, and T.M. Howard. “Efficient Grounding of Abstract Spatial Concepts for Natural Language Interaction with Robot Manipulators”. In: Proceedings of the 2016 Robotics: Science and Systems Conference. June 2016.
[4] J. Oh, et al. “Integrated Intelligence for Human-Robot Teams”. In: Proceedings of the 2016 International Symposium on Experimental Robotics. Oct. 2016.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.9 T2C2S2C – Collaborative Perception

PI: Tellex, Stefanie (Brown University), stefie10@cs.brown.edu, (617)383-4749

Objective and Benefits:

We will create new algorithms for integrating language and perception, which enable a robot to adapt its perceptual system in response to natural language input, leading to faster and more accurate responses to its human partners. Existing approaches to perception rely on fully autonomous algorithms that do not adapt to the mission and do not identify or recover from failure. Our approach will create a model that integrates language and perception so that the robot can adapt its perceptual approach to mission-relevant aspects of the environment and model and reason about its own understanding. By communicating this model to a person, the robot will establish a shared mental model to support error detection and recovery, increasing robustness.

Scientific approach:

Existing approaches to vision are unable to reason about transparent objects or localize previously unseen objects. Our approach will enable the robot to use a monocular camera that moves or array of cameras to integrate information into synthetic photographs useful for understanding the environment in human terms. Additionally, the robot can rapidly create a model for localizing a previously unknown object, integrated with language. As a result, if the person says "Pick up that screwdriver over there" and points, and the robot does not have a model for the word "screwdriver," the robot will interpret the pointing gesture to localize the object, then confirm "is this what you meant?" Next it will create a model for the screwdriver, so that next time it will be able to localize and pick the object without asking for confirmation. We are creating new techniques using light fields and synthetic photography to enable the robot to build robust models of objects by exploiting its ability to move its camera. Additionally, the person's language and gesture input will enable the robot to direct its movements and computation to just the locations and models necessary to satisfy the person's request. If the robot is looking for a screwdriver, there is no reason to search and localize other objects such as hammers. By integrating language and perception using light fields and a model factored around objects, this subtask will improve efficiency of human-robot collaboration. The robot will also collect data it needs to more robustly perform on object categories. If it lacks a screwdriver model, it can automatically acquire an instance based model that leads to robust manipulation, and also collect data that is necessary for learning models of object categories, to localize previously unseen screwdrivers.

The data to be collected consists of calibrated camera images. We need to know the image content, as well as the camera pose and camera calibration data, so that each pixel can be converted to a light ray in space. It is okay if these poses are in a local coordinate system as long as groups of images/poses are in the same system. Our prototype uses an arm with a camera mounted on a fixed base. This capability will rely on the robot's ability to localize itself in space over short periods of time.

We assume the world model represents the world in terms of objects. After modeling a new object, we will add a new object instance to the world model that can be tracked and detected. After groups of object instances are added, we will add object categories, such as "screwdrivers" in general, as opposed to "my red Phillips head screwdriver" in particular.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

This work complements T1C1S3A (Muelling, Anytime neural networks for semantic understanding on the move), which uses anytime methods to provide predictions. Based on human collaborative input, the system can dynamically decide where and when to invest computational resources to carry out perceptual tasks. We can also leverage T1C1S1D (Hebert, Rapid Learning of Recognition Models) to learn category-based models after acquiring instance-base data. This approach will allow us to generalize what the system has learned to new examples.

List of dependencies: This work is closely related to Nick Roy and Rohan Paul's work on language understanding (T2C2S2B).

Collaboration: T2C2S1C

Linkage to previous RCTA BPP sub-tasks: I5-5, P3-8
This work is connected to the previous subtask on Collaborative Perception. There we created POMDP models for interactive dialog and developed the light field models for perceiving the environment. Moving forward, we will work to handle unknown environments and integrate with complex language instructions by connecting with a robots' ability to perceive objects.

Expected products:
List verifiable products of this research effort by years (2017 and 2018): Examples are…
Publication at ICRA 2017 and RSS 2017
Demonstration on RCTA platform with the world model.
Ongoing collaboration with Rohan Paul and Nicholas Roy.
Workshop submissions; invited lectures at Georgia Tech, University of Rochester, and visit by ARL to Brown. CMU RI seminar in Fall 2016.

References:

John Oberlin and Stefanie Tellex. Autonomously Acquiring Instance-Based Object Models from Experience. In International Symposium on Robotics Research (ISRR), 2015.

John Oberlin, Maria Meier, Tim Kraska, and Stefanie Tellex. Acquiring Object Experiences at Scale. In AAAI-RSS Special Workshop on the 50th Anniversary of Shakey: The Role of AI to Harmonize Robots and Humans, 2015. Blue Sky Award.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.10 T2C2S2D - Transparency of Agent Reasoning and Operator Trust

PI: Chen, Jessie (ARL-HRED), yun-sheng.c.chen.civ@mail.mil, 407-384-5435

Objective and Benefits:

The informational needs of mutually interdependent human-agent teams must be anticipated if these teams are to act effectively. The objective of this research is to investigate the impact of using bidirectional communication to facilitate transparent interaction between a human and an agent, as well as the impact of role transparency on human-agent teaming effectiveness. The benefits of this research include the establishment of a compatible shared understanding between the human and the agent as well as determining how the presentation of information meant to support transparency can support effective human-agent team performance.

Scientific approach:

Since the Situation Awareness-based Agent Transparency (SAT) model was first proposed, progress in the field of artificial intelligence and in the capabilities of autonomous systems has advanced. To capitalize on these advances, we have sought to expand the SAT model. The goal of this expanded model is to better address joint action in human-agent teams, emphasizing the bidirectional communication of information and the way each team member apportions responsibilities in collaborative tasks.

The SAT model describes three independent levels, each of which denote information an agent needs to convey in order to maintain a transparent interaction with the human (Chen et al. 2014). Within this model, bidirectional communication between human and agent teammates can be used to make both a human's and an agent's reasoning available to the other party. Rather than resulting in agent transparency, interaction between the agent and the human can result in bidirectional transparency (Chen et al., 2017). This bidirectional transparency allows human and agent teammates to take each other into account when making decisions as well as allowing them to use each other's behavior and reasoning as a resource in their own decision making process (Fong, Thorpe, & Baur 2003; Kaupp, Makarenko, & Durrant-Whyte 2010).

Lyons and associates (2017) suggest that transparency is comprised of a number of different dimensions; in one of these dimensions, the *team* dimension, the agent supports the human's understanding of the division of labor between the human and the agent by conveying an understanding of which tasks fall under the human's purview, which tasks are the agent's responsibility, and which tasks require both humans' and agents' intervention (Bradshaw, Feltovich and and Johnson, 2011; Bruni et al., 2007; Lyons 2013). While overt communication of roles, responsibilities, and the fulfilment thereof is expressed in the human teamwork literature (Cannon-Bowers, Salas, and Converse 1993; DeChurch and Mesmer-Magnus 2010), discussion of these factors in the context of human-agent teams has only recently been receiving increased attention (Hayes and Scassellati 2013; Lyons 2013; Sycara and Sukthankar 2006).

In this subtask, the team dimension will be integrated with the SAT model and examined in a task where both the human and the agent communicate information to support transparency. This research seeks to granularly examine the influence these factors have on transparency and on a human-agent team. The results of this research effort will inform the development of guidelines for agent designs that can be used to to promote transparency in human agent teams.

Thrust and Capability Supported

This research will support the RCTA's Thrust 2, *Human-Robot Execution of Complex Missions*, Capability 2, *Distributed mission execution*. This subtask will examine transparency and bidirectional communication in the context of distributed interactions between human and agent team members carrying out complex missions.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

List of dependencies:
This subtasks depends on the use of the RCTA MMI and related software, so this research may depend on Dr. Daniel Barber's T4C3S1E, *HRI Driven Enhancements*.

Collaboration:
This subtask will have close ties with some members of Thrust 2, Human-Robot Execution of Complex Missions, and provide opportunities for collaboration. This subtasks will involve collaboration with Dr. Daniel Barber (T2C1S4B, *Adaptive multimodal communication* and T4C3S1E, *HRI Driven Enhancements*) due to his status as developer for the RCTA MMI. This collaboration will entail coordination with regard to experiment design and interface development. This subtask may also involve collaboration with Dr. Florian Jentsch (T2C2S3B, *Establishing Appropriate Human-Robot Trust through Empirically Validated Training Techniques: Investigating Mental Models, Situational Risk, and Trust* and T2C1S4A, *A user-centered design (UCD) approach to creating usable naturalistic communication interfaces for Soldier-robot teaming*), given his work with mental models and trust. Additionally, collaboration with Nicholas Roy (T2C2S2B, *Grounding Natural Language Instructions with Learnt Visual-linguistic Knowledge Representation* and T2C2S1C, *Learning Symbolic Planning Models*) is anticipated, due to his work with planning models and agent reasoning.

Linkage to previous RCTA BPP sub-tasks:
This subtask builds upon research performed in 2015-2016 under Task H2, Applying Models of Cognition to Soldier-Robot Teams, subtask 3, Transparency of Agent Reasoning and Operator Trust. The research focused on the transparency of an agent's reasoning process and how the availability of this information influenced a human teammate's engagement, trust, situation awareness, workload, and performance. The 2017-2018 research will continue to explore the impact of an agent's reasoning transparency on engagement, trust, situation awareness, workload, and performance. However, the 2017-2018 research will expand on the goals of the 2015-2016 research by including transparency of an agent's role and responsibilities as well as exploring bidirectional communication in transparent human-agent interaction.

Expected products:
Publication of this subtask's findings are anticipated in conferences (e.g., International Conference on Human-Computer Interaction, International Conference on Applied Human Factors & Ergonomics, and International Conference on Human-Robot Interaction) and journals specializing in human factors and/or human-robot interaction. By the end of FY2017, design guidelines for supporting transparency, based upon results of simulation experiments and assessments, will be developed. By the end of FY2018, refined guidelines will be provided, as well as a technical report detailing experimental procedure, findings, and how the guidelines were developed.

References:

1. Bradshaw, J., Feltovich, P., & Johnson, M. (2011). Human–agent interaction. In G. Boy (Ed). *Handbook of Human–Machine Interaction*, pp. 283–302. Farnham, U.K.: Ashgate.
2. Bruni, S., Marquez, J. J., Brzezinski, A., Nehmeand, C., Boussemart, Y. (2007). Introducing a Human-Automation Collaboration Taxonomy (HACT) in Command and Control Decision-Support Systems. in *12th International Command and Control Research and Technology Symposium*, Newport, RI.
3. Cannon-Bowers J., Salas E., Converse S., (1993). Shared mental models in expert team decision making. In N.J. Castellan Jr (Ed), *Individual and Group Decision Making: Current Issues*. pp. 221–46. Hillsdale, NJ: Erlbaum

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4. Chen, J. Y. C., Procci, K., Boyce, M., Wright, J.L., Garcia, A., & Barnes, M. J., (2014). SA-based agent transparency (ARL-TR-6905). Aberdeen Proving Ground, MD: US Army Research Laboratory.
5. Chen, J. Y. C., Lakhmani, S.G., Stowers, K., Selkowitz, A. R., Wright, J.L., & Barnes, M. (2017). Situation Awareness-based Agent Transparency and Human-Autonomy Teaming Effectiveness. Manuscript submitted for publication.
6. DeChurch, L. A., & Mesmer-Magnus. J. R. (2010). The Cognitive Underpinnings of Effective Teamwork: a Meta-Analysis. *Journal of Applied Psychology* 95 (1): 32-53.
7. Fong, T.W., Thorpe, C., & Baur, C. (2003). Robot, Asker of Questions. *Robotics and Autonomous Systems* 42 (3): 235-243.
8. Hayes, B., & Scassellati, B. (2013). Challenges in shared-environment human-robot collaboration. *Learning*, *8*, 9.
9. Kaupp, T., Makarenko, A., & Durrant-Whyte, H. (2010). Human–robot communication for collaborative decision making—A probabilistic approach. *Robotics and Autonomous Systems*, *58*(5), 444-456.
10. Lyons, J. B. (2013, March). Being transparent about transparency: A model for human-robot interaction. In *2013 AAAI Spring Symposium Series*.
11. Lyons, J. B., Sadler, G. G., Koltai, K., Battiste, H., Ho, N. T., Hoffmann, L. C., Smith, D., Johnson, W., & Shively, R. (2017). Shaping trust through transparent design: theoretical and experimental guidelines. In *Advances in Human Factors in Robots and Unmanned Systems* (pp. 127-136). Springer International Publishing.
12. Sycara, K., & Sukthankar, G. (2006). Literature review of teamwork models. *Robotics Institute, Carnegie Mellon University*, *31*.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.11 T2C2S2E – Semantic Framework for Concepts in Context

PI: Summers-Stay, Douglas (ARL), douglas.a.summers-stay.civ@mail.mil, (301)394-1990

Objective and Benefits:

This subtask involves developing a framework for understanding commands in the context of the immediate situation. This includes:

1. Mapping natural language commands into formal representations suitable for reasoning.
2. Mapping higher level statements about goals and intent into the same representation.
3. Mapping facts about the spatial arrangement of objects in the environment into the same representation.
4. Reasoning about the commands to interpret them in a way that makes sense in the context of all these other facts, and generating appropriate tactical behavior specifications.

Scientific approach:

The representations we are using for performing the reasoning are a hybrid of a knowledge base and a distributional semantic vector space (DSVS). Using the ability of a DSVS to discover true analogies, we can add flexibility and generalization capability to the knowledge base. [Summers-Stay 2016]

There are several areas of this technique that we are exploring how to improve. Which method of creating DSVS is the most effective for this task? Our experiments have shown that vectors which incorporate sentence word role information [Turney 2012] are superior to vectors derived from an indiscriminate window for analogy tasks.

In terms of the knowledge base, some of the questions we are trying to address are as follows: How should facts about intent be captured in the knowledge base in a way that allows us to reason about them effectively? How can we use reasoning to generate tactical behavior specifications [Oh 2015] that will be interpreted by lower level components of the system? How can we take turn information extracted from natural language into assertions in the knowledge base? How can we assert facts about the surrounding area such that we can reason about spatial properties [Boularias 2015] effectively? How can we extract information about visual scenes captured by the robot's cameras in a way that this information can be asserted into the knowledge base?

List of dependencies:

T3C1S1C – Discovering Task-Related Implications of Object Arrangements will be using some of the same semantic vector reasoning to guide exploration and scene interpretation. Our reasoning system will output tactical behavior specifications, a form which the R-CTA already has experience turning into robotic actions.

Collaboration:

The core team at ARL on this research will be Dr. Douglas Summers-Stay and Dr. Dandan Li, along with graduate students. We anticipate collaboration with with Michael Martin and Christian Lebierre of Carnegie Mellon on aspects of cognition and with Jean Oh's team on aspects of natural language understanding.
T2C2S1A, T3C1S1C

Linkage to previous RCTA BPP sub-tasks:

I1-2015-16-Framework for Collaborative Intelligence Level: 6.1 Basic Research
Subtask I1-1 Cognitive Architecture for Human-Robot Collaboration (CMU)
Subtask I1-3 Grounding Spatial Language in Spatial Representations
(ARL)

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Expected products:
Two conference publications
reasoning software

References:

Summers-Stay, D., Voss, C., & Cassidy, T. (2016). Using a distributional semantic vector space with a knowledge base for reasoning in uncertain conditions. Biologically Inspired Cognitive Architectures, 16, 34-44.

Turney, P.D. (2012), Domain and function: A dual-space model of semantic relations and compositions, Journal of Artificial Intelligence Research (JAIR), 44, 533-585.

Oh, J. H., Suppé, A., Duvallet, F., Boularias, A., Navarro-Serment, L. E., Hebert, M., & Dean, R. (2015, January). Toward Mobile Robots Reasoning Like Humans. In AAAI (pp. 1371-1379).

Boularias, A., Duvallet, F., Oh, J., & Stentz, A. (2015, May). Grounding spatial relations for outdoor robot navigation. In 2015 IEEE International Conference on Robotics and Automation (ICRA) (pp. 1976-1982). IEEE.

Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., & Parikh, D. (2015). Vqa: Visual question answering. In Proceedings of the IEEE International Conference on Computer Vision (pp. 2425-2433).

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.12 T2C2S3B - Establishing Appropriate Human-Robot Trust through Empirically Validated Training Techniques: Investigating Mental Models, Situational Risk, and Trust

PI: Jentsch, Florian (Ucf), Florian.Jentsch@ucf.edu, (407) 882-0304

Objectives and Benefits: The objective of this research is to enhance our understanding regarding how human-centric and environmental-centric factors impact trust in an autonomous robot team member with the goal of developing, and empirically validating, training techniques focused on calibrating human-robot trust. Specifically, this research will investigate how human-centric factors (e.g., training, mental models) and environmental factors (e.g., situational risk) can influence trust in a robot team member. Numerous researchers have emphasized the importance that accurate mental model formation and accurate perceptions of risk play in ensuring appropriately calibrated trust in a robot. However, little research has been conducted to empirically examine these claims. Given this, our primary interest is in examining the relationship between these three concepts (training, mental models, and situational risk). This research serves to benefit the RCTA and scientific community through the development of guidelines for (1) designing robots (e.g., what is the minimum reliability a robot needs to have in order for a human to trust the robot when the task the robot is responsible for involves considerable risk?) and (2) training interventions (e.g., how can training help human teammates develop an accurate mental model of a robot and appropriately calibrate their trust to reduce trust issues such as misuse and disuse?).

Scientific Approach: The scientific approach for this research will involve empirical studies with human participants that examine the effectiveness of various training approaches for building, and calibrating, mental models and trust in an autonomous robot teammate. Initial experiments (during 2017) will compare the effectiveness of different training approaches within a lab setting (e.g., interaction with a simulated robot). Results from the lab will serve as a baseline for comparison against follow-on field experiments (during 2018). We plan to utilize the Husky platform (either simulated or real) and the multi-model interface (MMI) for our studies, and thus, each study will involve coordination with the integration team (e.g., Daniel Barber) to ensure our products can be integrated into the 2019 Capstone. All studies will be conducted with novice robot users, as these are the users that are most likely to benefit from targeted training to dispel inaccurate expectations, help them to establish an accurate mental model of their robot teammate, and subsequently, to appropriately calibrate their trust in the robot.

This research is most closely linked to research being performed by Elizabeth Phillips (T2C2S3D *Appropriate calibrations of trust for supporting Soldier-robot teaming*) and Jessie Chen (T2C2S2D *Transparency of agent reasoning and operator trust*) during BPP 2017-18 as those subtasks seek to understand factors impacting human trust in one (or more) robot teammates. It also builds upon prior research conducted under H1.1 (Jentsch; *Determinants of Shared Mental Models for Soldier—Robot Teams*) regarding human mental models of, and trust in, robot teammates (see references).

This research primarily will support the RCTA's Thrust II/Capability 2 "Distributed Mission Execution" capability in that our experiments will be situated within the context of a human-robot team performing a mission that involves distributed interactions (e.g., non-line of sight communications). This research is also well suited to support other RCTA capabilities that involve distributed interactions/teaming between humans and robots such as Thrust II/Capability 1 "Situation awareness in unstructured environments." This research could also be adapted to support capabilities focused on non-distributed teaming situations (e.g., human and robot working within close proximity of one another). Furthermore, the research proposed under this subtask helps achieve the RCTA Capstone vision by helping to enhance the likelihood that human team members will appropriately trust their robot team members.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

List of Dependencies: To meet the goals of Thrust II/Capability 2 "Distributed Mission Execution," this subtask will be depend upon T4C3S1E *HRI Driven Enhancements* (Daniel Barber) as our experiments will involve teaming situations that require distributed communication with a robot team member. These communications will be facilitated through the use of a multi-modal interface (either simulated or real). We plan to use the multi-modal interface (MMI) being developed by Daniel Barber, and thus, will require his assistance in using the MMI and potentially creating new functionality needed for our research. In addition, the results of our research will filter back to T4C3S1E *HRI Driven Enhancements* in the form of design and training guidelines.

Collaboration: This subtask will involve collaboration with other RCTA researchers conducting human-robot trust research. One major collaborator will be Elizabeth Phillips, Brown University, (T2C2S3D *Appropriate calibrations of trust for supporting Soldier-robot teaming*) to consult on experimental designs, facilitate data collection, possible student exchanges, and produce joint papers/publications. In addition we plan on potentially collaborating/coordinating with Jessie Chen, ARL (T2C2S2D *Transparency of agent reasoning and operator trust*) given her previous and current work on trust and transparency. Further, we intend to collaborate with Daniel Barber, UCF (T2C1S4B *Adaptive Multimodal Communication* and T4C3S1E *HRI Driven Enhancements*) on experiment design, stimuli development, and MMI integration.

Linkage to Previous RCTA BPP Sub-tasks:

- H1-1: *Determinants of Shared Mental Models for Soldier—Robot Teams* (see references)

Expected Products:

- **End of BPP Year 2017:** Initial training and design guidelines will be provided to the integration team and program team as well as to key collaborators. These initial guidelines will be based upon results of our lab simulation experiments and integrated assessments. The results of our lab simulation experiment will be presented at an appropriate conference and made available via a technical report. Candidate conferences include the *Annual Meeting of the Human Factors and Ergonomics Society (HFES), ACM/IEEE International Conference on Human-Robot Interaction (HRI),* or *HCI International (HCII).*
- **End of BPP Year 2018:** Refined training guidelines will be provided to the integration team and program team as well as to key collaborators. These refined guidelines will be based upon results of a field study, possibly as part of a multi-component or integrated assessment. The results of the field study will be published in an appropriate journal. Candidate journals include *Journal of Human-Robot Interaction or Human Factors.*

References:

Previous H1.1 publications:

1. Ososky, S. J. (2013). *Influence of task-role mental models on human interpretation of robot motion behavior* (Doctoral dissertation). University of Central Florida, Orlando, Florida

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2. Ososky, S., Sanders, T., Jentsch, F., Hancock, P., & Chen, J. Y. (2014). Determinants of system transparency and its influence on trust in and reliance on unmanned robotic systems. In *SPIE Defense+ Security* (pp. 90840E-90840E). International Society for Optics and Photonics.

3. Ososky, S., Schuster, D., Phillips, E., & Jentsch, F. G. (2013). Building appropriate trust in human-robot teams. In *2013 AAAI Spring Symposium Series.*

4. Phillips, E., Ososky, S., Grove, J., & Jentsch, F. (2011). From tools to teammates toward the development of appropriate mental models for intelligent robots. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 55, No. 1, pp. 1491-1495). SAGE Publications.

5. Phillips, E., Schaefer, K. E., Billings, D. R., Jentsch, F., & Hancock, P. A. (2016). Human-Animal Teams as an Analog for Human-Robot Teams: Influencing Design and Fostering Trust. *Journal of Human Robot Interaction*, *5*(1), 100-125. doi: 10.5898/JHRI.5.1.Phillips.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.13 T2C2S3C - Investigating Multi-Modal Soldier-Robot Bidirectional Exchanges

PI: Elliot, Linda (ARL), linda.r.elliott.civ@mail.mil,

Objective and Benefits:

Researchers at the ARL HRED Fort Benning Field Element continue to examine issues regarding Soldier-robot multimodal and bidirectional communications to enhance performance and mission effectiveness, building upon on previous efforts (Redden, Elliott & Barnes, 2013; Elliott & Redden, 2013, Elliott, Mortimer, & Skinner, 2014; Elliott, L., Hill, S., Barnes, 2016; Elliott, L. 2016, Abich, Barber, & Elliott, 2016, Mortimer & Elliott, 2016; Hartnett et al 2015;Elliott et al, 2016, Pettitt et al., in press). During 2016 HRED researchers worked with ARL, UCF, and other researchers to investigate uni- and bidirectional communications between Soldiers and robots. 2016 efforts collected Soldier subjective and performance based data to evaluate individual and combined options for multimodal communications. Gesture-based controls using instrumented gloves were evaluated for robot control and manipulation (Anthrotronix phase 2 SBIR). Tactile salience, learning, and recall was further explored through manipulations of tactor and signal patterns. Soldier feedback was also collected to validate and refine map-based inferences in a route reconnaissance scenario

For 2017, research will further refine aspects regarding tactile HRI for robot-Soldier communications, to better measure and construct tactile signals that are easily felt and interpreted. In collaboration with SOARTECH (phase 2 SBIR), speech and gesture commands will be investigated through emerging technology developed for the SSID platform. Collaborations with UCF include two protocols (Abich et al, Talone et al). Deliverables include technical reports detailing empirical experiment-based results of mission performance using alternative bidirectional speech/gesture control approaches. The results will be used to-develop HRI design guidelines. This work is related to other on-going work in the Robotics CTA on multimodal, bi-directional communication between humans and robots. Progress will be reported through two invited panels, one for Applied Human Factors and Ergonomics Conference on multimodal tactile interfaces (June 2016) and one for International Human Computer Interface panel on wearable technology.

Technical Approach to Overcome Barriers and Advance SOA:

All technical approaches rely on the participation of Soldiers with experience and expertise. Interview and survey-based techniques will be used to capture and probe issues with regard to device use in various Soldier missions. Experiment- and simulation-based evaluations will focus on Soldier performance using the device(s) in mission-relevant operational tasks, such as navigation and route reconnaissance, that includes squad level robotic assets. The technology inherent in the displays and controls is advanced and emerging. . The approaches taken will capitalize on Soldier expertise, operational context, and experimental rigor, with regard experiment design and interpretation of Robotics CTA and advances our understanding of human-robot communications – a major thrust area of ARL HRI research.

Metrics & Progress Assessment:

Publications (ARL tech reports and other peer-reviewed publications) and transitions of research results (including user interface design recommendations) to Army robotics programs (e.g., RCTA and TARDEC programs) are expected.

Year Qtr Subtask H4.3 Quarterly Milestones and Deliverables

2017 Q1

•Submit protocol for SOARTECH examining advanced speech and gesture commands for soldier-robot bidirection communications

• Document 2016 experiment regarding Anthrotronix gesture and tactile displays for HRI.

2017 Q2

•Submit protocol for UCF collaboration (Squad Level Soldier-Robot Communication Exchanges)

• Conduct SOAR tech data collection

Conduct tactile EAI data collection

Document data collections

2017 Q3

•Conduct UCF data collection (Squad Level Soldier-Robot Communication Exchanges)

• TR from SOARTECH experiment

AHFE and HCI Invited Panels

2017 Q4

•TR from UCF data collection

References

Elliott, L., Hill, S., Barnes, M. (2016). Gesture-based controls for robots: Overview and implications for use by Soldiers. ARL Technical report 7715.

Elliott, L. A Vision for Future Soldier-Robot Teams. Panel Chair for session presented at the 7th International Conference on Applied Human Factors and Ergonomics 2016, Orlando.

Abich, Julian, Barber, Daniel. & Elliott Linda.: "An initial investigation of exogenous orienting visual display cues for dismounted Soldier-robot communication". Paper presented at the 7th International Conference on Applied Human Factors and Ergonomics 2016, Orlando.

Mortimer, Bruce & Elliott, Linda. "Context Sensitive Tactile Displays for Bidirectional HRI Communications. Paper presented at the 7th International Conference on Applied Human Factors and Ergonomics 2016

Elliott, L. CTA approaches to support scaled simulations and human performance in military scenarios Presentation to Human Factors and Ergonomics Society, HFES, September 2016.

Hancock, P., Elliott, L.,Cholewiak, R., van Erp, J., Mortimer, B., Rupert, A., Schmeisser, E. Redden, E. . Tactile cueing to augment multisensory human-machine interaction (2015). Ergonomics in Design.

Elliott, L., Mortimer, B., Hartnett-Pomranky, G., Zets, G., Mort, G. (2015). Augmenting Soldier situation awareness and navigation through tactile cueing. Proceedings of the HCI International 2015, the 17th International Conference on Human-Computer Interaction, The Westin Bonaventure Hotel, Los Angeles, CA, USA, 2-7 August 2015.

Swiecicki, C., Elliott, L., & Wooldridge, R. Squad-level Soldier-robot dynamics: Exploring future concepts involving intelligent autonomous robots (Feb, 2015). Technical Report No. 7215. Army Research Laboratory Human Research and Engineering Directorate.

Elliott, L., Mortimer, B., Cholewiak, R., Mort, G., Hartnett, G., Pettitt, R. (2015) Salience of tactile cues: An examination of tactor actuator and tactile cue characteristics. Technical Report ARL-TR-7392. Army Research Laboratory Human Research and Engineering Directorate.

Pomranky-Hartnett, G., Elliott, L., Mortimer, B., Mort, G., Pettitt, R. (2015). Soldier based assessment of dual-row tactor displays during simultaneous navigation and robot monitoring tasks. Technical Report ARL-TR-7397. Army Research Laboratory Human Research and Engineering Directorate.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.5.14 T2C2S3D - Appropriate calibrations of trust for supporting soldier-robot teaming.

PI: Phillips, Elizabeth (Brown University), elizabeth_phillips1@brown.edu, (401)863-1359
Additional PIs: Malle, Bertram (Brown University), Bertram_malle@brown.edu, (401)863-6820

Objective and Benefits: In future soldier-robot teams, human trust in the robotic teammate(s) will be needed to support successful peer-to-peer interactions. Complicating this picture, however, is the fact that humans tend to hold unrealistic expectations of robot performance and are sensitive to violations of these expectations (Phillips, Ososky, Grove, & Jentsch, 2011). As robotic teammates cannot yet perfectly replicate the capabilities of their human counterparts, expectations of robots are likely to be violated and, as a result, trust diminished or potentially lost. It is therefore important to investigate means by which the actions and capabilities of the robot can help to mitigate such violations and appropriately calibrate the human's level of trust in the robotic teammate(s) (Ososky, Schuster, Phillips, & Jentsch, 2013). The objective of this research is to investigate mechanisms, such as the sharing of a robot's norm awareness, to calibrate a soldier's trust in the robot. Appropriately calibrated human trust in robotic teammates will be especially important for helping to mitigate inappropriate reliance on the robot.

Scientific approach: This research will investigate means such as the robot's sharing of metacognitive information, including detected cues and detected norms (e.g., implicit, explicit, and/or derived from environmental perception) that can bring about appropriately calibrated human trust in a robot throughout the course of an OPTEMPO mission. For instance, a robot that communicates what it perceives to be the relevant norms in the situation may signal that the robot has critical coordinating mechanisms that underlie the team's goals across missions and tasks. Team actions, in particular, rely on norms to coordinate, streamline, and legitimize team behavior (Feldman, 1984); and team members who follow these norms may be trusted and relied on (Crisp & Jarvenpaa, 2013). It therefore stands to reason that humans who act jointly with robots in teams expect their robot partners to be aware of and follow many of the same norms that they themselves follow. If robots in fact do so, human team members are likely to trust them, and justifiably so. Evidence for the robot's norm awareness becomes a calibration device for human trust: trust is warranted in those specific contexts for which the robot abides by relevant norms. The robot's communication of its norm awareness may be one means to help soldiers maintain appropriate trust in the robot, and avoid under reliance (Disuse) when the team encounters novel situations in which the robot could be useful, and/or overreliance (Misuse) when the team assigns a difficult task to the robot that it is not capable of completing. The proposed research will utilize a combination of survey methods and laboratory experimentation. Survey methods will be used to identify the norms that govern team behavior, as well as those that specifically apply to robots operating in soldier-robot teams. Simulation methods will be used for laboratory experimentation, which will be used to determine if robots who indicate norm awareness can increase trust and whether differential robot norm awareness can instill calibrated human trust in the robot. The research will conclude with recommendations for implementing norm awareness in robot cognitive architectures. It is anticipated that recommendations for implementation of norm awareness via cognitive architectures in real robots will be platform independent. This research thus supports the overall research integration by providing technical guidance for implementation in cognitive architectures, and supports research conducted under Thrust 2, Capability 2, by addressing mechanisms that can be communicated to a human teammate by a robot, and can support appropriate calibrations of trust in the soldier-robot team.

List of dependencies: Implementation on a robot would require a dependency on the robot's perceptual understanding of the features, cues, and signals from the environment and the team, that convey context for norm

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

understanding. Subtasks under the Thrust 2, Capability 1 *Semantic Scene Understanding in Unstructured Environments* and *Dynamic Scene Understanding in Unstructured Environments* represent the most appropriate dependency mapping for this subtask. On the other hand, full use of the benefits of norm awareness in robots would require a mechanism for the robot to communicate its norm understanding to the members of the team. In addition, a proposed output of this subtask is technical guidance for implementing norm awareness in robotic cognitive architectures. Thus, subtasks under Thrust 2, Capability 1 (T2C1S3) *Shared Multi-Modal Dialog* as well as Thrust 1, Capability 1 (T1C1S2), *Recognize and Anticipate Motion and Activities, Robot Behavior Change Based on Social Context* would benefit from the outputs of this research.

Collaboration: Collaborations and sharing are planned with researchers both outside and within the consortium. Specifically, findings generated from research activities will be shared with outside collaborators experienced in developing cognitive architectures located at Tufts University (Matthias Scheutz). Within the consortium, we plan to share our insights and recommendations with the HRI researchers. Specifically, our findings are applicable to the researchers working on Thrust 2, Capability 1 (T2C1S3) *Shared Multi-Modal Dialog* as well as Thrust 1, Capability 1 (T1C1S2), *Recognize and Anticipate Motion and Activities, Robot Behavior Change Based on Social Context*.

Linkage to previous RCTA BPP sub-tasks:

This research extends the prior work conducted under RCTA BPP 2015-2016 *H1.3 Transparency and Trust in Soldier-Robot Teams,* by exploring signals and communications from the robot that can influence soldier trust that are different from, but build upon the transparency mechanisms (e.g., state projections, SA related information) investigated in the prior subtask. In addition, this work is complementary to the research on trust being conducted under the RCTA BPP 2017-2018 T2C2S3B *Establishing Appropriate Human-Robot Trust through Empirically Validated Training Techniques: Investigating Mental Models, Situational Risk, and Trust* which investigates the features of the task and training programs that influence trust.

Expected products: Several verifiable products and deliverables are planned. These include novel empirical surveys and survey methods, a conference paper submission and presentation to the Human Factors and Ergonomics Society Annual Meeting or the ACM IEEE Human-Robot Interaction Conference, technical reports of findings, as well as a paper providing technical guidance for implementing robot norm understanding in robotic cognitive architectures.

BPP 2017-2018		
2017	**Research activities**	**Deliverables**
Q1	Acquiring IRB protocol for research. Survey development	• Novel empirical surveys
Q2	Survey administration and data collection	• Paper submission (e.g., Human Factors and Ergonomics Society Annual Meeting (or comparable conference)
Q3	Data Analysis and follow-up surveys	• Preliminary technical report of findings from survey studies
Q4	Development of human-robot interaction (HRI) experiment and pilot testing	• Presentation at relevant conference • Technical report of findings • Preliminary report of findings from pilot study and technical recommendations for implementation in cognitive architecture

References:

Crisp, C. B., & Jarvenpaa, S. L. (2013). Swift trust in global virtual teams. *Journal of Personnel Psychology,12*(1), 45-56.
(
Feldman, D. C. (1984). The development and enforcement of group norms. *Academy of management review*, 9(1), 47-53.

Ososky, S., Schuster, D., Phillips, E., & Jentsch, F. G. (2013, March). Building appropriate trust in human-robot teams. In *2013 AAAI Spring Symposium Series*.

Phillips, E., Ososky, S., Grove, J., & Jentsch, F. (2011). From tools to teammates toward the development of appropriate mental models for intelligent robots. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *55*, (1), 1491-1495.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

3.6 Thrust 2/Capability 3: Trusted Execution of Verified Missions

3.6.1 Overview

A commander orders a robot teammate, along with another Soldier, to approach a specific intersection, traveling stealthily, and to observe a specific location and provide reconnaissance details. The Soldier acknowledges the command and projects that the task should be feasible, but the robot, based on the environment surrounding the specific intersection, reports back to the commander and to the accompanying Soldier that the task cannot be executed as requested. Utilizing knowledge obtained from its cognitive architecture through the capabilities described under Thrust 2 Capability 2 "Distributed Mission Execution" above, and combining it with software for plan verification developed under this capability, the robot explains to the Soldier the specific reasons of why it is unable to complete the task: A reason could be, for example, that there is an open field without cover or concealment leading to the intersection, thus limiting the robot's ability to plan a route that meets the requirements for moving "stealthily".

The explanation is provided through the multimodal interface via either speech-base communication, a visual display (e.g., chat display w/ visual representation of issues), a combination of the two, or some other form of multimodal communication. This exchange of relevant information is used by the robot so that its decision, plans, and limitations are transparent to the commander and the Soldier, thus, allowing both of them to better calibrate their trust. After providing the commander and Soldier with relevant information about why the task cannot be executed, the robot suggests other alternative paths via the multimodal interface, as well as information regarding why those particular paths were chosen (e.g., the robot associates stealthy with limiting visual exposure and thus suggests a plan to move quickly will also meet the objective of limiting visual exposure).

Based on the information provided by the robot, the commander selects an alternative path that will allow the robot to execute the task. The development of a solution may require that the Soldier selects from a list of potential options or provides clarification to the robot about how to prioritize mission parameters. This dialogue between robot and Soldier has the potential to significantly affect the overall efficiency of the team. By communicating about intended actions and potential issues and then collaborative developing solutions, Soldier-robot teams can perform tasks with greater efficiency, reducing the possibility of mission failures due to misunderstanding or miscommunication.

3.6.2 Relevance to the RCTA and the Army

It has been stated that trust is a prerequisite for autonomy. Certainly, for traditional Soldier teams, trust enables the confidence required for squad leaders to give commands and know that, with limited supervision, tasks will be completed as required by individuals or by, for example, fire teams. Given the parallels to human-only teams in both the dynamic environments and different mission types which Soldier-Robot teams will perform, it is important that trust between teammates is established and preserved to achieve mission success and efficiency. In many cases, trust is a transaction between two (or more) entities, in which the trustor puts him/herself at risk by putting his/her trust into the trustee. Trustworthiness is typically demonstrated, and trust is established, through "transparency" and through experience. In the absence of pertinent experience, transparency relates to providing the trustor with information on the trustee's capabilities, functioning, motivations, etc. that allow the trustor to make a better informed decision on the trustee's motive and competence. For instance, a low-transparency ("opaque") system does not provide the human team member(s) information about its inner workings, capabilities, limitations, operating status/progress, reasoning for decisions, or future projections. In contrast, a high-transparency ("transparent") system provides such information to the team member through documentation, demonstration, training, explanation/observation in example situations, or through direct communication/

conversation during an on-going task. In order to maintain an effective Soldier-Robot unit in which dialogue, collaboration, and assistance allow for a well-executed mission, trustworthiness of the robot is a necessity. By demonstrating that the robot has the capability to achieve high levels of transparency, the Soldier teammate will be more inclined to appropriately trust the robot's actions and decisions. This creates high efficiency interactions in which Soldier-robot teams have a clearly synchronized understanding of executed actions and decisions that have an impact on the overall progress of a specific mission. The following are some key reasons appropriately calibrated trust in robots is critically important to the Army:

- Lack of trust in robots will lead to Soldiers not using them (i.e., disuse) which will result in money, time, and resources lost.
- Too much trust in robots may results in utilizing them for purposes for which they are not well suited (i.e., misuse) which could result in harm to the Soldier, inefficient interactions, or failure to accomplish important missions.
- Appropriately calibrated trust will ensure that human-robot teams interact in the most efficient manner possible which will result in optimal team performance given each team members strengths and weaknesses.

3.6.3 Required research areas

To achieve this capability, research needs to focus on investigating how to establish high-transparency ("transparent") systems (e.g., by providing information to the team member) and verifying that internal algorithms operate reliably and validly. Providing relevant information (e.g., feasible trajectories) or querying the Soldier teammate for clarification about mission infeasibility and/or mission limitations, helps to make system limitations and processes transparent, thereby supporting trustworthiness and improving the outcome of the overall mission. In order to maintain trust during the execution of complex missions, this capability focuses on three distinct yet overlapping research areas: (a) verifiable planning, (b) maintaining correctness and trust through communication, and (c) system design to maintain trust. Specifically, research performed under "verifiable planning" aims to develop efficient planning and control algorithms that ensure that the probability of satisfying the mission is guaranteed or maximized. Research performed under "maintaining correctness and trust through communication" aims to develop predictable robots that are capable of describing and explaining their perception, decisions, and plans during normal task progress and unexpected events. Finally, research performed under "system design to maintain trust" aims to empirically validate means by which appropriately calibrated trust (and potentially mistrust) in the capabilities of a robot are fostered during a mission.

3.6.3.1 Verifiable Planning

Identifying and communicating mission planning options, discrepancies, and alterations are critical factors to successful teaming during complex tasks. This issue involves components related to semantic environmental understanding, natural language understanding, and collaborative planning. The first big challenge that we will pursue is how to perform higher-level missions in a formal, adaptive and computationally effect way in partially known, or unknown but learned, environments. Virtually all previous work in the area focuses on verifiable controller and supervisory logic design, assuming a known environment. Research will pursue the challenge of balancing learning partially known (or unknown environments) and mission planning while providing rigorous guarantees. This approach considers for the first time robot motion planning under temporal logic constraints in probabilistic maps obtained by semantic simultaneous localization and mapping (SLAM). The overall approach for temporal logic planning in an unknown environment needs to handle the following competing objectives during online

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

execution:

(i) explore the unknown environment to discover landmarks relevant to the mission and construct a map,

(ii) follow sensing trajectories in the known portion of the map to reduce uncertainty about landmarks and robot states, and

(iii) decide if and how the logic specification may be satisfied given the current information on the map.

Communicating the learned environment effectively can be achieved via natural language. Natural language interaction provides a rich, intuitive, and flexible medium for humans and robots to communicate mission objectives and situational awareness, facilitating teaming and peer-to-peer tactical behaviors. Significant progress has been made in the development of probabilistic models capable of grounding rich linguistic data (possessing composition, hierarchy, monologues etc.) with complex symbolic representations (like relations, abstraction, constraints etc.) describing the varied workspaces the robot may be operating in. Contemporary models can efficiently ground instructions like "put the pallet on the truck" or "pick up the nearest block on the table" or "go to the fire hydrant behind the traffic cone".

The grounding or interpretation of a language command is intimately tied to the robot's knowledge about the world. State of the art models embed the input instruction in a world representation that includes entities like object, regions, and constraints that can be currently perceived by the robot. This excludes knowledge related to behaviors/events that may have occurred in the past (interactions between objects and agents), as well as declarative facts that may be stated by the human ("Post A is guarded by a soldier", "this is my bag" etc.). Further, the symbol grounding process is treated independent of the process for acquiring the world model (typically delegated to a perception or semantic mapping system) is treated independent of the symbol grounding process. However, though the two problems are fundamentally coupled. For example, consider the instruction "pick up the item that the soldier put down" that includes both imperative and declarative components which contemporary approaches cannot model.

We will address this problem by formulating a unified probabilistic model that jointly expresses the robot's knowledge about the world and infers the groundings (objectives and constraints) for intended actions that it should execute. Research will expand the space of symbols to include temporal events (higher-order interactions) as observed by a visual perception system and incorporate factual information or mission relevant statements that may be communicated by human team mates via declarative language commands. Joint inference will allow the symbol grounding model to determine intended actions by appropriately accounting for uncertainty in its belief over the world knowledge, enabling grounding of language instructions that involve simultaneous grounding of imperative and declarative statements.

Finally, interactive mission planning research will integrate the cognitive models representing human task procedures in a computational form suitable for robots (but understandable by humans), with the mission planning work used to decompose mission-level tasks received from the Soldier into tactical behaviors. The former assumes a procedure communicated all at once, which then is gradually augmented with learned experiences to improve individual decisions. The latter assumes a self-contained generative process from a one-time user input, also improved over time by learning from the outcomes of generated plans. This work would integrate the two processes while transforming mission planning from a batch process to a highly interactive one in which the robot learns to decompose missions by combining its own reasoning and previous experience with explanatory input received as needed from human teammates, who would gradually teach and train the robot to fulfill the original "boot camp" vision.

3.6.3.2 Maintaining Correctness and Trust through Communication

Being able to explain one's own behavior is one of the key capabilities to be a competent teammate. For autonomous systems, this requires transparency and easily understood communications (i.e., visually salient or utilizing natural language techniques) about not only what the intended behaviors are, but also why those behaviors have been chosen. One of the tasks related to this will be to investigate an approach for a ground robot to generate explanations for its navigation paths. Current path planning algorithms focus on single factors for decision making. For example, minimizing travel distance can be expressed as finding a "shortest" path, whereas the objective of an algorithm for minimizing the probability of collisions can be labeled as finding a "safest" path. Instead, the proposed learning framework will learn to map descriptive and explanatory labels to a system's route choices. One of the difficulties in employing this method will be learning the mapping between language description and primitive concepts related to navigation behavior. The proposed approach to tackle this challenge builds on prior work on imitation learning for language understanding in outdoor navigation, where the focus was on learning the mapping from language to actions. In the proposed work, the focus will be on the reverse direction from actions to language. To generate explanation for a decision that has been made according to some black box algorithm, a discriminative approach will be used, i.e., to explain the chosen decision by comparing that to alternatives. Following the maximum margin principle, the problem can be formulated as an optimization problem where the loss is the difference between the costs of demonstration and generated paths. In the proposed work, we will collect a corpus describing the features in natural language to learn another mapping between language and features; such feature-level description in conjunction with behavior-level description will be used to generate explanations for chosen paths.

Another task is to determine the influence of Situation awareness-based Agent Transparency (SAT) on Soldier-Robot Team performance. More specifically, the SAT model has been expanded to incorporate Teamwork Transparency (robot's understanding of its roles and responsibilities). In the pursuit of this goal, bidirectional transparency is being prioritized to better anticipate the needs of mutually interdependent human-robot teams. This approach can be expanded upon in the future with the addition of controls, so overt collaboration, as well as mutual transparency, can be examined. This will address gaps in transparency research, including transparency into the thought process of a robot's human teammates and the typically unidirectional nature of communication. As such this research addresses the peer-to-peer tactical teaming thrust (specifically two-way dialogue with explanations) and the proactive, fine-grained contextual understanding thrust (fine grained recognition of humans and social cues).

3.6.3.3 System Design to Maintain Trust

In near-future Soldier-robot teams, human trust in the robotic teammate(s) will be needed to support successful peer-to-peer teaming. Complicating this picture, however, is the notion that humans tend to hold overly presumptuous expectations of robot performance, and trust in robotic partners is highly sensitive to violations of these expectations. As robotic teammates cannot yet perfectly replicate the capabilities of their human counterparts, expectations of robots are likely to be violated, and as a result, trust diminished or potentially lost. It is important to investigate means by which features and capabilities of the robot can help to mitigate such violations, and instill appropriate calibrations of trust in the robot. Because trust is an important predictor of reliance (and over reliance) on automated systems, appropriately calibrated human trust in robotic teammates will be especially important in situations that are novel, when team performance depends on developing an alternative plan, or when the robot is not able to carry out the mission as well as originally expected.

A goal of the research under this capability is to empirically investigate means by which appropriately calibrated trust (and mistrust) in capabilities of the robotic system can be fostered during the course of an op tempo team task. This research seeks to investigate what we term "trust signatures" (e.g., metacognitive

information sharing from the robot, contextual cues, communication of detected norms derived from environmental perception, estimations of reliability) that may convey trust and mistrust. These trust signatures may be able to help a human calibrate their trusting, and adjust their reliance on the robot in ways that appropriately support the team's completion of the mission.

We will examine the dimensions of trust, cognitive workload, and team performance as they apply to human-robot interaction (HRI) through the communication medium of mutual chat features. This effort will incorporate multiple users and multiple robots in both live and simulated mission-specified environments. The assessed chat communications will involve two-way dialog analysis. More specifically, we will examine expressed operator trust based on a sequence of problem-solving exercises, engaged with the robotic teammate.

The effect of chat content and frequency as a reflection of robot transparency level will also be evaluated. This series of experiments would provide a platform for collaborating with the on-going evolutions in robotic intelligence and perception by providing robot perception and response requirements for an operator. We seek to distill which elements of robot communication provide the best foundation for trust establishment and sustenance.

3.6.4 Demonstration plan

If funded, the experimental demonstrations related to 'Trusted Execution of Verifiable Missions' will center around integrating components of vehicle perception (i.e., gathering data about the environment), vehicle intelligence (i.e., interpreting sematic features), bi-directional communications (i.e., natural language grounding; effects of transparency), collaborative teaming (i.e., interactive mission planning), and calibrated trust (i.e., accurate team mental models). Demonstrations will focus on human-in-the-loop interactions, engaging in complex communication tasks relevant to mission success. Specific experimental demonstrations will show how algorithmic perception and mapping solutions provide input which can be transitioned to sematic labeling for collaborative planning tasks and the demonstration of multimodal communication techniques to develop shared situation awareness and increased system understanding. Calibrated trust research will provide input, into each of the communication techniques, designed to enable robust perceptions of trust that facilitate the efficiency in which the interactions take place in the above human-in-the-loop demonstrations.

We envision three milestones for integrating the technical elements for capability 3:

- Demonstrating Plan Verification (CY17): This first milestone for this capability will demonstrate the ability to take a fixed world model and a plan created by the cognitive architecture, and show that infeasible plans can be identified, and feasible plans confirmed by the verification system, increasing trust from the human team-mate.
- Demonstrating Verification and Trusted Replanning (C18): This second milestone builds on the ability to report whether a plan is feasible or not, and replan with the human team-mate in a transparent manner.
- Demonstration of System Designed to Maintain Trust (CY19): The final milestone demonstrates increased human team-mate trust as a result of both the underlying algorithmic plan verification and the overall system design for interaction.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.0 Thrust 3: Mobile Manipulation

4.1 Overview

Robots must be able to interact physically with the environment, as well as move through the environment. To access indoor, subsurface, and super surface locations, they need to ingress through various openings, which may involve opening doors, moving aside curtains, or climbing through windows or holes in walls. They may also need to negotiate stairs or ladders, which may involve grasping handrails or ladder rungs to help with balance or movement. Robot movement may be obstructed by debris, furniture, or other man-made objects, which therefore may have to be cleared out of the way. Mission functions may include manipulating objects, such as removing IEDs, retrieving objects, opening drawers, closets, or vehicle compartments to search or secure an area, handling supplies as part of automated logistics operations, or participating in various medical or casualty evacuation procedures. Many of these actions will require manipulating objects while the robot itself is moving, which entails coordinated mobile manipulation for systems with many degrees of freedom. Some actions will require moving or manipulating heavy objects, which may require actions that exploit bracing or dynamic behaviors.

The current state of the art in mobile manipulation involves using a large amount of prior knowledge of what objects are to be manipulated, where those objects are, what are their shapes and inertial properties, and how they should be grasped for a given task. Payloads often are relatively light (e.g. < 5 lbs), operations are not robust to failure, and specifying actions requires considerable input from human operators, hence large communication bandwidth to convey situational awareness to the operators. The scientific challenges involve relaxing the amount of prior knowledge required, enabling manipulation of heavier objects, reducing the amount of operator input required by increasing onboard autonomy in perception and planning, developing fault tolerant mobile manipulation behaviors that can recover from failures, and learning about properties of objects through experimentation, in order to better manipulate them.

We group the required research in this thrust into two areas of capability focus:

1. Manipulation in cluttered human spaces, where the general goals involve developing methods that cope with increasing uncertainty in the geometry of objects and their spatial arrangement;
2. Manipulation of unknown, heavy objects, with uncertain inertial properties, which requires more force control that the first capability.

The challenges and capability focus areas in this thrust are illustrated in Fig. 1.

Experimentation and demonstration in this thrust will use the RoMan platform, with two arms on a tracked mobility base, and the Robosimian platform, with four limbs that are used for mobility and manipulation, as well as other platforms available at member institutions for early development and testing of algorithms.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

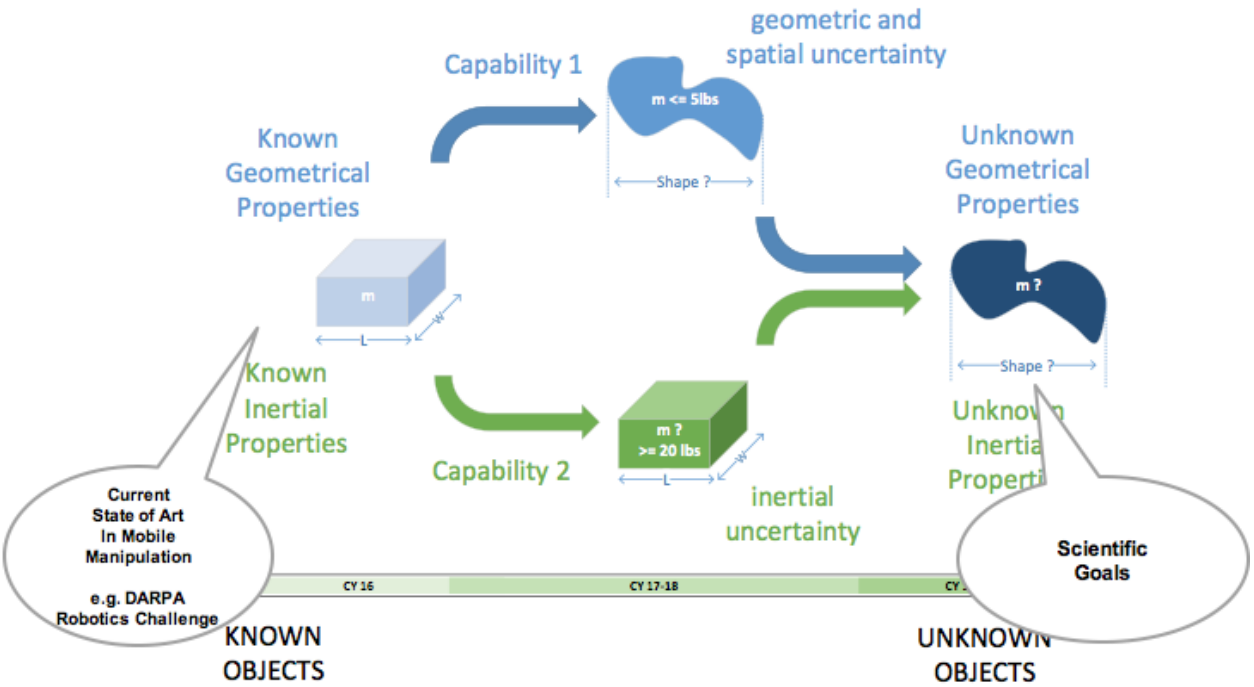

Figure 3 Problem Space and Focus

4.2 Thrust 3/Capability 1: Manipulation in cluttered human habitable spaces

4.2.1 Overview

This cluster is about manipulation in the presence of geometric and spatial uncertainty. The aim of this cluster is to develop a robot that reasons about unknown objects and their physical properties, and decides how to best interact with different objects to achieve a specific task. Given the task and a set of perceptual observations, the robot needs to identify potentially relevant objects, figure out how objects can be grasped (where, with what preshape), and decide what type of interaction to use (lifting/moving, pushing, pulling …) to achieve its task. In the process, the robot may ask a human teammate for further instructions and demonstrations. Planning, learning and integrating the two together cut across all these problems, enabling the robot to improve its performance over time and to adapt to new scenarios by building models on-the-fly while incorporating the power of model-based reasoning.

Cluster 1 will target the following capabilities:

- Placing objects such as sensors in confined spaces, possibly invisible to humans.
- Reliably identifying and localizing objects that are heavily occluded by clutter
- Recognizing and moving through ingress openings, including figuring how to manipulate and go through such openings as doors with curtains, sliding doors, push/pull doors with different sorts of handles.
- Identifying movable obstacles, then pushing, displacing, or walking over obstacles.
- Searching and retrieving specific items, including finding objects of a given class, and opening storage elements or containers (such as drawers and cabinets) to search them.

4.2.2 Relevance to the RCTA and the Army

Large human habitable spaces pose challenges for robots in terms of reaching a goal. While the Mobility clusters deal with moving through confined spaces, this cluster aims to support mobility in cluttered spaces by pushing or lifting objects out of the way, and identifying/actuating ingress openings. The robot will be able to manipulate objects commonly found in human habitable spaces, to perform behaviors related to manipulation such as rearranging objects to enable clear identification in a search. [May consider using standard issue equipment used in searches (chemical detector, sample swab, etc).] Other beneficial related search capabilities include opening cabinets and containers to find items of a specific class, such as keys, or electronic storage. The robot will be able to place and extract sensors to/from areas not easily accessible to humans, or provide instructions on how to retrieve the sensors it has placed.

4.2.3 Required research areas

The research areas are clustered around the three fundamental problems:

- Where are objects?

Brittleness of perception is currently a huge bottleneck for reliable manipulation in cluttered spaces. We seek to develop approaches that overcome this brittleness. Specifically, subtasks targeting this research question aim to develop algorithms that reason over multiple hypotheses from perception, develop approaches that are capable of finding and estimating the pose of objects via tactile sensing, and develop methods for incorporating context and object relationships to improve the quality of perception hypotheses.

- How to grasp objects

Grasping novel objects of arbitrary shape in a way that enables the robot to manipulate them according to a task is a challenge. In particular, the challenge is how to bridge the gap between multi-modal perception of the scene with the grasp reasoning. The subtasks targeting this research question will study how semantic understanding and perception of objects and of the scene more generally can be combined with the reasoning about the geometric constraints on the scene and the robot's arm and hand.

- How to manipulate objects?

Current state-of-the-art approaches to manipulating objects rely on having perfect, a-priori built, models of objects. This makes the application of autonomous mobile manipulators limited purely to the tasks that were pre-programmed and also makes it brittle to any uncertainties in the models of the objects the robot has to manipulate. The goal of the subtasks targeting this research question is to make autonomous mobile manipulation capable of going beyond what has been pre-programmed into them and robust with respect to uncertainties. The general theme is to integrate model learning with model-based reasoning. Specifically, these subtasks develop algorithms for planning manipulation with multiple hypotheses about object models, develop algorithms that learn manipulation policies on-the-fly via interaction, and develop methods for learning manipulation behaviors from human demonstrations.

4.2.3.1 Active tactile learning of complex objects

Robots encounter situations where objects of complex shape and perhaps unseen before have to be grasped, removed out of the way, or inspected. In many situations, visibility and illuminations conditions do not allow for a visual recognition of the object or even for applying visual affordances. A tactile interaction with the object will not only enable a recognition of the object but also directly facilitate grasping and manipulation instead of first seeing and then grasping. Preliminary results using a 3-finger hand and an invariant geometric representation has shown promising classification on objects learnt virtually in simulation and physically by touching. Because of the tediousness of the sensing, inference has to rely on the fewest probes possible. This can only happen by solving the following two fundamental challenges: i) learn a model for grasping (nor first classification and then grasping) using examples of dynamic simulation. We will apply a shallow convolutional network with input a 3D volume occupied with

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

touches and output heatmaps of the grasping points. ii) apply an optimization framework for the best pose and grasp, and iii) apply an active probing strategy by minimizing simultaneously the grasp failure and the number of probes during inference.

4.2.3.2 Geometry learning of rigid and deformable objects

Recent advances in 3D reconstruction and recognition have enabled the reconstruction of objects and humans by supervised detection and learning from landmark points. Today, we have access not only to millions of images of object classes but also to corresponding 3D models. However, even if we follow an end to end deep learning approach we need a supervised correspondence between 2D keypoints on the image and 3D keypoints on a geometric model. Such a correspondence is challenging even for a human and has limited existing datasets to simple furniture and vehicle classes.

In preliminary work [1,2], we have learnt to detect distributions of 2D keypoints and established their correspondence to 3D keypoints through a probabilistic approach that looks at the 2D positions as latent variables. But the challenge of whole possibly unstructured models defined by surfaces or volumes rather than keypoints remain open. We propose to study the learning of volumetric representations of objects given a prior probabilities about the context (an obstacle on the way of the robot) or the class of the object. This could be also achieved as in [3], we addressed the reconstruction as the consistency of multiple views of different instances of the same class. Rigid or deformable objects can be represented through a sparse representation of shape parameters that will be learnt directly on the volumetric space rather than on keypoints.

4.2.3.3 Planning with Uncertainty in the Model and Perception

This task will extend our current work on scalable planning techniques that deal with multiple hypotheses in perception of objects and in the models of the objects. In particular, our latest work was on Deliberative Perception that combines Search and Learning to deal effectively with multiple hypotheses in perception, searching for a plausible scene that best explains the observed data. For the coming years, the plan is: a) extend our work to explicitly reason about and search through multiple hypotheses about environmental conditions under which the objects are being detected (for example, using deliberative perception to understand that the strong pointed light source behind the robot is the reason why the object in a captured RGBD image looks so different from the learned model), b) develop planning methods that combine Deliberative Perception with taking actions that gather more information about the scene (re-positioning the robot, moving objects, etc.) to better understand the scene, c) continue to integrate these planning algorithms onto a Roman/RoboSimian platform (and later a legged platform) to provide the functionality of finding an object of interest, grasping and manipulating it.

The above work integrates with Sisir's work at JPL on the low-level control of Roman platform, with Kostas's work at UPenn on generating perception hypotheses about the object, with GDRS's work on the Roman platform itself (Terence) and overall integration (Ed Weller), and with RR's research (Lenny Sapronov) on navigation to the poses generated by our planner. The demo is: a Roman (or similar) platform opening a door into the room, searching the room to find the object of interest, opening previously unseen articulated objects such as cabinet doors and drawers and grasping the found object to bring it back.

4.2.3.4 Discovering Task-Related Implications of Object Arrangements

Considering context and the relationship between objects and events in the environment is a key component to developing a level of situational awareness necessary to accomplish abstract goals or effectively share information with human and non-human teammates. In this research effort we wish to extend our previous work on exploiting priming in cognitive robotics to enable the robot to use perception

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

and cognitive reasoning to understand the implications of spatial arrangements of objects and humans in the environment to a specific task. We focus on two tasks related to the 2019 capstone. The first task (Step 3) involves finding a secured observation position near a designated building. While state of the art perception algorithms can identify specific objects in the vicinity of the building, cognitive reasoning can identify ad-hoc and natural arrangements of those objects that could impact the choice of an observation position. Similarly, in Step 6d, identifying dangerous items in the building is a combination of state of the art perception supported by an efficient search strategy. Cognitive reasoning can contribute to this search strategy by identifying potential hiding spots for items relevant to the mission. We will leverage declarative and procedural memory and learning mechanisms in the ACT-R cognitive architecture to create a feedback loop between perceptual and cognitive systems. The cognitive system (ACT-R) – will encode environmental context (e.g., features, object/scene labels) provided by semantic perception and augment that context via contextual mechanisms (including goal-relevant information) and prior knowledge. The augmented context will be shared with perception to help with ambiguity resolution and/or goal-directed semantic mapping to support proactive, fine-grained contextual understanding. The input to perception could take the form of context-dependent expectancies, or of a probability distribution over multiple leading hypotheses. In turn, perception output will result in learning at the cognitive level, including addition of new knowledge and adaptive tuning of parameters involved in contextual inference.

4.2.3.5 Generalizing grasps across objects and tasks

This subtask addresses robot interaction with novel objects. We proceed by integrating (1) task constraints that emerge from an understanding of scene semantics with (2) gripper constraints that are defined by the geometry of the scene and the gripper and robot reachability.

Semantic scene understanding (1), through object classification (CNNs), helps the agent find grasps that are suitable to the task, by defining how a certain object class should be grasped to achieve a specific task. Semantic understanding also allows the robot to integrate information related to the weight, material properties (e.g. coefficient of friction) and rigidity of objects, which helps the agent compute how it can establish safe and robust hand-object bindings or how to push objects.

Geometric constraints (2 above) bias the robot towards gripper configurations that are compatible with the geometry of the scene. From the data coming from the robot's depth sensors, we identify points where the shape of the underlying object is compatible with the shape of the gripper. By the same process, we also exclude grasps that would collide with visible surface elements. We optimize motion with respect to the three constraints listed above, to take advantage of degrees of freedom offered at the object level (several tools available), grasp level (e.g. equivalence of grasps around symmetrical object), and embodiment level (in the robot's kinematic chain).

4.2.3.6 Lifelong Improvement of Robot Capabilities and Its Robustness

In this task, we will continue to work on planning methods such as Experience Graphs we developed in previous years of CTA that improve the performance of the robot behaviors via experience and demonstrations. In particular, we will continue on these methods in the context of complex full-body manipulation. In the coming years, we will be looking at how to combine Model-based Planning and Model-free Learning in a principal fashion. For example, imagine a manipulation platform that learns how to open a spring-loaded door from several demonstrations without having any prior model of a spring-loaded door. We will study how these skills can be incorporated with model-based planning for the platform to deal with situations such as a spring-loaded door in a cluttered space. Such scenarios require planning a feasible full-body motion through the clutter and combining it with learned skills to accomplish opening the door. On the applied side, we will demonstrate this work on a Roman platform and then on a

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

legged robot once it becomes available. For the legged robot, we will be looking at such tasks as teaching it how to go upstairs.

4.2.3.7 Techniques for fast human-like manipulation

In many tasks of relevance to the RCTA such as rummaging through a drawer or inspecting the contents of an automobile, the slow and robot-like way is to grasp each object one at a time, pick it up, and then put it down. The fast and human-like way is to probe, to push things around, to flip through things. This subtask will perform basic research to develop fast and effective techniques for such tasks. This includes some quick informal studies on how humans perform such tasks, followed by development of (1) pushing techniques for objects in clutter in a largely planar environment, and (2) probing techniques to move objects in clutter in 3D environments. Principled modeling based on applied mechanics, augmented by machine learning, is central to our approach.

4.2.3.8 Learning manipulation actions from demonstration

Generating motion plans to perform complex manipulation plans autonomously in real-world scenarios mandates a tightly coupled reasoning across symbolic states and continuous spaces. Most symbolic planners rely on the knowledge of pre-conditions and effects of individual manipulation actions that are manually defined and fixed. However the correct mapping from the continuous state space to the pre-conditions and effects of a manipulation action is unknown. Therefore, we seek to develop a learning algorithm that can learn these pre-conditions and effects given a few (sequences of) manipulation action demonstrations and the context in which they have been recorded. This representation should also allow for updating the pre-conditions on-the-fly through its experience, i.e., if the robot encounters a scenario in which the pre-condition of a manipulation action was faulty the robot needs to be able to update this pre-condition to reflect the encountered scenario. As a result, the robot will be able to expand its repertoire of manipulation actions through learning from demonstrations and to chain those using higher-level planning.

4.2.3.9 Dexterous Manipulation of Articulated Objects from Visual and Lingual Cues

If robots are to work effectively alongside soldiers in unstructured environments, they must be able to manipulate the articulated objects (e.g., doors, drawers, levers, car trunks and hoods, pliers, etc.) frequently found in environments built by and for humans with a level of dexterity that approaches that of their human partners. Traditionally, this interaction is predefined in the form of manipulation policies that must be manually specified for each object that the robot may interact with, thereby limiting the space of policies to simple objects drawn from a small set. These models fail to generalize and in practice, the robot typically regresses to remote teleoperation, whereby the soldier cautiously controls the low-level motion of each of the robot's degrees-of-freedom. Mobile manipulators would benefit from the ability to acquire the skills necessary to interact with articulated objects opportunistically from observation. Such a capability would facilitate the 2019 Capstone, by allowing robots to open doors that may block the entrances/exits of the target building and each of its rooms, inspect the contents of drawers, and open the doors, hood, and trunk of automobiles.

We propose to develop of a method that enables robots to learn to manipulate articulated objects from visual demonstrations and spoken descriptions of their motion given in situ (i.e., without the need for fiducial markers). Specifically, we will design a multimodal learning framework that exploits the complementary nature of visual and spoken observations to acquire the kinematic models that govern the motion of articulated objects. Visual observations of object motion can be used to learn models that express the rotational, prismatic, and rigid relationships that govern the motion of its individual parts. However, relying solely on visual cues makes the method sensitive to occlusion and errors in segmentation that occur as a

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

result of clutter or a lack of visual features. Natural language descriptions (e.g., accompanying the demonstration) offer a flexible, bandwidth-efficient means that soldiers can readily use to convey knowledge of an object's operation. These descriptions help to overcome the limitations of visual observations, e.g., by providing cues regarding the number of parts that comprise the object or the motion type (e.g., rotational) between a pair of parts. The learned models then allow the robot to manipulate the objects on-demand, e.g., given a spoken command issued by a soldier. The methods will be developed with and integrated on Walter's Rethink Baxter manipulator and as available, a RoboSimian or a Clearpath Husky UGV equipped with a Universal Robots UR5 arm (for which Walter has submitted a DURIP proposal).

This proposed subtask builds off of previous and ongoing work by Walter under the RCTA. H4-2: "Human-Facilitated Learning of Shared Semantic World Models" developed a method for learning semantic environment models from natural language descriptions. H3-3/4: "In-Situ Monologue Comprehension for Faster Mission Tempo" lead by Walter and Howard developed language understanding capability that will be used to interpret the soldier's descriptions of an object's motion and to then interpret his/her commands to manipulate objects. Capabilities developed under H7-4: "Learning to Handle Objects from Human Annotations", a new task that uses human supervision to learn grasping strategies, will be used to grasp objects. Meanwhile, the subtask complements previous and ongoing subtasks lead by RCTA members including Barber (multimodal display), Boularias (learning spatial relations for objects manipulation), Daniilidis (perceiving object geometry and semantics), Hebert (object recognition), and Karumanchi (RoboSimian support). The result of the subtask will be a demonstration of the articulation learning framework in the form of a video (involving Walter's Baxter or UR5-equipped Husky platform) or a physical demo on RoboSimian (as available). The subtask aligns with two technology thrusts: **Fast, Adaptive, Online & On-the-fly Learning** (i.e., Learning of techniques for dynamic locomotion, rapid grasping, and manipulation; Fast on-line interactive learning using all available sources) and **Interaction with Complex 3D Environments** (i.e., Human-like manipulation in complex situations).

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3 Thrust 3/Capability 1 Subtasks

4.3.1 T3C1S1A - Active tactile learning of complex objects

PI: Daniilidis, Kostas (UPenn), kostas@cis.upenn.edu, (215)898-8549

Overall objective/goal:

Robots encounter situations where objects of complex shape and perhaps unseen before have to be grasped, removed out of the way, or inspected. In many situations, visibility and illuminations conditions do not allow for a visual recognition of the object or even for applying visual affordances. A tactile interaction with the object will not only enable a recognition of the object but also directly facilitate grasping and manipulation instead of first seeing and then grasping. Preliminary results using a 3-finger hand and an invariant geometric representation has shown promising classification on objects learnt virtually in simulation and physically by touching. Because of the tediousness of the sensing, inference has to rely on the fewest probes possible. This can only happen by solving the following two fundamental challenges: i) learn a model for grasping (not first classification and then grasping) using examples of dynamic simulation. We will apply a shallow convolutional network with input a 3D volume occupied with touches and output heatmaps of the grasping points. ii) apply an optimization framework for the best pose and grasp, and iii) apply an active probing strategy by minimizing simultaneously the grasp failure and the number of probes during inference.

Proposed research:

We develop an active algorithm for selecting a sequence of end-effector (3-finger hand) poses for object class recognition, and shape estimation by touch only. The recognition part uses an invariant tactile object descriptor based on the distribution of triangles formed by probing points on the fingers. The main weakness of the previous work is that it requires hundreds of systematic contacts to classify an object. Here, instead of covering all parts of an object, we strategically select a sequence of end-effector poses that will increase the probability of recognition and decrease the error in shape estimation.

We formulate the problem as a Markov decision process (MDP). The states are observed using training data, which associate observations with actions, defined as relative transformations between pairs of end-effector poses. The objective is to find a sequence of actions that maximize the probability of recognition and minimize the movement cost. The optimal sequence is estimated by Monte Carlo tree search (MCTS). In the 2nd year, we will enhance our approach with actual grasps and not just poses so that we combine class/shape with graspability values.

At the beginning of the process, a random action is selected for an observation to initialize the object model. While advancing through the states, we do not make any robot movements to observe actual data, as it would make the search impractically slow. Instead, we rely on the training data and simulated grasps. Over time, the histogram becomes more filled and is expected to approach the true histogram computed at training.

To construct a tree from the MDP, each state in the MDP becomes a node in the tree, and each action becomes an edge. A sequence of actions is defined by a path from the root to the bottom of the tree. The maximum depth of the tree is the number of actions we wish to limit.

Besides the state, a reward is also stored in each node. In an MCTS, the reward is multi-armed, one reward for each available arm (action, or edge) at a node. Throughout the tree search, a node is likely traversed many times, each time taking a different edge with its own reward outcome. At the end, the multi-arm rewards are used to determine the most rewarding action from each node, recursively yielding the best path down the tree.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Our goal is to find a path of actions from the root of the tree, such that the cost of movement and the estimation error is minimized. The MCTS uses a tree policy to search for actions that maximize reward, or equivalently, minimize regret. For a multi-arm problem, there is an exploitation-exploration dilemma. In other words, if the arm with the maximum reward at a node is always chosen, then the other arms will never be given a chance to be explored. This can miss arms that start out with low reward but can potentially be an overall optimal arm at the end.

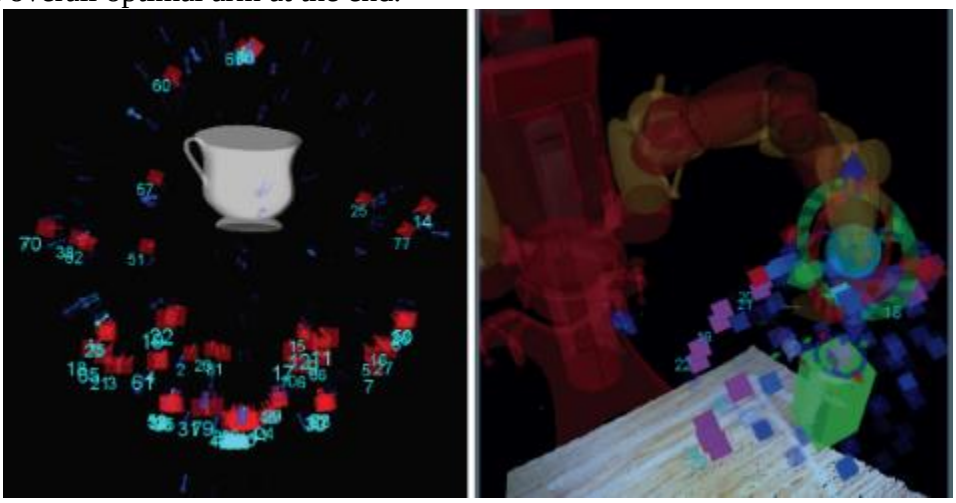

How the proposed research supports the thrust/capability:
Visibility conditions due to clutter or illumination decimate the performance of vision-based techniques in manipulation. In addition, grasp probes can help decide about the graspable parts of an object.

Scientific challenges:
The main scientific challenges addressed by the proposed research are in active perception: Tactile measurements are costly and have to be minimized by following an appropriate policy that at the same time minimizes the error of misclassification and pose and shape estimation. We study Markov Decision Processes and we will apply Monte Carlo Tree Search.

Collaboration: T3C1S1B

Connections to the past work under R-CTA:
P3-2, P3-5: We have demonstrated in our past work how to classify a 3D object in the dark and classify it given a number of class models.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.2 T3C1S1B - Planning with Uncertainty in the Model and Perception

PI: Likhachev, Maxim (CMU), maxim@cs.cmu.edu, (678)592-4601

Objective and Benefits:

The high-level goal of this subtask is to enable proactive understanding of the environment. In particular, it explores principal approaches to tight coupling of decision-making/planning with perception and world modeling. This coupling has inherently massive amounts of uncertainty that planning needs to reason over in a scalable way. Consequently, the subtask develops planning approaches that deal effectively with this uncertainty while providing real-time performance. The work is primarily done in the context of mobile manipulation under uncertainty in the model of objects that are being manipulated.

Specifically, the research we propose in this subtask aims to increase the robustness of the perception, which has frequently been one of the main culprits of brittleness in mobile manipulation, by explicitly reasoning over multiple hypotheses generated by perception to find the best explanation of the observed scene as well as taking actions to disambiguate the uncertainties that are critical. The benefit this research will provide is that it will aid in making it possible for the mobile manipulation platforms such as Roman platform to identify objects within clutter and to robustly manipulate them.

Scientific approach:

This task will extend our current work on scalable planning techniques that deal with multiple hypotheses in perception of objects and in the models of the objects. In particular, our latest work was on Deliberative Perception that combines Search and Learning to deal effectively with multiple hypotheses in perception, searching for a plausible scene that best explains the observed data [1, 2]. For the coming years, the plan is: a) to extend our work to explicitly reason about and search through multiple hypotheses about environmental conditions under which the objects are being detected (for example, using deliberative perception to understand that the strong pointed light source behind the robot is the reason why the object in a captured RGB-D image looks so different from the learned model), b) develop planning methods that combine Deliberative Perception with taking actions that gather more information about the scene (re-positioning the robot, moving objects, etc.) to better understand the scene, c) continue to integrate these planning algorithms onto a Roman platform to provide the functionality of finding an object of interest, grasping and manipulating it.

The main scientific challenges addressed by the proposed research are in investigating scalable algorithms that enable Deliberative Perception. Reasoning over multiple hypotheses leads to huge search-spaces. We have previously shown how reasoning over multiple hypotheses in perception can be formulated as a tree search [2] and how our recently developed Multi-Heuristic A* search [3] can be used to efficiently search this tree for close-to-optimal solutions. However, to support reasoning over environmental conditions, we need to be able to scale up the search dramatically with the size of the search tree. To this end, we propose to investigate how the tree search can be sped up without losing guarantees on near-optimality. One of the directions we propose to investigate is massive parallelization of Multi-heuristic A* search. The main challenge then is in developing parallelization techniques which don't sacrifice near optimality, which is non-trivial since search is inherently a sequential operation.

The proposed work contributes to the thrust 3, Mobile Manipulation, capability 1, Manipulation in cluttered human habitable spaces. It is one of the critical components in the ability to complete missions like retrieving an object, which is part of the Integrated Research Assessment and Capstone. Specifically, it is being used

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

to help with the manipulation of articulated and non-articulated objects, reasoning about how to position the platform in order to grasp an object and how to deal with uncertainty in perception. Some of the work developed in this subtask already ran on HDT/Husky platform in 2014/2015 and on the Roman platform in 2016. We will continue to port our work onto the Roman platform in 2017 and 2018. This research links to the research by others on perception for manipulation, and in particular, generating multiple hypotheses. This includes subtasks T3C1S2A by K. Daniilidis and T3C1S1C by M. Fields. The work also links to the development of the World Model that is supposed to maintain multiple hypotheses generated by perception. In addition, the Deliberative Perception we have developed has already been transitioned to other programs. In particular, a team participating in Amazon Picking Competition in July 2016 ran the Deliberative Perception framework for identifying rigid objects. We plan to transition the outcomes of the proposed research and use it again in 2017 Amazon Picking Competition

The proposed research is somewhat related to the work done in T1C1S4D by L. Navarro-Serment. However, the difference is that their work primarily targets the problem of *positioning* sensors to minimize the effect of environmental conditions such as getting sun into the camera as opposed to *reasoning* about the observed scene to parse through multiple perception hypotheses. Further, they primarily target the domain of perception and tracking of humans in outdoor environments and use human behavior examples for learning effective sensor placement strategies, whereas we target the problem of indoor object detection. To support the proposed research, the World Model needs to contain RGB-D data obtained from Kinect (or a similar device). This data is already being collected on the Roman platform for the purposes of perceiving an object. The Deliberative Perception module will need access to it. In addition, the current message from Perception module (developed by K. Daniilidis's subtask) will need to be extended to contain multiple hypotheses instead of a single, most likely one. The current pose estimation of the Roman platform is adequate for our purposes. Similarly, the feedback on the Roman arm is adequate based on our experiments in 2016. We don't anticipate that any other additional information from the World Model will be required.

List of dependencies: The proposed work integrates with S. Karumanchi's work at JPL on the low-level control of Roman platform (subtask T4C1S1B), with K. Daniilidis's work at UPenn on generating perception hypotheses about the object (subtasks T3C1S1A and T3C1S1C), with the work of M. Fields's group at ARL on generating perception hypotheses based on object relationships and context (subtask T3C1S1C), with GDLS's work on the world model that fuses data and will maintain hypotheses, Robotics Research's work on adaptive behaviors (subtask T4C3S1D), in particular navigation for search which consumes the poses generated by one of the planners that we built and the work by Ed Weller's group at GDLS on the overall software integration related to the planning for Roman platform (subtask T4C3S1C).

Collaboration: As part of this task, we have been collaborating with K. Daniilidis's (T1C1S1A) group on integrating perception with planning, with JPL on integrating planning with the control of the mobile manipulation platform, with RR on integrating planning for mobile manipulation with navigation and GDLS on integrating planning with the world model structure. As part of these collaborations, researchers from JPL, RR and GDLS have visited my lab several times, and we plan to continue this process.

Linkage to previous RCTA BPP sub-tasks:
This work relies heavily on the work we have been doing under the subtask I3-1 in the previous years of R-CTA (figure 1a). In particular, the concept of Deliberative Perception was developed in the last year of the program and shown to be effective on PR2 robot. In addition, under the same subtask, we have been

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

developing techniques for planning with actions that disambiguate uncertainty and have shown some of these algorithms in the context of "Retrieve Object" mission in the Integrated Research Assessment (figure 1b).

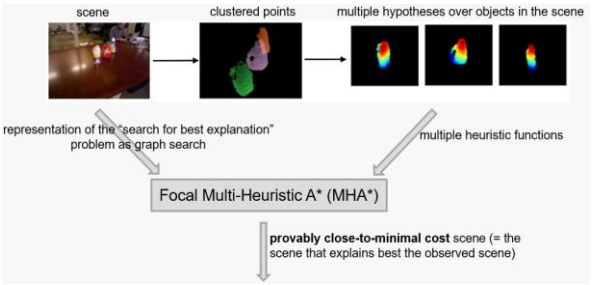

(a) Deliberative Perception framework

(b) "Retrieve Object" scenario integrates two of our planners

Figure 1.

Expected products:

The deliverables include papers at conferences such as RSS, AAAI, ICAPS, etc. The researched topic of Deliberative Perception is also a topic of PhD thesis by a PhD student. Consequently, a thesis document and defense on this topic is also an anticipated deliverable. We also plan to hold a workshop related to Deliberative Perception in 2017 and hope to co-organize it with other R-CTA researchers working on perception such as Kostas Daniilidis. Finally, we plan to run the proposed Deliberative Perception on the Roman platform in the context of "Retrieve Object" mission in the Integrated Research Assessment and Capstone.

Year	Qtr	Subtask T3C1S1B. Quarterly Milestones and Deliverables
2017	**Q1**	Formalize the framework of Deliberative Perception to reasoning about environmental conditions.
2017	**Q2**	Develop a parallelized version of Multi-heuristic A* (and/or another way to speed up the search), provide its theoretical analysis & port it into Deliberative Perception framework
2017	**Q3**	Transition the Deliberative Perception framework onto the Roman platform and integrate it with existing perception modules generating multiple hypotheses
2017	**Q4**	Fine-tune the framework and perform experimental evaluation on the Roman platform

Year	Qtr	Subtask T3C1S1B. Quarterly Milestones and Deliverables
2018	**Q1**	Extend Deliberative Perception to Active Deliberative Perception by reasoning over actions that the robot can take to disambiguate critical uncertainties
2018	**Q2**	Extend planning algorithm to provide scalable performance and provide its theoretical analysis
2018	**Q3**	Transition the Active Deliberative Perception framework onto the Roman platform and integrate it with existing perception modules
2018	**Q4**	Fine-tune the framework and perform experimental evaluation on the Roman platform

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

References:

[1] *Venkatraman Narayanan and Maxim Likhachev, "**Discriminatively-guided Deliberative Perception for Pose Estimation of Multiple 3D Object Instances,**" Proceedings of the Robotics: Science and Systems Conference (RSS), 2016.*

[2] *Venkatraman Narayanan and Maxim Likhachev, "**PERCH: Perception via Search for Multi-Object Recognition and Localization,**" Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2016.*

[3] Sandip Aine, Siddharth Swaminathan, Venkatraman Narayanan, Victor Hwang and Maxim Likhachev, ***"Multi-Heuristic A*,"*** International Journal of Robotics Research (IJRR), 20154.3.3 T3C1S1C – Discovering Task-Related Implications of Object Arrangements

MaryAnne Fields, Army Research Laboratory, mary.a.fields22.civ@mail.mil, 410-278-6675

Michael Martin, Carnegie Mellon University, mkm@andrew.cmu.edu, 724-799-0788

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.3 T3C1S1C – Discovering Task-Related Implications of Object Arrangements

PI: Fields, MaryAnne (ARL), mary.a.fields22.civ@mail.mil, (410)278-6675
Additional PI: Michael Martin (CMU), mkm@andrew.cmu.edu, (724)799-0788

Task Description and Capabilities Supported

This subtask extends our previous work on exploiting cognitive priming to augment semantic perception to the discovery of task-related implications of object arrangements. Our goal is to create a feedback loop between cognitive reasoning and active perception by building a computational model of comprehension that learns about and assesses the relevance of spatial arrangements of objects and humans in the environment. The comprehension model will support missions where task performance depends on knowledge about how the actions of other goal-driven agents (past, present or future) could alter the environment, and how their actions might impact search (e.g., 2019 capstone, Step 6d), observation (e.g., 2019 capstone, Step 3), perception, movement, manipulation, etc. Our work leverages research tasks in three thrusts, including Semantic Framework for Concepts in Context (T2C1S1D), Understanding of Environment at Contextual and More Generalized Levels (T1C1S1B), Multi-modal, Multi-scale Reasoning for Scene Understanding (T1C1S1D), Grounding Natural Language Instructions with Learnt Visual-linguistic Knowledge Representation (T2C2S2B), Interactive Multi-Agent Planning (T2C2S1A), Object Manipulation in Cluttered Spaces (T3C1), Where Objects Are (T3C1S1). It also provides information to support downstream algorithms (e.g., task and path-planners, goal-driven perception-action cycles).

Scientific Challenges and Technical Approach

Our approach combines a cognitive architecture (ACT-R) with computational lexicons and ontologies (e.g., WordNet, FrameNet, Displacer) to construct situation models of functional relations among objects in the environment. The main challenges include: (1) mapping architectural mechanisms (focus on process) to lexicons/ontologies (focus on content) so that general world knowledge can be bootstrapped, (2) learning from sparse data in high dimensional problem spaces, and (3) selecting task-relevant features (i.e., affordances) so that downstream algorithms can be informed.

Affordances can be selected by making inferences about task goals, environmental context (observed scenes) and likely manifestations of agent activities (expected scenes), and then examining the overlap among the elaborated representations. This process can be implemented in ACT-R as a combination of established discourse comprehension models (e.g., Kintch's construction-integration model, Zwaan's event-indexing model). Parsing of natural language can be avoided by using information placed in the World Model by other modules. The general world knowledge required for elaborative inferences can be obtained from lexicons like WordNet (for synonyms, meronyms, associations, etc.) and FrameNet (for abstract activity structures). The similarity between bootstrapped concepts can be obtained from ontologies based on distributional semantics (e.g., Displacer). Similarity, in turn, underlies the spread of activation among related concepts and influences the retrieval of contextually relevant concepts. Thus contextually constrained inferences can be made and inferences concerned with functional arrangements of objects can be included in a situation model. As the situation model evolves, activation will accrue on a coherent set of task-relevant features and object relations, which can be used to inform downstream algorithms such as planners for search and observation tasks and goal-driven perception/action cycles to resolve perceptual ambiguities.

We focus on two tasks related to the 2019 capstone. The first task (Step 3) involves finding a secured observation position near a designated building. While state of the art perception algorithms can identify

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

specific objects in the vicinity of the building, cognitive reasoning can identify ad-hoc and natural arrangements of those objects that could impact the choice of an observation position. The second task (Step 6d) - identifying dangerous items in the building - is a combination of state of the art perception supported by an efficient search strategy. Cognitive reasoning can contribute to this search strategy by identifying potential hiding spots for items relevant to the mission.

Metrics for assessing the comprehension model include: efficiency of search, judgments of observation choice, generalization across tasks and activities, etc.

Connection to Previous Tasks
Our work in this subtask will extend our previous work on the simultaneous learning of object classes and the selection of relevant features (Task I4-5), including our Gazebo tools for public spaces and aspects of the feature-selection model.

Collaborations
The core research team for this effort will be Michael Martin and Christian Lebiere of CMU; and MaryAnne Fields and Craig Lennon of ARL. In addition, we plan to collaborate with Doug Summers-Stay (ARL) regarding the contribution of distributional semantic vector spaces and Jianbo Shi (UPenn) to enable us to focus on task-relevant perception.

Deliverable Products
We will demonstrate our early results (FY17) on a simulated robot emphasizing our ability to reason with semantically labeled objects and events. Later demonstrations (FY18-19) will utilize a physical robot operating near (or inside) a building of interest.

2017	• A methodology to relate objects and their arrangements to specific tasks and activities (task/scene comprehension model) • A methodology to use the comprehension model to task the perception system to collect information to support the overall robotic behavior • Research to incorporate semantic vector spaces into comprehension model • Simulated scenario demonstrating improved search strategy based on ACT-R comprehension model for relatively simple task (find object left behind by another entity working on a known task) • Conference paper(s)
2018	• Integration of Y1 algorithms on RCTA platform, demonstration in FTIG scenario • Extension of comprehension model to more complex tasks • Conference paper(s)

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.4 T3C1S1D - Lifelong Improvement of Robot Capabilities and Its Robustness

PI: Likhachev, Maxim (CMU), maxim@cs.cmu.edu, (678)592-4601

Objective and Benefits:

The objective of the task is to study how a robot can learn to perform and adapt on the fly complex multi-step tasks involving challenging manipulation based on experience and human demonstrations. If successful, this research will enable naïve users (e.g., soldiers) to teach robots with manipulators to perform complex tasks such as disassembling a weapon, disabling an IED and opening door locks.

Scientific approach:

We propose to continue our work on planning methods such as Experience Graphs that we have developed in the previous years of CTA [1, 2]. These methods improve the performance of the robot behaviors via experience and demonstrations. In particular, we will continue working on these methods in the context of complex full-body manipulation and will be looking at how to combine Model-based Planning and Model-free Learning in a principal fashion. For example, imagine a manipulation platform that learns how to open a spring-loaded door from several demonstrations without having any prior model of a spring-loaded door. We will study how these skills can be incorporated with model-based planning for the platform to deal with situations such as a spring-loaded door in a cluttered space. Such scenarios require planning a feasible full-body motion through the clutter and combining it with learned skills to accomplish opening the door. On the applied side, we will demonstrate this work on a Roman platform. (We will also demonstrate it on a legged robot once/if it becomes available. For the legged robot, we will be looking at such tasks as teaching it how to go upstairs.)

The specific scientific challenge addressed by this subtask is in researching principal approaches to combining learning and planning. Learning typically relies on building a mapping from sensory information onto actions. Planning typically relies on searching through a world and robot model for a plan that achieves the desired task. This subtasks studies how to formalize the problem of combining the two problems together and investigating principal approaches that scale to real-world problems yet provide rigorous guarantees on performance. We anticipate for the proposed research to bridge the gap between classical AI tools and modern ML techniques.

A big part of Thrust 3 is the ability of the robot to manipulate environment around it such as opening drawers, doors, cabinets of various shapes and forms without having to hardcode perfect models of all of these articulated objects a-priori. This capability is critical in such missions as search and retrieve an object which requires the robot to be able to enter rooms through sliding, push, pull and other kinds of doors and to open drawers, cabinets, refrigerator, and other containers that may contain the object of interest. The proposed research explores how such manipulation capability can be learned by the robot through experience and demonstrations. We aim to transition this capability onto the Roman platform.

This research links to the research by others on perception for manipulation, and in particular finding objects that are to be manipulated. This includes subtask T3C1S2A by K. Daniilidis and our own subtask T3C1S1B. The work also links to the development of the World Model that is supposed to be extended to maintain enough data for us to decide on the right set of experiences to use. In addition, our prior work on Experience Graphs has already been transitioned into several projects outside of R-CTA including commercialization in the context of building industrial robots for autonomous paint stripping. We are also in the process of transitioning this work under a project dedicated to building a Humanoid robot jointly with Mitsubishi.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Experiences will be stored within the Experience-based Planning module. While the world model itself will not need to store and maintain experiences, it will need to fuse perception data and pass it to the planner, so that the planner is able to extract the right set of experiences to use as well as to use the data for planning. In its current form, the skills will be specific to manipulation and will be planned over by the planner. In the future, we may investigate how the approach to combining learning skills and model-based planning can be extended to navigation, especially navigation tasks requiring dynamics. The current pose estimation of the Roman platform is adequate for our purposes. Similarly, the feedback on the Roman arm is adequate based on our experiments in 2016. We don't anticipate that any other additional information from the World Model will be required.

List of dependencies:
This subtask depends on S. Karumanchi's work at JPL on the low-level control of Roman platform, and in particular full-body controller for the Roman platform (subtask T4C1S1B), Daniilidis's work at UPenn on generating perception for manipulation (subtasks T3C1S1A and T3C1S1C), GDLS's work on the world model that fuses perception data required for planning manipulation and extracting the right set of experiences, our own work on reasoning of multiple perception hypotheses in the context of manipulation (subtask T3C1S1B) and the work by Ed Weller's group at GDLS on the overall software integration related to planning for Roman platform (subtask T4C3S1C).

Collaboration: As part of this task, we have been collaborating with JPL on integrating planning with the control of the mobile manipulation platform, with RR on integrating planning for mobile manipulation with navigation and GDLS on integrating planning with the world model structure. As part of these collaborations, researchers from JPL, RR and GDLS have visited my lab several times, and we plan to continue this process. T4C3S1A, T3C1S3C, T3C1S2A.

Linkage to previous RCTA BPP sub-tasks:
In the previous several years of R-CTA, the subtask I4-2 introduced the concept of Experience Graphs [1] and has shown how, in the context of Heuristic Search-based planning, Experience Graphs can be used to improve the performance of planning through experience [1] and demonstrations [2]. These approaches have been used in the context of mobile manipulation (figure 1a). The proposed efforts partially build on this work. In addition, the proposed work ties into our earlier work done jointly with V. Kumar under M-1 subtask (within DMUM) on Mobile Manipulation. That subtask addressed model-based planning techniques for full-body mobile manipulation required for such tasks as door opening (figure 1b).

(a) (b)

Figure 1. *(a): Our prior work on using Experience Graphs in planning that learns how to plan articulation of objects from demonstrations; (b): Our prior work with V. Kumar at UPenn in M-1 subtask (within DMUM) on Model-based planning for mobile manipulation tasks such as door opening*

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Expected products:
The deliverables include papers at conferences such as RSS, AAAI, ICAPS, etc. In addition, the proposed work is a critical component in the "Retrieve Object" task in the Integrated Research Assessment and Capstone. The outcome of the proposed research will be the planning/learning module that will enable the Roman platform to operate articulated objects such as doors, drawers, cabinets, etc. without their full model a-priori. We anticipate to have this module running on the Roman platform in 2018.

Year	Qtr	Subtask T3C1S1D. Quarterly Milestones and Deliverables
2017	**Q1**	Formalize the problem of integrating Model-based planning with Model-free learning in the context of mobile manipulation tasks.
2017	**Q2**	Develop a scalable approach and provide its theoretical analysis.
2017	**Q3**	Prototype the approach on the Roman platform
2017	**Q4**	Evaluate the approach experimentally on a simple manipulation task.

Year	Qtr	Subtask T3C1S1D. Quarterly Milestones and Deliverables
2018	**Q1**	Extend the approach to more complex manipulation skills such as opening a spring-loaded door with a fully unknown model
2018	**Q2**	Extend the approach to more complex manipulation skills such as opening a spring-loaded door with a fully unknown model (cont'd)
2018	**Q3**	Transition the extended approach onto the Roman platform.
2018	**Q4**	Evaluate the extended approach in the context of the "Retrieve Object" scenario within IRA

References:
[1] *Mike Phillips, Benjamin Cohen, Sachin Chitta and Maxim Likhachev, "**E-Graphs: Bootstrapping Planning with Experience Graphs,**" Proceedings of the Robotics: Science and Systems Conference (RSS), 2012.*
[2] *Mike Phillips, Victor Hwang, Sachin Chitta and Maxim Likhachev, "**Learning to Plan for Constrained Manipulation from Demonstrations,**" Autonomous Robots (AURO), 2015.*

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.5 T3C1S1E – Discovering Task-Related Implications of Object Arrangements

PI: Martin, Michael (CMU), mkm@andrew.cmu.edu, (724)799-0788

Objective and Benefits:

This subtask extends our previous work on exploiting cognitive priming to augment semantic perception to the discovery of task-related implications of object arrangements. Our goal is to create a feedback loop between cognitive reasoning and active perception by building a computational model of comprehension that learns about and assesses the relevance of spatial arrangements of objects and humans in the environment. The comprehension model will support missions where task performance depends on knowledge about how the actions of other goal-driven agents (past, present or future) could alter the environment, and how their actions might impact search (e.g., 2019 capstone, Step 6d), observation (e.g., 2019 capstone, Step 3), perception, movement, manipulation, etc.

Scientific approach:

Our approach combines a cognitive architecture (ACT-R) with computational lexicons and ontologies (e.g., WordNet, FrameNet, Displacer) to construct situation models of functional relations among objects in the environment. The main challenges include: (1) mapping architectural mechanisms (focus on process) to lexicons/ontologies (focus on content) so that general world knowledge can be bootstrapped, (2) learning from sparse data in high dimensional problem spaces, and (3) selecting task-relevant features (i.e., affordances) so that downstream algorithms can be informed.

Affordances can be selected by making inferences about task goals, environmental context (observed scenes) and likely manifestations of agent activities (expected scenes), and then examining the overlap among the elaborated representations. This process can be implemented in ACT-R as a combination of established discourse comprehension models (e.g., Kintch's construction-integration model, Zwaan's event-indexing model). Parsing of natural language can be avoided by using information placed in the World Model by other modules. The general world knowledge required for elaborative inferences can be obtained from lexicons like WordNet (for synonyms, meronyms, associations, etc.) and FrameNet (for abstract activity structures). The similarity between bootstrapped concepts can be obtained from ontologies based on distributional semantics (e.g., Displacer). Similarity, in turn, underlies the spread of activation among related concepts and influences the retrieval of contextually relevant concepts. Thus contextually constrained inferences can be made and inferences concerned with functional arrangements of objects can be included in a situation model. As the situation model evolves, activation will accrue on a coherent set of task-relevant features and object relations, which can be used to inform downstream algorithms such as planners for search and observation tasks and goal-driven perception/action cycles to resolve perceptual ambiguities.

We focus on two tasks related to the 2019 capstone. The first task (Step 3) involves finding a secured observation position near a designated building. While state of the art perception algorithms can identify specific objects in the vicinity of the building, cognitive reasoning can identify ad-hoc and natural arrangements of those objects that could impact the choice of an observation position. The second task (Step 6d) - identifying dangerous items in the building - is a combination of state of the art perception supported by an efficient search strategy. Cognitive reasoning can contribute to this search strategy by identifying potential hiding spots for items relevant to the mission.

Metrics for assessing the comprehension model include: efficiency of search, judgments of observation choice, generalization across tasks and activities, etc.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

List of dependencies: Our work leverages research tasks in three thrusts, including Semantic Framework for Concepts in Context (T2C1S1D), Understanding of Environment at Contextual and More Generalized Levels (T1C1S1B), Multi-modal, Multi-scale Reasoning for Scene Understanding (T1C1S1D), Grounding Natural Language Instructions with Learnt Visual-linguistic Knowledge Representation (T2C2S2B), Interactive Multi-Agent Planning (T2C2S1A), Object Manipulation in Cluttered Spaces (T3C1), Where Objects Are (T3C1S1). It also provides information to support downstream algorithms (e.g., task and path-planners, goal-driven perception-action cycles). No subtask is on our critical path.

Collaboration: with T3C1S1C and T2C2S1A
The core research team for this effort will be Michael Martin and Christian Lebiere of CMU; and MaryAnne Fields and Craig Lennon of ARL. In addition, we plan to collaborate with Doug Summers-Stay (ARL) regarding the contribution of distributional semantic vector spaces and Jianbo Shi (UPenn) to enable us to focus on task-relevant perception.

Linkage to previous RCTA BPP sub-tasks:
Our work in this subtask will extend our previous work on the simultaneous learning of object classes and the selection of relevant features (Task I4-5, Exploiting Cognitive Context in Autonomous Perception), including our Gazebo tools for public spaces and aspects of the feature-selection model.

Expected products:
We will demonstrate our early results (CY17) on a simulated robot emphasizing our ability to reason with semantically labeled objects and events. Later demonstrations (CY18-19) will utilize a physical robot operating near (or inside) a building of interest.

2017	• A methodology to relate objects and their arrangements to specific tasks and activities (task/scene comprehension model) • A methodology to use the comprehension model to task the perception system to collect information to support the overall robotic behavior • Research to incorporate semantic vector spaces into comprehension model • Simulated scenario demonstrating improved search strategy based on ACT-R comprehension model for relatively simple task (find object left behind by another entity working on a known task) • **Publication**: academic paper submission targeting AAAI, IROS or ICRA
2018	• Integration of Y1 algorithms on RCTA platform, demonstration in FTIG scenario • Extension of comprehension model to more complex tasks • **Publication**: academic paper submission targeting AAAI, IROS or, or ICRA

References:

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.6 T3C1S2A - Geometry learning of rigid and deformable objects

PI: Daniilidis, Kostas (UPenn), kostas@cis.upenn.edu, (215)898-8549

Overall objective/goal:

Recent advances in 3D reconstruction and recognition have enabled the reconstruction of objects and humans by supervised detection and learning from landmark points. Today, we have access not only to millions of images of object classes but also to corresponding 3D models. However, even if we follow an end to end deep learning approach we need a supervised correspondence between 2D keypoints on the image and 3D keypoints on a geometric model. Such a correspondence is challenging even for a human and has limited existing datasets to simple furniture and vehicle classes.

In preliminary work, we have learnt to detect distributions of 2D keypoints and established their correspondence to 3D keypoints through a probabilistic approach that looks at the 2D positions as latent variables. But the challenge of whole possibly unstructured models defined by surfaces or volumes rather than keypoints remain open. We propose to study the learning of volumetric representations of objects given prior probabilities about the context (an obstacle on the way of the robot) or the class of the object. We further addressed the reconstruction as the consistency of multiple views of different instances of the same class. Rigid or deformable objects can be represented through a sparse representation of shape parameters that will be learnt directly on the volumetric space rather than on keypoints.

Proposed research:

We propose two approaches that will be studied in 2017 and 2018 respectively. In the first approach we combine statistical models of appearance and the 3D shape layout of objects for pose estimation. The approach consists of two stages: we first reason about the 2D projected shape of an object captured by a set of 2D semantic keypoints and then we estimate the 3D shape and camera viewpoint consistent with the keypoints. In the first stage, we use a high capacity convolutional network (convnet) to predict a set of semantic keypoints. The network takes advantage of its ability to aggregate appearance information over a wide-field of view to make reliable predictions of the semantic keypoints. In the second stage, the semantic keypoint predictions are used to explicitly reason about the 3D shape of an object and the camera pose modeled by a weak or full perspective camera model. Shape and pose estimates are realized by maximizing the geometric consistency between the parameterized shape model and the 2D semantic keypoints.

In our second approach, we will try a 3D pose and shape estimation with an end-to-end architecture based again on semantic keypoints. we cast 3D pose estimation as a keypoint localization problem in a discretized 3D space.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Instead of directly regressing the coordinates of the joints we train a ConvNet to predict per voxel likelihoods for each keypoint in this volume. This volumetric representation, is much more sensible for the 3D nature of our problem and improves learning. Effectively, for every keypoint, the volumetric supervision provides the network with ground truth for each voxel in the 3D space. This provides much richer information than a set of world coordinates. We plan to extend this approach to boundaries of objects and multiple views.

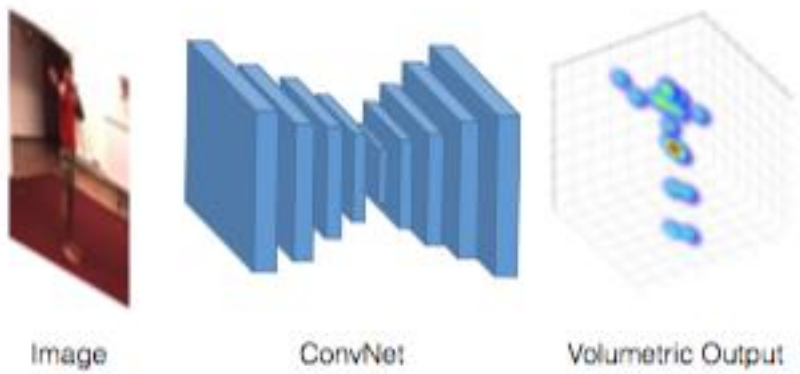

How the proposed research supports the thrust/capability:
Pose and shape of objects that belong to a category even if their instances have not been seen before is crucial for manipulation purposes. Imagine, for example, that the robot has to use a drill that has not seen before. Our approach will establish the pose and the shape of the drill. In addition graspable points can be predicted based on the existing semantic keypoints.

Scientific challenges:
The main scientific challenges addressed by the proposed research are in harnessing the power of convolutional networks for geometric problems. At the end of these two years we will have increased our understanding on whether category-based 3D reconstruction is a problem that can be solved end-to-end with a neural network or we rely on the neural network only on detection and perform the inference of geometry analytically.

Connections to the past work under R-CTA: P3-2
We have demonstrated in our past work how to search for a 3D object and if found estimate its accurate pose. The shape was assumed to be known and only the color/appearance could vary.

Collaborations:
The proposed work integrates with Likhachev's and Papon's work at CMU and JPL, respectively. Jason Owens at ARL is working for his PhD on this topic with Daniilidis. T3C1S1D, T3C1S1B, T3C1S2B.

Transitions:
Our object pose estimation has been used in the DARPA Robotics Challenge as well as in the Fast Light Autonomy program.

Deliverable products:
The deliverables include papers at conferences such as CVPR, ICCV, ECCV, RSS, ICRA. George Pavlakos is dedicating his PhD on this topic and works closely with CMU and GDLS colleagues.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Year	Qtr	Subtask I3.1 Quarterly Milestones and Deliverables
2017	**Q4**	Assessment of CNN based keypoint detection and geometric reconstruction
2018	**Q4**	Assessment of end-to-end geometry learning.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.7 T3C1S2B – Transferring Manipulation Strategies to New Objects

PI: Detry, Renaud (JPL/Caltech), detry@jpl.nasa.gov, (818)354-1068
Additional PIs: Papon, Jeremie (JPL/Caltech), papon@jpl.nasa.gov, (818)354-4729
Matthies, Larry (JPL/Caltech), lhm @jpl.nasa.gov, (818)393-5007

Objective and Benefits:

The primary objective of this task is to augment the autonomous manipulation capabilities of RCTA's robot platforms, to the end of lightening the workload of their human operators and supporting low-bandwidth–communication operations. By contrast to T3C1S3C which enhances autonomy via force sensing, this subtask leverages visual perception and learning to generalize manipulation plans across objects and tools that differ in shape and appearance. Specifically, this subtask enables the transfer of manual tasks to new tools and new objects, by providing means of generating grasping plans that are geared to the task. In this domain, our community has focused on grasp models and task models. Grasp models [1, 2] determine grasping points that are suitable for picking up an object, while task models [3, 4] assume the pre-existence of a satisfactory grip on the object and focus on modeling the motion that realizes the task. The objective of this work is to bridge the gap between these two domains, i.e., grasping objects to the end of completing a task that imposes constraints on the grip configuration. This subtask complements T3C1S3C (Karumanchi, *Dual Arm Contact Driven Mobile Manipulation Behaviors on the Roman Platform*), by allowing the transfer of S3C force-controlled behaviors to previously-unseen objects.

This task will endow the RoMan platform with means of adapting a set of pre-programmed behaviors to the new tools or objects that it is expected to encounter in T3C1's "Manipulation in cluttered spaces" scenario. For instance, it will allow RoMan to compute how to grasp a stretcher of unfamiliar geometry in order to lift it up, or how to pick up pieces of debris to move them out of its way.

Scientific approach:

This subtask aims to establish a task-oriented grasp model, that encodes grasps whose placement on an object enable a given task. For instance, if the task is to hand over an object to an operator, the model will encode grasps that leave part of the object's surface available for the operator to secure his own grip. The model will consist of two independent agents: First, a geometric grasp model that computes, from an RGBD image, a distribution of 6D grasp poses for which the shape of the gripper matches the shape of the underlying surface. The model will rely on a dictionary of geometric object parts annotated with workable gripper poses and preshape parameters, learned from experience via kinesthetic teaching, building on Detry et al. [5]. The second agent will consist of a semantic model that encodes task-compatible grasping regions. It will rely on a CNN that parses a scene into a set of object classes, building on Papon et al. [6]. On top of this classifier, we will train a qualitative statistical model of object-gripper relations that enable a specific task. The semantic model will allow us to encode relationships such as *grasp from the top*. The product of the geometric and semantic agents will allow us to initiate manipulative tasks on previously-unseen objects by identifying grasping regions where the shape of the gripper fits the shape of the tool or object, and where the positioning of the gripper allows the robot to perform the intended task. This work advances the state-of-the art by leveraging data-driven semantic scene understanding and combining approximate semantic constraints to geometric gripper constraints, thereby providing solutions that are both contextually relevant and physically (mechanically) realizable.

Metrics and Progress Assessment: We will evaluate progress based on two experiments. Experiment 1 (E1) will evaluate the performance of our approach on a manipulation platform composed of an industrial arm, a

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

three-finger hand and a 3D camera. This platform will allow us to perform hundreds of trials and collect statistically-significant characterizations. In E1, the metric will correspond to the success rate of grasp executions, and a qualitative evaluation of the suitability of grasp positioning for the intended task. The second experiment (E2) will validate the work on the Surrogate/Robosimian platform. The task will be to extract items from a gym bag suspected of containing an explosive device (or an analogous task). The task constraints developed in this subtask will allow the robot to extract items from the bag while avoiding grasps near parts that exhibit sensitive elements such as electrical wires. We will measure success by counting the number of objects that are extracted from the bag without contacting sensitive elements.

Transition Opportunity: The second experiment (E2) above will assess the applicability of this work to autonomous inspection and disposal of suspected explosive devices. It will provide preliminary results to help planning future research on this problem.

List of dependencies: The success of this task does not depend on other subtasks, nor do other subtasks depend on this one. However, this task is synergetic with several other tasks (listed as collaborations below), and integration will provide capabilities that no task would be capable of providing individually.

This task does not rely on externally-produced prior information, it will not need a map, and it will not need anything atypical from the world model. We will be using raw data coming from cameras or other vision sensors directly. This task does not require reliable pose estimates. Velocity estimates would be helpful to compute small pose deltas (meter-scale) but not critical.

Collaboration:

- JPL, Karumanchi, T3C1S3C: ongoing year-long JPL collaboration. The work from this task will be integrated with S3C to allow the transfer of S3C force-controlled behaviors to previously-unseen objects.
- Rutgers, Boularias, T3C1S2C: research integration for MCA, student exchange. We will address the problem of grasping objects that are unreachable due to clutter. S2C will allow the robot to push objects to a pose from which they can be grasped.
- UPenn, Daniilidis, T3C1S1A: student exchange.
  T3C1S1A and T3C1S2B address tactile-based and vision-based grasp planning respectively. The two tasks complement each other: in high-clutter or low-light environments, a vision-based planner may fail to provide a workable solution, and pass control over to a tactile-based planner. This collaboration will integrate S1A and S2B to provide a system that gives precedence to either vision- or tactile-based planning based on environmental conditions.

Linkage to previous RCTA BPP sub-tasks:

- This task is a continuation of P2-3 under past RCTA work.
- This task is also tightly coupled to M3 tasks (M3-1 specifically) under past RCTA work.

Expected products:
2017:

- **MCA at JPL:** Experiments on **unknown objects (geometry and inertia)** at JPL, integrating the geometric grasp model with S. Karumanchi (JPL) and Abdeslam Boularias from Rutgers.

- Delivery of a **software package** for scene-wide grasp point detection. Package applicable to object manipulation (supporting this subtask), but also to generic grasp point detection, for instance for facilitating the identification of bracing or climbing fixtures or debris removal.
- One workshop presentation conceptualizing task-oriented grasp selection. Tentatively ICRA2017 workshops.
- One conference publication demonstrating task-oriented grasp selection on simulated data. Submission to IROS2017 or ICRA2018.

2018:

- **MCA at GD:** Specific Task is TBD. Tech transfer to Cognitive architecture on Roman Platform with GD, Max Likhachev at CMU, Abdeslam Boularias from Rutgers, Marc Killpack at BYU and S. Karumanchi at JPL.
- **MCA at JPL**: Evaluate task-oriented grasping with Experiment 1 and 2 (E1/E2 above), including evaluation on RoMan/Surrogate platform.
- One conference publication demonstrating training of multiple grasp types within a single network, and transfer of simulation-trained network to real tasks. Submission to ICRA2018.
- One journal publication coalescing E1 and E2. Submission to T-RO, or RA-L with presentation at ICRA2019.

References:

[1] A. Saxena, J. Driemeyer, and A. Y. Ng. Robotic Grasping of Novel Objects using Vision. International Journal of Robotics Research, 27(2):157, 2008.
[2] O. Kroemer, E. Ugur, E. Oztop, and J. Peters. A kernel-based approach to direct action perception. In IEEE International Conference on Robotics and Automation, 2012.
[3] S. Calinon, I. Sardellitti, and D. G. Caldwell. Learning-based control strategy for safe human-robot interaction exploiting task and robot redundancies. In IEEE/RSJ International Conference on Intelligent Robots and Systems, 2010.
[4] J. Peters, J. Kober, K. Mlling, O. Krmer, and G. Neumann. Towards robot skill learning: From simple skills to table tennis. In Machine Learning and Knowledge Discovery in Databases, pages 627–631. Springer, 2013.
[5] R. Detry, C. H. Ek, M. Madry, and D. Kragic. Learning a dictionary of prototypical grasp-predicting parts from grasping experience. In IEEE International Conference on Robotics and Automation, 2013.
[6] J. Papon and M. Schoeler. Semantic pose using deep networks trained on synthetic rgb-d. In Proceedings of the IEEE International Conference on Computer Vision, pages 774–782, 2015.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.8 T3C1S2C - Learning Mechanical and Geometric Models of Unknown Objects Online

PI: Boularias, Abdeslam Rutgers University, Boularias@gmail.com, (848)445-8848

Objective and Benefits:

We consider the problem of robot manipulation in an unstructured, rigid environment with multiple rigid objects. We aim to manipulate the objects with a redundant manipulator, while accounting for how objects may physically interact with each other and the robotic arm. Manipulation tasks may include pushing and reorienting objects to enable grasping in a cluttered scene. The kinematic properties of the objects (geometric information, pose) are estimated using a vision-based sensing solution [1]. The physical properties of the objects (e.g., mass, surface friction coefficients etc.) may not be known a priori and need to be learned.

The developed tools will be evaluated in the context of robot manipulation challenges. Some of the more interesting problems in this domain involve cluttered environments, where multiple contacts can arise between objects and the robot, as well as the use of non-prehensile manipulation operation, like pushing, which are highly affected by physical parameters of objects, such as friction and mass distribution. A major limitation of existing solutions in this domain is their inability to reason about complex interactions between a target object and its environment, which can influence the capability of a robot to manipulate it. Progress in this area can have significant impact in the capability of robots to achieve dexterous manipulation and their deployment in unknown environments, which are frequently cluttered by objects with a variety of physical properties.

Technical Approach to Overcome Barriers and Advance SOA:

Physics engines typically have two components: a collision detector and a dynamic simulator. While collision detection requires only geometric models of objects, such as meshes or shapes, dynamic simulations additionally require mechanical models. Mechanical models are mass, friction, and elasticity parameters of objects. In this project, we focus only on rigid objects as they are generally easier to model. Therefore, we do not consider elasticity parameters. If the object has a non-homogeneous mass distribution or friction parameters, then it can be represented as a set of objects tied together, where each new object corresponds to a subpart of the original object with different mass and friction coefficient.

To solve the problem of modeling mechanical properties of objects, we propose an online learning approach to identify mass and sliding models of objects as a Bayesian optimization problem. The goal is to allow the robot to use predefined models of objects, in the form of prior distributions, and to improve the accuracy of these models on the fly by interacting with the objects. The learning process should be in real time because it takes place simultaneously to the physical interaction.

Figure 1 shows an overview of the proposed approach. The first step consists in using a pre-trained object detector to detect the different objects present in the scene. We are currently using the Fast Region-based Convolutional Neural Network (Fast R-CNN) for object detection [1]. We are also using the 4PCS algorithm to estimate the pose of the detected objects by mapping them to a knowledge base of preexisting 3D mesh models. We propose to augment the 3D mesh models with mechanical properties of objects. The mechanical properties correspond to the mass and the static and kinetic friction coefficients of each subpart of a given object. These properties are represented as an d-dimensional vector θ. A prior distribution P_0 on θ is used instead of a single value of θ, because different instances of the same category of objects usually have different mechanical properties.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

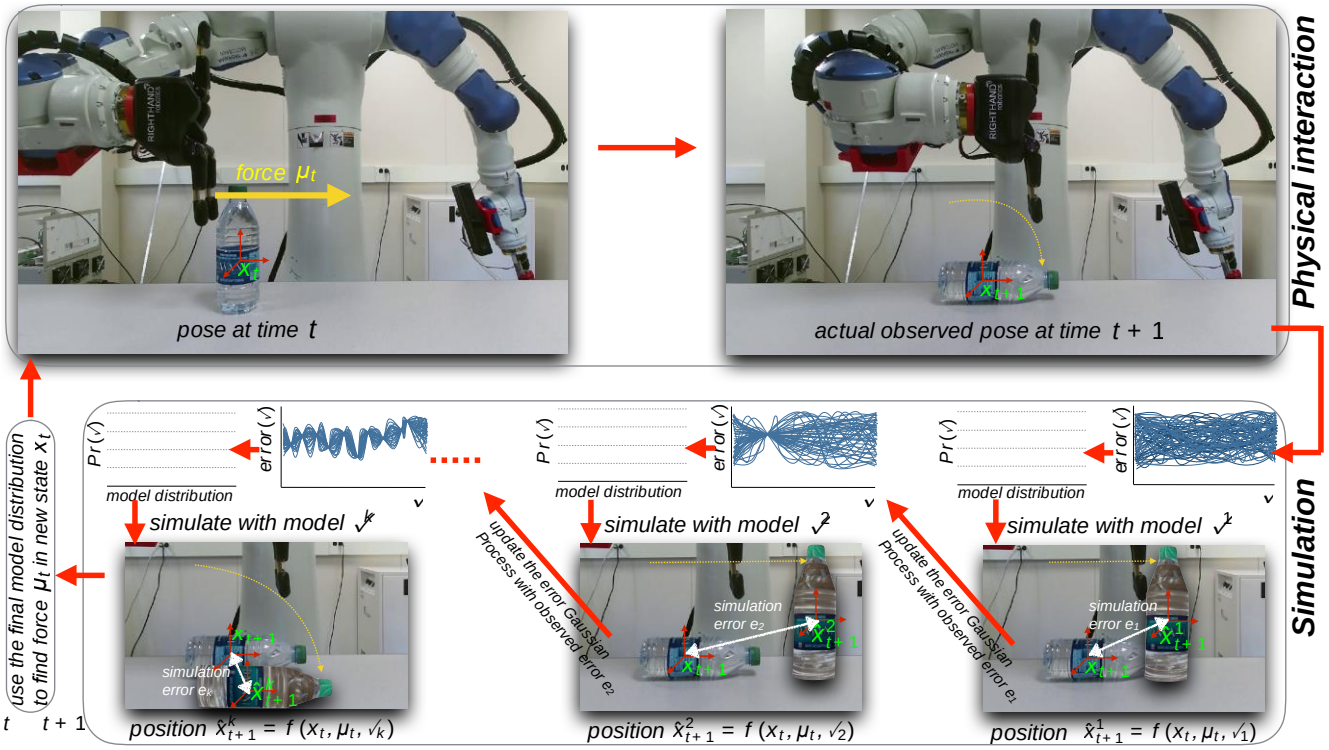

Figure 4: Overview of the proposed approach for learning object models with a physics engine. Small objects are used in the preliminary experiments, future experiments will focus on heavy objects.

The proposed online learning algorithm takes as input a prior distribution P_t on model parameters θ. Distribution P_t is calculated based on an initial distribution P_0 and a sequence $(x_0, \mu_0, x_1, \mu_1, \ldots, x_{t-1}, \mu_{t-1}, x_t)$, wherein x_t is the observed 6D pose (position and orientation) of the manipulated object at time t and μ_t is a vector describing a force applied by the robot's fingertip on the object at time t. Applying a force results in changing the object's pose from x_t to x_{t-1}. Forces are chosen according to a policy that maps a pose x_t and a parameter distribution P_t into a force μ_t.

Given a prior distribution Pt and a new observation (x_t, μ_t, x_{t+1}), we show how to use a physics engine to estimate a posterior distribution P_{t+1} on the models. We used the *Bullet* physics engine in our preliminary experiments. The posterior is obtained by simulating the effect of force μ_t on the object under various values of parameters θ and observing the resulting poses $\hat{x}_{t+1}$. The goal is to identify the model parameters that make the outcome $\hat{x}_{t+1}$ of the simulation as close as possible to the actual observed outcome x_{t+1}. In other terms, we solve the following black-box optimization problem:

$$\theta^* = \arg\min_\theta \text{Error}(\theta) = \arg\min_\theta \| x_{t+1} - f(x_t, \mu_t, \theta) \|_2,$$

wherein x_t and x_{t+1} are the observed poses of the object at times t and $t+1$, μ_t is the force that moved the object from x_t to x_{t+1}, and $f(x_t, \mu_t, \theta) = \hat{x}_t$ is the simulated pose at time $t+1$ after applying force μ_t in pose x_t. The simulator of the physics engine is modeled as an unknown function f, that is computationally costly to evaluate.

We formulate this problem in the Bayesian optimization framework, and we use the Entropy Search technique presented in [3]. We specifically use a greedy-search variant that we proposed in [4] for learning to grasp objects and that is more efficient computationally. This technique consists in representing the simulation error function (Error(θ) in the equation) as a Gaussian process that is learned from observed data and online simulations. To reduce the number of simulations needed to find θ^*, we simulate using models θ that contributes the most to the entropy of a distribution on θ^*, in an attempt to reduce this entropy as quickly

as possible and to converge to the true model θ^* in real time.

In a preliminary experiment, we learn the mass and the friction coefficient of a simple rigid box (an Expo eraser). The robot plays with the object by poking it in random directions with random forces and observes how it moves. We use the learned model to predict how the object would move. We uploaded on http://bit.ly/2dRgut1a video that shows the real motion of the object next to its predicted motion in Bullet using a model that was learned with the proposed approach.

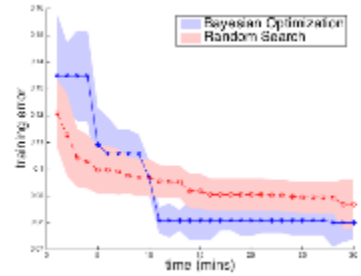

Figure 5: Difference between learned models and the best model as a function of the learning time

Metrics & Progress Assessment:

(a) Prediction error: In our preliminary experiment, we used the learned model to predict the motion of an object, and we computed the distance between the predicted pose of the object (3D position and 3D orientation) and the real observed pose. Our results indicate that this prediction error diminishes quickly as the robot interacts more with the object. This will be a primary measure of performance in our experiments. We will also report multi-step prediction errors, where we assess the capability of the learned model to predict the pose of an object multiple steps ahead.

(b) Model accuracy: We measure the difference between learned models and hand-tuned models as a function of the learning time. Figure 2 shows this difference as in the preliminary experiment.

(c) Time: We report the total time for each task, which includes planning, execution, and learning.

(d) Task performance: Given a certain predefined time window (e.g, 10 minutes), we report the overall average performance of the system within this time window. For example, the task performance is the average number of objects picked from a box in a given time window.

(e) Adaptation to changes: After the robot learns to perform a task on given objects, we replace all the objects with different ones from the same categories. We assess the system's ability to learn models of the new objects and to solve the tasks, using the above metrics.

How the research supports RCTA Thrust/Capability?

Dexterous manipulation of objects requires highly accurate models of the objects. This is particularly true for precise grasping and pushing of objects. Most objects encountered in the field are unknown, i.e. they do not have any prior model. Our research will result in a new technique for quickly identifying models of objects online and on-the-fly, using a small number of interactions between the robot and the objects. This component will be useful for all other tasks of RCTA that currently assume known models of the object, and use them for planning.

What scientific challenges are being addressed?

(1) Learning from few data points: while most of the research in machine learning had shifted to problems of dealing with large amounts of data (big data), there has been a very little focus on the problem of learning with a small number of data points. This problem is encountered whenever a robot needs to learn a new skill or to learn a model of a new object in real-time while operating in the field. People, for instance, are extremely good at learning from only one or two examples. Kernel and instance-based methods are typically preferred when the number of data points is too small. However,

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

deep architectures seem to perform significantly better with large data sets. It would be interesting to find a method that combines the best of the two worlds.

(2) Model selection: we are currently fixing the friction and inertial models and only trying to figure out some free parameters online. It would be interesting to also learn to select the model that would better fit a given unknown object (e.g., rigid vs non-rigid model selection). We are also interested in learning the kernel (i.e. metric) function that is more appropriate for our application.

(3) Learning with noisy data: we are currently using an object detector (a neural net) and a tracker to estimate the position and orientation of an object in a given image. The sequence of images taken after pushing the object in multiple directions is used to estimate a friction and inertial model of the object. However, the estimated positions and orientations of the objects are subject to a certain error, we are looking into ways to deal with this error in a principled fashion.

(4) Online learning: this type of learning is fundamentally different from offline learning. Labeling errors have immediate impact on the performance of the task. Online learning should also be real-time and incremental (i.e. sequential).

(5) Exploration vs Exploitation: this is the fundamental challenge in reinforcement learning We need to balance the time spent by the robot trying to learn a model of the object (i.e., playing with the object) and the time spent on actually performing the given task (e.g, grasping or pushing the object).

(6) Scalability: so far, we have been using simple models of the object (uniform friction and mass distribution). In reality, objects are made of several connected parts that have different properties. A major challenge is to scale our approach up in order to handle spaces with higher dimensions.

(7) Planning under uncertainty: our proposed approach is Bayesian, and returns a belief on all possible models of a given object. This is preferred to returning a single model because it is extremely difficult to identify an accurate model of unknown object with a small number of data points. The returned belief will be used for planning. This is known as belief space panning. It is still unclear how this can be done in real-time.

Who you will be collaborating with in consortium/outside?

The teams at Rutgers University (Abdeslam Boularias) and at JPL (Renaud Detry and Sisir Karumanchi) have elaborated a plan for collaboration. The team at JPL is developing, among other things, an approach for grasping unknown objects. In some situations, it is necessary to push objects before being able to grasp them. This is known as pre-grasping. To efficiently push an object without causing much damages, it is important to quickly build an approximate model (or a distribution of models) on-the-fly. The team at Rutgers will provide a package for learning such models and also for pushing objects. The team at JPL will develop a method for deciding when to grasp and when to push, and will use the component developed at Rutgers for the latter action.

This collaboration will be achieved through student exchange. A graduate student of Abdeslam Boularias will spend a few weeks of Fall 2017 at JPL and deploy the component developed at Rutgers on the robotic platform used at JPL.

Collaborations: T3C1S2B, T3C1S3C

Previous RCTA Subtask: M2-3

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Deliverable products:

	Subtask Quarterly Milestones and Deliverables *(Deliverables will include all data, software, evaluation analyses and documents such as source code algorithm flow charts, algorithm parameters, software documentation, research reports and publications.)*
Q1	Continue the development of the online model learning technique
Q2	Online simultaneous model learning and control (model-based reinforcement learning)
Q3	Experiments on table top objects using the manipulators currently available at Rutgers
Q4	Experiments on **heavy unknown objects** at JPL, integrating the proposed manipulation learning method with Renaud Detry and Sisir Karumanchi's work on grasping.

Publications: two conference papers, one submitted to R:SS and one to NIPS, in addition to a journal article prepared at the end of the year 2017.

References:

[1] Girshick, R. B. (2015). Fast R-CNN. CoRR, abs/1504.08083

[2] Coumans, E. Bullet physics engine. [Online]. Available: www.bulletphysics.org

[3] Hennig, P. and Schuler, C. J. (2012). Entropy Search for Information-Efficient Global Optimization. Journal of Machine Learning Research, 13, 1809–1837.

[4] Boularias, A., Bagnell, J. A., and Stentz, A. (2014). Efficient optimization for autonomous robotic manipulation of natural objects. In Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, pages 2520–2526.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.9 T3C1S3A - Learning manipulation actions from demonstration

PI: Muelling, Katharina (CMU), kmuelling@nrec.ri.cmu.edu, (412)692-1097

Objective and Benefits

Generating motion plans to perform complex manipulation plans autonomously in real-world scenarios mandates a tightly coupled reasoning across symbolic states and continuous spaces. Most symbolic planners rely on the knowledge of pre-conditions and effects of individual manipulation actions that are manually defined and fixed. However, the correct mapping from the continuous state space to the pre-conditions and effects of a manipulation action is unknown. Therefore, we seek in this subtask to develop a learning algorithm that can learn these pre-conditions and effects given a few (sequences of) manipulation action demonstrations and the context in which they have been recorded. Learning pre-conditions is particularly useful for this Capability, as it enables robot manipulation algorithms to move beyond pre-programmed behavior and, consequently, to enable the robot to autonomously interact with complex cluttered real-world environments. The developed capability can also reduce the number of manipulation actions needed to perform a task successfully.

Scientific approach

Our previous work in the area of manipulation actions focused on integrating TBS into the area of manipulation on a Barrett Wam platform in cooperation with Jean Oh (I1). Here, we achieved a task-level speech mapping to low level manipulation actions which allowed a manipulation platform to perform table-top manipulation task by following the speech instruction of a human collaborator. However, this work did not enable the system to solve higher level instructions which require to sequence multiple lower-level manipulation actions to achieve the goal.

In this task, we focus on developing the ability to learn pre-conditions and effects of manipulation action from demonstrated manipulation action sequences applicable for continuous real-robot applications. These manipulation actions can be understood as motor policies that map the state of the system to a control vector that can be executed by a robotic system [1,2]. Examples of such manipulation actions are reaching for and grasping of an object. Although, we consider motor policies specific to manipulation, the general approach is not constraint to this context. The proposed research is based on two assumptions: (1) complex motor tasks can be frequently solved using a small number of motor policies which are sequenced with respect to the goal and the environmental context and (2) that the environmental context in which the robot operates is available to the system. In this project, we first focus on learning the pre-conditions of manipulation primitives based on the context and goal in which they have been demonstrated. The demonstrations can be either given by kinesthetic teach-in or observation of human demonstrations. The individual manipulation actions are given to the system in form of a motor policy library. The technical approach includes the analysis of the performance of existing approaches and development of a hierarchical (deep) inverse reinforcement learning framework to infer the relative contextual information relevant to identify and extend the preconditions. We will further investigate planning representations that allow for updating the pre-conditions on-the-fly through its experience, i.e., if the robot encounters a scenario in which the pre-condition of a manipulation action was faulty the robot needs to be able to update this pre-condition to reflect the encountered scenario.

The environmental context considered in this task will be defined by the pose of the objects in the robots proximity relative to its base. Other information necessary, include the current robot configuration. As in most real-robot applications, we do not assume perfect knowledge of the object positions and instead intend to enable the algorithm to deal with uncertain information. However, we do assume that those information are within a reasonable amount.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Collaboration:
Planned collaborations include: Jean Oh (CMU) for creating multi-modal interfaces and Renaud Detry (JPL) for the detection of environmental clues and possible grasp affordances.

Linkage to previous RCTA BPP sub-tasks:
I1-1 Framework for Collaborative Intelligence and I6-1 Learning from Human Input, Data Mining, and Experience

Expected Product
The subtask will generate the following products:

Deliv. by end of	Products
2017	- Develop of hierarchical (deep) learning framework for inferring conditions and context of manipulation actions from a few demonstrated sequences - Design of framework for planning using the inferred pre-conditions
2018	- Integration of an online update algorithm that allows the agent to update the pre-conditions on-the-fly through its experience

Target publication venues for the suggested work includes AAAI, RSS, ICRA, and IROS.

References:
[1] A.D. Dragan, K. Muelling, J.A. Bagnell, S.S. Srinivasa: Movement primitives via optimization. ICRA 2015.
[2] K. Muelling, J. Kober, O. Kroemer, J. Peters. Learning to select and generalize striking movements in robot table tennis. The International Journal of Robotics Research. 2013.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.10 T3C1S3C– Dual Arm Contact Driven Mobile Manipulation Behaviors on the Roman Platform

PI: Karumanchi, Sisir (JPL/Caltech), Sisir.B.Karumanchi@jpl.nasa.gov, (818)356-3867

Research Description:

Objective and Benefits:
The primary objective of this subtask is to enhance the interaction capabilities of a robotic platform as a team member. Such a robot will has two properties that are complementary

1) Will be mobile and maneuverable in complex 3D conditions where it can maneuver via physical interaction
2) Will have manipulation capabilities that will enable it to perform useful work once it reaches its destination

Current **state of art** in mobile manipulation **requires highly precise geometric sensing** of the world, **avoids all contact during mobility** about **can only manipulate payloads less than 5 lbs**. This task will aim to 1) reduce the reliance of accurate geometric knowledge of the world 2) extend its mobility capabilities to leverage contact to maneuver in challenging conditions and 3) extend manipulation capabilities to payloads > 20lbs by leveraging force feedback.

Specifically, the core objectives of this subtask is to research and develop next–gen behaviors that

1. Minimize the need to re-program for each object or obstacle type
2. Minimize reliance on accurate geometric knowledge of the world and the need for prior models of objects.
3. Leverage force feedback to manage reaction forces from the environment effectively,
4. Bridge the gap between mobility and manipulation.

The end goal is to achieve autonomous behaviors at the capstone that i) do not require prior models of objects or obstacles, ii) are robust to imperfect geometric sensors (exteroception) by performing manipulation and mobility via feel (proprioception) and iii) have the ability to learn from past experience (by associating proprioception with exteroception).

Specific Contributions to the Cognitive Architecture:
This task will leverage the Roman Platform within RCTA. The following contributions will be made to cognitive architecture

1) An architecture that generalizes for both mobility and manipulation. (locomotion vs manipulation)
2) An architecture that generalizes for both perception driven and contact driven behaviors (exteroception vs. proprioception)
3) An architecture that generalizes to contact driven physical human-robot interactions in addition to semantic commands.

Technical Approach to Overcome Barriers and Advance SOA:
The problem of using contact to supplement manipulation and mobility behaviors is complementary. Force feedback at the limbs and current feedback in the base form inputs into whole body planning that enable robots to generate contact based behaviors.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

This task will leverage the heavily geared RoboSimian limbs and talon base on the Roman Platform. These are especially well suited for tasks that require high torque requirements as those proposed here.

The ability to manipulate heavy (>20lbs) objects and manage high reaction forces will be tackled via four different yet complimentary technical strategies.

a) Whole body motions with current feedback at the tracks and force feedback at the limbs.
b) Multi-limbed grasps
c) Multi-robot cooperative behaviors
d) Human-robot cooperative behaviors

Metrics & Progress Assessment:
Three metrics will be used assess progress

- **Metric 1: No. of steps in behavior:** *Simple behaviors generalize better.* The no of steps in a manipulation behavior is an indication. The best behavior of a given task will be the one that can get the job done with the least no. of steps. When manipulating objects of high inertial uncertainty, a robot must be able to recover from failures and learn from experience. With more steps in a behavior more possible failure cases need to be handled which makes complex behaviors intractable and ineffective.

- **Metric 2: No. of objects per behavior:** *Interactive behaviors can generalize better.* By enabling a robot to physically inspect and reason about subsequent steps. One could use ***a single behavior to manipulate/maneuver many objects/obstacles*** instead of re-programming for every possible object/obstacle a robot might encounter.

- **Metric 3: Size of payload:** Assign dull/heavy work to robots and intense/finesse work to soldiers. By increasing the performance envelope of robotic manipulation beyond those of human limits. A robotic team mate can provide complimentary support to a soldier. Instead of trying to replicate human level dexterous performance it may be more useful to perform simple behaviors that operate beyond human levels.

Connection to past RCTA work:

- This task is a continuation of M3-1 under DMUM thrust under past RCTA work.
- This task was also tightly coupled with M2 and P2 tasks under past RCTA work.

Collaborations:

- *Ongoing year-long JPL/Caltech collaborations*
  - Burdick - T3C1S3D Mobility-aided manipulation on the Roman platform
  - Papon/Detry - T3C1S2B - Generalizing grasps across objects and tasks

- *MCA's Student exchanges via JPL summer internship programs*
  - Boularias, Rutgers - T3C1S2C - Learning Mechanical and Geometric Models of Unknown Objects Online
  - Goodrich & Killpack, BYU - T1C2S4C - Practical human-robot interaction for navigation and manipulation

- *IRA's and tech transfer to Roman Platform*

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

- Likhachev, CMU - T3C1S1B - Planning with Uncertainty in the Model and Perception
- Daniilidis - T3C1S2A - Geometry learning of rigid and deformable objects

Deliverables

	Subtask Quarterly Milestones and Deliverables
Q1-2017	Develop tracked base behaviors force contact-driven mobility in 2D cluttered environments
Q2-2017	Whole body mobile manipulation behaviors for contact-driven mobility in 2D cluttered environments
Q3-2017	**MCA at JPL:** Experiments on **geometrically and inertialy unknown objects** at JPL, integrating the proposed manipulation learning method with Renaud Detry and Abdeslam Boularias from Rutgers.
Q4-2017	**Tech transfer to Roman Platform**
Q1-2018	**MCA at JPL:** Whole body mobile manipulation behaviors with contact driven HRI. Aiming for **multi-robot/human-robot coordinated control** with Marc Killpack at BYU.
Q2-2018	Tight integration of perception for mobile manipulation behaviors with Renaud Detry
Q3-2018	**MCA at GD:** Specific Task is TBD. Tech transfer to Cognitive architecture on Roman Platform with GD, Max Likhachev at CMU, Abdeslam Boularias from Rutgers, Marc Killpack at BYU and Renaud Detry at JPL.
Q4-2018	Publish a Journal article on **"Scalable contact driven mobile manipulation behaviors for cluttered environments"**

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.11 T3C1S3D – Mobility-aided manipulation on the Roman platform

PI: Burdick, Joel (JPL/Caltech), jwb@robotics.caltech.edu, (626) 395-4139

Research Description:

Objective and Benefits:
This work is aimed to develop, implement, and test a new class of computationally efficient local motion planning algorithms which will allow mobile manipulation platforms such as the RoboSimian and SURROGATE platforms to engage in sophisticated manipulation tasks on uneven terrain, while ensuring platform stability. In this upcoming period, the goal is to extend these algorithms to handle heavy loads, loads with large uncertainties in inertial properties, and tasks involving large interaction forces. Additionally, this effort will seek to better integrate proprioceptive information in mobile manipulation planning and control.

Additionally, generalization of the above algorithms to handle actuators that address high torques and high bandwidths should also contribute to ***Optempo Mobility in Rough Terrain,*** as robots constructed with these actuators should be able to generate the higher forces and faster reactions needed to enable optempo operation.

***Technical Approach to Overcome Barriers and Advance SOA*:**
We have previously developed an efficient quadratic programming algorithm for local *kinematic* planning of mobile manipulation. Preliminary investigations support the idea that the class of quadratic planning algorithms can be extended to dynamic models, which will allow us to model and incorporate uncertain load inertia characteristics into the approach.

Metrics & Progress Assessment:

(a) Degree of uncertainty in load characteristics which can be handled while maintaining stance stability

(b) Magnitude of forces which can be controlled during mobile manipulation tasks

(c) Size of object which can dragged under proposed control framework

Connection to past RCTA work:

- This task is a continuation of M2 under DMUM thrust under past RCTA work.
- This task was also tightly coupled with M3-1 and P2 tasks under past RCTA work.

Collaborations:

- *Ongoing year-long JPL/Caltech collaborations*
  - Karumanchi - T3C1S3C Dual Arm Contact Driven Mobile Manipulation Behaviors on the Roman Platform
  - Papon/Detry - T3C1S2B - Generalizing grasps across objects and tasks

- *MCA's Student exchanges via JPL summer internship programs*
  - Boularias, Rutgers - T3C1S2C - Learning Mechanical and Geometric Models of Unknown Objects Online

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

- Goodrich & Killpack, BYU - T1C2S4C - Practical human-robot interaction for navigation and manipulation

- *IRA's and tech transfer to Roman Platform*
  - Likhachev, CMU - T3C1S1B - Planning with Uncertainty in the Model and Perception
  - Daniilidis - T3C1S2A - Geometry learning of rigid and deformable objects

Deliverables

	Quarterly Milestones and Deliverables for CY 2017
Q1	Analysis of mobile manipulator control under uncertain load characteristics in context of quadratic programming paradigm
Q2	Simulation of mobile manipulator handling a heavy load using quadratic programming approach
Q3	Prototype implementation of approach on SURROGATE or Robosimian platforms at JPL. Publish conference papers.
Q4	Evaluation of framework against proposed metrics and SOA.

	Quarterly Milestones and Deliverables for CY 2018
Q1	Generalize quadratic programming for locomotion.
Q2	Simulation of locomotion using quadratic programming approach
Q3	Prototype implementation of approach on JPL's Llama platform. Publish conference papers.
Q4	Evaluation of framework against proposed metrics and SOA

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.3.12 T3C1S3F - Learning novel force-based mobile manipulation – Navigation in Confined Spaces

PI: Yim, Mark (UPenn), yim@seas.upenn.edu, (215)898-5269
Additional PIs:

Name	*Org.*	*Phone*	*Email*
Dan Lee	Univ. Penn	215-898-8112	ddlee@seas.upenn.edu
James Dotterweich	ARL	410-278-9835	james.m.dotteerwich.ctr@mail.mil
Harris Edge	ARL	410-278-4317	harris.l.edge.civ@mail.mil
Chad Kessens	ARL	410-278-4452	chad.c.kessens.civ@mail.mil
Ray Von Wahlde	ARL	410-278-9738	raymond.vonwahlde.civ@mail.mil

This is research for flight platforms and their components which employ human-scale forces and moments to interact with the environment. Validate dynamic simulations of airborne platform interactions with the environment with real-world data.

How will the research support RCTA Thrust/Capability? Robotic manipulation is the ability to move, arrange or control objects in an environment. In a broad sense, mobility is the ability to manipulate a robot itself. The long-term objective of this work is to have the ability to accurately model robots interacting with the environment and objects in the environment in order to better understand how to design platforms that are safe and efficient and can operate at mission tempo or faster. Previous AMM team efforts in this area have been for ground systems. In addition to ground systems, this line of research will be expanded to flight systems. Typically, one does not consider forces for interaction with objects in the environment for a flight vehicle. Providing a means to accurately predict the forces experienced and that need to be applied by a flight vehicle performing work will be a great aid to future research.

Intelligent machine learning will also be applied in this context. How can lightweight robots learn to apply large forces on the environment in the right places to effectively leverage their interactions? Lightweight robots cannot simply apply large forces at a single point as they cannot use their inertia for reaction forces and therefore need to brace themselves against the environment. But how can a robot learn how to effectively brace itself using the environment?

Platforms will be investigated that can navigate in confined areas with the use of onboard sensors and visual cues from perception arrays such as onboard computer vision systems. This includes research into appendages and mechanisms that may facilitate mobility in constrained space. The work from this project will extend into designing air and ground robotic systems in the future that can increase the safety and range of capability for autonomous systems used by the military, first responders and DOE hazardous material operations. This research has great potential for making difficult operations in megacity environments feasible.

Scientific challenges: This project will develop the capability for flight vehicle interacting with solid objects in the environment. This involves computing flight and object interaction forces during simulation. In addition, a very unique flight vehicle will be employed for the airborne manipulation that will have the ability perform a variety of thrust vectoring which enables aircraft attitude adjustment independent to the propulsion vector force. Also, the vehicle will have telescoping arms that will allow the vehicle to perform

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

work and move objects. It should be noted that there are collaborative and parallel research efforts that will be leveraged. The first is the manipulator research in which RCTA researchers from UPenn will collaborate with ARL researchers to study generation and control of human-scale forces and moments by a UAV in order to perform work. The second leveraged research project is the subtask "Grasping to Perform Work", in which much of the basic theory and application of grasping is studied.

Part of our proposed approach for semi-autonomous environment interaction is to learn features of the environment that a novel tubular/cable robot can use for leverage and to sustain strength/weight ratios more than an order of magnitude better than traditional articulated robot arms. This will enhance capabilities for mobile, lightweight robots to manipulate objects in the environment, in order to open passageways or optimally place sensor payloads.

Measuring the aerodynamic interactions of these platforms with local environments has consisted of small scale bench tests of isolated rotors. Data was not linked to any control architecture or handling qualities studies. One area of research that may be explored is the use of touch or manipulation to aid airborne maneuvers. Little research has been done in this area.

Link to past RCTA subtasks: This work builds on *Intelligence I5-2 Online Learning of features for interaction tasks* and *DMUM M1.1 Mobile Manipulation with Large* Forces. This included work on developing the expanding tube light weight long reach arm for mounting on high agility mobile platforms.

Collaborations and Transition: This work joins two formerly separate RCTA groups, one in Intelligence and one in DMUM. Outside of Academia this work couples very tightly with the ARL staff James Dotterweich, Harris Edge, Chad Kessens and Ray Von Wahlde. Chad Kessens serves a dual role as both student exchange (visiting UPenn while a U.Md phD student) as well as collaborator at ARL. The group has already demonstrated collaborative work and everyone is excited to continue this work both for its interest to the research direction and the direct application to Army needs. In particular, promising opportunities have been communicated by ARL staff for potential transition outside of the RCTA.

Metrics & Progress Assessment:

Year	Qtr	Quarterly Milestones and Deliverables
FY17	**Q1**	Flight experiments of propulsion core.
FY17	**Q2**	Integration of ARL propulsion core and UPenn arm. Document findings.
FY17	**Q3**	Addition of ARL grasper. Characterization of wall-propulsion core interaction.
FY17	**Q4**	Flight experiments with wall interaction. Document findings.
FY18	**Q1**	Characterization of force and moment generation of flight vehicle. Approaches
FY18	**Q2**	Validation of simulation for flying limbed vehicle interacting with an object.
FY18	**Q3**	Demonstrate ability to open door with learned environmental features.
FY18	**Q4**	Document findings.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.4 Thrust 3/Capability 2: Manipulation for unknown heavy objects

4.4.1 Overview

This cluster is about manipulation in the presence of high inertial uncertainty. The end goal is to increase the performance envelope of robotic manipulation to objects heavier than 20lbs with the ability to recover/adapt from mid-manipulation failures due to unknown inertial properties. Compared to the previous cluster in manipulation, the focus is not on geometric perception of objects instead tactical simplifications (e.g. visual tags) will be made to aid experimentation.
This cluster will target the following capabilities:

- **Ability to manipulate heavy (> 20 lbs) objects.** Reliable behaviors that leverage whole body motions to manipulate objects beyond envelopes of human performance. Some sample strategies to investigate include push vs. pull of payloads, multi limbed whole body leverage, and multi – robot/human-robot coordinated lifting.
- **Ability to deal with high amount of uncertainty on inertial properties**. No access to prior models. Will use active perturbation to study manipulands before adaptively generating grasping strategies.
- **Ability to leverage proprioceptive feedback.** Exteroceptive sensors (camera, lidars etc.) are not well suited to reason about inertial properties, one has to leverage proprioceptive feedback.
- **Ability to learn from past experience to recognize / re-recognize inertial properties.** Can associate measured proprioceptive feedback with exteroceptive sensing for predictive understanding.
- **Ability to leverage reaction forces from the environment.** Relevant for Self-manipulation within the mega cities paradigm. Will enable robots to leverage environment fixtures to achieve tactical performance.

4.4.2 Army relevance

Robots cannot rely on perfect world/object knowledge from exteroceptive sensing when operating in real-world environments. Challenges include imperfect sensing and self-occlusions of manipulands during mid-operation. The ability to leverage proprioceptive sensing (e.g. force torque sensors, actuator torque/current feedback) will enable robots to operate via feel when necessary and will also lead to reduced resolution requirements on the perception suite which will reduce platform cost. Related capabilities include maneuvering heavy debris to clear access and to stack heavy inventory for logistics. Other beneficial capabilities include the ability to use force feedback as an information channel for physical human robot interaction (e.g. carrying a stretcher).

4.4.3 Required research areas

The research under this capability are clustered under three fundamental questions:

1. How to perceive and maneuver heavy objects?
2. How much force can be applied to or expected from the environment?
3. How to enable physical interactions via force feedback?

The following subsections describes the various flavors of research groupings proposed under this capability. Sections 4.5.3.1-4.5.3.3 discuss proposed research that addresses the first question. 4.5.3.4 discusses proposed research that addresses the second. Final1y, 4.5.3.5 discusses research on co-manipulation that addresses the third question of enabling physical interactions.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

4.4.3.1 Next Gen Mobile Manipulation Behaviors to Maneuver Heavy Objects via Force Feedback and Whole Body Planning.

The primary objective of this work is to enhance the interaction capabilities of a robotic platform as a team member. Such a robot will not just move from point A to B but once it reaches its destination it performs useful work. Current state of art in mobile manipulation requires highly precise geometric sensing of the world about can only maneuver payloads less than 5 lbs. This task will aim to 1) reduce the reliance of accurate geometric knowledge of the world 2) extend its capabilities to payloads > 20lbs by leveraging force feedback.

The end goal is to achieve autonomous behaviors at the capstone that i) do not require prior models of objects, ii) are robust to imperfect geometric sensors (exteroception) by performing manipulation via feel (proprioception) and iii) have the ability to learn from past experience (by associating proprioception with exteroception).

The ability to manipulate heavy (>20lbs) objects and manage high reaction forces will be tackled via recent advancements in four complementary technical strategies a) Whole body motions via local optimization based motion planning algorithms b) Multi-limbed force controlled grasps c) Multi-robot distributed force control based cooperative behaviors d) Human-robot cooperative behaviors via shared force feedback.

4.4.3.2 Novel Robot Actuators for Next-Gen Mobile Manipulation

Typical robotic actuators are either chosen for torque or speed performance which currently is a fundamental tradeoff in robotics. For example RoboSimian actuators were chosen for high torque density and can handle heavy loads but are slow. On the other hand, walking robots have compliance, back-drivability and high torque bandwidth which enables them to move and adapt much more quickly.
Under this research problem, design of novel mechanisms are being considered that could provide high torque density, compliance, back-drivability, and good torque bandwidth, all of which are needed to enable next generation human-scale mobile manipulation platforms. Such proposed mechanism will retain the torque density of RoboSimian actuators while achieving compliance and back-drivability that is required for more dynamic platforms. We have proposed a novel actuator architecture involving two outrunner motors, dual planetary gear systems, and a novel series-elastic actuator. Preliminary modeling supports our hopes for high torque density and high bandwidth.

Over time, if this actuator concept is successful, it will contribute to ***Manipulation of heavy objects.*** Additionally, the high torques and high bandwidths of these actuators should also contribute to ***Optempo Mobility in Rough Terrain,*** as robots constructed with these actuators should be able to generate the higher forces and faster reactions needed to enable optempo operation.

4.4.3.3 Learning Mechanical and Geometric Models of Unknown Objects Online

We consider the problem of robot manipulation in an unstructured, rigid environment with multiple rigid objects. We aim to manipulate the objects with a redundant manipulator, while accounting for how objects may physically interact with each other and the robotic arm.

Manipulation tasks may include pushing and reorienting objects to enable grasping in a cluttered scene. The kinematic properties of the objects (geometric information, pose) are estimated using a vision-based sensing solution [1]. The physical properties of the objects (e.g., mass, surface friction coefficients etc.) may not be known a priori and need to be learned.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

To solve the problem of modeling mechanical properties of objects, we propose an online learning approach to identify mass and sliding models of objects as a Bayesian optimization problem. The goal is to allow the robot to use predefined models of objects, in the form of prior distributions, and to improve the accuracy of these models on the fly by interacting with the objects. The learning process should be in real time because it takes place simultaneously to the physical interaction.

4.4.3.4 Navigation and Work in Confined Space

Robotic manipulation is the ability to move, arrange or control objects in an environment by mechanical means. In a broad sense, mobility is the ability to manipulate a robot itself. Previous efforts in this area have been confined for ground systems. This line of research will be expanded to flight systems. Typically one does not consider forces for interaction with solid objects in an environment in conjunction with a flight vehicle. This will be a fairly unique situation that is experienced by flight systems that perform manipulation or work. Providing a means to accurately predict the forces experienced and that need to be applied by a flight vehicle performing work will be a great aid to future research in this area.

The objectives of this work is to enable Lightweight robots to apply large forces by bracing against the environment. Our proposed approach is to learn features of the environment that a novel tubular/cable robot can use for leverage and to sustain strength/weight ratios more than an order of magnitude better than traditional articulated robot arms. This will enhance capabilities for mobile, lightweight robots to manipulate objects in the environment, in order to open passageways or optimally place sensor payloads.

The long-term objective of this work is to have the ability to accurately model robots interacting with the environment and objects in the environment in order to better understand how to design platforms that are safe and efficient and can operate at mission tempo or faster.

4.4.3.5 Practical Human-Robot Interaction for Navigation and Manipulation:

The research cluster aims to leverage force feedback for physical interactions between humans and robots. As a sub-focus of human-robot interaction for co-manipulation this research will build on RCTA work on predicting human intent from haptic data. Data from over 40 human participants has allowed the development of models produced by Deep Neural Nets to predict human intent based on haptic (force-torque) data. For the end of the 2016-2017 BPP, we are implementing a controller that enables a human and robot to co-manipulate an object, which in a practical scenario neither could move by themselves. Most current robot platforms are fairly slow, and as robot hardware continues to evolve the ability for a robot to physically interact with a human as a teammate at optempo will be important. This is especially true as lighter, faster robots are developed that don't have the same payload capacity as (heavy, power-consuming) highly geared or hydraulic robots have. For this reason, we propose to develop whole-body interactive control for HRI. By predicting human intent for specific objects being manipulated, dynamic tasks that two humans can already naturally do (such as pushing or tossing rubble to clear an area, blocking or unblocking a door, breaking down a door) can be done by a human with a robot teammate. We will also extend our preliminary work on manipulating rigid objects to collaboratively manipulating non-rigid objects such as moving a wounded soldier or boosting another soldier up to be able to climb a wall or look through a window.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.0 Integrated Research

5.1 Overview

Integrated Research provides standard platforms and a software environment to demonstrate research results on robots. Integration is an ongoing, iterative process where the goal is to efficiently capture implemented capabilities using a common development environment and repository. We leverage continuous integration software development processes to reduce integration risk by early identification of interface and dependency issues. The continuous integration process and common development environment should reduce rework and code refactoring and lessen the software integration cycle time.

In addition, we have defined and maintain code and capability module standards. These efforts will make it easier for new researchers to pick-up and advance previous research or to more quickly add needed capabilities. The end result is to provide the Army with a living, maintainable, software capability that can be readily used for future Army needs.

There is an immediate impact to the researchers since they will need to learn the new software development and configuration control environment, but the required short term cost will be small compared to the bigger long term payoff.

Integrated Research is the area which defines the RCTA as a robotics program, rather than a collection of research in AI, machine vision, mechanisms, and team cognitive psychology. The requirement that systems operate on robotic platforms in challenging environments forces the consideration of problems such as interfaces, uncertainty, real-time performance, non-ideal conditions, and interactions between the collected research components. These are problems that are sometimes seen as uninteresting in a detached research environment, but are critical to that research being relevant to autonomous robotics.

State of the Art and Technical Barriers:

Robotic platforms for research and development are becoming widespread. Many such platforms provide a mobile base that can be controlled via a software API. Some come with robotic manipulators, or can easily be connected to them. However, few of these platforms provide support for high-level intelligent mobility, perception, or manipulation. RCTA researchers need a state of the art software development environment that provides a foundation that supports multiple computer, multiple process, real time, inter- and intra-process communications, with software tools that support debugging, profiling, logging, and visualizing the data being collected.

In addition the robot platforms need to have computational density for the very computational intensive algorithms for example perception processing, path planning in highly dynamic environments, and adaptive mission planning.

Moving Beyond the State of the Art:

Our approach in Integrated Research is to bring the technology together on common platforms, , using a common Intelligence Architecture and Open Architecture Infrastructure. This infrastructure supports a common software development environment and collaboration tools, to supply a common module structure based on common coding standards.

This architecture has been developed by the research tasks in I1 and I2, to provide a framework that supports advanced semantic reasoning, metric robotic behaviors, a rich world model containing metric and semantic data, and control algorithms for manipulation and mobility. This architecture is designed to be compatible with multiple protocols used by researchers (such as ROS, NML, LUMP and DDS), and to support the Open Architectures initiative with its exhortations against closed, proprietary systems. We developed the architecture to run on a freely available operating system (Ubuntu Linux). It runs on the Husky, ROMAN, ROBOSIMIAN, and will also run on the LAMA platform.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Integrated Research Assessments (IRAs) are designed to evaluate our technologies in a relevant environment. This means that the subsystem has to operate outside laboratory settings; indoors or outdoors, on restricted computing hardware (that may be shared with other tasks), under real-time constraints, and subject to interactions from other parts of the system. These additional difficulties often illustrate ways in which algorithms and hardware can be usefully improved, and provide insight into the limits of the system as a whole, as well as its various subsystems. The process of integrating these technologies is inherently collaborative, and gets our researchers working together outside the typical meeting and workshop environments.

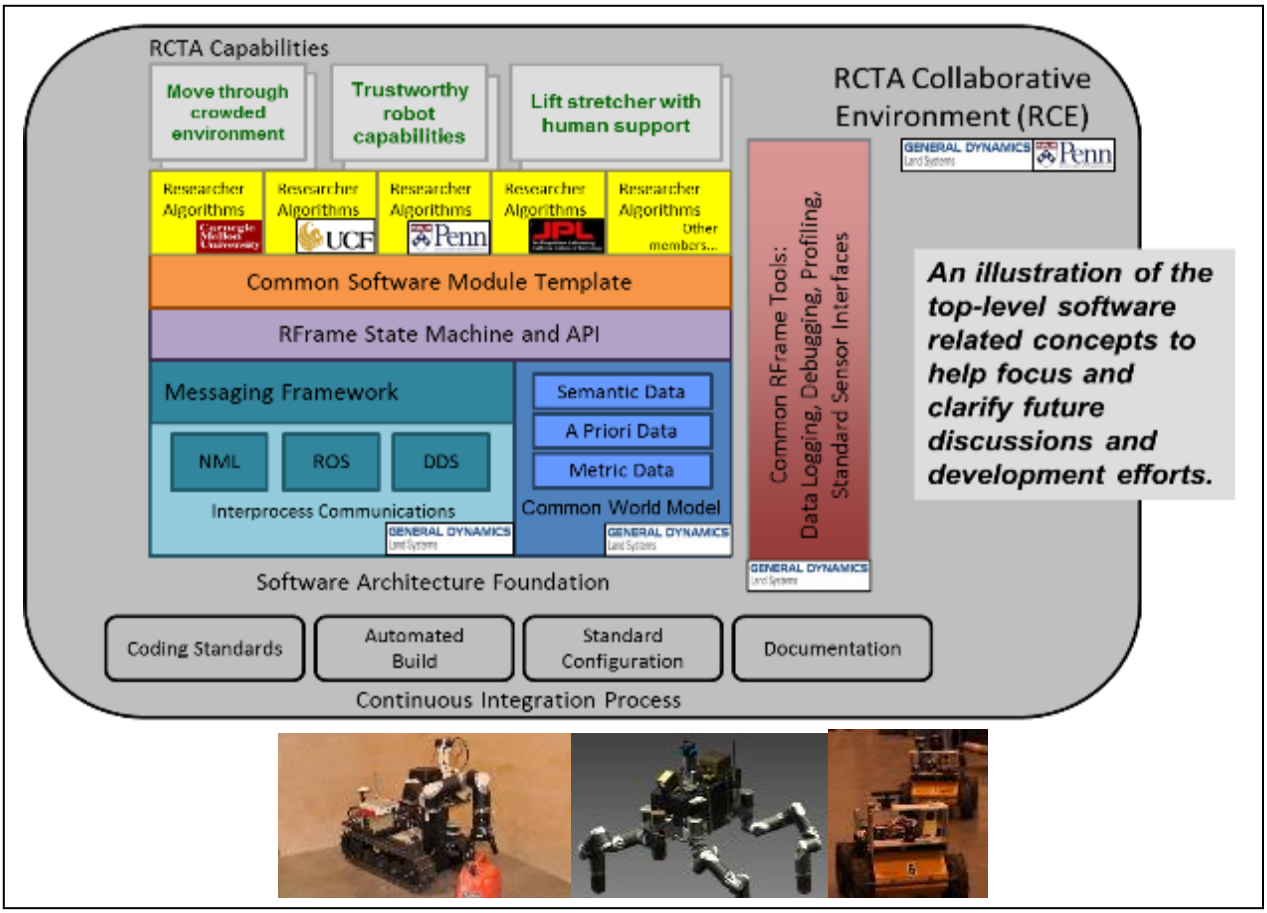

Figure 6 RCTA Integrated Research provides a state-of-the-art environment to facilitate researchers to rapidly develop algorithms and capabilities and implementations and test them on highly capable platforms.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Integrated Research – Key Tasks

RCTA Collaborative Environment

GDLS Mike Henke, UPenn IT

As part of our efforts to adopt a continuous integration philosophy, the RCTA Collaborative Environment (RCE) has been established. Based on the collaborative software tools of Atlassian, the RCE infrastructure is hosted at UPenn. UPenn will perform the hardware and software maintenance, system backup, security, and system access control. GDLS will configure, manage, organize, and provide access control for the software development tools and process.

A key task in this effort is to establish a process workflow to transition research from laboratory into the RCE common development environment. The process will allow the RCTA to take full advantage of the functionality provided by the Atlassian tools, and will guide researchers through software development steps, algorithm and code documentation, and the cyclical continuous integration process. Initially, the effort will be centered on the transition of existing systems from each institution to the centralized RCE. Once these transitions have taken place, the automated build processes and unit and integration tests can be applied to their full advantage. Documentation of assessments, including data sets and experimental results, will also be captured in the RCE. The major portion of this subtask is in designing this software

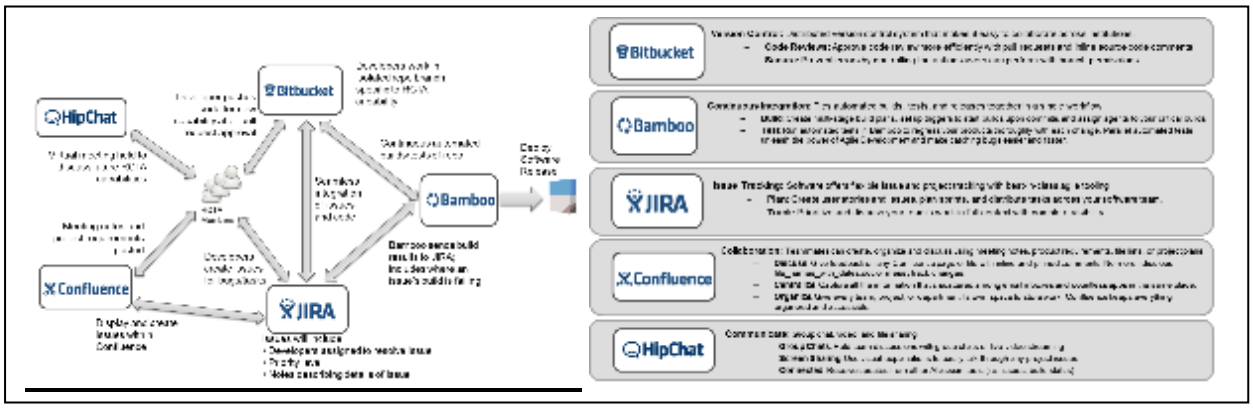

delivery, documentation, and test process.

Figure 7 RCTA Collaboration Environment - provide common database, configuration control, build tools, issue tracking.

An important aspect of this task that should be emphasized is the user support requirements. Researchers will need help transitioning from existing systems as well as in moving new algorithms into the RCE framework. Neglecting this aspect of the task will result in a collaborative environment in name only that produces few of the benefits described.

Related subtasks that stand to benefit from this work: This will provide support for all thrusts and capabilities being developed on the program.

Software Infrastructure Support

GDLS/ARL

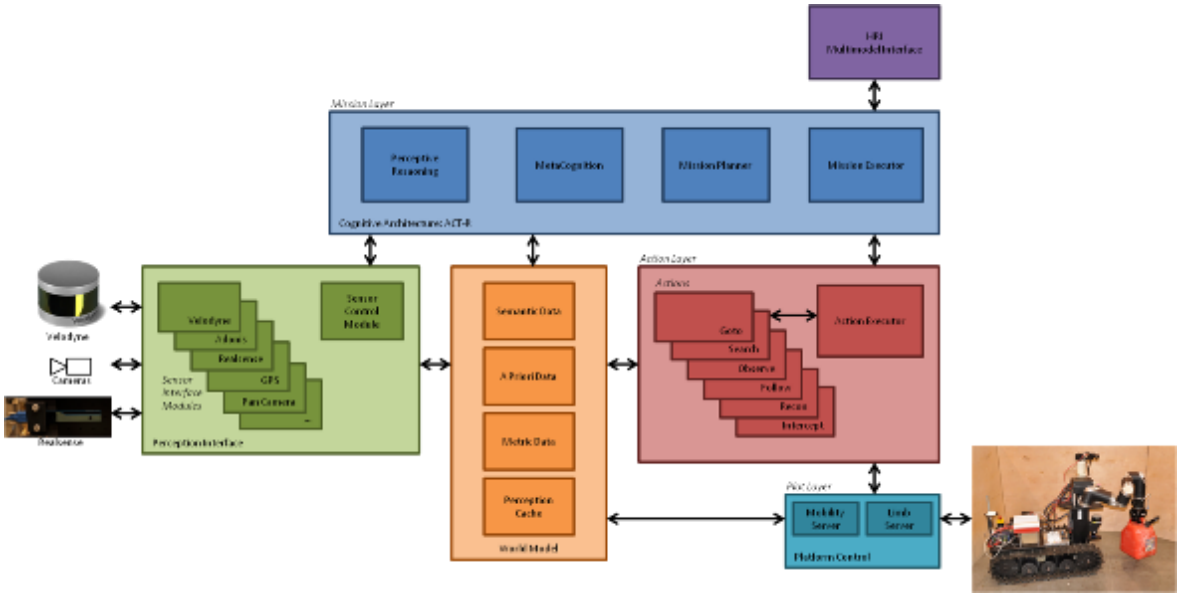

Figure 8 Software support efforts are identifiable in this high level picture of the RCTA Intelligence Architecture a) Perception Sensor Interfaces b) World Model Infrastructure Enhancements, c) Robotic Actions, d) Platform and Limb Control Enhancement

Perception Sensor Interfaces

When new and better sensors become available and are applicable to the RCTA capabilities being developed the software infrastructure support team will work to develop the sensor interfaces required and integrate the sensor onto the platform. Ensuring that the mechanical, electrical and software interfaces are implemented will remove the burden from the researchers. In addition as new perception capabilities become available the integration team will fold those capabilities into the system. Such as advancement in simultaneous mapping and localization, visual odometry, structure from motion, object detection and segmentation, optical flow, outlier detection and handling.

World Model Infrastructure Enhancements

There are several aspects related to the world model that need to be researched including supporting more types of representations and better ways to abstract the world model data so it can be shared more efficiently, better ways to support a distributed world model and identifying the relevant pieces of the world models that need sharing. The world model management implementation needs to support multi-threaded parallel processing for updates so that access to the world model can quickly handle multiple sensors updates, data and world model alignment, and data abstraction, data storage, and support high level analysis of the world context.

The data abstraction needs to include pushable and deformable object labeling along with labeling to support tall grass and shrubs that are, and are not, traversable.

Robotic Actions

The software infrastructure and support team will continue to develop new actions or skills that can then can be sequenced and utilized by the mission layer as needed to accomplish the mission.

Platform and Limb Control Enhancement

The software infrastructure and support team will continue to develop and refine the platform and limb control. Enhancing the robustness, speed, controllability of the platform and limb actuators.

Related subtasks that stand to benefit from this work: This will provide support for all thrusts and capabilities being developed on the program.

Modeling and Simulation Support

The objective of this subtask is to enable the use of simulation to decrease the iteration time, facilitate repeatable testing, and overcome the limited availability of robotic platforms and sensors. The thrust of this work is to increase the usage and capabilities of simulation for development and integration throughout RCTA. This will be accomplished by the addition of new platforms, improved sensor models, and additional human/machine interfaces and controls. Additionally, the simulation team will work closely with developers and the integrated research team to provide user support required for development and testing.

Historically, our approach to simulation has been exemplified by the RIVET simulation environment. Designed for robotic vehicle simulation, and based on the Torque 3D game engine, the RIVET environment has many advantages, including several years of user experience; detailed CACTF terrain models; multiple vehicle, sensor, and robot models developed especially for RCTA; and realistic human non-player character models. Although RIVET has been a workhorse simulation for developing sensor perception algorithms and testing and integration of planners, as well as in the area of human-robot interaction, the growing emphasis on team interaction and manipulation has started to reveal its limitations. The problem is one of scale: the fundamental scenario assumptions of a vehicle driving in a fairly large area begin to break down when the scale of the scenario shifts to identifying gestures and grasping small articles like door knobs. The RIVET environment is based on modeling large drivable areas, and introducing smaller, high-fidelity details required by these new tasks is difficult given the constraints of the underlying game engine. Rather than try to make RIVET a universal tool, we will examine other simulation frameworks in order to identify if these might be of use in RCTA and determine how these other tools might provide the necessary functionality. Notable candidates include Gazebo and MATLAB/Simulink. In particular, we will provide tools to support integration and regression testing to deploy in a continuous integration environment.

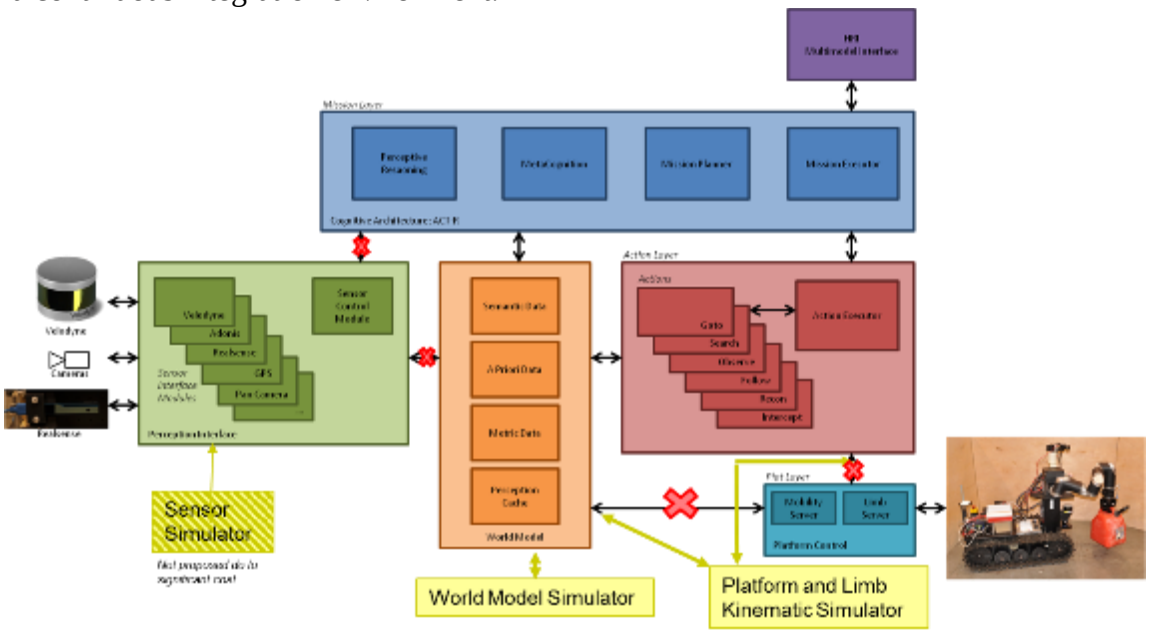

Figure 9 The RCTA modeling and simulation capability would enable the ability to provide a powerful development tool and provide some automated regression testing capability.

Related subtasks that stand to benefit from this work: This will provide support for all thrusts and capabilities being developed on the program.

Software Integration

Description: This subtask is to support researchers in developing their software, delivering their software to the central RCTA repository, and integrating it into the system. This means providing platforms for testing when appropriate, setting up and verifying a test bed that can be duplicated by researchers for their testing, and supporting organized integration and assessment activities. This sub-task provides support to the researchers for carrying out their particular sub-tasks, including for Task Based Assessments and/or Integrated Research Assessments (IRA), as appropriate. Task Based Assessments will be used as entry points for higher level integration activities and IRAs. It is important to make sure that technology is functioning at a reasonable level to avoid wasting researchers' time due to immature technology.
This effort will use the RCTA Collaborative Environment (RCE) to integrate software capabilities, verify interface compliance, document test procedures and expected operating parameters and limitations, and perform system regression testing. The major portion of this sub-task, however, is in aiding researchers in delivering, integrating, and testing the resulting software.
As part of this effort, and leveraging the development and results of the simulation sub-task, an integrated test-bed for use at researcher's desks and as an automated regression test environment will be developed. The integrated test-bed environment, based on the RCTA software architecture and framework will support verification of intelligence and behavior algorithms using simulated or recorded data. It is expected that RCTA software that enters into this process for incorporation in the assessment code base will have undergone basic researcher level or task-based assessment using this test-bed prior to being provided for integration. Basic integration problems, including interface, communication, data interpretation, execution timing, and numerical or algorithmic problems can be detected and isolated easily in this environment. This checkout process will occur based on collaboration with researchers providing technology to be integrated and in concert with the ongoing field integration events.

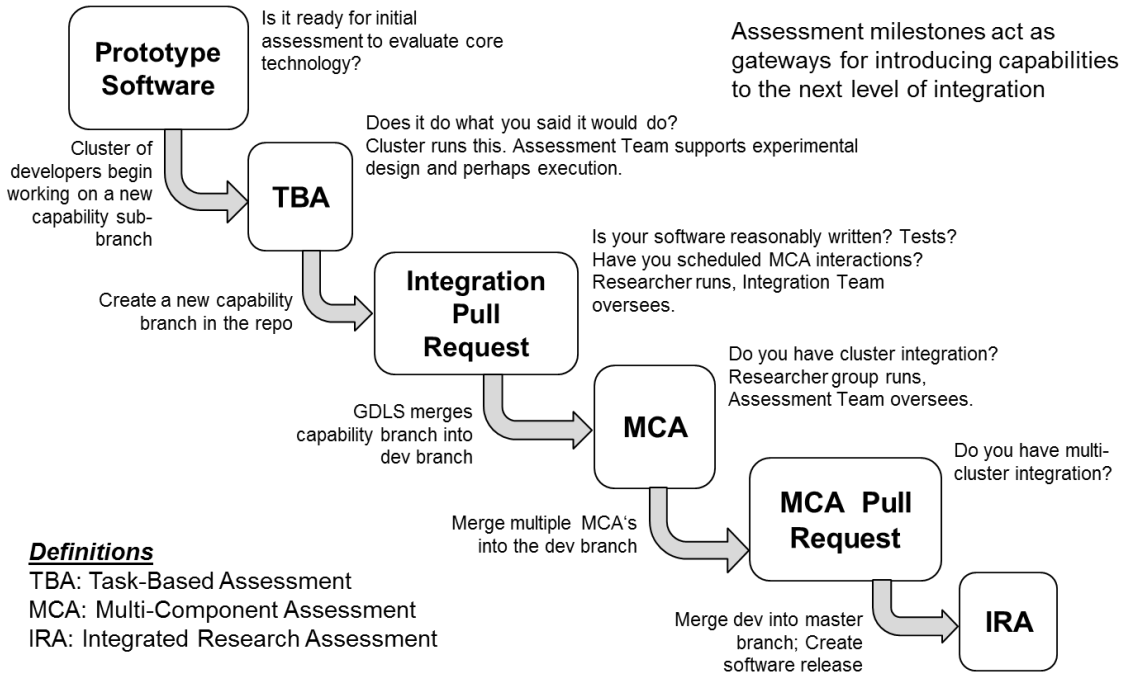

Figure 10 Software Integration Process

Related subtasks that stand to benefit from this work: This will provide support for all thrusts and capabilities being developed on the program.

Platform Enhancements and Support Logistics

The RCTA has 10 Husky and 4 ROMAN platforms that are being used by the RCTA research community. The effort is to continue to upgrade the existing platforms as needed to enhance the sensor capability, increase the computational ability, and increase the platform power storage to enable longer run times. In addition there will be additional ROMAN, ROBOSIMIAN, and LAMA platforms developed in the future. The task will focus on keeping the vehicles in working order and making sure that they are systematically upgraded.

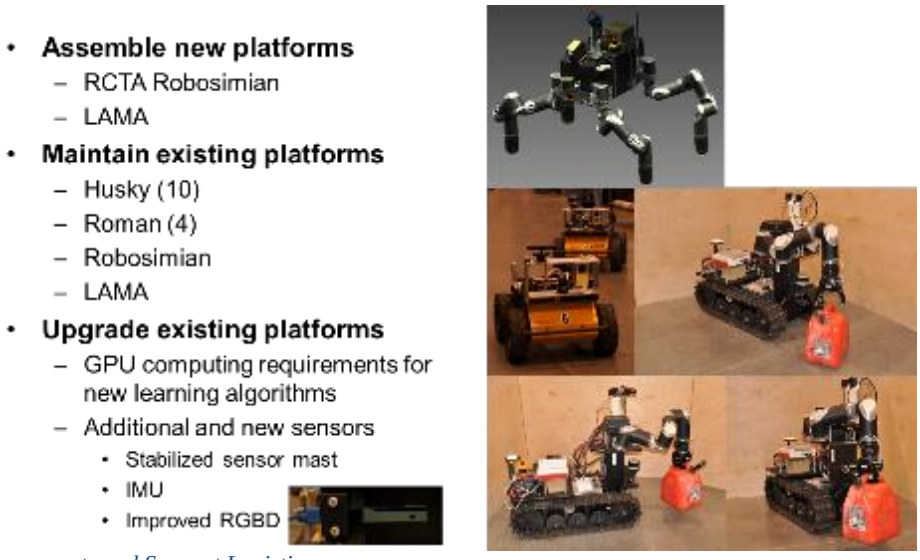

Figure 11 Platform Enhancements and Support Logistics

Related subtasks that stand to benefit from this work: This will provide support for all thrusts and capabilities being developed on the program.

5.2 Thrust 4 Subtasks

5.2.1 T4C1S1A - Maintain current platforms – Husky, Roman, Robosimian

PI: Mark Gonzalez, GDLS, gonzalm2@gdls.com, (410)386-1741

Objective and Benefits:

The RCTA has 10 Husky and 4 Roman platforms that are being used by the RCTA research community. The effort is to continue to upgrade the existing platforms as needed to enhance the sensor capability, increase the computational ability, and increase the platform power storage to enable longer run times. In addition there will be additional Roman, Robosimian, and LAMA platforms developed in the future. The task will focus on keeping the vehicles in working order and making sure that they are systematically upgraded.

Connection to past RCTA work:

This is an ongoing RCTA task that was previously funded under subtask IR3, Platforms and Testbeds for Integrated Research.

Collaborations:

- *Tech transfer cross- collaboration*
  - Karumanchi, JPL - T4C1S1B - Maintain current platforms - Robosimian
  - Taylor, QNA - T4C1S1C - Maintain current platforms - TL1 – Roman
  - ARL, T4C1S1D - Platform maintenance support
  - Karumanchi, JPL - T3C1S3C Dual Arm Contact Driven Mobile Manipulation Behaviors on the Roman Platform

Deliverables

Year	Goal	Benefits
2017	• Get Robosimian platform built, tested, and integrated with RCTA system • Upgrade Husky, Roman platforms with better sensors and GPU processing to support DNN processing. • Add gimbal sensor stabilization to a Roman platform and test and measure performance. • Test various limb configurations to determine appropriate motor, controller, and gearing selection and sizing. • Add audio sensing and output to the platforms.	Husky, Roman, and Robosimian platforms with an enhanced common computer and sensor systems with integrated RCTA software capability to support RCTA researchers.
2018	• Fully functional LLAMA platform with integrated RCTA computer system. • Perform Husky, Roman, Robosimian computer power and sensor upgrades as needed.	Husky, Roman, Robosimian, LLAMA platforms with an enhanced common computer and sensor systems with integrated RCTA software capability to support RCTA researchers.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

2019	• Maintain the platforms make upgrades as needed. Increasing computational power and sensor performance while reducing size weight and power usage.	Husky, Roman, Robosimian, LLAMA platforms with an enhanced common computer and sensor systems with integrated RCTA software capability to support RCTA researchers.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.2 T4C1S1B - Maintain current platforms – RoboSimian

PI: Karumanchi, Sisir (JPL/Caltech), Sisir.B.Karumanchi@jpl.nasa.gov, (818)356-3867
Additional PI: Bergh, Charles (JPL/Caltech), cfb@jpl.nasa.gov, (818)354-5567

Research Description:

Objective and Benefits:
T4C1S1B will provide hardware design, integration support for the RCTA RoMan and RoboSimian platforms. In addition, ongoing support will be provided for software development and integration support for current and future RCTA integration research platforms.

Metrics & Progress Assessment:

(a) Code delivery, Completion of software and hardware integration for the Fall IRA.

(c) Energy source, actuation and systemic design specifications for next gen robot.

(d) Access to JPL's expertise in creating the materializing the next generation platform designs.

Specific Contributions to the Cognitive Architecture:
This task will provide ongoing support towards the Roman Platform within RCTA. The efforts under this task along with research performed under Subtask T3C1S3C will make the following contributions to the cognitive architecture. Further, efforts made under task T1C2S5A for Llama legged platform development will help us generalize architectural design to be consistent across locomotion and mobile manipulation.

Connection to past RCTA work:

This is an ongoing RCTA task that was previously funded under subtask IR3.7 for the Roman Platform design and development.

Collaborations:

- *Leveraging JPL/Caltech subtasks for software support*
  - Karumanchi - T3C1S3C Dual Arm Contact Driven Mobile Manipulation Behaviors on the Roman Platform
  - Burdick - T3C1S3D Mobility-aided manipulation on the Roman platform
  - Matthies - T1C2S2A - Perception for Maneuver in Complex Terrains
  - Papon/Detry - T3C1S2B - Generalizing grasps across objects and tasks

- *Tech transfer cross- collaboration*
  - Gonzalez, GDLS - T4C1S1A - Maintain current platforms - Husky, RoMan, RoboSimian
  - Taylor, QNA - T4C1S1C - Maintain current platforms - TL1 – RoMan
  - ARL, T4C1S1D - Platform maintenance support

Deliverables

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

	2017 Quarterly Milestones and Deliverables
Q1	Travel to Train GD personnel in Actuator Build, Deliver Planning and Control Software for RoMan platform. Maintain a mirror RoMan-like /Surrogate Platform or RoboSimian Platforms at JPL for mirrored testing.
Q2	Travel for MCA at GD.
Q3	Ongoing support of RoMan platform prior for Fall IRA.
Q4	Travel for IRA. Mirrored testing with JPL's Surrogate/RoboSimian platform to complement FITG testing and testing at GDLS.

	2018 Quarterly Milestones and Deliverables
Q1	Travel to Train GD personnel in Actuator Build, Deliver Planning and Control Software for RoMan platform. Maintain a mirror RoMan-like /Surrogate Platfrom or RoboSimian Platforms at JPL for mirrored testing.
Q2	Travel for MCA at GD.
Q3	Ongoing support of RoMan platform prior for Fall IRA.
Q4	Travel for IRA. Mirrored testing with JPL's Surrogate/RoboSimian platform to complement FITG testing and testing at GDLS.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.3 T4C1S1C – Support and Maintain TL1/ROMAN

PI: Taylor, Thaddeus (QNA), thaddeus.taylor@qinetiq-na.com, (412)449-1518

This task is a collaboration with GDLS that supports the TL1 platform. This support includes assistance integrating various sensors and supporting equipment on to the TL1, redesigning and upgrading the platform as requested by various stakeholders, and providing repair and replace as needed and requested by GLDS.

Relevance

The TL1 is the current non-ITAR platform of choice for RCTA given its mobility profile, its ability to carry weight, and its power. Supporting all RCTA tasks that utilize the TL1 helps ensure the success of the program.

Previous Subtasks

- IR2-2 – Collaboration Integration Support

MILESTONES

Quarter	Milestones
2017 Q1	Continued technical support of the TL1/ROMAN platform as requested by GDLS, funding dependent
2017 Q2	Continued technical support of the TL1/ROMAN platform as requested by GDLS, funding dependent
2017 Q3	Continued technical support of the TL1/ROMAN platform as requested by GDLS, funding dependent
2017 Q4	Continued technical support of the TL1/ROMAN platform as requested by GDLS, funding dependent
2018 Q1	Continued technical support of the TL1/ROMAN platform as requested by GDLS, funding dependent
2018 Q2	Continued technical support of the TL1/ROMAN platform as requested by GDLS, funding dependent
2018 Q3	Continued technical support of the TL1/ROMAN platform as requested by GDLS, funding dependent
2018 Q4	Continued technical support of the TL1/ROMAN platform as requested by GDLS, funding dependent

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.4 T4C1S1D – Maintain Current Platforms

PI: Westrich, Jeff (ARL), jeffrey.r.westrich.ctr@mail.mil, (410)278-0541
Additional PIs: Carpenter, Howard (ARL), howard.l.carpenter.ctr@mail.mil, (410)278-3942

Objective and Benefits:

The RCTA has 10 Husky and 4 Roman platforms that are being used by the RCTA research community. The effort is to continue to upgrade the existing platforms as needed to enhance the sensor capability, increase the computational ability, and increase the platform power storage to enable longer run times. In addition there will be additional Roman, Robosimian, and LAMA platforms developed in the future. The task will focus on keeping the vehicles in working order and making sure that they are systematically upgraded.

Connection to past RCTA work:

This is an ongoing RCTA task that was previously funded under subtask IR3, Platforms and Testbeds for Integrated Research.

Collaborations:

- *Tech transfer cross- collaboration*
  - Gonzalez, GDLS - T4C1S1A - Maintain current platforms - Husky, Roman, Robosimian
  - Karumanchi, JPL - T4C1S1B - Maintain current platforms - Robosimian
  - Taylor, QNA - T4C1S1C - Maintain current platforms - TL1 – Roman
  - Karumanchi, JPL - T3C1S3C Dual Arm Contact Driven Mobile Manipulation Behaviors on the Roman Platform

Deliverables

Year	Goal	Benefits
2017	• Get Robosimian platform built, tested, and integrated with RCTA system • Upgrade Husky, Roman platforms with better sensors and GPU processing to support DNN processing. • Add gimbal sensor stabilization to a Roman platform and test and measure performance. • Test various limb configurations to determine appropriate motor, controller, and gearing selection and sizing. • Add audio sensing and output to the platforms.	Husky, Roman, and Robosimian platforms with an enhanced common computer and sensor systems with integrated RCTA software capability to support RCTA researchers.
2018	• Fully functional LAMA platform with integrated RCTA computer system. • Perform Husky, Roman, Robosimian computer power and sensor upgrades as needed.	Husky, Roman, Robosimian, LAMA platforms with an enhanced common computer and sensor systems with integrated RCTA software capability to support RCTA researchers.
2019	• Maintain the platforms make upgrades as needed. Increasing computational	Husky, Roman, Robosimian, LAMA platforms with an enhanced common

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

	power and sensor performance while reducing size weight and power usage.	computer and sensor systems with integrated RCTA software capability to support RCTA researchers.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.5 T4C2S1A – RCTA Collaborative Environment (RCE) and SW Management

PI: Henke, Mike (GDLS), henkem@gdls.com, (240)446-9745

Objective and Benefits:

As part of our efforts to adopt a continuous integration philosophy, the RCTA Collaborative Environment (RCE) has been established. Based on the collaborative software tools from Atlassian, the RCE infrastructure will be hosted at UPenn. UPenn will be providing the infrastructure: hardware and software maintenance, system backup, security, and system access control. GDLS will configure, manage, organize, and provide access control for the software development tools and process.

SubTask Description

A key task in this effort is to establish a process workflow to transition research from laboratory into the RCE common development environment. The process will allow the RCTA to take full advantage of the functionality provided by the Atlassian tools, and will guide researchers through software development steps, algorithm and code documentation, and the cyclical continuous integration process. Initially, the effort will be centered on the transition of existing systems from each institution to the centralized RCE. Once these transitions have taken place, the automated build processes and unit and integration tests can be applied to their full advantage. Documentation of assessments, including data sets and experimental results, will also be captured in the RCE. The major portion of this subtask is in designing this software

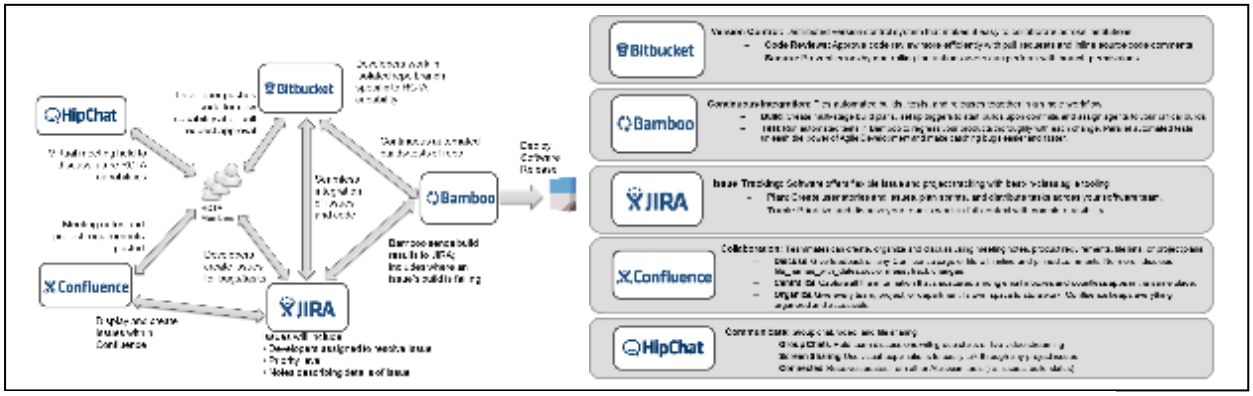

Figure: RCTA Collaboration Environment - provide common database, configuration control, build tools, and issue tracking.

delivery, documentation, and test process.

An important aspect of this task that should be emphasized is the user support requirements. Researchers will need help transitioning from existing systems as well as in moving new algorithms into the RCE framework. Neglecting this aspect of the task will result in a collaborative environment in name only that produces few of the benefits described.

Connection to past RCTA work:

IR5-6, IR5-1

Collaborations:

- Widyono, UPenn – T4C2S1B RCE Server and Application Maintenance
- Fink, ARL – T4C2S1C RCE and SW Management

Deliverables

RCE Process and Workflow Workshop
Software Integration Framework and Workshop
Framework and Architecture Upgrade Support
RCE Test Environment Role Out
RCE Integration Delivery Support for TBA, MRA, IRA

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.6 T4C2S1B - RCE Maintenance and Integration Support

PI: Kostas Daniilidis (UPenn), kostas@cis.upenn.edu, (215)898-8549
Additional PI: Widyono, Daniel, widyono@seas.upenn.edu, (215)898-4707

Objective and Benefits:
UPenn integration team has two objectives: (1) maintaining and monitoring the RCTA Collaborative Environment (RCE), and (2) Provide **integration support** for the UPenn research.

RCTA Collaborative Environment (RCE)
UPenn has purchased server hardware, non-cloud based storage media, and the Atlassian framework. It now hosts the RCTA Collaboration Environment (RCE) system. UPenn will be responsible for installing applications as listed below and maintaining the environment for the alliance.
The Atlassian based collaboration applications selected by the alliance are:

- Confluence (Knowledge Management)
- Bamboo (Continous Integration)
- Bitbucket (Development Workflow)
- Hipchat (Chat tool)
- Low level help desk support

UPenn will manage this environment under the guidance of the RCTA Alliance and Consortium lead managers. These managers will determine RCE access, best use practice, and assist UPenn is resolving pertinent problems.

Integration Support:
In addition, UPenn will designate an engineer or researcher to be the integration support liaison between UPenn researchers and the RCTA integration team.

Connection to past RCTA work: IR5-5

Collaborations:

- The entire consortium will be using the RCE but specifically the UPenn Team works with GDLS.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.7 T4C3S1A - World Model Maintenance and Enhancements

PI: Patel, Jaymit (GDLS), patelj1@gdls.com, 410-386-1762

Objective and Benefits:

There are several aspects related to the world model that need to be researched including:

- Supporting more types of representations
- Better ways to abstract the world model data so it can be shared more efficiently
- Better ways to support a distributed world model
- Identifying the relevant pieces of the world models that need sharing

To support more types of representations, voxels will be integrated at the metric data level to directly enable additional perception algorithms. The world model management implementation needs to support multi-threaded parallel processing for updates so that access to the world model can quickly handle multiple sensor updates, data and world model alignment, as well as data abstraction, data storage, and support high level analysis of the world context.

The data abstraction needs to include pushable and deformable object labeling along with labeling to support tall grass and shrubs that are, and are not, traversable.

Additionally, the software infrastructure and support team will continue to develop new actions or skills that can then can be sequenced and utilized by the mission layer as needed to accomplish the mission. Likewise, the team will continue to develop and refine the platform and limb control in order to enhance the robustness, speed, and controllability of the platform and limb actuators.

1. **Connection to past RCTA work:**
   This is an ongoing RCTA task that was previously funded under subtask IR5-3.

2. **Collaborations:**
   - Matthies, JPL – T1C2S2A – Perception for maneuver in complex terrain
   - Henke, GDLS – T1C2S2B – Perception for optempo mobility in rough terrain
   - Gonzalez, GDLS - T4C1S1A - Maintain current platforms - Husky, Roman, Robosimian
   - Taylor, QNA - T4C1S1C - Maintain current platforms - TL1 – Roman
   - ARL, T4C1S1D - Platform maintenance support

3. **Deliverables**
   a. World model is realized by using it in a multi-component and/or IRA
   b. Mission layer can functionally enable actions through the architecture, as exemplified in a multi-component and/or IRA.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.8 T4C3S1B – Perception Sensor and Computing Interfaces

PI: Henke, Mike (GDLS), henkem@gdls.com, (240)446-9745
Additional PI: Quang, Long (GDLS), quang@gdls.com, (410)386-1763

Objective and Benefits:

The integration team will work on finalizing sensor selection, developing and supporting a robust software infrastructure for a standard platform sensor suite. In addition, the integration team will work with and provide support to researchers for their sensor interface needs. Ensuring that the mechanical, electrical and software interfaces are implemented will remove the burden from the researchers. In addition, as new perception capabilities become available, the integration team will fold those capabilities into the system and maintain them throughout the program. Projected new capabilities are advancement in simultaneous mapping and localization, visual odometry, structure from motion, object detection and segmentation, optical flow, and outlier detection and handling.

Likewise, the computing interfaces of both perception and detection algorithms must be upgraded and improved as more sophisticated algorithms are deployed. Typical image processing solutions lean heavily on the availability of GPUs, so this subtask will investigate appropriate GPU computing solutions that meet the SWAP requirements for deploying on a robotic platform.

Connection to past RCTA work:

This is an ongoing RCTA task that was previously funded under subtask IR3-6, On Platform Perception, Communication, and Computing

Collaborations:

- Gonzalez, GDLS - T4C1S1A - Maintain current platforms - Husky, RoMan, RoboSimian
- Taylor, QNA - T4C1S1C - Maintain current platforms - TL1 – RoMan
- ARL, T4C1S1D - Platform maintenance support

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.9 T4C3S1C – Capability Integration

PI: Weller, Ed (GDLS), weller@gdls.com, (410)386-1778

Objective and Benefits:

This subtask is to support researchers in developing their software, delivering their software to the central RCTA repository, and integrating it into the system. This means providing platforms for testing when appropriate, setting up and verifying a test bed that can be duplicated by researchers for their testing, and supporting organized integration and assessment activities. This sub-task provides support to the researchers for carrying out their particular sub-tasks, including for Task Based Assessments and/or Integrated Research Assessments (IRA), as appropriate. Task Based Assessments will be used as entry points for higher level integration activities and IRAs. It is important to make sure that technology is functioning at a reasonable level to avoid wasting researchers' time due to immature technology.

This effort will use the RCTA Collaborative Environment (RCE) to integrate software capabilities, verify interface compliance, document test procedures and expected operating parameters and limitations, and perform system regression testing. The major portion of this sub-task, however, is in aiding researchers in delivering, integrating, and testing the resulting software.

As part of this effort, and leveraging the development and results of the simulation sub-task, an integrated test-bed for use at researcher's desks and as an automated regression test environment will be developed. The integrated test-bed environment, based on the RCTA software architecture and framework will support verification of intelligence and behavior algorithms using simulated or recorded data. It is expected that RCTA software that enters into this process for incorporation in the assessment code base will have undergone basic researcher level or task-based assessment using this test-bed prior to being provided for integration. Basic integration problems, including interface, communication, data interpretation, execution timing, and numerical or algorithmic problems can be detected and isolated easily in this environment. This checkout process will occur based on collaboration with researchers providing technology to be integrated and in concert with the ongoing field integration events.

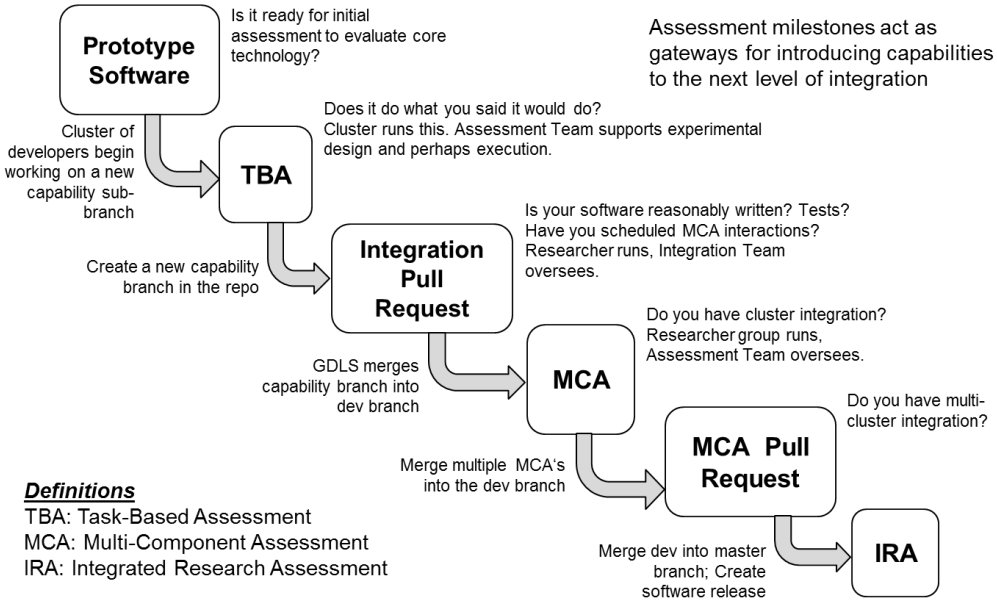

Figure: Software Integration Process

Connection to past RCTA work: IR2-1

Collaborations:

All capability subtasks with integration deliverables will benefit from this subtask.

Deliverables

	2017 Quarterly Milestones and Deliverables
Q1	RCE and Process Workflow Workshop
Q2	Software Integration Framework
Q3	RCE Test Environment Complete
Q4	RCE Integration Delivery for Initial Capability IRA

	2018 Quarterly Milestones and Deliverables
Q1	Software Integration Workshop
Q2	Interim Framework and Architecture Upgrade
Q3	RCE Integration Test Environment Complete
Q4	RCE Integration Delivery for Interim Capability IRA

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.10 T4C3S1D – Human Location and Platform Interfaces

PI: Sapronov, Lenny (Robotic Research), sapronov@roboticresearch.com, (240)631-0008

Objective and Benefits:

This subtask is aimed at developing and integrating several new capabilities for the RCTA common platforms. The first is WarLoc, a foot-worn device for localizing human assets. Another is a common platform web interface to facilitate configuration and testing of the Husky and Roman platforms. Other work that may fall under this subtask includes development of new software modules to support intelligent behaviors for the IRA and Capstone demos (e.g. coordinated planning).

WarLoc integration will enable testing of human detection algorithms and tactical teaming behaviors, specifically by allowing comparison of detection algorithm output with ground truth locations determined through WarLoc.

The platform interface aspect of this subtask includes supporting general hardware and software integration on RCTA common platforms (Husky and TL1/Roman). This includes software drivers, autonomous mobility software, hardware configuration (nav units, lidars), as well as any improvements to these modules over the course of the platforms' lifetimes. The goal of this effort is to ensure proper operation of the platforms, adequate mobility and reliable operation in support of tests and assessments conducted by other consortium members.

Connection to past RCTA work:

This is an ongoing RCTA task that was previously funded under subtask IR3-4, IR5-2.

Collaborations:

- *WarLoc*
  - Navarro-Serment – T1C1S2A: Modeling Interactions and Actions
  - Oh – T1C1S4B: Navigation among Dynamic, Reactive Agents
  - Navarro-Serment – T1C1S4D: Modeling Trajectories for Robust Perception
- *Platform Interfaces*
  - Gonzalez, GDLS - T4C1S1A - Maintain current platforms - Husky, Roman, Robosimian
  - Taylor, QNA - T4C1S1C - Maintain current platforms - TL1 – Roman
  - ARL, T4C1S1D - Platform maintenance support

Deliverables

- Support for ongoing integration efforts, which includes software updates, troubleshooting, new sensor integration
- ROS interfaces for any existing software modules that don't currently support ROS
- WarLoc ground nodes for more accurate localization. Currently, the WarLoc system relies only on vehicle-mounted ranging radios for trilateration. Addition of stationary ground nodes will improve position accuracy for applications where a high degree of accuracy is important. The exact number of ground nodes needed is TBD.
- Updated common platform website with more vehicle configuration and visualization options

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.11 T4C3S1E - HRI Driven Enhancements

PI: Barber, Daniel (UCF), dbarber@ist.ucf.edu, (407)882-1128

Task Description

Proposed Work (6.2): The goal for this task is to support development of prototype software and hardware that facilitate human-in-the loop HRI research with platforms demonstrating RCTA capabilities. Specific objectives include (1) supporting researchers use of RCTA platforms (e.g. Husky) within human-in-the-loop experiments by establishing baseline capabilities, runtime support, and team performance logging, (2) transitioning results from basic research into working software and physical products like the multimodal interface that interact with integrated platforms (3), aiding in the development of human-in-the-loop experiments incorporating integrated platforms in terms of capabilities available and ensuring IRB compliance, (4) assisting in collaborations with other thrust areas to add requirements to research to further advance the science human-robot collaboration, and finally (5) modifications to Multimodal Display application to incorporate dialogue and robot state management data to support integration within a team and delivery of new scene and mission reasoning information.

Thrust and Capability Supported

This research primarily will support the RCTA's Thrust IV/Capability 3: RCTA SW Infrastructure and Integration.

Relation to Previous RCTA 2015-2016 Research

This subtask primarily builds upon research performed in 2015-2016 under subtask IR3-3 *Multimodal Display.*

Collaborations

This subtask will involve extensive collaboration with HRI related subtasks, including:

- Barber (T2C1S4B *Adaptive multimodal communication)*
- Florian Jentsch (T2C1S4A: *A user-centered design (UCD) approach to creating usable naturalistic communication interfaces)*
- Beth Philips (T2C2S3B *Appropriate calibrations of trust for supporting Soldier-robot teaming*)
- Linda Elliott (T2C2S3C *Multimodal comms research*)
- Tal-Oron Gilad (T2C1S3C *Investigating Tactical Multi-modal Soldier-Robot Exchanges*)
- Michael Goodrich (T1C2S4C *Practical human-robot interaction for navigation and manipulation*)
- Jean Oh (T2C1S3A *Two-way translation between vision and language*)
- Matthew Walter (T2C1S3B *Generating natural language scene descriptions for shared situation awareness*)
- Hebert (T1C2S1D reasoning for scene understanding)

Scientific Approach

This task will:

- Support setup and running of integrated platforms in experiments for human-in-the-loop studies
- Transition HRI research into software and physical capabilities for integrated platforms
- Maintain HRI capabilities across platforms to support researchers
- Allocate runtime support for integrated platform(s)
- Transition basic research into working software on the physical platform
- Add support for performance logging of overall robot and human interactions/teaming

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

- Support additional dialogue management for integration with new perception and intelligence capabilities generating additional state and scene understanding during mission/task execution.
- Software prototypes will support all platforms supporting the RCTA architecture (e.g. Husky, Roman).

Product/Outcomes of this Research

- End of BPP Year 2017: Integrate HRI research results and recommendations across the RCTA program to develop human in-the-loop experimentation and facilitate transition of basic research into hardware and software platforms for successful demonstration of the RCTA's program end-goal.
- End of BPP Year 2018: Integrate HRI research results and recommendations across the RCTA program to develop human in-the-loop experimentation and facilitate transition of basic research into hardware and software platforms for successful demonstration of the RCTA's program end-goal.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.12 T4C3S1F – Capability Integration Support

PI: Childers, Marshal (ARL), marshal.a.childers.civ@mail.mil, (410)278-7996
Additional PI: Kaplan, Matthew (ARL), matthew.a.kaplan2.ctr@mail.mil,

Objective and Benefits:

This subtask is to support researchers in developing their software, delivering their software to the central RCTA repository, and integrating it into the system. This means providing platforms for testing when appropriate, setting up and verifying a test bed that can be duplicated by researchers for their testing, and supporting organized integration and assessment activities. This sub-task provides support to the researchers for carrying out their particular sub-tasks, including for Task Based Assessments and/or Integrated Research Assessments (IRA), as appropriate. Task Based Assessments will be used as entry points for higher level integration activities and IRAs. It is important to make sure that technology is functioning at a reasonable level to avoid wasting researchers' time due to immature technology.

This effort will use the RCTA Collaborative Environment (RCE) to integrate software capabilities, verify interface compliance, document test procedures and expected operating parameters and limitations, and perform system regression testing. The major portion of this sub-task, however, is in aiding researchers in delivering, integrating, and testing the resulting software.

As part of this effort, and leveraging the development and results of the simulation sub-task, an integrated test-bed for use at researcher's desks and as an automated regression test environment will be developed. The integrated test-bed environment, based on the RCTA software architecture and framework will support verification of intelligence and behavior algorithms using simulated or recorded data. It is expected that RCTA software that enters into this process for incorporation in the assessment code base will have undergone basic researcher level or task-based assessment using this test-bed prior to being provided for integration. Basic integration problems, including interface, communication, data interpretation, execution timing, and numerical or algorithmic problems can be detected and isolated easily in this environment. This checkout process will occur based on collaboration with researchers providing technology to be integrated and in concert with the ongoing field integration events.

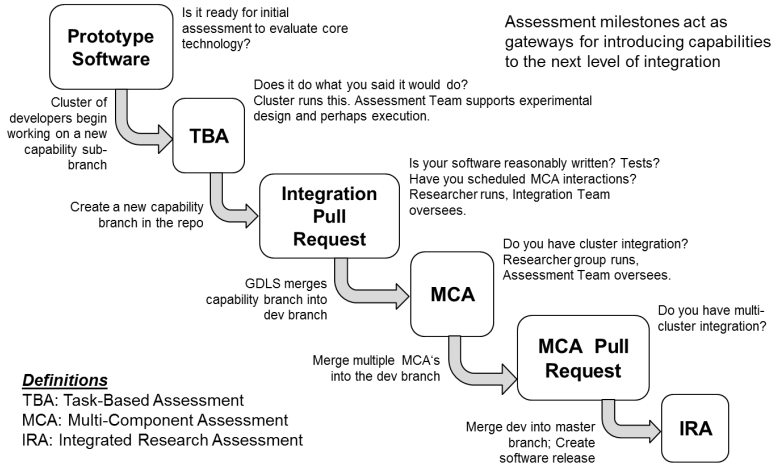

Figure: Software Integration Process

Connection to past RCTA work:

This is a new RCTA task but is related to earlier efforts funded under IR5, Collaborative Software Integration.

Collaborations:

All capability subtasks with integration deliverables will benefit from this subtask.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Deliverables

	2017 Quarterly Milestones and Deliverables
Q1	RCE and Process Workflow Workshop
Q2	Software Integration Framework
Q3	RCE Test Environment Complete
Q4	RCE Integration Delivery for Initial Capability IRA

	2018 Quarterly Milestones and Deliverables
Q1	Software Integration Workshop
Q2	Interim Framework and Architecture Upgrade
Q3	RCE Integration Test Environment Complete
Q4	RCE Integration Delivery for Interim Capability IRA

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.13 T4C3S1G – Capability Integration Support

PI: Roy, Nicholas (MIT), nickroy@csail.mit.edu, +1-617-875-2506

Objective and Benefits:
As the integration efforts of the RCTA expand to include more tasks, there will be substantial effort required to work with the world model to ensure support for relevant data types, and to assist with the interfaces between the perception, planning and natural language subtasks. The world model and RFrame will still reside with GDLS, but this subtask will involve working with the GDLS performers and the integration leads at the other institutions to ensure that the data and interfaces are consistent across the RCTA.

The primary responsibilities will include

- Ensuring that perception data can be transitioned from the perception subtasks in each capability (T1C1S1A, T1C1S1C, T1C1S3A, T1C1S1D, T1C1S2A, T2C1S2B, T1C2S2A, T1C2S3A, T3C1S1A, T3C1S2A, T3C1S2C, T2C1S2A, T2C1S2D) and incorporated into the planning models used by the planning subtasks (T1C1S4B, T1C1S4D, T3C1S1B, T1C2S1A, T1C2S4C, T2C2S1A, T2C2S1C, T2C2S1G, T2C2S1J, T3C1S1D, T3C1S2B, T3C1S3A, T3C1S3C, T3C1S3D)
- Ensuring that the perception data can be transitioned from the perception subtasks in each capability (T1C1S1A, T1C1S1C, T1C1S3A, T1C1S1D, T1C1S2A, T2C1S2B, T1C2S2A, T1C2S3A, T3C1S1A, T3C1S2A, T3C1S2C, T2C1S2A, T2C1S2D) and incorporated into the language understanding models (T2C1S3A, T2C1S3B, T2C1S4B, T2C2S2A, T2C2S1D, T2C2S2B, T2C2S2C
- Working with the other integration leads to develop simulation-based approximations to the RCTA capabilities for testing
- Working with the other integration leads to ensure appropriate acceptance and regression tests for the higher-level computational processes
- Ensuring data collection and integrity exists for the learning tasks.

List of dependencies:
The dependencies are listed above.

Collaboration:
This subtask will support a collaboration with Jon Fink and Ethan Stump at ARL on extending the world model.

Linkage to previous RCTA BPP sub-tasks: I5-3, H1-4

Expected products:

5. Improved documentation of cross-RCTA data types and interfaces, including process flow
6. Improved simulation capabilities
7. Improved acceptance and regression tests for higher-level computational processes

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.14 T4C3S1H - Capability Integration Support

PI: Luis Navarro-Serment (CMU), lenscmu@ri.cmu.edu, (412) 268-6034.

Description: The objective of this task is to guide and advise CMU researchers with the design, development, and integration of their software into the RCTA Collaborative Environment (RCE). After researchers have developed and tested their algorithms using a variety of tools and environments (e.g. Matlab, OpenCV, custom ray tracers, Scilab, etc.), they need to integrate their newly created capabilities using the RCE, and deploy them for testing during Multi-Component and Integrated Research Assessments (MCA and IRA, respectively). However, this integration must follow the guidelines for incorporation in the assessment code base. For example, it is expected that the RCTA software that will be added to the code base will have undergone basic researcher level or task-based assessment before being released for integration. Moreover, the software should be designed in a way that maximizes the use of resources already available in RCTA platforms and configurations, to avoid unnecessary repetition of effort and capabilities. Finally, the software must be designed according to the RCE architecture to guarantee that new software modules coexist in the RCE without disrupting the performance of preexisting modules.

This effort will facilitate the transition from existing individual systems and repositories created by CMU researchers to the RCTA RCE. It will focus on guiding the software design to integrate software capabilities; to verify compliance with message and interfaces; to supervise the creation and documentation of software operation and testing procedures, and the selection and tuning of parameters. Finally, it will provide assistance to CMU researchers in delivering, integrating, and testing the resulting software using the target platforms.

Linkage to previous sub-tasks: I1-1.

Collaborations: This sub-task will collaborate with T4C3S1F to ensure that compatibility with the RCE is maintained; that new modules pass the integration requirements; and that the corresponding documentation is added to the repository.

Deliverables:

End of CY2017	• Participation in RCE Workshop. • Familiarization with RCE Test Environment. • Internal CMU RCE integration workshop.
End of CY2018	• Baseline testing of new capabilities. • Support for MCA • Support for IRA.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.15 T4C4S1A – Task Based Assessment (Local and Field)

PI: Lennon, Craig (ARL), craig.t.lennon.civ@mail.mil, (240)760-0667

Objective and Benefits:

This subtask will provide for review of task level technology assessments. Before a technology is ready for integration onto an existing platform, the researcher developing it should conduct experimentation to establish the capabilities of the technology. This experimentation may be at the sub-task level for some technologies, and as a combination of sub-tasks for others, depending on how interconnected the technologies are. Researchers should provide to the integration and assessment team a description of the capabilities their technology will provide, and the experimentation they intend to perform, or have performed, to demonstrate its capabilities. The assessment team will review this experimentation, and will provide assistance with the experimentation if requested by the researcher.

Metrics & Progress Assessment:

This task is assessed by how useful the integration team found the input from the task based assessment.

Connection to past RCTA work:

This is an ongoing RCTA task that was previously funded under subtask IR2 for Integration, Experimental Design, and Assessment.

Collaborations:

- Childers, ARL Assessment Lead
- ARL, T4C1S1D - Platform maintenance support
- Weller, GDLS – T4C3S1C: Capability Integration
- Capability leads

Deliverables

FY 17 and FY 18: Reviews of task level assessments and assistance with experimentation as requested.

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.16 T4C4S1B – Integrated Assessment (Local and Field)

PI: Lennon, Craig (ARL), craig.t.lennon.civ@mail.mil, (240)760-0667

Objective and Benefits:

This subtask will provide 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 RCTA investigators providing iterative feedback on the experimental design, data capture, and analysis tools.

These technology assessments prepare the researchers, integrators, and technology for participation in the 2019 Capstone evaluations, in which technologies can show how they contribute to the capabilities the RCTA is trying to develop.

Metrics & Progress Assessment:

When technology has reached a level of maturity sufficient for integration onto one of the RCTA platforms, the integration team and the assessment team will collaborate with the researchers to determine what benefits the new technology might provide, what limitations it will operate under, and in what relevant environment the technology might be appropriately evaluated. The assessment team will then design and conduct an experiment in collaboration with the researchers and the integration team. After the evaluation, the assessment team will provide feedback to the researchers and integration team on the performance of the technology as part of the platform.

Connection to past RCTA work:

This is an ongoing RCTA task that was previously funded under subtask IR2 for Integration, Experimental Design, and Assessment.

Collaborations:

- Childers – ARL Assessment lead
- Weller, Cullop – T4C4S1D: Integrated Assessment (Local and Field)
- Gonzalez, GDLS - T4C1S1A - Maintain current platforms - Husky, RoMan, RoboSimian
- Taylor, QNA - T4C1S1C - Maintain current platforms - TL1 – RoMan
- ARL, T4C1S1D - Platform maintenance support
- Weller, GDLS – T4C3S1C: Capability Integration

Deliverables

FY 17 and FY 18: Integrated research assessments as determined by the integration schedule

FY 19: Capstone evaluation of technologies as contributing to capabilities

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

5.2.17 T4C4S1D – Integrated Assessment (Local and Field)

PI: Weller, Ed (GDLS), weller@gdls.com, (410)386-1778
Additional PI: Cullop, Barry, cullopb@gdls.com, (410)386-1793

Objective and Benefits:

This subtask is to support field testing and assessment of research results within the program. The task will focus on integration of subsets of capabilities, expanding towards complete systems as capabilities mature. The assessment sequences will focus on key portions of the individual portions of the overall RCTA missions, rather than end-to-end scenarios.

This includes helping develop test plans and metrics; identifying the required data outputs and ensuring that the data is made available and captured during experiments; and providing tech and safety support for the field range test activities. Additionally, this includes coordinating test activities with researchers, customer, and other stake holders.

Metrics & Progress Assessment:

(a) Local integrated assessments as scheduled by researchers.

(b) Pre-IRA trips to FTIG for IRA preparation

(c) Support for Fall IRA at FTIG.

Connection to past RCTA work:

This is an ongoing RCTA task that was previously funded under subtask IR2-1 and IR5-3

Collaborations:

- Lennon – T4C4S1A, T4C4S1B: Task Based and Integrated Assessments (Local and Field)
- Gonzalez, GDLS - T4C1S1A - Maintain current platforms - Husky, Roman, Robosimian
- Taylor, QNA - T4C1S1C - Maintain current platforms - TL1 – Roman
- ARL, T4C1S1D - Platform maintenance support
- Weller, GDLS – T4C3S1C: Capability Integration

Deliverables

	2017 Quarterly Milestones and Deliverables
Q1	Local Integrated Assessments (to be scheduled)
Q2	Local Integrated Assessments (to be scheduled)
Q3	Preparation for Fall IRA
Q4	Initial Capability IRA (FTIG)

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

	2018 Quarterly Milestones and Deliverables
Q1	Local Integrated Assessments (to be scheduled)
Q2	Local Integrated Assessments (to be scheduled)
Q3	Preparation for Fall IRA
Q4	Interim Capability IRA (FTIG)

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.

Appendix A Quarterly Reports

All funded subtask researchers/engineers will provide a quarterly report that summarize research/engineering activities for the reporting quarter. The report should include:

Subtask Number & Title:
PI Name, Institution Name, email, and phone:
Quarter/Year:

Research Issues Addressed:

General Progress:

Detailed Progress:

Collaboration and Integration Activities:

Plans for Next Quarter:

Other Significant activities: (Examples: Conference, Presentation, Publication, Workshop, Student Exchange, Internship, Demo, Experiment, Test, Transition, etc.)

DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited.
The Government shall have the right to duplicate, use, or disclose the data to the extent provided in the Robotics CTA Cooperative Research Agreement W911NF-10-2-0016.