

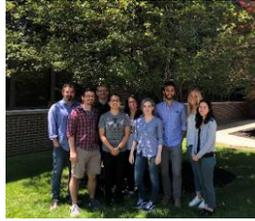
STRONG CYCLE 2
PROPOSED EFFORTS ~~TO BE~~ AWARDED
COOPERATIVE AGREEMENTS
MAY 2020



Jasmin Cloutier



Jennifer Kubota



Neural Network Models of Competency Status to Enhance Human-Agent Team Performance

University of Delaware

Research Team:

Jasmin Cloutier: University of Delaware; Social Neuroscientist and Social Status; Senior Investigator

Jennifer Kubota: University of Delaware; Social Neuroscientist and Decision-Making; Senior Investigator

Denise Barth: University of Delaware, Graduate Student, Junior Investigator

Summary:

The changing landscape of technology enabling interaction among human-agent teams represents a new challenge for many social organizations in the 21st century, including the military. Despite progress in neuroscience and psychology, it remains unclear what factors affect team performance, and whether we can utilize neuroscientific methodologies to both predict and enhance teaming. Social status is a critical factor that fundamentally shapes how humans interact, influencing our attention, evaluations, and trust/cooperation with others. Within military rank hierarchies, differential social status can be based on perceived competence of a skillset. Aligned with high-priority basic research questions tied to the strengthening of teamwork for robust operations in novel heterogeneous human-intelligent agent teams (STRONG), this multidisciplinary, collaborative seedling effort investigates how the ascribed and perceived competency status of human and non-human teammates impacts attention and evaluations; how these processes predict the coordination and cooperation of teams performing tasks; and whether we can predict team performance based on neural networks relevant to status processing.

To model how individual differences in neural responses during impression formation of human and non-human team members predicts the emergence of team performance, we will

focus on the implementation of a novel multivariate brain network approach. We will complement this network analysis with Q-learning computational reinforcement modeling. To investigate these questions, individuals will first learn the rank performance of human (low or high, based on skills) and intelligent agents (low or high, based on amount of artificial intelligence training) in efficiently solving an escape room scenario. Individuals will then take part in an fMRI session during which they will form impressions of human and intelligent agents to assess variability in their recruitment of brain networks supporting attention and evaluation and also complete a reinforcement learning task to model the value associated with low and high competency status human and non-human agents. Following scanning, individuals will complete an escape room challenge with a subset of the humans and agents previously encountered during the impression formation task. Individual differences in the recruitment of brain networks and reinforcement learning parameters, in combination with a series of behavioral and physiological measures, will then be used to predict the performance of the team during the escape room challenge.

The ultimate goal following this seedling effort is performance optimization through explicitly influencing brain network states derived from this multidisciplinary approach that are sensitive to the status of others. This research is significant because status-based attentional and evaluative biases impact how individuals cooperate and coordinate with humans and may also impact how individuals respond to non-human agents. A better understanding of the mechanisms underlying competency-based attention and evaluations as a function of agency (humans versus non-humans) may be critical for predicting and improving team performance in the 21st century.



Stephen Gordon



Kevin King

Temporal Aspects of Sustained Attention

DCS Corporation

Research Team:

Stephen Gordon: DCS Corporation; signal processing and ML approaches for neurophysiological analysis and state detection; Senior Investigator

Kevin King: DCS Corporation; developing tools to enable the intuitive use of AI to improve team performance and outcomes; Junior Investigator

Summary:

Many tasks require an individual to sustain attention for extended periods of time to achieve optimal performance. Although there has been much research on the physiological and performance markers of sustained attention, hence SAT, in individuals [Oken, et al., 2006], less is known about the specific dynamic aspects of SAT. Thus, while prior work has revealed that 1) maintaining high levels of SAT is difficult, if not impossible, for long periods of time [Mackworth, 1964] and 2) SAT tends to decrease over time [Parasuraman, et al., 1987], the specific temporal fluctuations in SAT during long but only moderately demanding tasks remains a mystery. For instance, in real world settings, self-reports indicate that SAT waxes and wanes, often conspicuously, in response to task demands [Oken, et al., 2006; O'Connell, et al., 2008]. It is an open question how the waxing/waning of SAT is strategically employed by individuals in naturalistic settings in the face of fluctuating task demands, limited cognitive resources, and competing internal factors (e.g. mood, circadian rhythm, etc.) to ensure primary task performance.

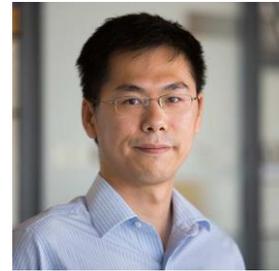
We argue that the answer to this question has strong implications for heterogeneous team design since a core part of teaming is interaction and information exchange. During critical moments the human team members will likely be in various levels of SAT. As a result, the ability of those team members to digest and respond appropriately to new information will vary and the consequences of inappropriately disrupting said individuals will also vary. Furthermore, coping strategies employed by people, such as mind wandering and background chatter during periods of low task demands [Thomson, et al., 2015], may be more difficult to perform with nonhuman teammates or be otherwise perturbed by nonhuman interactions.

The goal of the proposed research effort is to investigate how SAT fluctuates in naturalistic settings and during group interactions to inform the design of policies governing how nonhuman agents integrate within heterogeneous teams. While there are multiple roles that nonhuman agents may take within a heterogeneous team, we are considering two specific roles: first, as information delivery systems and second, as potential actors in the background plays used by their human teammates to cope with task demands. To accomplish our goal we will focus during the seedling effort on applying existing models to analyze the moment-to-moment fluctuations in multiple underlying dimensions believed to index SAT. This will include, but not be limited to: 1) neural function assessed using EEG markers and previously trained deep learning models, 2) autonomic nervous function using pupillary and heart rate measures, 3) behavioral measures such as eye and head movements, and 4) task performance. Using these dimensions, we will define a state space that can be analyzed using Markovian approaches. We will apply this method to both data sets involving brief, highly demanding tasks as well as less demanding, longer duration tasks. Finally, we will analyze existing data sets involving dyadic human interactions in order to investigate how said interactions perturb, or otherwise, modify the underlying state space model for the individual.

If successful, this research would 1) improve our overall understanding of the temporal fluctuations of SAT in humans and, thus, what strategies humans use to maintain performance on tasks requiring SAT, 2) provide initial results describing the extent to which those strategies require the presence, or absence, of other individuals, and 3) most importantly provide the framework for enabling nonhuman agents to learn how to better integrate with human teams performing cognitively demanding tasks.



X. Jessie Yang



Cong Shi



Yaohui Guo



Jose Sanchez

Trust-Driven Human-Agent Teaming: Modeling and Predicting Trust Dynamics

University of Michigan

Research Team:

X. Jessie Yang: University of Michigan; Human factors, Human-autonomy interaction, Human-autonomy teaming; Senior Investigator

Cong Shi: University of Michigan; Design and analysis of efficient algorithms for stochastic systems, Approximation and online algorithms; Senior Investigator

Yaohui Guo: PhD student, Junior Investigator

Jose Sanchez: Master's student, Junior Investigator

Summary:

Human-agent teaming is a major emphasis in the ongoing transformation of Army operations wherein human agents and autonomous agents are expected to work as a team in a multi-tasking environment subject to uncertainty and dynamic changes. To enable effective teaming, trust has been identified as one central factor. The majority of existing literature on trust in automation/autonomy adopted a "snapshot" view of trust and measured trust once, usually at the end of an experiment. This STRONG proposal is focused on trust dynamics - how trust forms and evolves based on moment-to-moment interactions in the human-agent team. In our prior work, we formulated the trust evolution process as a Bayesian inference process and developed a personalized trust estimator. In the proposed effort, we aim to enhance the trust estimator and model trust dynamics in a human-agent team wherein a human interacts with a highly capable autonomous agent, and to develop computational models enabling the autonomous agent to make optimal decisions by explicitly considering trust dynamics in the

team. We will conduct a series of human-in-the-loop experiments, using a simulated intelligence, surveillance and reconnaissance (ISR) task. We will focus on dyadic human-agent teams in Year 1. If successful, the following years will be devoted to extending the methods and models developed in Year 1 to multi-human-multi-autonomy teams.



Monica Nicolescu



Mircea Nicolescu



David Feil-Seifer



Andrew Palmer



University of Nevada, Reno



Bashira Akter Anima

From Heterogeneous Individual Capabilities to Emerging Teamwork: An Architecture for Effective Human-Agent Teams

University of Nevada, Reno

Research Team:

Monica Nicolescu: University of Nevada, Reno; Multi-robot systems, human-robot interaction and learning; Senior Investigator

Mircea Nicolescu: University of Nevada, Reno; Visual motion analysis, perceptual organization, vision-based surveillance and activity recognition; Senior Investigator

David Feil-Seifer: University of Nevada, Reno; Socially assistive robotics, unmanned autonomous systems; Senior Investigator

Two graduate students: Andrew Palmer, Bashira Anima; will perform research on human-robot teamwork performance; Junior Investigators

Summary:

In this project we address the problem of improving human-robot team (HR-Team) performance through a novel approach that takes into account differences in physical and cognitive skills among the individual teammates, as well as the variability of these skills over time. The performance of the HR-Team is highly influenced by the members' ability to distribute responsibilities in the most efficient way possible. The task allocation process of the HR-Team must therefore take into account several key aspects that impact the effectiveness of workload distribution in the team.

In this project, we propose a novel approach for dynamic task allocation in an HR-Team that models dynamic heterogeneity and human agent preferences and provides the ability to evaluate team performance based on the known capabilities and on the interactions of the

individual agents in the team. We plan to build our solution on our team's prior research in heterogeneous multi-human, multi-robot teams, intent recognition and natural language processing. This work provides several key capabilities to be used as our stepping stones: 1) a distributed method for task allocation in human-robot teams, which uses a theory of mind approach to maintain a shared understanding of the mission progress, 2) an approach for understanding human intentions, which enables more effective human-robot collaboration, and 3) a framework for translating complex verbal instructions into hierarchical task representations, which encode multiple types of execution constraints. We will build on our prior work and design a new, team-aware approach for dynamic task allocation that seamlessly adapts to changes in the environment and in agent capabilities. In particular we will work toward addressing the following research objectives:

RO-1. Define dynamic models of human and robot capabilities that encode the most up to date evaluation of their perceived performance.

RO-2. Develop approaches for representing and learning preferences of the human teammates.

RO-3. Design a new approach for task allocation within the team that utilizes the above dynamic models in order to best coordinate the actions of a self-organizing heterogeneous team.

RO-4. Design and evaluate a model of HR-Team performance that emerges from a hierarchical representation of the mission and the dynamic models of individual teammates.

This project addresses the following key question identified in the STRONG Cycle 2 FOA: "How can technologies assist in the dynamic allocation of tasks to individual team members to appropriately balance the variability in both physical and cognitive skills and capacity across the team in order to maximize team performance? In addition, by incorporating dynamical models of human/robot capabilities into the task allocation process, it enables the evaluation of emergent team performance.



Modeling the Relationship between Performance and Momentum within Individuals and Teams

University of Cal - San Diego

Research Team:

Ying Choon Wu: UC San Diego; higher-order cognitive functions, including multi-modal discourse processing and problem solving; Senior Investigator

Chi-Yuan Chang: UC San Diego; artifact removal for EEG signal, signal processing, machine learning, brain-computer interfaces, and mental health assessment; Junior Investigator

Tzyy-Ping Jung: Swartz Center for Computational Neuroscience; cognitive sciences and engineering in service of insights into functions of the brain, cognition, and behavior; Senior Investigator

Summary:

Psychological momentum (PM) is a concept that captures the dynamics of cognition, affect, motivation, and behavior in face of either perceived progress or inertia with respect to an objective [1]. This project seeks to model individual differences in resilience and vulnerability to negative momentum and to understand how the distributed effects of PM across team members influence emergent team states and processes. A multi-tiered model is hypothesized, wherein for some people, inertia toward or regression from a desired outcome is followed by changes in brain activity related to motivation and/or changes in autonomic activity related to stress – and subsequently, by subtle changes in allocation of attention and performance of task work. On the other hand, for individuals who are resilient to or even boosted by negative trajectories of

events, physiological stress responses are expected to be uncoupled from adverse motivational states and behavioral trajectories. At the team level, it will be examined how such anticipated patterns of vulnerability and resilience to periods of negative momentum relates to diminished or enhanced team co-ordination, and ultimately to the effectiveness of action processes.

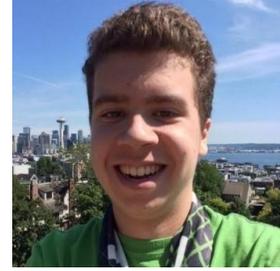
To explore these questions, individuals will partake in a virtual human-agent team tasked with collecting targets distributed across a large space organized in sectors. Within each active sector, targets that are found within a three-minute time window will result in points that translate into monetary benefit, whereas undiscovered targets will expire within that window, leading to a loss of points and money. Periods of progress are operationalized as any three-minute segment that ends in a net gain in points and follows a preceding segment that was also characterized by net gain. Reciprocally, periods of regression are defined as any segment of net-loss that is preceded by a corresponding net-loss segment. The success (or failure) rate of each segment will be artificially controlled so that pseudo-randomly interleaved periods of progress or regression are achieved. During participation in the virtual Human-Agent team, players' electroencephalographic (EEG) and electrocardiographic (ECG) data will be recorded via wireless and wearable biosensing systems, while eye movements will be recorded using an integrated eye tracker within the VR headset. Through this work, it is hoped that physiological markers of sensitivity to negative momentum can be uncovered, giving rise to the possibility of diagnostic tools for evaluating the degree to which each team member's performance is likely to decline versus improve during periods of prolonged impeded progress toward a mission objective.



Bradley Hayes



University of Colorado
Boulder



Matthew Luebbers



Christine Chang

Opportunistic Planning for Emergent Capability in Dynamic Human-Robot Teams

University of Colorado, Boulder

Research Team:

Bradley Hayes: University of Colorado Boulder; Explainable AI, Human-Robot Teaming; Junior Investigator

Christine Chang: University of Colorado Boulder; Human-Robot Communication, Multi-agent Planning and Coordination; Junior Investigator

Matthew Luebbers: University of Colorado Boulder; Explainable AI, Multi-agent Planning and Coordination under Uncertainty; Junior Investigator

Summary:

Planning algorithms that can react online to changing team dynamics while encouraging the utilization of emergent capabilities within multi-agent cooperative scenarios represents a major milestone for human-machine teaming; this proposal introduces new methods for modeling and leveraging the abilities of agent sub-groups for enabling mixed human-robot teams with dynamically shifting availability to more rapidly and efficiently execute challenging missions.

The proposed work targets two complementary objectives: (1) Developing online planning algorithms for mixed human-robot teams that can leverage predictive models of emergent capabilities, namely, models that are able to estimate capability augmentation that is possible only when groups of agents work together; and (2) Adapt established work in explainable planning techniques to efficiently convey these plans to human teammates,

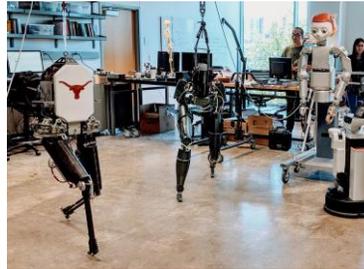
enabling more capable and agile human-robot teaming in the field. To focus and motivate our efforts, we propose to evaluate our work within the domain of alpine avalanche rescue operations: a time-sensitive and hazardous domain frequently operated in by dynamically changing, ad-hoc teams that necessitates expert coordination to achieve positive results.



Efstathios Bakolas



Luis Sentis



Emergent cooperative behaviors in human-robot teams induced by individualized and adaptive decision making, active sensing, and task allocation

University of Texas – Austin

Research Team:

Efstathios Bakolas: University of Texas at Austin; problems related to multi-agent control, decision making, and optimization; Senior Investigator

Luis Sentis: University of Texas at Austin; control, embodiment and intelligence of humanoid robots as well as strength augmentation exoskeletons; Senior Investigator

Graduate Research Assistant: Junior Investigator (TBA)

Summary:

The goal of this project is twofold. The first goal is to develop algorithms for adaptive decision making, active sensing and dynamic task allocation for mixed human-robot teams that improve team performance via both the individual actions of the team members (or agents) as well as their local interactions (coordination of actions). The second goal is to evaluate these algorithms through realistic experimental simulations and validate our hypothesis that the proposed theoretical and algorithmic tools will be applicable to different classes of mission scenarios. We propose to address the following tasks:

Task 1: Devise a theoretical and algorithmic framework for adaptive decision making and dynamic task allocation in human-robot teams that promote emergent cooperative processes via optimized and adaptive individualized behaviors.

Task 2: Construct a sufficiently rich set of admissible tactics and utility functions for theoretical analysis based on the framework of potential games.

Task 3: The objective of this task is to create an environment for principled experimentation and a set of metrics for measuring teaming performance for different types of “missions”.

The PI of the project is Dr. Efstathios (Stathis) Bakolas who is an Associate Professor in the Department of Aerospace Engineering at the University of Texas at Austin. He received his Diploma in Mechanical Engineering with highest honors from the National Technical University of Athens, Greece, in 2004 and his MS. and Ph.D. degrees in Aerospace Engineering from the Georgia Institute of Technology, Atlanta, in 2007 and 2011, respectively. Dr. Bakolas research interests include distributed control and game-theoretic decision-making for multi-agent networks and autonomous systems, optimal control and optimization-based control for uncertain systems, control of complex systems, and motion planning and path planning for robotic systems. He will lead the efforts for the design and implementation of the proposed tools of Task 1 and Task 2.

Dr. Luis Sentis is the General Dynamics Endowed Associate Professor in Aerospace Engineering at the University of Texas at Austin and co-founder of Apptronik Systems. He received a Ph.D. in Electrical Engineering from Stanford University and was a La Caixa Foundation Fellow. He leads the Human Centered Robotics Laboratory, a computational and experimental facility focusing on control, embodiment and intelligence of humanoid robots as well as strength augmentation exoskeletons. His research interests include real-time decision and control of human-centered robots, hardware design of high performance humanoid robots, and safety capabilities for mobile robots. He was awarded the NASA Elite Team Award for his contributions to NASA's Johnson Space Center Software Robotics and Simulation Division. He has been a speaker in popular events such as SXSW and NASA Thoughts Leader Seminar. He will lead the efforts for the design and implementation of the proposed experimental demonstrations of Task 3.



Kshitij Jerath



Paul Robinette



Reza Ahmadzadeh

Emergence of Trust Clusters in Human-Agent Teams Operating under Resource Constraints

University of Massachusetts – Lowell

Research Team:

Kshitij Jerath: University of Massachusetts Lowell; modeling of collective emergent behaviors, influence in multi-agent systems, self-organizing networked systems; Senior Investigator

Paul Robinette: University of Massachusetts Lowell; human-robot trust in time-critical situations in simulation, the lab, and the field; Senior Investigator

Reza Ahmadzadeh: University of Massachusetts Lowell; robot learning, reinforcement learning, imitation learning, learning from demonstration, and optimization; Senior Investigator

Ashwin Nair: University of Massachusetts Lowell; Junior Investigator

Summary:

Teams succeed because of the network of relationships they possess, and the emergent behaviors this network facilitates. These behaviors arise due to three constructs: (a) multiple agents in the team capable of taking actions, (b) interactions between the agents, and (c) emergence of global-scale patterns due to the interactions. The overall research strategy and objective tackles each of these constructs as separate tasks in the context of a search and rescue mission. Search and rescue operations are often severely resource constrained in terms of time, energy, and information organization. Operating in such resource-constrained scenarios can impact the ability of human-agent teams to tackle complex problems, resulting in sub-optimal outputs.

The research objective of the proposed work is to demonstrate how varying resource constraints applied to individual learning entities (governed by microscopic level model parameters of agents or humans) impacts the macroscopic level emergence of trust clusters in teams, and eventually team performance and cohesion. The proposed framework builds on the theories of Commons [1], Graves [2], and Jerath [3][4], to create a modeling, reinforcement learning, and adaptation framework that is functionally identical for both humans and agents, and incorporates dynamic models of trust. We solve three tasks in the proposed work to achieve the research objective:

- i. Task 1: Microscopic-scale task learning in a team environment that are governed by an individual's resource constraints (encoded as a 'complexity profile')
- ii. Task 2: Modeling interaction between entities (humans or agents) that depend on dynamic models of trust and are based on quality of shared information (encoded as a 'trust profile')
- iii. Task 3: Experimentally validating the relationship between microscopic-scale individual resource constraint profile and emergence of trust clusters via design of experiments approach

In the proposed work, we use the notion of 'complexity profile' to identify the level of complexity that an entity (or team) can continue to learn at or manage, under varying levels of resource constraints. We implement novel multi-agent reinforcement learning strategies to achieve proficiency in taskwork at fast timescales (i.e., learning to complete the search and rescue task), while leveraging dynamical models of trust based on quality of shared information to produce emergent trust clusters at slower time scales (i.e., building team cohesion and confidence in team members abilities). Online search and rescue simulations with human participants will be used to experimentally validate the impact of complexity profiles on the emergence of trust clusters and team performance metrics using a design of experiments approach.

The long-term goal of the proposed work is to leverage advances made during the seedling grant to answer these broader questions in the potential 3-year follow-on project period. This will also lays the foundations to pursue the long-term goals and broader needs of the STRONG project, such as training agents and humans to adapt their microscopic-scale, individual-level trust and complexity profiles in a way that optimizes macroscopic-scale, emergent team processes. The research team consists of three PIs across Computer Science, Mechanical Engineering, and Computer Engineering, and a graduate student, and brings over 30 years of collective experience in the field of emergence, complex adaptive systems, reinforcement learning, and human-robot trust.



Katia Sycara



Michael Lewis



Dana Hughes



University of
Pittsburgh



Huao Li

Effective Individualized Dynamics and Emergent Team Behaviors in Human Agent Teams

Carnegie Mellon University, University of Pittsburgh

Research Team:

Katia Sycara: Carnegie Mellon University; Multi-agent Systems, Human-Agent Teaming, Multi-agent Control and Learning; Senior Investigator

Michael Lewis: University of Pittsburgh; Human Factors and Trust in Human-Agent Teams; Senior Investigator

Dana Hughes: Carnegie Mellon University; Multi-agent Systems, Reinforcement Learning; Junior Investigator

Huao Li: University of Pittsburgh; Human Factors and Trust in Human-Agent Teams; Junior Investigator

Summary:

We propose to develop models and algorithms of adaptive coordination in human-agent teams considering individualized dynamics of team mates and differential skills, capabilities and preferences. A crucial objective will be to (a) measure and evaluate how the developed learning processes and models enable and influence team states and processes and (b) predict human-agent team performance. To achieve the above objectives, our methods include theory formulation, scenario and testbed development, reinforcement learning and lab experiments with teams comprised of only humans, only agents and human-agents.

The agents will learn to share information through human- and agent-intelligible targeted communication with other team members to establish and maintain a shared mental model, and

coordinate with other team members to improve team synchrony and coherence. Our model will learn several key components to achieve these goals. Agents will learn to allocate tasks, based on team mate skills, capabilities and task requirements. They will also learn to communicate grounded messages to targeted team mates, essentially learning what message to send to whom. Agents will also learn to model other team members' behaviors, in order to learn which messages and actions improve or maintain team coordination. We will train our agents using several intrinsic motivations that encourage effective communication and actions for information sharing and coordination. Agents will perform online adaptation to their models of human team mates, allowing their behavior to be individualized to specific team members. We will evaluate our research using a Team Disaster Relief testbed (TDR), where 3-5 member human-agent teams perform multiple concurrent tasks (e.g., searching an area, locating victims, treating already found victims, putting out fires) in a dynamic and resource limited environment. We will test the performance of agent-only teams and human-agent teams in TDR to identify potential agent characteristics that would further benefit human-agent team performance. Team coordination and coherence will be evaluated based on observed communication between team members; agent performance will be additionally evaluated based on their ability to correctly predict and adaptively coordinate with team members.

In the Cycle 1 seedling and follow-on option, we focused on (1) developing effective and adaptive agents via imitation learning, reinforcement and perturbation learning for human-agent teams in role-based tasks that require tight coordination between team members, in variants of the Team Space Fortress (TSF) simulation and (2) incorporate physiological measurements to better understand how individual human characteristics, such as attitude toward risk and ambiguity affect human-agent team performance (ARL collaboration). In Cycle 2, we will build upon these results to delve deeper and provide validated models of human-team performance that demonstrate how individual dynamics of humans and agents enable and influence emergent properties, states and behaviors of human agent teams. Supporting human-intelligible human-agent adaptive team models and communications is particularly challenging because they rely on what are commonly presumed to be human cognitive abilities. Replicating and evaluating these team processes in human agent teams would be a significant step in understanding and optimizing team performance.