

STRONG CYCLE 1
PROPOSED EFFORTS AWARDED
MAY 2019



Noshir Contractor



Northwestern
University



Cornell University



Aaron Schecter



Malte F. Jung



UNIVERSITY OF GEORGIA



Leslie DeChurch

The Signatures of Success in Human-Agent Teams

Project Investigators:

Noshir S. Contractor: Northwestern University; Industrial Engineering, Management, & Communication; Senior Investigator

Malte F. Jung, Cornell University; Computer & Information Science; Junior Investigator

Aaron Schecter, University of Georgia; Management Information Systems; Junior Investigator **Leslie**

A. DeChurch, Northwestern University; Communication & Psychology; Senior Investigator

Summary:

In the US Military, Army squads continue to be the decisive force (Brown, 2011). Squads are cohesive groups of 6 to 10 soldiers tackling focused engagements from peacekeeping to provincial reconstruction. The central role of the squad to a variety of military engagements has made teamwork capabilities a cornerstone of military science for decades (National Research Council, 2014). Now, as the digital era transforms the basic nature of work, nowhere are the possibilities of human-agent collaboration greater than in military squads (DeCostanza, Marathe, Bohannon, Evans, Palazzolo, Metcalfe, & McDowell, 2018). Rapid technological advances offer new ways to augment human collaboration to make military engagements safer and more effective. Realizing this vision will require foundational advances to understand the affective, cognitive, and behavioral underpinnings of effective human-agent teams. Simply put, how do individuals and agents interact effectively in teams? To answer this question, we must understand both human-agent interaction, and also the nature of interpersonal interactions when agents join teams. Our project discovers the signatures of success in human-agent teams using advances in event network analysis.

Squads are essentially teams, allowing us to understand them using a large body of work that identifies conditions that enable team effectiveness (Kozlowski & Ilgen, 2006). These factors can be grouped into three higher order categories: affective, cognitive, and behavioral. Team affective processes include, for example, trust, motivation, and cohesion; affect is what holds a team together and makes its members want to contribute their effort to team goals. Team cognitive processes are organized knowledge systems that enable members to anticipate one another's needs and to coordinate their actions. The two fundamental cognitive processes being shared mental models and transactive memory systems. Behavioral processes like mission analysis, coordination, information sharing, and backup behavior, describe the verbal and behavioral interaction mechanisms through which individuals combine their efforts toward the pursuit of team goals. Effective teams require all three types of processes.

A recent advance in studying team processes finds there are subtle cues in the patterns of interaction over time that can signal the extent to which a team exhibits functional affective, cognitive, and/or behavioral processes (Leenders, Contractor, & DeChurch, 2016; Pilny, Schechter, Poole, & Contractor, 2016). This relational event perspective leverages the content of the big three processes, but finds innate tendencies, or signatures, that characterize team interaction as it unfolds over time, can also distinguish effective and ineffective teams. Early research on these signatures finds they are more predictive of team performance than are traditional quality, quantity, or pattern approaches to team process (Schechter, 2017). Can this approach to uncover subtle interaction patterns that foretell success be used to discover markers of successful human-agent teams? That is the goal of our STRONG project.

Our research focuses on agents that are robots, and we use the terms agent and robot interchangeably in the text that follows. The basic insights that come from human-robot teams can provide a valuable basis for work more generally on human-agent teams where agents consist of, for example, sensors, autonomous vehicles, and intelligent assistants. The defining aspect of human-agent teams being that the technology “can perform taskwork as part of the larger team (DeCostanza et al., 2018, p.3).” In human-agent teams, the technology is a team member, rather than a resource or material feature of the context.

Research on human-robot interaction finds the same social processes critical to interpersonal interaction are also essential for human-robot interaction (Nass & Moon, 2000). For example, successful interactions, require the ability to detect emotion and tune messages accordingly (affective) and to anticipate what information is needed and how it is used to make decisions (cognitive). In this project, we leverage the relational event perspective, using it to discover the subtle tendencies of human-agent teams that best predict performance.

A growing literature on human-robot teaming develops important insights towards understanding and enabling teamwork with embodied autonomous robots (Hancock, Billings, & Schaefer, 2011). Human-robot collaboration, or teaming, refers to collaborative partnerships between humans and robots aimed at completing tasks; this work focuses on coordination between a human and a robot (see review by Ajoudani et al., 2018). A central premise of this work is that robots and humans each bring distinctive capabilities, making them ideal teammates (Shah, 2017). Humans are superior at critical thinking, creativity, and managing people, whereas robots are superior at coping with harsh environments and performing highly repetitive work, for example. Adding robots to teams holds the potential to create dream teams, capitalizing on the diverse strengths of each. However, realizing this potential will first require an understanding of the processes that enable performance.

Research on human teams provides a solid foundation for understanding the functional needs of human-robot teams, but it requires extending the paradigm to consider the social dynamics among humans and how they are affected by the addition of robotic teammates. Human teams require affective (e.g., cohesion, trust), cognitive (e.g., transactive memory), and behavioral (e.g., coordination, information sharing) processes to succeed. These processes have extensions to human-robot teams. In the affective domain, members in human teams must share a common desire to succeed to motivate their efforts, and they must trust one another's intentions. In human-robot teams, humans need to trust the robot's intentions, and be able to understand and predict the behavior of the robot in different situations (Hayes & Shah, 2017).

To use a cognitive example, human teams perform best when they create transactive memory systems whereby team members specialize their expertise and maintain a directory of who knows what. Humans and robots need similar knowledge structures, holding a shared understanding of the distinctive capabilities of the human and the robot. The presence of robot teammates not only requires effective human-robot interactions, it also conditions the nature of human-human interactions as well. Human teammates must share an understanding of the robot's capabilities and how to effectively incorporate robots into the teams. This suggests a new form of team mental model specifying members' schematic representations of the nature of interaction between humans and robots - a human-robot interaction mental model. The human-robot interaction model is one example of how teams theories need to be usefully extended to explain effective human-robot interaction as well as human-human interaction in the presence of a robot.

Researchers focused on the HRI interface are growing increasingly aware of the need to understand human-human interactions when robots join teams. Three sets of studies find embodied autonomous robots affect interpersonal relations among humans. First, the mere presence of a robot can affect how individuals interact with each other (Dole, 2017; Riether et al., 2012). Second, research shows robots can be designed to mediate interpersonal interaction. For example, the addition of a robot to a human dyad can improve conflict dynamics (Jung, Martelaro, & Hinds, 2015), problem solving (Short et al., 2017), group emotion (Correia, et al., 2018), and/or coordination (Iqbal, Rack, & Riek, 2016; Iqbal & Riek, 2017). Third, even robots that are not intentionally designed to shape interpersonal interactions can do so by reshaping social norms (Lee et al, 2012). Given these findings, building effective human-agent teams will require a detailed understanding of the mechanisms through which robots influence social processes.

Creating effective human-robot teams requires advances in two adjacent areas. In the teams domain, theories need to be extended to understand how human social interaction can be functionally extended when robots join teams. In the robot domain, theories need to extend beyond the "one human interacting with one robot" paradigm (Jung & Hinds, 2018). Furthermore, advanced analytic approaches to capture dynamic interaction process are needed for human-agent teams. The current state of research on HRI focuses on when and how people work with robots. We extend this work to consider different patterns of interaction among human-robot teams, and to ultimately link these patterns to performance differences.

To do this, we leverage relational event networks (Butts, 2008). A relational event is a concrete behavior or action engaged in by one social actor that is directed toward another social actor. A social actor could be a robot or a human. The sequence of these behaviors or events hold valuable insights into the functioning of social groups. Returning to the big three processes, we use relational events to characterize functional team behaviors such as coordinating, information sharing, monitoring, and

backing up. In addition, we leverage a coevolution perspective to consider the affective and cognitive states that shape and are shaped by teaming behaviors. Thus, we can discover the signatures of affective, cognitive, and behavioral processes commencing over time in human-human-robot teams. We identify signatures of human-agent team interaction to answer three questions.

- RQ1: How do different types of human and agent actions shape the emergence of human-agent interaction signatures?
- RQ2: Which interaction signatures are most likely to emerge in human-agent teams?
- RQ3: Which interaction signatures are most functional to the performance of human-agent teams?

Answering these three provides the basis for programmatic applied research to uncover the likely tendencies of human-agent teams and also the functional, or optimal signatures of them. This is a necessary first step for developing individualized adaptive technologies to support human-agent teams.

In terms of the STRONG objectives, this work provides important building blocks for the first two objectives set out in W911NF-19-S-001: “Theories linking individual dynamics to emergent processes in human-agent teams” and “Models bridging dynamic, continuous individual and group attributes, emergent team behaviors, and team performance.”



Jamie C. Gorman



John M. Flach

Polycentric Control and Team Resilience: A Dynamic Systems Approach

Project Investigators:

Jamie C. Gorman: Georgia Institute of Technology; Teams, Dynamical Systems; Senior Investigator

John M. Flach: Mile Two LLC; Cognitive Systems Engineering; Senior Investigator Two

Georgia Tech graduate research assistants TBD; Junior Investigators

Summary:

Technological advances associated with communications (e.g., global networking capabilities) and computations (e.g., autonomous AI agents) have increased the complexity (i.e., requisite variety) of cognitive work. On the positive side, this has led to new affordances and opportunities, but on the negative side this has created new threats and risks to stability, sustainability, and resilience. This increased requisite variety creates new challenges that typically exceed the capacity of unaided individuals and thus, increases the demands for coordinated action of teams of humans and automated agents. The goal of this effort is to explore the implications of this challenge in terms of alternative principles and forms of coordination and control (e.g., subsidiarity; command intent; polycentric control; self-organization). The short-term goal is to understand and model how high-functioning teams and organizations achieve stable solutions that balance higher-order constraints (i.e., shared intentions and values), while maintaining a capacity to flexibly adapt to local changes and disturbances. The long-range goal is to set a foundation for making training and design decisions that enhance human-technology teaming and system performance.

This challenge is illustrated by constructs such as command intent and subsidiarity. Both of these suggest that taking advantage of networked communications and distributing control responsibility to subsystems in an organization (e.g., allowing increased autonomy or discretion to junior officers to adapt to local situations) increases flexibility and stability. However, there is no clear principle to specify how much autonomy is enough. That is, complete autonomy will result in a system where the elements (e.g., teammates) will often be working at cross purposes, so that the global intent will not be satisfied and the overall system will not be sustainable (as illustrated by the tragedy of the commons). Conversely, local (operator; agent) constraints might not be anticipated a priori, and guidance from a

centralized authority would take too long. One goal of this project is to examine a range of organizational structures to identify alternative control strategies for managing complexity in distributed human-technological systems.

We propose to develop a common information language to track system coordination dynamics across different control strategies, with the ability to communicate how human and technological subsystems contribute to emergent system states. Using this modeling approach, the dynamic properties of system state are represented as information flows within and across subsystems in order to visualize and diagnose threats to achieving and maintaining a sustainable system equilibrium. We will attempt to codify and formalize these attributes in high performing teams and organizations in order to explore trade-offs and implications for stability and resilience. This effort will begin to tackle the challenge of providing representational tools and displays that allow agents in these systems to deal with complexity.

The theory and modeling technologies envisioned under the proposed project will provide the groundwork to amplify polycentric control by facilitating direct communication of system state across human and technological subsystems. A longer-range goal of the proposed effort is aimed at translating the insights gained through field observations (Aim 1) and analytical modeling (Aim 2) into guidance for training and interface solutions for shaping resilient team performance (Aim 3). Ultimately, the dynamic representations and theories that will result from this project directly address several STRONG program goals: *establishing theories that link individual dynamics to emergent, system-level processes in humanagent teams; developing models that bridge dynamic, continuous individual and system-level attributes, emergent team behaviors, and team performance; and predicting and leveraging individual (subsystem) fluctuations for improved overall (system) performance.*



Tinoosh Mohsenin



Tim Oates

Dynamic Graphs for Learning to Optimize Human-Agent Teams

Project Investigators:

Tinoosh Mohsenin: University of Maryland Baltimore County; implementation of computational intensive machine learning techniques to meet the real-time, power consumption and noise constraints to integrate the systems in real-world environment; Senior Investigator

Tim Oates: University of Maryland, Baltimore County; graph grammars, time series analysis, natural language processing, relational learning, and social media analysis; Senior Investigator Graduate Student: TDB; Junior Investigator

Summary:

Teams and the work that they do are inherently relational and dynamic. Said differently, they can be effectively represented by graphs that change over time. Nodes correspond to team members, domain objects with which they work, properties of objects, tasks, task inputs and outputs, roles and responsibilities, capabilities of team members, and so on. Edges correspond to communication, actions that affect state, input/output relationships, organizational structure, and a wide variety of other kinds of interactions among nodes, including state, property, and role changes through time. As the team's work progresses, this graph is elaborated, and at various points in time can be decorated with measures of performance.

We claim that team states and dynamics can both be represented as subgraphs, that machine learning methods can be used to identify states and processes that are predictive of team performance, and that biasing the learning methods can effectively expose the role of agents in the performance of human/agent teams. In particular, we propose to use graph embedding methods to convert teamwork graphs into vectors, and use bias injection to focus the learned embeddings on agents and human/agent interactions.

Our two hypotheses are: (1) Specifying edge weighting functions to focus graph walks on agents and human/agent interactions is an effective way of understanding the idiosyncrasies of what makes human/agent teams work well or perform poorly. (2) Machine learning applied to agent/human teamwork graphs will yield substantial insights into what makes such teams work well or poorly, and

how to improve their performance dynamically. We will test these hypotheses using a publicly available dataset of 1,144 human participants in a game simulating complex management processes in the presence of automated participants (agents).



Christoph Riedl



Northeastern
University



Brooke Foucault Welles



Richard Radke



Rensselaer



Kate Coronges



Michael Foley



Indrani Bhattacharya

Strengthening Teamwork for Robust Operations in Novel Groups (STRONG)

Project Investigators:

Christoph Riedl: Northeastern University; Collective Intelligence, Collaboration in Teams; Senior Investigator

Brooke Foucault Welles: Northeastern University; Computational Communication; Senior Investigator

Richard J. Radke: Rensselaer Polytechnic Institute; computer vision problems related to human-scale, occupant-aware environments; Senior Collaborator

Kate Coronges: Northeastern University; social structures and their dynamics on communication patterns, behaviors, and performance; Senior Collaborator

Michael Foley: PhD student at Northeastern University; agent-based modeling and statistical analysis, modeling of social network processes; Junior Investigator

Indrani Bhattacharya: PhD student at Rensselaer Polytechnic Institute; machine learning algorithms for multi-modal sensor analysis; Junior Investigator

Summary:

Building on several years of interdisciplinary, collaborative research, the proposed work will develop social scientific models of dynamic processes of human-agent teams to inform the design of physical and virtual adaptive agents able to enhance team performance. We will leverage a cutting-edge suite of technological resources that we developed to study emergent processes in human-agent teams, including:

- A sensor-equipped Smart Conference Room that captures and dynamically processes verbal and non-verbal team communication through multi-modal sensor data;
- A suite of state-of-the-art machine learning algorithms that turn raw sensor data into meaningful higher-level measures including automated speaker detection, head pose estimation, and visual focus of attention estimation;
- A toolbox for agent-based modeling that supports adaptive behavioral strategies and dynamic networks using reinforcement learning;
- An online experimentation laboratory to study group decision-making and collaboration;
- The ability to integrate the same intelligent adaptive agents from the agent-based modeling framework as team members in online experiments; and
- A method for efficient automation and optimization of adaptive experiments that allows us to increase the gained information and to conduct experiments with many treatment arms.

Together, these technological advances allow us to undertake unprecedented research on complex collective behavior that emerges in human-agent teams. Our work seeks to understand emergent team communication with two distinct, yet complementary, approaches that will identify points where adaptive cognitive intelligent agents can intervene to improve performance in human-agent teams. In the seedling year, we will develop protocols for adaptive intelligent cognitive agents that will achieve the following: (Theme 1) identify dynamical stages when adaptive intelligent cognitive agents can intervene, nudge, or inform human teams to improve performance by eliminating bias and exclusion in team deliberation, and (Theme 2) improve collective decision making by inserting agents as team members designed to rewire communication network structures.

These protocols and technology demonstrations will serve as a foundation for our participation at the Summer Innovation Summit, where we will develop collaborations to build and test physical and virtual agents that can participate in face-to-face, online, and hybrid human-agent teams and improve their performance.



Paul Sajda



James McIntosh

A Physiologically-Informed Artificial Intelligence (PI-AI) Framework for Human-Agent Teaming

Project Investigators:

Paul Sajda: Columbia University; rapid decision-making; Senior Investigator

James McIntosh: Columbia University; decision making with respect to choice behavior, recurrent neural networks; Junior Investigator

Summary:

In this STRONG project we propose to develop and investigate a new model and framework, which we term Physiologically-Informed Artificial Intelligence (PI-AI; pronounced pie-AI) for enabling robust and adaptive human-agent teams. We define physiologically-informed broadly, meaning utilizing neural and systemic physiological information together with affective cues from facial features and posture/gait to infer latent cognitive and emotional states of a human(s) interacting with AI(s) in a collaborative decision-making task. Our approach builds upon, in a novel way, the concept of human-in-the-loop (HITL) learning. Specifically we build upon HITL by not only utilizing explicit behaviors and actions of the human to train and adapt the AI, but also latent cognitive and emotional states which we infer from neural, physiological and behavioral cues. Specifically the one year seedling project we propose looks to develop and provide a proof-of-concept for an important part of the PI-AI model, specifically investigating the hypothesis that "physiological signals can be used to infer human cognitive state and that these can in turn be used to modulate reinforcement signals for an AI-agent." We believe that confirmation of such a hypothesis will be critical for enabling a "theory of mind" or "shared mental model" between human and agent, which is ultimately critical for enabling complex collaborative decision-making in human-agent teams. To make the proposed seedling project tractable, we will leverage several current experimental paradigms and data, specifically data streams that have been previously recorded for a boundary avoidance task (BAT) implemented as a flight-based navigation task. In this previous experiment we simultaneously collected multiple physiological signals including electroencephalography (EEG), electrodermal activity (EDA), heart rate (HR), pupil dilation and eye

movements (EYE) in over 30 subjects in an open-loop mode while they performed the task in virtual reality. Here we plan to investigate whether a deep reinforcement learning (DRL) agent can learn a policy to assist the human in the task when they are in most need, namely when the human's confidence is low and/or their arousal is high. Our specific research focus will be to investigate models for inferring these cognitive state variables (e.g. confidence and arousal), given the physiological inputs, and comparing how inferences for different classes of models affects the learned policy of the DRL and ultimately human-agent performance. Our long-term goal is to team with ARL researchers and other collaborators to investigate these cues, as well other cues related to social signaling, such as facial expression, gait and speech, to develop a PI-AI framework for optimizing human-agent teams for complex and collaborative decision making scenarios.



Julie A. Shah



**Massachusetts
Institute of
Technology**

Human-Robot Co-Training for Unpredicted Tasks

Project Investigators:

Julie A. Shah: Massachusetts Institute of Technology; artificial intelligence, human factors, and systems engineering to develop interactive robots that emulate the qualities of effective human team members; Senior Investigator

Graduate Student TBD; Junior Investigator

Summary:

The aim of this project is to envision new types of human-robot training procedures that would enable robots and humans to accomplish tasks more efficiently together than they do separately. The focus is specifically on enabling humans and robots to work together in unpredictable situations. In our prior investigation, we formalized perturbation training as a computational problem and provided the first clear pathway for incorporating an intelligent robot, with algorithms for learning and communication, to participate in the training process alongside a human team member. Specifically we developed the first end-to-end computational framework for perturbation training that enables a robot to co-train with a person during live interactions.

The proposed project will augment the co-training model with a new communication component that enables the human-robot team to leverage human insight when confronted with a new situation, to rapidly come to agreement on acceptable and unacceptable strategies for working together. The effort will utilize PI Shah's new computational models for transfer learning that enable humans and machines to draw from their shared experiences in a compatible manner.

Our specific objective is to augment our Perturbation Training (PT) model to support bidirectional communication of acceptable and unacceptable teaming strategies using temporal logic specifications. We will assess the utility of the communication component in terms of its ability to quicken team learning of a repetitive task, improve the team's success rates on novel tasks, and promote human-robot shared understanding of the teaming strategy in laboratory studies. If successful, the following years will be devoted to operationalizing the training process for Army-relevant tasks and demonstrating impact of the approach in Army-relevant simulation environments.



Amelia Solon



Stephen Gordon

Preemptive and Adaptive Group Design

Project Investigators:

Amelia Solon: DCS Corporation; ML (deep learning) and bridging the gap between the lab and the real world in the areas of neural signal analysis and brain-computer interfaces (BCI); Junior Investigator

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

Summary:

The introduction of intelligent agents into operational environments has the potential to amplify their dynamic and complex nature. Effective human-agent teaming will thus hinge on the identification of capable and adaptable humans and the real-time evaluation of these humans as they perform these critical roles. Current aptitude testing methods cannot replicate real-world stress, nor can they gauge the real-time mental responses during the performance of a particular task. As such, there is a need to develop tools to identify, during training, those individuals who are better able to cope with the influx of real-world stress, and during task execution, real-time changes in individual and team state.

The goal of the proposed effort is to understand performance variability across and within individuals operating in a human-agent teaming environment with the aim of improving teamwork outcomes through task pre-allocation and reallocation. We believe that optimizing initial team composition as well as enabling contingencies to adapt to inevitable team changes are integral to optimizing teamwork. Thus, this effort is hierarchical in its approach.

First, we investigate methods for “Preemptive Group Design” that will expand upon standard “stable ability detection” approaches, such as subjective self-measures (i.e., Five Factor Model [Judge and Ilies, 2002]) and objective tests of aptitude (i.e., Armed Services Vocational Aptitude Battery [ASVAB]). In other words, we seek to better predict, prior to team formation, an individual’s strengths so that his/her tasks make the best use of the individual’s existing and future abilities, while complementing the strengths and tasks of other teammates.

For this, we will first investigate and verify predictors of task performance on existing data sets. Then, we will study changes in task performance as an individual learns a task, as models of individual abilities will need to be able to adapt over time as the individual improves in a task. This study will allow us to test the extent of our ability to predict an individual’s learning trajectory, as well as learn how

certain individual traits evolve over time and how these individual states change with individual performance.

Second, we will investigate “Adaptive Group Design” to reallocate the tasks of human agents in response to anticipated and actual performance fluctuations. This effort will lay the groundwork for closed-loop systems that are capable of detecting individual- and team-state changes and for investigation of the ideal system responses to said state changes.

Both efforts will hinge on passive physiological/behavioral sensing, using generalized machine learning (ML) approaches to detect and integrate the states of interest within and across teammates. While this proposal focuses on the optimal combination of human teammates, the methods by which the teammates will be sensed are expected to be used as input to intelligent agents in follow-on work.



Katia Sycara



Michael Lewis



Dana Hughes



University of
Pittsburgh



Huao Li

Robust Teamwork in Human Agent Teams

Project Investigators:

Katia Sycara: Carnegie Mellon University; artificial intelligence, machine learning, multi-agent systems and robotics; Senior Investigator

Michael Lewis: University of Pittsburgh; human factors and human-agent interaction; Senior Investigator

Dana Hughes: Carnegie Mellon University; artificial intelligence, machine learning and distributed systems; Junior Investigator

Huao Li: University of Pittsburgh; cognitive science and human factors; Junior Investigator

Summary:

Teamwork is a set of interrelated reasoning, actions and behaviors of each team member that adaptively combine to fulfill team objectives [21]. Human only teamwork has been extensively investigated in the literature. Drivers of team effectiveness have been identified to be: team leadership, mutual performance monitoring, backup behaviors, adaptability, team orientation, shared mental models closed loop communication and mutual trust [24, 25]. As AI and Machine Learning technologies enable development of increasingly intelligent agents, these agents are envisioned to play important roles as team mates of humans in military teams. This would require re-consideration and reconceptualization of human roles and team interactions. This entails re-examination of theories of teamwork, team processes and states that have been identified for human-only teams.

There are many challenges in this area. One of the largest one is human variability that could reduce predictability in human-agent teams. Additional challenges and elements with deleterious effects on teamwork, where human variability also plays a role, are miscalibrated trust, complacency, automation bias, degraded situational awareness, inadequate system-specific knowledge. Many different approaches have been proposed to mitigate these deleterious effects, primarily investigating different schemes for better functional allocation between humans and machine agents [18, 19, 20, 27,28] Unfortunately, most of these approaches suffer from limitations that make them applicable only in controlled and limited contexts [18]. Recent innovative approaches have been the privileged sensor framework [18] that treats the human as a special type of sensor and [9] Both call for integrating humans and autonomous agents based on their individual characteristics to exploit their strengths and mitigate their weaknesses to achieve enhanced human-agent team performance. This is at the crux of the STRONG program and of our proposed research effort.

Our overall research goal is to develop new theoretical principles that would characterize the processes and states of high performance human-agent teams in the overall context of individualized and adaptive technologies. To fulfill the overall research goal, our research objectives in Cycle 1 are: (i) investigate methods for individualized accounting in real-time of strengths and weaknesses of team mates in human-agent teams, (ii) develop efficient computational learning techniques for training and adaptation of agents to human roles, behaviors (policies), strengths and weaknesses, (iii) use the above individualized adaptation to investigate influencing individual team members (e.g. influence role allocation) as individual dynamics change in order to increase team performance, (iv) investigate team stability, rapport and adaptation of humans and agents, (v) evaluate the role of the gap in human and agent's speed of reasoning and acting, and its effect on team synchronization.

These fundamental research endeavors and results will constitute building blocks for follow-on options and later Cycles.

Our methods include theory formation, lab experiments, simulations and, in the outyears, field data collection in experiments with robotic platforms. These methods will support our overall research strategy which is to start by using a relatively simple game testbed, because it offers experimental control but also flexibility and scalability in complexity of agents, environments and teamwork scenarios. In subsequent follow-on options, we will extend the testbed in multiple ways including addition of multiple agent teammates, making the environment partially observable, and making the agents heterogeneous in terms of sensing and acting. These extensions reflect more realistic contexts from teamwork and will enable us to explore additional complexities and arising theoretical issues in human agent team processes. Extending the testbed and scenarios in these ways is one avenue for future scientific inquiry. In parallel, we will work closely with ARL researchers and other collaborators to identify more realistic scenarios, simulators and platforms to be used in follow-on options in later years of the grant.

Our proposed approach is three-pronged. First, we will develop a testbed called Team Space Fortress, by leveraging the Space Fortress testbed, that is open source and has been used for many years to perform psychology experiments. Space Fortress consists of a fortress whose location is static in the middle of the screen. The fortress can turn and orient itself to shoot missiles at an aircraft, controlled by a human that must also orient itself to align with the fortress and shoot missiles to destroy the fortress. There are several variants of this game. Space Fortress satisfies multiple characteristics of a test environment to facilitate and ground empirical and computational basic research including: (a) simple enough to be experimentally suitable, (b) tractable enough to allow easy implementation and

modification, (c) challenging yet learnable in experimental settings, and (d) flexible enough to investigate multiple scientific questions in human-agent team processes. In Team Space Fortress we will develop a team of two agents engaged in interdependent activities where one agent, the bait, establishes preconditions for the actions of the other agent, the shooter, with high risk to itself. In various experimental conditions, the players would be humans or agents. These humans and agents will have different levels of skills that will need to be exploited and integrated in appropriate ways to enhance team performance.

Second, we will develop a training and deployment approach that allows agents to learn complex behaviors based on human examples, and improve on them through extended self-practice and reinforcement learning. The agents will thus be able to dynamically adapt their behavior in real-time and also influence individual human team mates during deployment to enhance team performance. In particular, the agents will aim at identifying the strengths and weaknesses of their human team mate, so they can adapt (a) their execution policy to productively complement human capabilities, and (b) change their role, if they have learned through experience that the human may not have the requisite skills for good performance in her current role.

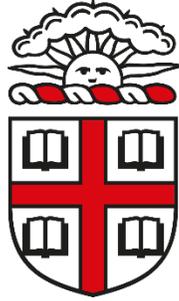
Third, we will design and run human experiments to test experimental hypotheses regarding (1) underpinnings of teamwork (rather than taskwork) in human-agent teams, (2) effectiveness of agent adaptation strategies for enhancing team performance, (3) team stability and possible human confusion as a result of the agent changing its role, and (4) the effect on teamwork of the gap in human and agent speed of reasoning and acting.

The research will deliver: (i) validated theoretical findings regarding teamwork processes and states, (ii) a teamwork testbed, Team Space Fortress, that we will make available to the research community, and (iii) experimental scenarios, results and training data that could also be available.

We will participate fully in the Summer Innovative Summit and will collaborate with ARL researchers and other collaborators to plan and formulate collaborative follow-on options.



Stefanie Tellex



BROWN



Ellie Pavlick



Peter Haas

Human-Robot Teaming with Natural Language and Mixed Reality

Project Investigators:

Stefanie Tellex: Brown University; robots that seamlessly collaborate with people to meet their needs using language, gesture, and probabilistic inference; Senior Investigator

Ellie Pavlick: Brown University; computational semantics and natural language inference; Junior Investigator

Peter Haas: Brown University; LI-DAR mapping and autonomous navigation; Associate Director of the Humanity Centered Robotics Initiative; Program Management / Research Technician

Graduate Student TBD; Research Assistant Undergraduate

Student TBD; Research Assistant

Summary:

The aim of this proposal is to test the hypothesis that an object-oriented Partially Observable Markov Decision Process (POMDP) framework can enable a shared mental model for a heterogeneous team of humans and robots using language, gesture, and mixed reality. This proposal will focus on building a mixed reality (MR) and language based interface for a human-robot team, driven by a POMDP framework for modeling each team member's mental state. Each human in the team will be equipped with a mixed reality headset that allows the robot to express its state and future actions. Additionally the human team can interact with the robot using language and gesture. We will create an initial prototype of this interface and assess its quality with a pilot user study. Our proposed effort will explore design concepts for MR interfaces for team formation and establishing shared mental models. By using a

POMDP framework, we aim to allow the robot to explicitly model the person's mental state and reason about information-gathering actions such as asking a question with natural language or showing extra information with MR to reduce uncertainty without overloading the human team members with too much information.

Technical Approach: We define a novel model, the Human/Agent Teaming Partially Observable Markov Decision Process (HAT-POMDP) to intelligently select when and what modality a robot should provide feedback based on its model of other team members. A Markov Decision Process (MDP) is a decision problem formalism in which an agent observes the state of the environment and takes actions in discrete time steps. It is defined by the tuple $\{S, A, R, T\}$, where S is the set of environment states, A is the set of actions, $R(s, a)$ is a reward function that specifies how much instantaneous reward is received for taking action a in state s , and $T(s, a, s')$ is the transition function that defines the probability of the environment transitioning to state s' after the agent takes action a in state s . The goal of the agent is to find an action in any given state (a policy) that maximizes the expected future reward. In an MDP, it is assumed the agent knows the true state at each timestep. For many problems, this assumption is invalid. Partially Observable Markov

Decision Processes (POMDPs) extend MDPs to describe the case when the agent can only indirectly observe the underlying state at each time step from a set of observations Ω . These observations are modeled as conditionally dependent on the true hidden state by the observation function $O(s, o)$, which defines the probability that the agent will observe observation o in state s .

Our HAT-POMDP will augment the POMDP framework by explicitly including a model of the person's mental state as part of the model, and model the effects of augmented reality interventions when choosing what to display to the person. Our previous work has shown that for an object fetching task, a robot can improve speed and accuracy by modeling a person's mental state compared to never asking a question or always asking a question. Our other previous work showed that we could improve the speed and accuracy of a person's collaboration with a robot by using augmented reality [20]. However this work used a fixed intervention of showing the robot's future path. This intervention provides information to the user, but might be overwhelming if every team member showed this information at all times.

Our proposed approach will selectively display intention information to teammates who are grouped with the robot in order to establish common ground. The system will decide what to show (in language, gesture, or mixed reality); when to show it; and to whom to show it based on its model of task performance and teaming. Mixed reality will allow us to quickly and efficiently ask teaming questions to resolve ambiguities. Figure 1 shows a prototype from one of our preliminary works demonstrating a mixed reality and language interface for commanding a robot drone. Here the human issues a command in language, and uses landmarks present in mixed reality. This demonstration shows one human and one robot. In the proposed work, we will expand this to a human-robot team with an air and ground robot.

Our proposed design will include the ability to visualize team membership and teammate goals in mixed reality and language. The HAT-POMDP will selectively choose what to show the human teammate to develop, manage, and communicate shared mental models of the problem, environment, and states and intentions of other team members. Mixed reality will provide a concrete, grounded framework to explicitly show intentions overlaid on the actual physical environment, eliminating many sorts of ambiguity. Language output can provide high-level descriptions of goals and intentions, and allow a person to communicate shared goals and intentions. Performing all this in a belief-space driven

way will allow the robot to model its own uncertainty and perform information gathering actions, including social actions or mixed reality actions to reduce uncertainty.

For example, it may first ask a question using language since it distracts the user less and does not require them to use their hands or a menu to reply. But if this method fails, it can fall back to mixed reality, which is less ambiguous but may take longer.

Assessment: We will perform two major assessments to test our hypotheses that the HATPOMDP combined with an MR interface improves the speed and accuracy of a human-robot team compared to baselines. Our proposed approach will use not only a drone interface as in our preliminary work, but add an additional ground robot, as well as two human teammates. The team will collaborate to collect information about an environment, using mixed reality to coordinate their activities, specifically to map and find targets in an environment with hybrid sensors, in coordination with the person. Our first assessment will consist of a design review and critique conducted by implementing our proposed approach and collecting feedback. We will carry this assessment out at the Summer Innovation Summit half-way through the program, by contributing a live demonstration of our prototype system for attendees of the summit to participate in. The second assessment will consist of a user study to assess the system quantitatively and qualitatively at a human-robot teaming task. Our task will be a patrol and surveillance task with a hybrid ground robot and aerial robot and lead to a publication at a top robotics conference. This task will set the groundwork for future phases of the program in which we will develop more sophisticated interfaces on more complex tasks with larger teams.

Research Team: The research team will consist of PIs Stefanie Tellex and Ellie Pavlick, as well as Peter Haas who will provide additional direction and participation at the Summer Innovation Summit. PI Tellex has expertise in collaborative robotics and language, as well as significant prior work in mixed reality applications in robotics. PI Pavlick has expertise in user modeling as well as compositional semantics and language processing. Haas has aerial robotics experience and will provide additional project management.