

Integrating Context into Artificial Intelligence: Research from the Robotics Collaborative Technology Alliance

Kristin E. Schaefer, Jean Oh, Derya Aksaray, Daniel Barber

■ *Applying context to a situation, task, or system state provides meaning and advances understanding that can affect future decisions or actions. Although people are naturally good at perceiving contextual understanding and inferring missing pieces of information using various alternative sources, this process is difficult for AI systems or robots, especially in high-uncertainty and unstructured operations. Integration of context-driven AI is important for future robotic capabilities to support the development of situation awareness, calibrate appropriate trust, and improve team performance in collaborative human-robot teams. This article highlights advances in context-driven AI for human-robot teaming by the Army Research Laboratory's Robotics Collaborative Technology Alliance. Avenues of research discussed include how context enables robots to fill in the gaps to make effective decisions more quickly, supports more robust behaviors, and augments robot communications to suit the needs of the team under a variety of environments and team organizations and across missions.*

Integrating context to support AI development provides a number of potential benefits for efficient teaming and collaborative task accomplishment for human-robot teams. For military teams in particular, integration of context into AI architectures is essential to facilitate collaboration and successful operation in complex and dynamic environments. Take, for example, when a soldier reports that a hostile threat is in a target area. Given this information, a robot could be expected to change how it navigates to

the target environment, make its primary objective enemy detection, and provide guidance for the movements and future actions of both friendly and adversarial human counterparts so that team members can remain undetected. However, a human teammate's interpretation of the robot's behaviors is directly influenced by the robot's ability to adequately communicate reasoning for its own previous and current actions. Otherwise, its behavior may appear ambiguous or incorrect from a human perspective. Therefore, the robot needs to understand both how context will or could affect its own decisions as well as how it could affect team members' decisions. Integrating contextual understanding allows shared situation awareness and shared mental model development, improves joint decision making and categorization of data, provides better processing times, and enhances learning both online and offline for the team.

Challenges in Developing Context-Driven AI

Although our early work identified the potential benefits of incorporating context in inference and planning models to enhance shared situation awareness and improve intent-based communication of human-robot teams, we have also identified a number of open research challenges in making these types of advances (Schaefer et al. 2017; Schaefer et al. 2019). These identified gaps in the scientific research community include a large number of unknowns about what constitutes context, differing opinions on how to reason about context, differing recommendations for how to acquire new contextual knowledge, and even the specific means to transfer or communicate the teammate's knowledge of context to other members of the team. We explain these open challenges in detail in the following subsections and then describe how the current work by the Army Research Laboratory's Robotics Collaborative Technology Alliance (RCTA) has been addressing and advancing research in this area.

Representing Context

Context means different things to different people. One well-cited, yet broad, definition suggests that context is "any information that can be used to characterize the situation of an entity [whereby an] entity is a user, a place, or a physical or computational object that is considered relevant to the interaction between a user and an application, including the user and application themselves" (Dey, Abowd, and Salber 2009, p. 106). Long, all-encompassing definitions like this one, with a

lack of consensus for an agreed upon operational definition, have resulted in large variability in the representations of context across the scientific community. Depending on the problem, context might refer to a relevant part of the state-space (Xiong and Huber 2010), a probability distribution over the concepts in an environment (Singhal, Luo, and Zhu 2003), a set of relationships between objects (Rabinovich et al. 2007), logical statements that represent cause and effect (Zettlemyer and Collins 2009), or a function to select relevant features for object recognition (Heitz and Koller 2008). Although these representations can capture some aspects of contextual knowledge, none of them are applicable to general settings because of their strong dependence on domain-specific knowledge. Overall, the key gaps are linked to reconciling different views of context and developing a coherent representation that can be used in various ways.

Inferring Context

In the current literature, inferring context from the world occurs primarily through use of visual and language-based sensor data. However, limitations in current perception systems can inhibit inference capabilities. For example, the presence of a highly cluttered environment, dynamic objects, and even changes in the image resolution can cause failures in accurate scene detection and inference processing. Similarly, inference models for natural language understanding might struggle in the presence of arbitrary sentences with complex grammatical structure and untrained words. Therefore, a future goal is to extend the application areas of existing inference methods and models for generalization. A second goal is to use sensor data other than vision or language, such as temperature, humidity, smell, and audio, to provide additional, and possibly redundant, sources of information to support these inference models. In a mixed-initiative team, these factors (such as heat exhaustion, distraction, inability to hear) may influence performance of human team members, requiring adaptations to robot behavior for optimal collaboration. Moreover, temperature and humidity support inferences about potential changes in weather conditions (for example, fog, wet roadways), smells indicate the presence of chemicals affecting safety, and auditory sensing can be crucial in inferring changes in the scene, such as approaching vehicles or people or imminent danger. Overall, the challenge is in the identification of which combination of sensory information gives an efficient way for better understanding of context to improve decision making. A major limitation to this area of research is

the limited knowledge on how to derive the required types of sensor data in terms of mission specifications that may make for an efficient context inference.

Learning New Classification Schemes

Given a finite set of features of a context model (for example, indoor versus outdoor spatial areas), it is possible to learn classification schemes to determine the current context. However, prespecified contextual features are unlikely to be sufficient for long-duration robots operating in populated environments. Because people naturally grow their sense of context over time, it is desired for robots to have the same capability for resilient team-based operation. It is not realistic to assume that all possible contextual features are known before a mission. In this sense, mining new contextual knowledge online and incorporating it into context models are two crucial issues for improving the performance of robots that are likely to collaborate with humans. Researchers are developing algorithms to acquire new knowledge from the perceived world and to reconstruct the inference model as new information is added (Tucker et al. 2017). However, new methods are needed to address the problem of updating the current context model in a scalable way for incrementally discovered information. The main challenge here is that there is not a common solution to this problem because the context models are diverse, application dependent, and incapable of accommodating all types of contextual knowledge.

Communicating Contextual Information

In team collaboration, any contextual variable that drives the inference process must be transparent to team members to build trust and to explain future decisions based on context derived from sensor data. People reason about the world in a way that incorporates diverse contextual information about logic concepts, prior sensor data, or time histories of state estimates. It is possible to use this information in their inference and reasoning processes to make decisions. When sharing that level of contextual understanding with robotic team members, it is important to enable direct (for example, language, text) or indirect (for example, gestures, emotions, posture) communication methods. Similarly, not all information that has been formed through AI is transparent to humans. For instance, graphical models and neural networks are powerful representations in a robot's inference mechanisms that they can use to instantiate contextual knowledge from a set of sensor

data. However, these mathematical representations without annotated explanations may not necessarily be meaningful from the perspective of a human. Thus, similarly to human-to-robot communication, it is crucial to develop a communication architecture that can also support information transfer from robot to human in a transparent fashion. In other words, how do we transform what is in the black box of machine learning or other algorithms into something a human can easily understand? Furthermore, how do we represent this information through visual or other modes so that comprehension of the data is quick without overloading human cognition?

Current Work

This article reviews work on developing context-driven AI for human-robot teaming conducted by the RCTA that shows advances to the research gaps noted. Within the RCTA efforts, context in support of human-robot teaming is defined as any available information that can fill gaps, addressing uncertainty, to enable shared understanding and team collaboration. When integrated appropriately, context supports teaming initiatives for collaborative interactions including planning and prediction, communication, advanced mission goals, and independent as well as collaborative robot decision-making capabilities.

The following sections of this article describe advances in theoretical and applied contributions, as well as research and development efforts, to advance context-driven AI in support of advanced collaborative human-robot teaming. It begins with a general overview of the research goals of RCTA and how advancing context-driven AI is an integral component of that research. The importance of environmental, mission-specific, and social context for advanced human-robot teaming is discussed. This discussion is followed by specific RCTA research efforts advancing the research and development associated with the aforementioned challenges for developing context-driven AI. Our work describes the development of a multimodal interface (MMI) that supports advances in context-driven AI related to natural language, world modeling, and novel concept acquisition. All these technical advances benefit collaborative human-robot teaming.

RCTA

RCTA is a large multidisciplinary research program that includes a consortia of industry and academic partners working directly with government organizations to advance the state of the art in human-robot teaming or, in other words, to revolutionize

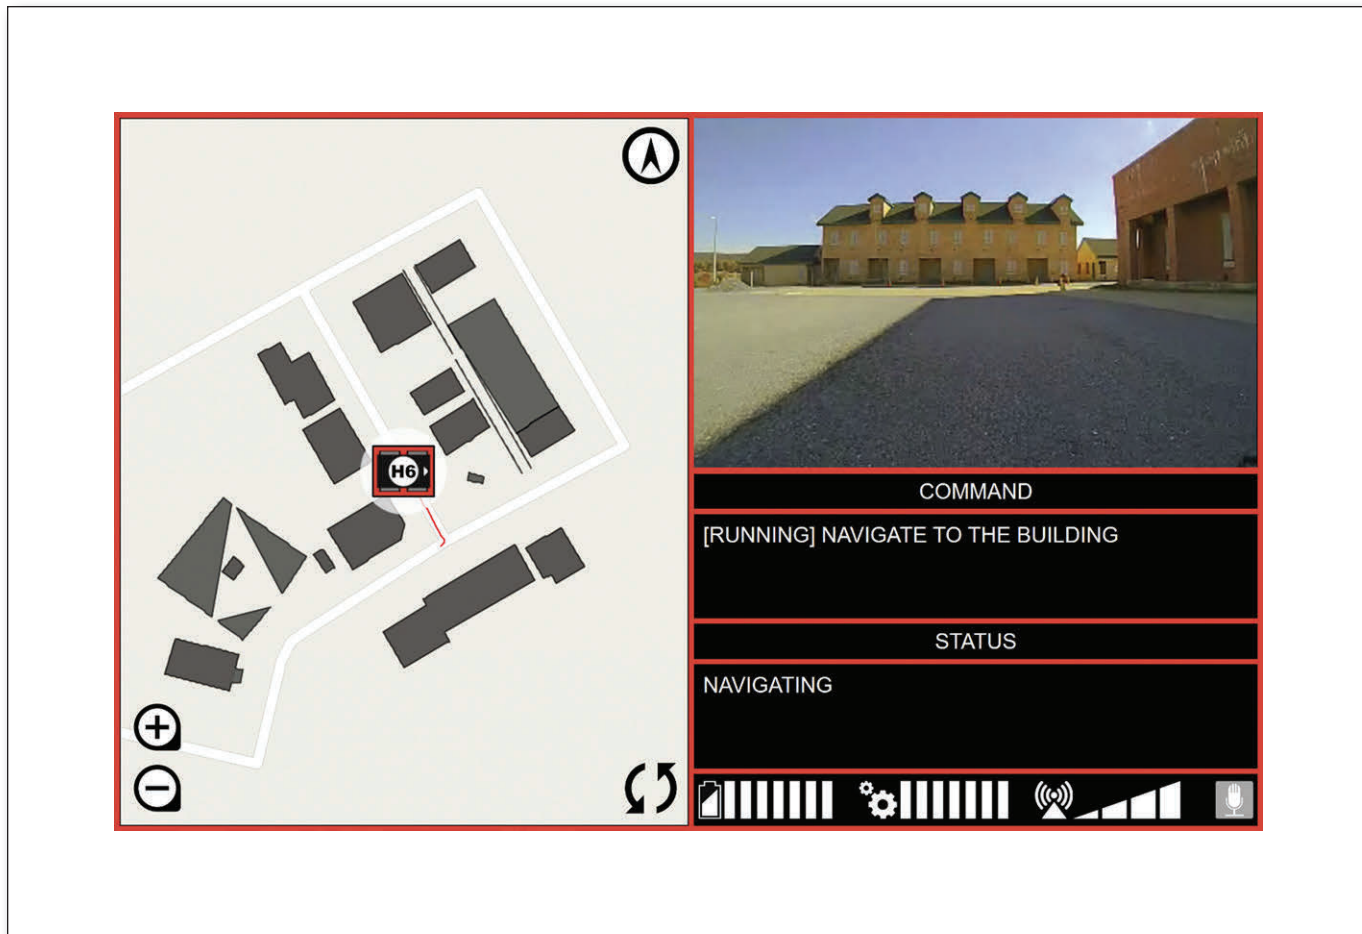


Figure 1. RCTA Multimodal Interface Visual Interface.

The display includes a semantic map (left), video from the robot's perspective or other imagery data (top right), and the robot's action and health status (bottom right).

robots from tools to teammates (Phillips et al. 2011). There are four main research areas and required capabilities: (1) OPTEMPO maneuvers¹ in unstructured environments, including mobility in dynamic scenes and across rough terrain, (2) human-robot execution of complex missions requiring situation awareness of unstructured environments and distributed mission execution, (3) mobile manipulation in cluttered spaces, and (4) integrated research that combines and assesses capabilities delivered from the other thrusts on multiple robotic platforms. Although a number of research interests are being addressed to advance teaming, context plays an important role in all of these areas. In particular, subcategories of RCTA research that drive the advancements in context-driven AI include advancements in semantic perception, adaptive behavior generation, metacognition, machine learning, and a hybrid cognitive and metric world model.

Semantic perception moves robotic perception beyond simply detecting what is or is not an obstacle toward semantic understanding of an environment in a way similar to how human team members would perceive or reason about the environment, for example, by recognizing the types of objects and terrains of interest for a specific task, such as navigation (Oh et al. 2015a; Oh et al. 2015b; Oh et al. 2016; Shiang et al. 2017). Adaptive behavior generation combines previously developed robotic planning algorithms, machine learning techniques, and semantic understanding of an environment within the context of a high-level task. This enables robots to generate effective mission plans in partially known and unstructured environments and to compute these plans online whenever necessary by following natural language commands (Boularias et al. 2015; Boularias et al. 2016; Paul et al. 2016; Paul et al. 2017; Paul et al. 2018; Tucker et al. 2017) or navigating while



Figure 2. Natural Language Understanding Integrated in Various Robot Platforms.

Barret robot arms (left), Clearpath Husky robot (center), and CMU Ballbot (right; see Lauwers, Kantor, and Hollis [2005]). The difficulty in integrating context-driven AI is in relation to perspective. In the image on the left, perspective is in the word *left*. What does left of the green block mean — my left or the robot's left? In the center image, the key word is *back*. What is the back of a building? In the image on the right, there are questions as to which laboratory or which table and what should be told to the human.

adapting to social context (Vemula, Muelling, and Oh 2017; Vemula, Muelling, and Oh 2018).

Metacognition enables the use of intuitive, human-level commands for soldier-robot communication to facilitate the creation of shared mental models and the development of shared situation awareness (Ososky et al. 2013). The world model spans a range including traditional metric data to an associated semantic understanding to underpin the cognitive levels of reasoning (Dean 2013). Within the RCTA goal of advancing robots from tools to teammates, there are three different but interrelated contexts that are directly relevant to improve teaming capabilities: an environmental context related to the physical world, a mission-specific context derived from the task criteria and goals, and a social context related to the agents and interaction of agents within the environment.

Environmental Context

Collaborative team operations occur within an environment that directly influences the perception of team members' actions and actual decision-making behaviors. The main difficulty is that humans are very quick to infer meaning from their surroundings. However, the process that robots use for inferring meaning from the world often does not match the process or capabilities of the human, leading to different planning and reason-based decision making (Perelman et al. 2017; Perelman et al. 2018; Schaefer et al. 2018). Although difference is not necessarily negative, it does affect the development of shared situation awareness and shared understanding among the team (Chen et al. 2014; Wright et al. 2017).

Context from the environment is semantic information perceived dynamically or provided a priori (for example, terrain maps). Traditionally, the information extracted from the environment is limited to physical objects and used for navigation and mobility — for example, to go from point A to point B as quickly as possible. However, the semantic environment includes much more than physics, such as additional data derived and inferred from the aggregation of world model information. It can be attributes associated with semantic objects in a robot's world model, part of short- and long-term memory, or temporary variables computed in decision-making algorithms. Hence, context is important to understanding the scene in terms of functionality and affordances, inferred relationships between objects, and influences of the world on the people in that world. Functionality and affordances provide contextual understanding through interpretation of a scene. For example, a person may be sitting on a box, but a box also holds things. What is in the box could change the amount of risk associated with operating in that particular area of the environment. Inferred relationships between objects can also directly affect navigation and mobility. For example, a tree is usually an object that should be avoided during path planning, but trees provide cover; operating within close proximity to a tree line could improve stealthy maneuver. It is also important for a robot to be able to interpret and infer how environmental context can influence people (that is, psychophysiological states including stress, fatigue, and workload). Overall, this type of data supports priming of perception

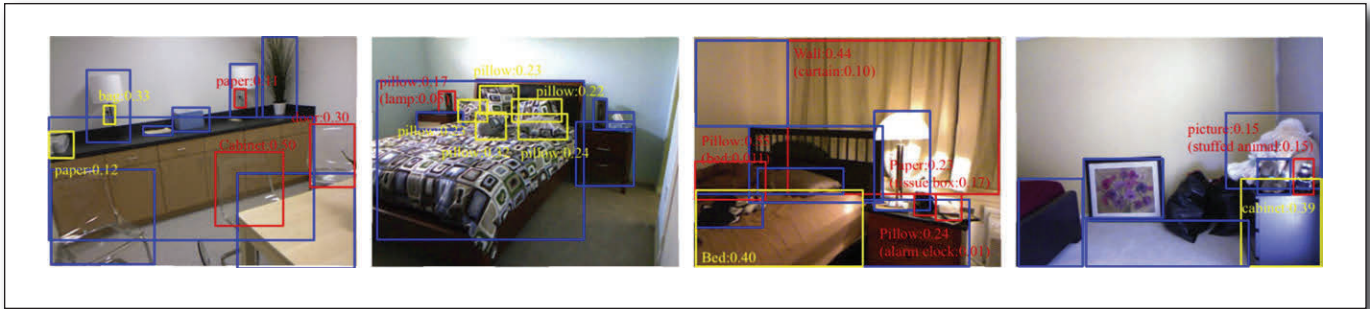


Figure 3. How Fusing Vision and Language Can Improve Object Recognition Results.

The blue boxes represent the correct classification by vision, the red boxes represent misclassification by vision, and the yellow boxes represent those that have been corrected by using language with one to three spatial descriptions (for example, the chair is near the table).

systems to better focus on what should be versus what might be in a scene.

Advancement of environmental context-driven AI will produce a better understanding of social situations, as well as improved detection of specific environmental features that enhance scene understanding and reasoning for observed behaviors. The RCTA research is specifically looking at how the integration of AI can incorporate elements of environmental context into the decision-making process, which would offer two major benefits. First, it would allow the robot to reason about the world, which could affect its subsequent decision-making process, and second, it would support natural language and bidirectional communication between the robot and other team members to facilitate collaborative decision making for mission execution.

Mission-Specific Context

A task or complex mission goal given to a robot is a rich source of information to be included in AI, and it must be able to support decision making that accounts for rules of engagement, social norms, and the prioritization of objectives. This context can therefore reduce the solution space of actions related for a given decision with heuristics or optimize selections within the space with improved weighting of variables. In line with the earlier example of an enemy reported in an area, the type of mission underway will dictate the robot's goal prioritization and behaviors. For example, should the robot try to scout ahead, stay in formation with the team, or take steps to avoid detection? Thus, context influences the robot's selection of an actions execution. From a team dynamics perspective, to accomplish a task, each member has its own set of objectives and goals that may be shared or may be independent of each other and the robot. Task or mission context directly tells the robot what role it should be in and what roles others will take, such that each individual

is doing what is best for overall team performance, especially if the team is distributed throughout an area and not collocated. The RCTA research is specifically looking at how the decomposition of the overall mission goal into tactical behaviors is sensitive to the context of the instruction given by a human teammate related to both the events in the overall environment and to the specific goals of all the team members.

Social Context

Social context is important for teaming because it supports collaborative decision making, including decision making associated with navigation and how both robots and humans can operate in a shared space (Schaefer et al. 2016; Schaefer et al. 2017; Straub and Schaefer 2018), advanced robot awareness and its understanding of human behavior such as inferring human intent through pose, location, and social signals (Fiore et al. 2013; MacArthur, Stowers, and Hancock 2017; Schaefer, Hill, and Jentsch 2018; Wiltshire et al. 2016), and prediction of future activities based on context related to the scene, objects, and people. Social context guides the robot's interactions with ambient society where people are not bound to mission roles or context and supports detection of team member cues that may result in behaviors outside of mission norms. Each member of a team navigating through crowded environments is required 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 recognizing and anticipating the motion of humans, social signals, and the activities in which they are engaged. RCTA efforts are developing approaches to recognize and make inferences about people's actions from the robot's visual data streams to track and predict future actions of

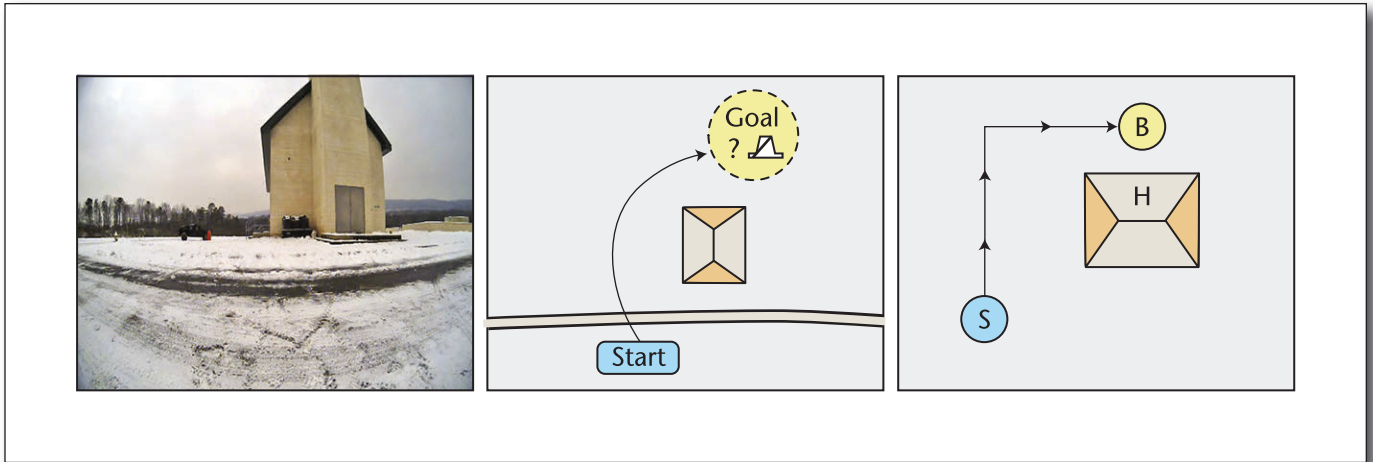


Figure 4. An Example of How People Plan in Context.

This figure illustrates how people plan for the command, “Stay to the left of the building; then go to the barrel behind the building,” given the leftmost image showing an outdoor environment. On the right are two examples of plans drawn by human subjects for the given command, depicting the start position of the person, the assumption of where the barrel will be located behind the building, and the person’s planned path for attempting to reach the barrel.

people and other dynamic objects. Here, human pose features can be used to provide context to the human’s current or potential future activities. Understanding social interactions includes not only being able to predict the actions of other agents but also reasoning about how one’s own actions can affect those of other agents. This idea has led the RCTA investigation to socially compliant navigation planning and computational models for representing such social interactions, for example, using a variation of Gaussian processes (Vemula et al. 2017) and a deep learning model known as social attention (Vemula et al. 2018). These models use context to better support the effectiveness of team communications, minimize the associated cognitive burdens for each teammate, and improve synchronization of human-robot team member tasks.

Research Advances in Context-Driven AI

Today’s soldiers are required to operate within inherently complex, dynamic environments. Working within teams, regardless of the presence of a robot, includes multiple activities during which these soldiers must pay attention to their own task execution and their teammates. Inclusion of robots can easily result in teammate overload for a number of reasons. First, robots do not communicate or think like human teammates. Second, while they are equipped with advanced sensors capable of streaming large amounts of data, the data still must be parsed for human consumption. Third, there is a degree of

required oversight and additional attention for operations based on the capabilities of the autonomy. Research conducted through RCTA has sought to address these issues through natural and intuitive bidirectional communication. These efforts have led to the development of an MMI (Barber et al. 2015; Barber, Howard, and Walter 2016) that facilitates the exchange of environmental, mission, and social context to AI efforts (see figure 1). Specifically, this technology and associated research supports the development of a common ground where shared understanding allows for joint decision making and collaborative operations.

The RCTA MMI has leveraged work in speech recognition, natural language understanding, gesture recognition, synthetic speech, tactile displays, and visual displays that model human-robot communication after human-human communications, providing natural and redundant communication channels (see, for example, the work of Duvallet et al. 2016; Elliott, Hill, and Barnes 2016; Mortimer and Elliott 2017). These modalities (audio, visual, and tactile) provide additional ways to tailor the MMI to adapt the communication of information between teammates to ensure message delivery, robust (fault-tolerant) communication, and shared situation awareness. For example, if a robot determines a soldier is fatigued or under high workload, the MMI could enable haptic feedback and increase the frequency of radio communications to ensure a message receipt. Moreover, understanding where a soldier is looking or pointing can add context to the robot’s understanding to resolve an ambiguous command, such as to go behind a building,

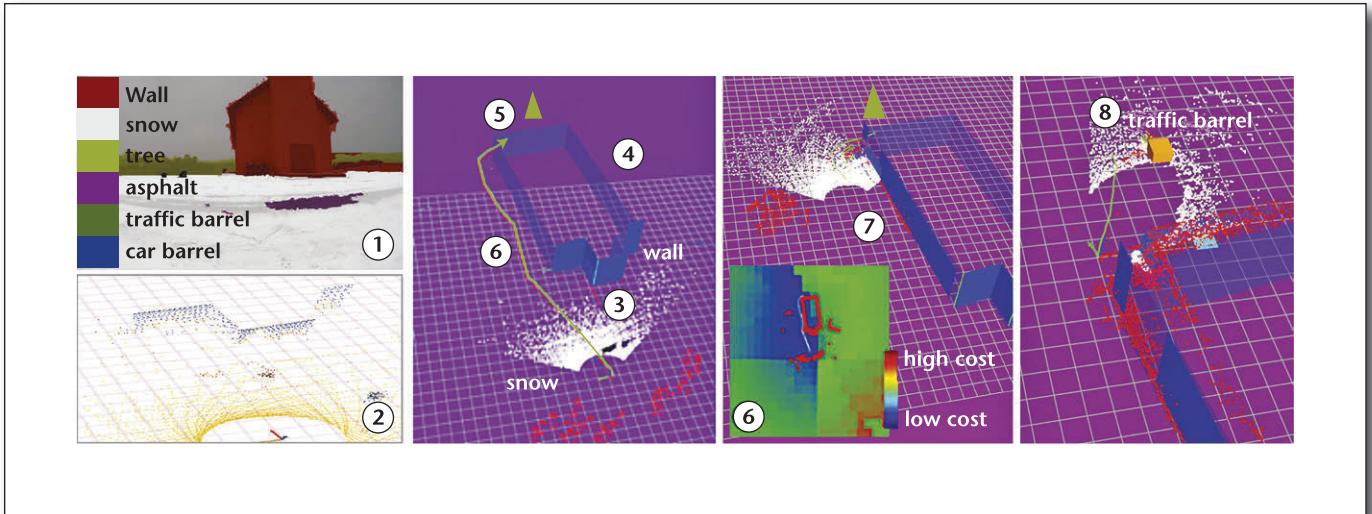


Figure 5. Assumptive Planning Approach.

This is an example illustration of the assumptive planning approach for a robot given the command, Stay to the left of the building; then go to the barrel behind the building (Oh et al. 2015b). Steps 1 and 2 show the camera and LIDAR sensor data. Steps 3 to 6 represent the hypothesized space for robot reasoning about the spatial constraint behind the building needed to locate the hypothesized barrel as a target goal and generate a plan. Steps 7 and 8 demonstrate the robot's ability to continuously update its world model of the environment and its subsequent plan as it perceives more information about the actual environment.

by denoting a specific building or what structures in the environment the team considers buildings.

The MMI is a gateway device toward a shared context between robot and human teammates. It provides flexible methods for a soldier to give high-level commands to a robot, such as, "Screen the back of the building," down to specific goal-driven semantic navigation, such as, "Go to the town square." It also enables bidirectional communication of a robot's state on a continual basis or through speech dialogues to request scene descriptions or explanations of its behavior (for example, "Where are you going?" or "What do you see?"). Beyond the provision of dialogue with the robot and a command input, the MMI further attempts to classify the human teammates' states to contribute to the others' environmental context. Information about what is in the environment informs the MMI about factors that may be changing the soldier's decision-making behaviors. Hence, the MMI is not just a portal into the robot for human teammates but also a sensor about the humans for the robot. This sensor facilitates the acquisition of information for all three categories of context, capturing what each soldier is doing, where they are, physiologic and cognitive capacity, and what information they are communicating to all supporting actors. These specific technical advances to context-driven AI support natural language communication, world model development, and novel concept acquisition.

Context-Driven AI and Natural Language

One important component of the MMI is development that could support natural language communication. Natural language is a capability critical to facilitating direct human-to-robot mission-specific communication. Speech is the most commonly used method of interaction among human teammates. When a team of agents (human and robot) is performing a shared task, the clarity of the communication and how the context is understood are crucial for the team's success. How language is understood directly affects development of the shared context, that is, whether teammates interpret the task and the environment in the same way such that they can perform the task as one cohesive team. However, our work in developing AI for natural language processing brought to light a number of challenges with integration of robot perception and associated cognition for interpreting human-to-robot communication and performing associated actions that support the team leader's intent. The key finding in this research was that the addition of visual descriptors alone does not provide enough contextual understanding to initiate appropriate robot response (figure 2).

Natural language processing can also support reduction of ambiguity and improved shared situation awareness by leveraging teammates' inferences of environmental or social context. This added capability can support current technical limitations

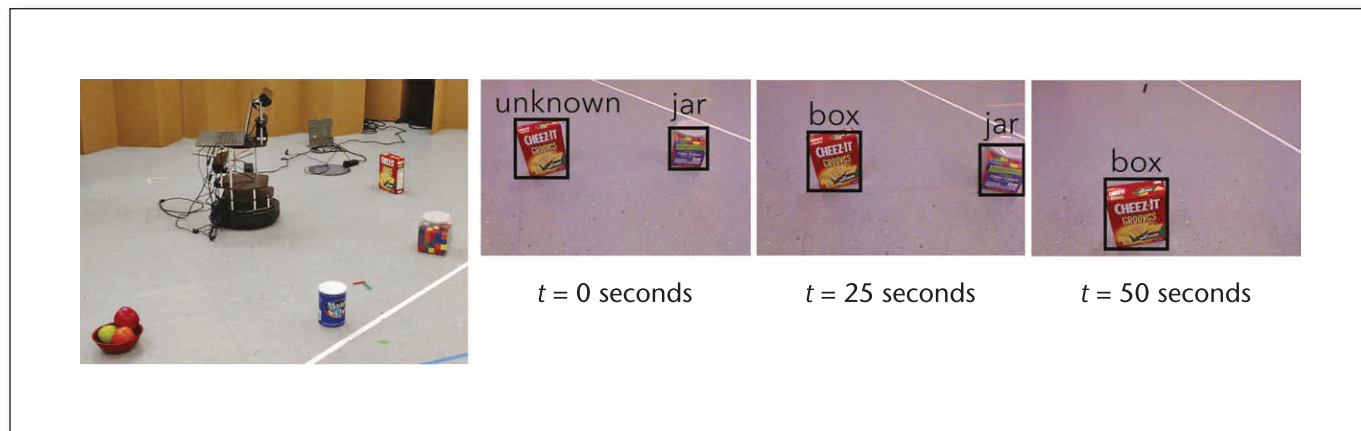


Figure 6. Online Concept Acquisition in Complex Partially Known Workspaces.

This figure demonstrates a grounding scenario (left image) with the TurtleBot mobile platform in an environment populated with objects (that is, box, jar, can, and fruits in clockwise order) from the YCB data set (Calli et al. 2015). The robot is equipped with a Kinect sensor with a limited field of view ($62^\circ \times 48.6^\circ$). The goal is to acquire new grounding symbols (right images; see Tucker et al. [2017]). For example, the robot has a model for grounding a jar object but was not trained to recognize or ground a box object. The robot receives a command *move to the box*. Due to the presence of an unknown object in its perceived world, the model grounded the unknown phrase *the box* to the unknown object and drove to the box. Online retraining was performed with the acquired set of visual observations and the lexical token *the box*. Inference occurred in 2.34 seconds.

in robot perception. Although recent advances in computer vision show a promising future, to date, robotic perception in real-life environments remains a difficult challenge for context because of variations in lighting conditions, illumination, weather, seasons, and more. For instance, recognizing objects in an outdoor environment based only on computer vision can include both false positive and false negative errors that can result in serious misinterpretation of higher-level tasks. Leveraging language-understanding skills, robots may harvest additional context information to improve their understanding, such as through listening to human teammates in the same environment. For example, consider the navigation task in figure 2, in which a person commands a robot to navigate to the building that is behind a vehicle. Even if the two team members are in the same environment, the way they see the environment may be different. For instance, because of its imperfect perception, the robot may misperceive the camouflaged vehicle for part of the background environment, such as a tree. From the robot's perspective, the command to go to the back of the vehicle, which does not exist in its world model, does not make sense because it is unlikely that the human teammate would use a landmark not close to the robot. Within the context of this task, the robot may update its world model to be more consistent with the command so that the object the robot first recognizes as a tree may be the vehicle of interest.

To account for these issues, the RCTA has looked into how language understanding improves vision-

based object recognition on an open data set (NYU Depth Dataset V2, [Silberman et al. 2012]). Context-based reasoning is modeled as multimodal understanding whereby the robot continuously updates its task context by fusing new information from vision and language understanding (figure 3). Graphical model approaches such as random walk (Shiang et al. 2017) and conditional random fields (Shiang, Gershman, and Oh 2017) use context for object recognition, whereby the nodes and the edges of a graph represent objects and their relationships, respectively. These models represent the robot's world model, where the contextual relationships among object types can be updated using both vision and language. This general idea has been applied to various human-robot teaming problems to enable RCTA robots to perform complex tasks in natural language (Boularias et al. 2015; Oh et al. 2015b; Oh et al. 2016). Our intelligence system for multimodal understanding has been integrated on several physical robot platforms with varying capabilities (figure 2). These examples demonstrate the use of domain-specific contextual information for various tasks, including manipulating tabletop objects and navigating in outdoor environments, and in indoor perception.

Context-Driven AI and World Modeling

Context can be useful when robots operate in an unknown or partially known environment. Consider

a navigation command, “Go to the barrel behind the building,” when an agent is situated as pictured in figure 4 (left image). When this task is given to humans, most people come up with a plan in which they hypothesize the rear side of the building not shown in the picture, the space behind the building, and a path toward an imaginary goal.

The RCTA robot planning approach follows this humanlike planning concept, known as assumptive planning, whereby the robot must hypothesize the unseen part of an environment to fill the gap between the given command and its world model. The idea of assumptive planning enables the robot to plan and execute commands that require environmental context reasoning in an unknown or partially known environment. Figure 5 illustrates how the robot solves the same navigation command using camera and three-dimensional LIDAR sensors. The semantic analysis is based solely on what has been sensed. The robot then gradually propagates the information to include the hypothesized space. At the same time, the robot computes a metric cost map representing the commander’s preference — stay to the left of the building — with the knowledge it has gained from having been trained offline via a machine-learning technique called imitation learning, in which experts demonstrate desired behaviors (Boularias et al. 2015; Boularias et al. 2016). The robot continuously updates its world model of the environment, known as the world model, and its subsequent plan as it perceives more information. Using context, the robot is able to accomplish tasks that it had previously been unable to complete because of missing information. In two sets of extensive outdoor experiments during the RCTA assessments, the task-level performance for navigation improved from 50 percent to 75 to 93 percent under conditions with varying difficulties (Oh et al. 2016).

Context-Driven AI and Novel Concept Acquisition

A robot must be able to communicate known and unknown information. Quantifying what is unknown makes it possible to dictate needs for contextual understanding to fill the gaps in a robot’s knowledge and reasoning process. A model was developed under the RCTA research efforts to enable learning the meaning of a large variety of phrases in complex environments, facilitating learning new words and objects online. The model contributed to language-guided models that enabled online concept acquisition in complex partially known work spaces (Tucker et al. 2017). For example, the leftmost picture in figure 6 illustrates a robot and four objects in the environment. Initially,

the robot perceives an unknown object and a jar and does not know the box. When the robot is given a command such as “Go to the box,” it incorporates environmental context into its decision making (that is, reasoning about the existence of a box object around itself) and infers that the perceived unknown object is supposed to be a box, then approaches. Moreover, as it approaches, it collects various visual images of the object and uses this information in its model to recognize boxes in future instances. This approach aids reasoning about a world that contains known and unknown objects and gaining new knowledge through human-robot communication.

Humans are good at understanding abstract notions such as a group of cars, a row of blocks in front, or the nearest door of a row of three doors. Recent work has extended existing grounding models to accommodate abstract notions (Paul et al. 2016) or acquire factual knowledge (for example, lift the box that I put down, or, this is my helmet, pick it up) from natural language instructions (Paul et al. 2017). These works significantly extend the space of commands that robots can understand and enable a better understanding of temporal and spatial context for more effective and efficient human-robot interaction.

Conclusions

The RCTA research efforts have demonstrated that environmental, task, and social context can significantly support human-robot teaming through the development of effective bidirectional communication. To date, the primary work has been in improving robot perception and learning to support tasks related to navigation. The development of an MMI has facilitated major improvements in the development of context-driven AI through natural language understanding, updating the world model, and learning through identification of unknown objects. Additional investigations are underway to understand how to improve the robot’s understanding of context through effective bidirectional dialogues. These dialogues include natural language understanding capabilities to better categorize speech into commands, queries, and reports for the joint construction of context using the research we described here.

Future human-robot teams are envisioned to operate in high-risk OPTEMPO and cluttered environments. Therefore, the robot must be able to autonomously adapt its communication modalities and the types of information it possesses to better support team member needs in times of high workload or stress. One of the ongoing research

efforts within RCTA aims to integrate a better contextual understanding of the human team member through the real-time classification of a human's state using wearable technologies (for example, heart rate sensors) to infer human work load, stress, and possibly even trust. This source of information will allow the MMI to adapt a relevant modality of communication, frequency of communication, and types of information transmitted among human and robot teammates while also supplying additional input to AI decision-making modules to maximize effectiveness of interactions and decision making.

Acknowledgments

This research was sponsored by the US Army CCDC Army Research Laboratory and was accomplished under Cooperative Agreement W911NF-10-2-0016. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army CCDC Army Research Laboratory or the US Government. The US Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. We thank Dr. Nicholas Roy, MIT, and Dr. Susan G. Hill, ARL, for reviews and guidance on this article. We also acknowledge the members of the RCTA Consortium from the CCDC Army Research Laboratory, Carnegie Mellon University, Florida State University, General Dynamics Land Systems, Jet Propulsion Laboratory/California Institute of Technology, Massachusetts Institute of Technology, QinetiQ North America, University of Central Florida, and University of Pennsylvania who have supported these research efforts.

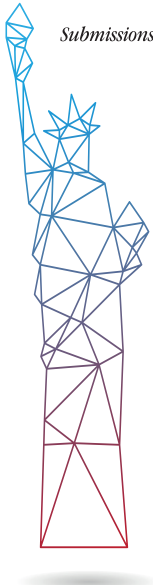
Note

1. Operations tempo for pace of an operation, its planning, and resupply (Castro and Adler 1999).

References

- Barber, D.; Abich, I. V. J.; Phillips, E.; Talone, A.; Jentsch, F.; and Hill, S. 2015. Field Assessment of Multimodal Communication for Dismounted Human-Robot Teams. In *Proceedings of the 59th Human Factors and Ergonomics Society Annual Meeting* 59(1):921–5. Los Angeles: SAGE Publications.
- Barber, D.; Howard, T.; and Walter, M. 2016. A Multimodal Interface for Real-Time Soldier-Robot Teaming. In *Proceedings of SPIE Defense and Security: Unmanned Systems Technology XVIII* 9837: 98370M. Bellingham, WA: SPIE.
- Boularias, A.; Duvallet, F.; Oh, J.; and Stentz, A. 2015. Grounding Spatial Relations for Outdoor Robot Navigation. In *Proceedings of the IEEE Conference on Robotics and Automation*, 1976–82. Piscataway, NJ: IEEE. doi.org/10.1109/ICRA.2015.7139457
- Boularias, A.; Duvallet, F.; Oh, J.; and Stentz, A. 2016. Learning Qualitative Spatial Relations for Robotic Navigation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 4130–4. Palo Alto, CA: AAAI Press.
- Calli, B.; Singh, A.; Walsman, A.; Srinivasa, S.; Abbeel, P.; and Dollar, A. M. 2015. The YCB Object and Model Set: Towards Common Benchmarks for Manipulation Research. In *Proceedings of the International Conference on Advanced Robotics*, 510–7. Piscataway, NJ: IEEE. doi.org/10.1109/ICAR.2015.7251504
- Castro, C. A., and Adler, A. B. 1999. OPTEMPO: Effects on Soldier and Unit Readiness. Parameters. *Journal of the US Army War College* 29: 86–95.
- Chen, J. Y. C.; Procci, K.; Boyce, M.; Wright, J.; Garcia, A.; and Barnes, M. 2014. Situation Awareness Based Agent Transparency, Technical Report ARL-TR-6905. Aberdeen Proving Ground, MD: US Army Research Laboratory.
- Dean, R. M. 2013. Common World Model for Unmanned Systems. In *Proceedings of SPIE Defense, Security, and Sensing: Unmanned Systems Technology XV* 8741: 87410O. Bellingham, WA: SPIE. doi.org/10.1117/12.2016606
- Dey, A. K.; Abowd, G. D.; and Salber, D. 2009. A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications. *Human-Computer Interaction* 16(2–4): 97–166. doi.org/10.1207/S15327051HCI16234_02
- Duvallet, F.; Walter, M. R.; Howard, T.; Hemachandra, S.; Oh, J.; Teller, W.; Roy, N.; and Stentz, A. 2016. Inferring Maps and Behaviors From Natural Language Instructions. In *Experimental Robotics*, edited by M. Hsieh, O. Khatib, and V. Kumar 373–88. Vol. 109. *Springer Tracts in Advanced Robotics*. Berlin: Springer. doi.org/10.1007/978-3-319-23778-7_25
- Elliott, L. R.; Hill, S. G.; and Barnes, M. 2016. Gesture-Based Controls for Robots: Overview and Implications for Use by Soldiers, Technical Report, ARL-TR-7715. Aberdeen Proving Ground, MD: US Army Research Laboratory doi.org/10.21236/AD1011904
- Fiore, S. M.; Wiltshire, T. J.; Lobato, E. J.; Jentsch, F. G.; Huang, W. H.; and Axelrod, B. 2013. Toward Understanding Social Cues and Signals in Human-Robot Interaction: Effects of Robot Gaze and Proxemics Behavior. *Frontiers in Psychology* 4(859): 1–15. doi:10.3389/fpsyg.2013.00859
- Heitz, G., and Koller, D. 2008. Learning Spatial Context: Using Stuff to Find Things. In *European Conference on Computer Vision*, edited by D. Forsyth, P. Torr, and A. Zisserman, editors, 30–43. Vol. 5302. *Lecture Notes in Computer Science*. Berlin: Springer. doi.org/10.1007/978-3-540-88682-2_4
- Lauwers, T.; Kantor, G.; and Hollis, R. 2005. One Is Enough! In *Robotics Research: Results of the 12th International Symposium ISRR*, edited by S. Thrun, R. A. Brooks, and H. Durrant-Whyte, 327–36. Berlin: Springer.
- MacArthur, K. R.; Stowers, K.; and Hancock, P. A. 2017. Human-Robot Interaction: Proximity and Speed — Slowly Back Away from the Robot. In *Advances in Human Factors in Robots and Unmanned Systems*, edited by P. Savage-Knepshield and J. Y. C. Chen, 365–74. Vol. 299 *Advances in Intelligent Systems and Computing*. New York: Springer. doi.org/10.1007/978-3-319-41959-6_30
- Mortimer, B. J. P., and Elliott, L. R. 2017. Information Transfer within Human Robot Teams: Multimodal Attention Management in Human-Robot Interaction. In *Proceedings of the*

- 2017 IEEE Conference on Cognitive and Computational Aspects of Situation Management, 1–3. Piscataway, NJ: IEEE. doi.org/10.1109/COGSIMA.2017.7929608
- Oh, J.; Navarro-Serment, L.; Suppe, A.; Stentz, A.; and Herbert, M. 2015a. Inferring Door Locations from a Teammate's Trajectory in Stealth Human-Robot Team Operations. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5315–20, Hamburg, Germany. Piscataway, NJ: IEEE. doi.org/10.1109/IROS.2015.7354127
- Oh, J.; Suppe, A.; Duvallet, F.; Boularias, A.; Vinokurov, J.; Navarro-Serment, L.; Romero, O.; Dean, R.; Lebiere, C.; Hebert, M.; and Stentz, A. 2015b. Toward Mobile Robots Reasoning Like Humans. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1371–9. Palo Alto, CA: AAAI Press.
- Oh, J.; Zhu, M.; Park, S.; Howard, T. M.; Walter, M. R.; Barber, D.; Romero, O.; Suppe, A.; Navarro-Serment, L.; Duvallet, F., et al. 2016. Integrated Intelligence for Human-Robot Teams. In *2016 International Symposium on Experimental Robotics*, edited by D. Kulić, Y. Nakamura, O. Khatib, and G. Venture, 309–22. New York: Springer.
- Osofsky, S.; Phillips, E.; Schuster, D.; and Jentsch, F. G. 2013. A Picture Is Worth a Thousand Mental Models. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 57(1): 1298–302. doi.org/10.1177/1541931213571287
- Paul, R.; Arkin, J.; Aksaray, D.; Roy, N.; and Howard, T. M. 2018. Efficient Grounding of Abstract Spatial Concepts for Natural Language Interaction with Robot Platforms. *International Journal of Robotics Research* 37(10): 1269–99. doi.org/10.1177/0278364918777627
- Paul, R.; Arkin, J.; Roy, N.; and Howard, T. M. 2016. Efficient Grounding of Abstract Spatial Concepts for Natural Language Interaction with Robot Manipulators. In *Proceedings of Robotics: Science and Systems*. doi.org/10.15607/RSS.2016.XII.037
- Paul, R.; Barbu, A.; Felshin, S.; Katz, B.; and Roy, N. 2017. Temporal Grounding Graphs for Language Understanding with Accrued Visual-Linguistic Context. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 4506–14. Palo Alto, CA: AAAI Press. hdoi.org/10.24963/ijcai.2017/629
- Perelman, B. S.; Evans, A. W., III; and Schaefer, K. E. 2017. Mental Model Consensus and Shifts During Navigation System-Assisted Route Planning. *Proceedings of the Human Factors and Ergonomics Society* 61(1): 1183–7. doi.org/10.1177/1541931213601779
- Perelman, B. S.; Evans, A. W., III; and Schaefer, K. E. 2018. Attitudes Toward Risk and Effort Tradeoffs in Human-Robot Heterogeneous Teams. *Proceedings of the Human Factors and Ergonomics Society* 62(1): 1098–102. doi.org/10.1177/1541931218621252
- Phillips, E.; Osofsky, S.; Grove, J.; and Jentsch, F. 2011. From Tools to Teammates: Toward the Development of Appropriate Mental Models for Intelligent Robots. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 55(1): 1491–5. doi.org/10.1177/1071181311551310
- Rabinovich, A.; Vedaldi, A.; Galleguillos, C.; Wiewiora, E.; and Belongie, S. 2007. Objects in Context. In *Proceedings of the 11th International Conference on Computer Vision*. Piscataway, NJ: IEEE. doi:10.1109/ICCV.2007.4408986
- Schaefer, K. E.; Aksaray, D.; Wright, J. L.; and Roy, N. 2019. Challenges with Addressing the Issue of Context within AI and Human-Robot Teaming. In *Computational Context: The Value, Theory and Application of Context with AI*, edited by W. Lawless, R. Mittu, and D. Sofge. Boca Raton, FL: CRC Press.
- Schaefer, K. E.; Brewer, R.; Putney, J.; Mottern, E.; Barghout, J.; and Straub, E. R. 2016. Relinquishing Manual Control: Collaboration Requires the Capability to Understand Robot Intent. In *Proceedings of the International Conference on Collaboration Technologies and Systems*, 359–66. Piscataway, NJ: IEEE. doi.org/10.1109/CTS.2016.0071
- Schaefer, K. E.; Hill, S. G.; and Jentsch, F. G. 2018. Trust in Human-Autonomy Teaming: A Review of Trust Research from the US Army Research Laboratory Robotics Collaborative Technology Alliance. In *Advances in Human Factors in Robots and Unmanned Systems*, edited by P. Savage-Knepshield and J. Chen, 102–14. Vol. 784. Advances in Intelligent Systems and Computing. New York: Springer.
- Schaefer, K. E.; Perelman, B.; Brewer, R. W.; Wright, J. L.; Roy, N.; and Aksaray, D. 2018. Quantifying Human Decision-Making: Implications for Bidirectional Communication in Human-Robot Teams. In *Virtual, Augmented and Mixed Reality: Interaction, Navigation, Visualization, Embodiment, and Simulation*, edited by J. Y. C. Chen and G. Fragomeni, 361–79. Vol. 10909. Lecture Notes in Computer Science. Berlin: Springer. doi.org/10.1007/978-3-319-91581-4_27
- Schaefer, K. E.; Straub, E. R.; Chen, J. Y. C.; Putney, J.; and Evans, A. W., III. 2017. Communicating Intent to Develop Shared Situation Awareness and Engender Trust in Human-Agent Teams. *Cognitive Systems Research* 46: 26–39. doi.org/10.1016/j.cogsys.2017.02.002
- Shiang, S.; German, A.; and Oh, J. 2017. A Generalized Model for Multimodal Perception. Paper presented at the AAAI Fall Symposium, Natural Communication for Human-Robot Collaboration. Arlington, VA, November 9–11.
- Shiang, S.; Rosenthal, S.; Gershman, A.; Carbonell, J.; and Oh, J. 2017. Vision-Language Fusion for Object Recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 4603–10. Palo Alto, CA: AAAI Press.
- Silberman, N.; Hoiem, D.; Kohli, P.; and Fergus, R. 2012. Indoor Segmentation and Support Inference from RGBD Images. In *Computer Vision*, edited by A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, 746–60. Vol. 7576. Lecture Notes in Computer Science. Berlin: Springer. doi.org/10.1007/978-3-642-33715-4_54
- Singhal, A.; Luo, J.; and Zhu, W. 2003. Probabilistic Spatial Context Models for Scene Content Understanding. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 1: 235–41. Piscataway, NJ: IEEE.
- Straub, E. R., and Schaefer, K. E. 2018. It Takes Two to Tango: Automated Vehicles and Human Beings Do the Dance of Driving – Four Social Considerations for Policy. *Transportation Research Part A: Policy and Practice*. doi.org/10.1016/j.tra.2018.03.005
- Tucker, M.; Aksaray, D.; Paul, R.; Stein, G. J.; and Roy, N. 2017. Learning Unknown Groundings for Natural Language Interaction with Mobile Robots. In *Proceedings of the International Symposium on Robotics Research*. New York: Springer.
- Vemula, A.; Muelling, K.; and Oh, J. 2017. Modeling Cooperative Navigation in Dense Human Crowds. In *Proceedings of the IEEE Conference on Robotics and Automation*, 1685–92. Piscataway, NJ: IEEE. doi.org/10.1109/ICRA.2017.7989199
- Vemula, A.; Muelling, K.; and Oh, J. 2018. Social Attention: Modeling Attention in Human Crowds. In *Proceedings of the*



Submissions Due November 15!

AAAI-20 WORKSHOP PROGRAM

Sponsored by the Association for the Advancement of Artificial Intelligence

<ul style="list-style-type: none"> Affective Response Artificial Intelligence for Cyber Security Artificial Intelligence for Education Artificial Intelligence in Team Sports Artificial Intelligence of Things Artificial Intelligence Safety Cloud Intelligence: AI/ML for Efficient and Manageable Cloud Services Deep Learning on Graphs: Methodologies and Applications The Eighth Dialog System Technology Challenge Engineering Dependable and Secure Machine Learning Systems Generalization in Planning 	<ul style="list-style-type: none"> Health Intelligence Intelligent Process Automation — RPA Meets AI Interactive and Conversational Recommendation Systems Knowledge Extraction from Unstructured Data in Financial Services Plan, Activity, and Intent Recognition Privacy-Preserving Artificial Intelligence Reasoning and Learning for Human-Machine Dialogues Reasoning for Complex Question Answering Reinforcement Learning in Games Reproducibility in AI — Future Direction and Reproducibility Challenge Statistical Relational AI
---	--

February 7–8, 2020 ☺ New York, New York USA

aaai.org/Conferences/AAAI-20/ws20

Submission requirements vary for each workshop. Please consult the individual workshop description for complete information about where to submit your paper and a link to the workshop supplementary website, where more detailed information will be available.

IEEE Conference on Robotics and Automation. Piscataway, NJ: IEEE. doi.org/10.1109/ICRA.2018.8460504

Wiltshire, T. J.; Lobato, E. J.; Garcia, D. R.; Fiore, S. M.; Jentsch, F. G.; Huang, W. H.; and Axelrod, B. 2016. Effects of Robotic Social Cues on Interpersonal Attributions and Assessments of Robot Interaction Behaviors. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 59(1): 801–5. doi.org/10.1177/1541931215591245

Wright, J. L.; Chen, J. Y. C.; Barnes, M.; and Hancock, P. A. 2017. Agent Reasoning Transparency: The Influence of Information Level on Automation-Induced Complacency, Technical Report ARL-TR-8044. Aberdeen Proving Ground, MD: US Army Research Laboratory.

Xiong, X., and Huber, D. 2010. Using Context to Create Semantic 3D Models of Indoor Environments. In *Proceedings of the British Machine Vision Conference*. Surrey, UK: BMVA Press. doi.org/10.5244/C.24.45

Zettlemoyer, L. S., and Collins, M. 2009. Learning Context-Dependent Mappings from Sentences to Logical Form. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP* 2: 976–84. New York: ACM. doi.org/10.3115/1690219.1690283

Kristin E. Schaefer is an engineer with the CCDC Army Research Laboratory. She received her PhD in modeling and simulation from the University of Central Florida in 2013. Her research interests lie primarily in the areas of artificial intelligence and modeling and simulation approaches

to enhance the development of bidirectional communication and trust in human-robot teams.

Jean Oh is a systems scientist (research faculty) at the Robotics Institute at Carnegie Mellon University. Her current research is focused on the intersection among vision, language, and planning in robotics. She is passionate about creating persistent robots that can coexist with humans in shared environments, learning to improve themselves over time through continuous training, exploration, and interactions.

Derya Aksaray is an assistant professor in the Aerospace Engineering and Mechanics Department at the University of Minnesota. She received her PhD degree in aerospace engineering from the Georgia Institute of Technology in 2014. She then held post-doctoral researcher positions at Boston University from 2014–2016 and at the Massachusetts Institute of Technology from 2016–2017. Her research interests lie primarily in the areas of control theory, formal methods, and machine learning with applications to autonomous systems, robotics, and human-robot teaming.

Daniel Barber is a research assistant professor at the University of Central Florida. He has extensive experience in the field of robotics, simulation development, training environments, and human state assessment. His current research focus is on human system interaction and training assessment including multimodal communication, user interaction devices, teaming, physiological assessment, and adaptive systems.