ESCAPE with PARTNER: Experimental Suite for Cooperatively Achieved Puzzle Exercises with the Platform Assessing Risk and Trust in Nonexclusively Economic Relationships

Julie L. Marble, Ariel M. Greenberg, Sean M. Kain, Brandon J. Scott, and Mary E. Luongo

AR VR MR

ABSTRACT

Trust is a socio-emotional construct used to explain why one is willing to be vulnerable to the variable and unpredicted actions of another independent actor. Virtual reality (VR) provides an excellent test platform for allowing researchers to assess trust, because it provides a safe environment yet participants can be made to feel vulnerable. In the research described in this article, we developed a VR-based game to assess trust between humans and robots in a collaborative task. This article describes the development of and first experiments with this experimental platform developed to explore humans' trust of bots.

INTRODUCTION

With the increasing capability and autonomy of intelligent systems, trust in or of autonomous systems has moved to the forefront of research on the adoption of technology. When a person does not trust a system, this lack of trust inhibits adoption of the technology. However, few attempts have explored trust separately from the perceived reliability of the technology. Virtual reality (VR) provides an excellent platform to explore these issues in systematic ways.

Previous research has attempted to characterize trust of autonomous systems, whether these systems have been embodied robots, decision aids,[1] or cyber systems. Much of this research has conflated the concept of perceived reliability with trust. While reliability may support the development of trust, other factors, such as workload, situational awareness, and learning, have also been shown to affect subjective measures of trust.[2,3] Trust is an internal state of the one who lends trust and is also dependent on both the agent receiving trust and the context; trust itself cannot be measured directly, although the dispositions that induce it and the behavior consequent to it may be. (Perkins et al.[4] included situational risk as a key factor in the operator's trust of the system.) We hypothesize that behaviors indicative of trust can be objectively measured and defined in the context of interest.

Several factors have been proposed as drivers of humans' trust in autonomous systems. One is transparency, which has been proposed as a predicate for trust, but it is not clear that understanding the reasoning of a teammate is necessary for trust. For example, Chen et al.[5] developed a model, Situation Awareness-based Agent Transparency (SAT), that explores requirements for the human's awareness of an agent's (1) current actions and plans, (2) reasoning process, and (3) outcome predictions. However, within humans the precise internal cognitive state of a teammate is hidden; even when both teammates are human, as one cannot truly know what the other is thinking.

The study of these phenomena requires an environment where the situational risk can be manipulated, workload can be assessed, and teammates' actions can be recorded. VR allows for manipulation of perceived risk, with little actual risk to the performer. VR is also a unique environment in that the identity of a collaborator is not immediately apparent; a person performing a task may not be able to immediately determine whether their collaborator is a human or a bot if key precautions (such as limited communications) are taken. Additionally, VR allows for creation of an artificial environment in which the experimenter can create the rules and gain more control of what occurs in the environment than in the real world. Finally, recent advances in VR technologies have simplified collection of human performance data. In the commercial world, we see the second wave of commodity VR, with VR systems now costing significantly less than when they first emerged. Many devices are now completely stand-alone, and controllers are high quality, offering six degrees of freedom and accuracy measured in millimeters. Systems that became available in the last year do not require controllers to track users' hands and have built-in eye tracking and voice recognition, allowing for assessment of the user's workload and stress.

ESCAPE with PARTNER (Experimental Suite for Cooperatively Achieved Puzzle Exercises with the Platform Assessing Risk and Trust in Nonexclusively Economic Relationships) was developed to provide an environment for objectively assessing trust behavior and decision-making between humans and autonomous systems. This article discusses its development and the preliminary results from experimentation, using the term *bots* to refer to virtual agents,[6] instantiated in the game environment but functioning as an independent teammate.

We are specifically interested in the types of collaboration that we predict will soon become commonplace: agents that take initiative at critical times (e.g., when the human is overloaded or too slow to react or is potentially in danger) to achieve shared goals. As in human–human teams, this would entail that roles of leader–follower become fluid, such that the robot take over the implementation of a plan or implement prespecified algorithms to achieve the objective. This concept expands beyond the highest level of self-governing autonomy proposed by Allen, Guinn, and Horvtz in 1999.[7] In this case, the robot can act as a collaborative peer, proposing alternative actions or taking action when the other teammate is unable to act (see Marble et al.[8] for a discussion of peer–peer interaction in human–autonomous systems).

PREVIOUS RESEARCH

Lee and See[9] defined trust as "the attitude that an agent will help achieve an individual's goals in a situation characterized by uncertainty and vulnerability." Rousseau et al.[10] defined trust as "a psychological state comprising the intention to accept vulnerability based upon positive expectations of the intentions or behavior of another." Finally, Hoff and Bashir[2] reviewed recent empirical research on factors that influence trust in automation to present a three-layered trust model that synthesizes existing knowledge. Their model blends preexisting knowledge, attitudes, and experience with dynamic understanding of the system that is built up through interaction; it implies that trust is essentially a reliance strategy that changes dynamically with the task and context.

Because poorly calibrated trust can result in catastrophic failures due to operator misuse or disuse of the automated systems, there has been a recent emphasis on trust in autonomy. Research on human–automation interaction has explored how people use, misuse, or disuse automation.[11,12] Previous research has often focused on system reliability.[13] However, this research conflates perceived reliability with trust in that people will demonstrate trust for partners that are somewhat unpredictable or not fully capable.[13,14]

Others[15] have focused on task and team characteristics, including task complexity and relationship equity.[16] Much of this research was synthesized by Hancock et al.[17] This meta-analysis underlies the development of Schaefer Trust Index questionnaire.[18] However, subjective questionnaires such as the Schaefer Trust Index or the Muir trust questionnaire[19] provide only the subject's perception of their trust of the robot. However, perception does not match honest measures of trust. These tools do not give insight into behaviors indicative of trust during use. In addition, few of these studies and tools attempt to define trust such that it can be modeled and tested.

We argue that trust is a socio-emotional construct to explain why trustees would choose to make themselves vulnerable to potentially faulty, unpredictable systems in dynamic environments. One implication of the definitions posed by Lee and See,[9] Rousseau et al.,[10] and Hoff and Bashir[2] is that many factors drive the development of trust between teammates: vulnerability to the actions of another, variability and unpredictability of their actions, uncertainty of context and events, and beneficence and capability or perceived capability of the trusted.

Our long-term goal is to understand how different factors impact trust and control allocation and, based on this information, to build a model that can estimate an operator's current level of trust so that the system can adjust in ways to move the operator's trust toward a more accurately calibrated position, to prevent inappropriate usage of the autonomy. To that end, we developed ESCAPE with PARTNER to provide a platform to assess the development of humans' trust of autonomous systems.

PARTNER DEVELOPMENT

PARTNER is a VR-based research platform developed in Unity and run on Oculus Rift–enabled systems.

ESCAPE refers to the four virtual escape room puzzles within PARTNER. In ESCAPE with PARTNER, the test participant teams with either a human confederate or an agent to identify the solution that will allow the two to escape the room. To solve the puzzle, teammates must actuate levers, platforms, and buttons; use "tractor beams" and "repulsor beams"; or virtually lift, carry, and push obstacles to create a route by which both teammates can reach the exit. The exit door activates only when both teammates are present.

At each decision point, one teammate must perform a risky action (such as jump from a high vantage point to a platform), while the second teammate takes action to enable the first teammate. This second action also comes with implicit risk. For example, to obtain a tool the first teammate must jump from platform to platform at risk of falling, while the second teammate enables the teammate's action by standing on a lever. The second teammate is also at risk because standing on the lever ensures that the platforms holding the first teammate remain, but the floor surrounding the second teammate slowly fills with deadly "acid," which will "kill" their avatar. The player standing on the lever could easily jump to safety, but doing so would jeopardize or "kill" their teammate. Avatar death in the game is typically due to a simulated fast fall from height in VR. While there is no risk to the participant, the height exposure and fall is mildly unpleasant. Examples of the PARTNER environment are shown in Figures 1 and 2. VR emulates vulnerability to teammate actions, which is considered critical to the development of trust between teammates.

In application, the test participant could team with either a human (a confederate on the experiment team) or an agent (the bot). The bot is a finite-state machine scripted to play the game with complete information on the game puzzle. Each step in the puzzle is indexed and associated with the relevant game elements to complete the step. Combined interaction events with these game elements would either advance or regress the puzzle step. Information on the association of game elements with a puzzle step as well as on the current puzzle step is not available to human players. The autonomous agent would use this information in combination with the teammate's current position to determine the transition to the next appropriate state. Each state is a heuristic routine connected via defined transition rules. The combination of a human teammate moving in three-dimensional space with multistep puzzles created a large state space that a finite-state machine would likely be unable to accommodate; while the finite-state machine was not comprehensive, it was competent over the majority of the space. To mitigate this risk, a human could override the autonomous player in real time, "nudging" the player into an appropriate state. When needed, the study team confederate who also played as the human teammate took this action.

Every game with a human teammate was played by the same confederate. The experiment confederate had extensive gaming experience. He controlled his avatar via standard gaming keyboard conventions. The confederate, when playing as the human teammate, was scripted to attempt to play each game as similarly as possible, including allowing the test participant time to attempt to identify the solution or explore the puzzle room. While the test participant played in VR, the confederate played on a flat screen using keyboard controls in a room separate from the test participant.

Teammates were not able to communicate with each other verbally or via text messaging. This constraint was enforced to disguise the identity of the human and bot teammates. Rather, teammates (including the bot) were equipped with pointer beams, sprites, and timers common to several games to attract the teammate's attention to objects and coordinate actions, as illustrated in Figures 1 and 2.

Teams had 7 minutes, a time frame selected to avoid potential VR sickness, to identify and implement the solution to each puzzle. Participants solved four different puzzles, two while teamed with a human and two while teamed with a bot. Each participant performed the puzzles in the same order, although the order of playing with a human or bot teammate was randomized for the first two and then the second two puzzles. After completing (or failing to complete) the puzzle, participants were placed in a virtual waiting room where they were asked a series

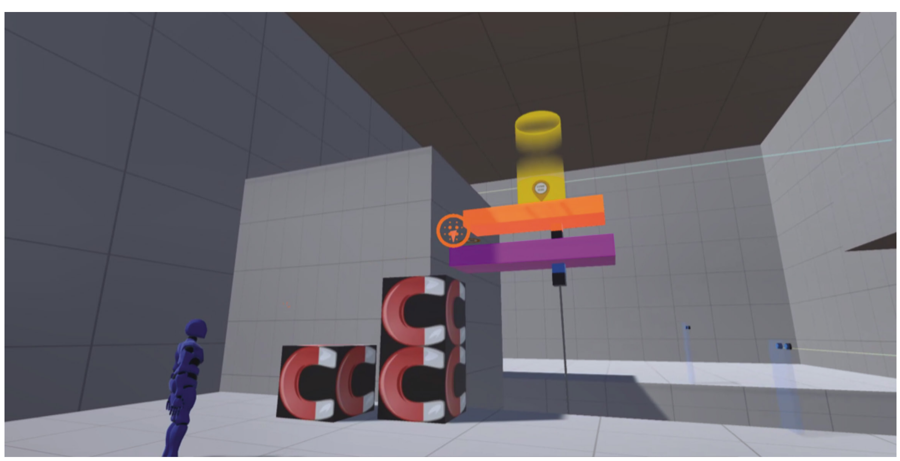

Figure 1. Still from ESCAPE with PARTNER. The image shows the player avatar and communication sprite.

Figure 2. Another still from ESCAPE with PARTNER. This image shows the teammate avatar gesturing to the control panel, showing a chasm to be navigated.

of questions about their performance, their teammate's performance, and the team's performance on the previous puzzle. To maintain players' immersion, questions were presented virtually, and players selected answers using the Oculus Rift hand controllers. Responses were collected using Qualtrics survey software.

Our goal was to attempt to identify a set of objective measures that would be indicative of trust between teammates. In this article we discuss the results around the most basic of these measures—performance assessment and teammate selection. If a person trusts their teammate, we hypothesized that they would choose that teammate over a less trusted teammate during higher-stakes events (such as winning the championship). We also wanted to explore the relationships between perceived capability and trust as well as actual performance and trust.

METHODS

Participants

Thirty-one participants were recruited from the population of APL staff members who self-reported enjoyment of video games and experience playing the game *Portal 2*. Data from 3 participants were lost or removed because of problems with the software. The data from the remaining 28 participants were used in the study. Participants ranged in age from 18 to 45; all reported normal or corrected-to-normal vision. Participants reported having previously used a VR system with no adverse effects.

Equipment and Setup

ESCAPE with PARTNER was played on an Oculus Rift headset with handheld controllers. Participants played in individual test rooms. The floor of each test room was demarcated with a prickly tape strip to prevent participants from bumping or walking into walls. Similarly, streamers were hung from the ceiling to mark boundaries of movement. The confederate, in the test control room, played on a standard laptop monitor using keyboard controls. To ensure their safety, participants were observed from the control room via video and audio feeds from the test rooms; however, the data were not recorded.

Procedures

At the start of the experiment, participants were given an overview of the experiment. They were told that they would be asked to solve a series of four escape room puzzles with the assistance of a human teammate or a bot teammate. It was implied that the bot and human were equally skilled at solving the puzzles. (The human confederate's actions were partially scripted to prevent his skill and speed from increasing far beyond the level of the bot.) Participants were told that they would have 7 minutes to solve each puzzle and that they could display a timer on the screen by activating a button on the controller.

Participants were fitted with the Oculus headset and handheld controllers. They were told that they could end the test at any time. The VR environment was started, and participants began the test in a practice room, which allowed them to learn how to use the handheld controllers, jump from object to object, and manipulate objects by directly contacting them or by using the tractor beams. When participants were comfortable with the environments and had been exposed to all the object types they would experience in the game, they moved on to the game itself.

Participants were asked in the VR environment which teammate they would like to play the first game with, and their response was recorded. Participants were then assigned a teammate independent of their preference. After each puzzle, participants were asked a series of questions about the previous game, such as how well they felt they played, how well they felt their teammate played, and how well they felt the team played. After answering these questions, participants entered the next puzzle, which they played with the other teammate type—that is, if they first played with a human, they next played with a bot. They were told that the next teammate was different from the first. As a result, participants played each round in one of the following conditions, where the boldface word indicates the actual identity and the italic word is the assumed identity:

Teaming conditions (**actual**/*assumed*)	Actual identity: **human**	Actual identity: **bot**
Assumed identity: *human*	**Human**/*human*	**Bot**/*human*
Assumed identity: *bot*	**Human**/*bot*	**Bot**/*bot*

After the first two puzzles were completed, participants were told that they had played one round each with the bot and the human, and would now play two more puzzles. They were again asked whether they would like to play with a human or a bot. They were then assigned a teammate independent of their indicated preference. After each puzzle, they were again asked to rate their own performance, their teammate's performance, and the team's performance in the game.

Questionnaire

When all four puzzles were completed, participants were told that they had "done very well" and were currently ranked in first place. They were asked whether, in the event of a tie for first place, they would come back to play a final round, and if so, with which teammate they would want to partner. Finally, participants completed a last questionnaire asking them in which puzzles they had teamed with a human versus a bot, and at which puzzles they thought they had performed the best (in which their avatars had died the fewest times and they had escaped most quickly).

After completing all puzzles and questionnaires, participants were debriefed about the purpose of the game and in which puzzles they had teamed with a human and with a bot. It was strongly recommended that participants not drive for at least 30 minutes after the game.

RESULTS

Participants were able to complete the puzzles in the time allotted, except puzzle 3, which was surprisingly more difficult to solve than the other three puzzles. While the majority of participants came very close and could likely have escaped had they had an additional 30 seconds, only one participant was able to escape. Based on the performance of the participants, we realized that the first two puzzles were significantly easier than the second two puzzles. Given this inconsistency in task difficulty resulting in undersampling across the relatively small number of participants, the following data are to be considered merely preliminary.

Figure 3 shows participants' choice of a human teammate versus a bot teammate. At the start of the puzzles, when asked which teammate they wished to start with, half of the participants desired to play first with a bot and half desired to play with a human. After the first two puzzles, there was a shift toward a preference to play with a human, with roughly 68% of participants indicating a preference for a human. When asked if they would come back to play a final round in the event of a tiebreaker, 100% said that they would. Additionally, roughly 90% expressed a preference for playing that final tiebreaker round with a human.

In general, however, participants were not able to accurately distinguish between playing with a human and playing with a bot (Figure 4). It is interesting to note that several participants indicated that they played at least three puzzles with a bot, while others indicated that

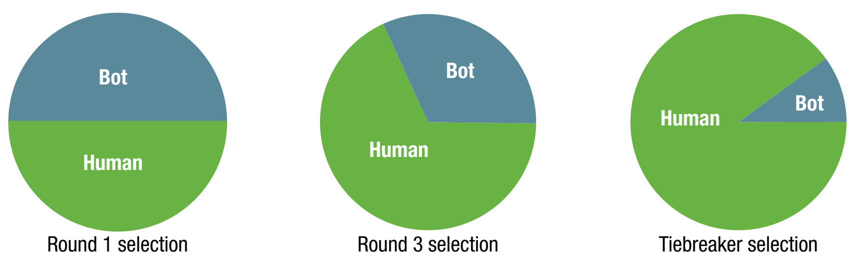

Figure 3. Participants' teammate selection (human vs. bot). Charts show participants' choices before the first, third, and hypothetical tiebreaker rounds. Blue denotes a choice of bot; green denotes a choice of human.

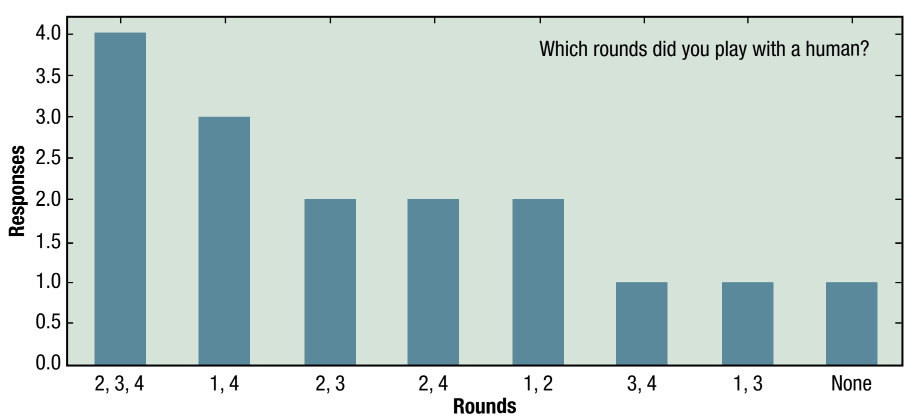

Figure 4. Responses to question about which puzzles were played with a human (vs. a bot). All participants played two puzzles with a human and two with a bot.

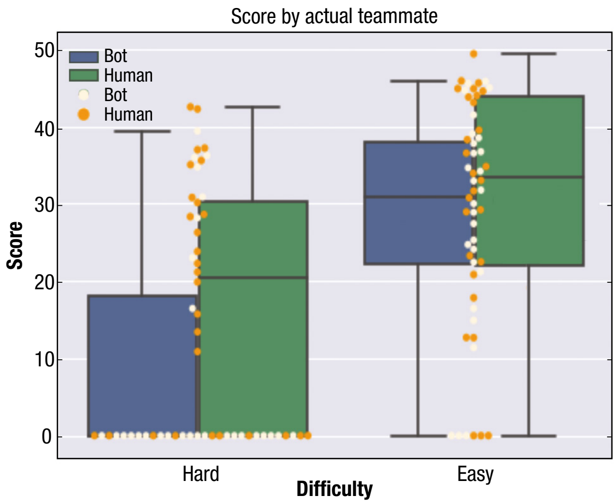

Figure 5. Distribution of scores as a function of true teammate. The scores for the more difficult puzzles are shown with the left bars, and those for the easier puzzles are shown with the right bars.

they did not play any puzzles with a bot. All participants played two puzzles with a human and two puzzles with a bot.

We then explored performance as a function of teammate across the puzzles. Performance was a hybrid score based on the time to complete the room and the number of times the participant avatar or bot avatar died. For the first two puzzles, as shown in Figure 5, the score was slightly higher when the teammate was human than when it was a bot. As shown on the left side of Figure 5, scores were noticeably higher on the first two puzzles than the second two. (We cannot assess this for significance because of the difficulty of puzzle 3, which only one participant escaped.)

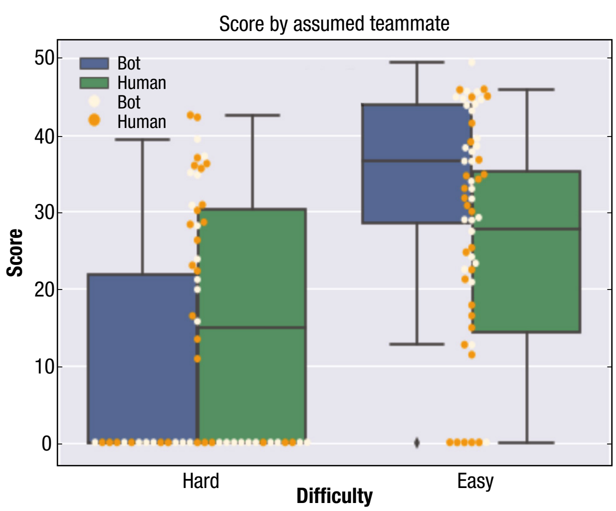

Figure 6. Distribution of scores as a function of assumed teammate. For the first two easier puzzles, scores were higher when the participant believed they were playing with a bot. This relationship may not be true for the third and fourth more difficult puzzles, but it is difficult to say given the lack of data on puzzle 3.

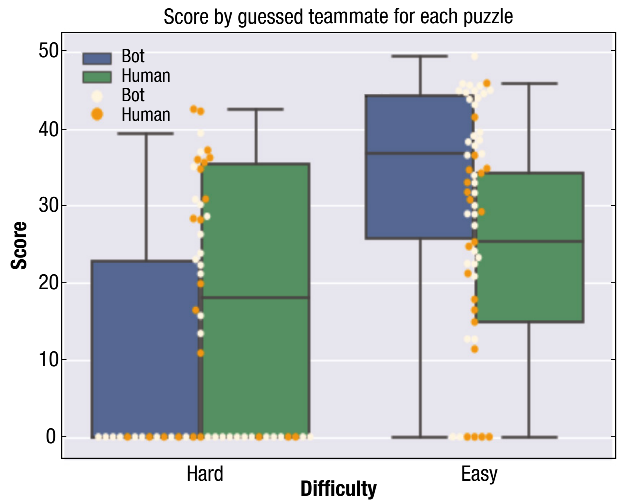

Figure 7. Distribution of scores as a function of participant's guess at teammate's identity for each puzzle. Results were very similar to the scores as a function of assumed teammate.

We also explored participant score as a function of assumed teammate. This is shown in Figure 6. For the first two easier puzzles, scores were higher when the participant believed they were playing with a bot than when they believed they were playing with a human. This relationship may not be true for the more difficult puzzles (puzzles 3 and 4), but it is difficult to say given the lack of data on puzzle 3.

We also explored score as a function of participants' guesses of which puzzle had been played with which partner (Figure 7). Intriguingly, the results were very similar.

To better understand these findings, we explored responses to the question of how well the teammate performed after each puzzle. When the participant assumed the teammate was a bot, the average score on a scale of 1–7, where 1 was bad and 7 was extremely good, was 4.14 out of 7, or moderate, and the modal score was 4. When the participant assumed the teammate was a human, the average performance score was 4.8, but the modal score was 6 out of 7.

CONCLUSIONS

Several conclusions can be tentatively drawn from these results. Trust in autonomous systems is often conflated with capability. From an engineering perspective, there is a belief that if the system is capable of performing the action, the human operator will trust it. In our research, this belief was not supported. While the bot was able to perform as well as the human, people developed a preference for the human teammate, even when they did not know on which puzzles they had teamed with a human to solve.

In addition, this preference developed after the first two puzzles. This is intriguing given the finding that performance on the first two puzzles was better (in terms of escape speed and fewer deaths of avatars) when the teammate was assumed to be a bot. When combined with the finding that participants were not good at distinguishing a bot from a human, this suggests that other factors are in play in the development of trust. This could be indicative of a predisposition to trust or simply a desire to press all the buttons and try all the game modes. In addition, while participants stated that the puzzles on which they performed best were those in which they thought they were paired with a bot (for the easy puzzles at least), they rated the bot as performing worse than the human teammate. While this finding is intriguing, it is difficult to expand on it given the difficulty of puzzle 3.

A number of explanations for these findings cannot be assessed here because of limitations of the data. Follow-on work will seek to rectify the deficiencies in the data and clarify the source of these findings.

ACKNOWLEDGMENTS: We thank Evan Fuller for his contributions to analysis.

REFERENCES

[1]M. S. Cohen, R. Parasuraman, and J. T. Freeman, "Trust in decision aids: A model and its training implications," in *Proc. Command and Control Res. Technol. Symp.*, 1998, pp. 1–37.

[2]K. A. Hoff and M. Bashir, "Trust in automation: Integrating empirical evidence on factors that influence trust," *Human Factors*, vol. 57, no. 3, pp. 407–434, 2015, https://doi.org/10.1177/0018720814547570.

[3]J. A. Colquitt, B. A. Scott, and J. A. LePine, "Trust, trustworthiness, and trust propensity: A meta-analytic test of their unique relationships with risk taking and job performance," *J. Appl. Psychol.*, vol. 92, no. 4, pp. 909–927, 2007, https://doi.org/10.1037/0021-9010.92.4.909.

[4]L. Perkins, J. E Miller, A. Hashemi, and G. Burns, "Designing for human-centered systems: Situational risk as a factor of trust in automation," in *Proc. Human Factors and Ergonomics Soc. Annu. Mtg.*, vol. 54, no. 25, pp. 2130–2134, 2010, https://doi.org/10.1177/154193121005402502.

[5]J. Y. Chen, S. G. Lakhmani, K. Stowers, A. R. Selkowitz, J. L. Wright, and M. Barnes, "Situation awareness-based agent transparency and human-autonomy teaming effectiveness," *Theor. Issues Ergonomics Sci.*, vol. 19, no. 3, pp. 259–282, 2018, https://doi.org/10.1080/1463922X.2017.1315750.

[6]Defense Science Board, *Defense Science Board Summer Study on Autonomy*, Washington, DC: Department of Defense, 2016.

[7]J. E. Allen, C. I. Guinn, and E. Horvtz, "Mixed-initiative interaction," *IEEE Intell. Syst. Appl.*, vol. 14, no. 5, pp. 14–23, 1999, https://doi.org/10.1109/5254.796083.

[8]J. L. Marble, D. J. Bruemmer, D. A Few, and D. D. Dudenhoeffer, "Evaluation of supervisory vs. peer-peer interaction with human-robot teams," in *Proc. 37th Annu. Hawaii Int. Conf. Syst. Sci.*, 2004, pp. 1–9, https://doi.org/10.1109/HICSS.2004.1265326.

[9]J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *Human Factors*, vol. 46, no. 1, pp. 50–80, 2004, https://doi.org/10.1518/hfes.46.1.50_30392.

[10]D. Rousseau, S. Sitkin, R. Burt, and C. Camerer, "Not so different after all: A cross-discipline view of trust," *Acad. Manage. Rev.*, vol. 23, no. 3, pp. 393–404, 1998, https://doi.org/10.5465/amr.1998.926617.

[11]R. Parasuraman and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," *Human Factors*, vol. 39, no. 2, pp. 230–253, 1997, https://doi.org/10.1518/001872097778543886.

[12]S. Dixon and C. Wickens, "Automation reliability in unmanned aerial vehicle control: A reliance-compliance model of automation dependence in high workload," *Human Factors*, vol. 48, no. 3, pp. 474–486, 2006, https://doi.org/10.1518/001872006778606822.

[13]C. D. Wickens and X. Xu, "How does automation reliability influence workload," in *Proc. 1st Annu. Robot. Consortium*, US Army Research Laboratory Collaborative Technology Alliance, 2003.

[14]M. Desai, M Medvedev, M. Vazquez, S. McSheehy, S. Gadea-Omelchenko, et al., "Effects of changing reliability on trust of robot systems," in *Proc. 7th Annu. ACM/IEEE Int. Conf. on Human-Robot Interact.*, 2012, pp. 73–80, https://doi.org/10.1145/2157689.2157702.

[15]J. D. Lee and N. Moray, "Trust, self-confidence, and operators' adaptation to automation," *Int. J. Human Comput. Stud.*, vol. 40, no. 1, pp. 153–184, 1994, https://doi.org/10.1006/ijhc.1994.1007.

[16]E. J. de Visser, M. M. Peeters, M. F. Jung, S. Kohn, T. H. Shaw, et al., "Towards a theory of longitudinal trust calibration in human–robot teams," *Int. J. Social Robot.*, vol., 12, pp. 459–478, 2020, https://doi.org/10.1007/s12369-019-00596-x.

[17]P. Hancock, D. Billings, K. Schaefer, J. Y. C. Chen, E. de Visser, and R. Parasuraman, "A meta-analysis of factors affecting trust in human-robot interaction," *Human Factors*, vol. 53, no. 5, pp, 517–527. 2011, https://doi.org/10.1177/0018720811417254.

[18]K. E. Schaefer, S. G. Hill, and F. G. Jentsch, "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*, J. Chen (ed.), Cham: Springer, pp. 102–114, 2019, https://doi.org/10.1007/978-3-319-94346-6_10.

[19]B. M. Muir, "Trust between humans and machines, and the design of decision aids," *Int. J. Man-Machine Stud.*, vol. 27, nos. 5–6, pp. 527–539, 1987, https://doi.org/10.1016/S0020-7373(87)80013-5.

Julie L. Marble, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Julie L. Marble, internationally recognized for her work in human reliability analysis, leads research in human–machine teaming in APL's Research and Exploratory Development Department. She has a PhD in human factors/cognitive psychology from Purdue University. Before joining APL in 2015, she was a program officer at the Office of Naval Research, with programs focused on cognitively compatible human–machine teaming and human factors considerations of cybersecurity. She started work in robotics at the Idaho National Laboratory. She was also president of Sentient Corporation and a program manager in research at the US Nuclear Regulatory Commission. Julie is coauthor of SPAR-H, the most commonly used method of human reliability analysis in the United States. She contributed a chapter called "Human Behaviour, Cognition and Decision Making" as part of the NATO Human Factors and Medicine Research Task Group's work in 2014, and she continues working with this task group today. Her email address is julie.marble@jhuapl.edu.

Ariel M. Greenberg, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Ariel M. Greenberg is a senior research scientist in APL's Intelligent Systems Center; assistant editor-in-chief of the *Johns Hopkins APL Technical Digest*; steering committee co-chair of the International Conference on Social Computing, Behavioral-Cultural Modeling & Prediction and Behavior Representation in Modeling and Simulation (http://sbp-brims.org/); and an instructor in APL's Strategic Education Program. He received degrees in biology and electrical engineering from University of Maryland, College Park. In the spirit of cells-to-societies transdisciplinary research, Ariel's work addresses topics in psychophysiology and computational social science, toward applications in defense, intelligence, and health. His e-mail address is ariel.greenberg@jhuapl.edu.

Sean M. Kain, Asymmetric Operations Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Sean M. Kain is a software engineer in APL's Asymmetric Operations Sector. He has a BS in geography and an MPS in geographic information systems, both from the University of Maryland, Baltimore County. Sean has a geoinformatics-focused skill set, and his research interests include biomimetic robotics, agent-based simulation, artificial general intelligence, and machine learning. His email address is sean.kain@jhuapl.edu.

Brandon J. Scott, Asymmetric Operations Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Brandon J. Scott is a software engineer in APL's Asymmetric Operations Sector. He has a BS in engineering mathematics from Saint Augustine University and an MS in computer science from North Carolina Agricultural and Technical State University. Brandon's background is in information assurance, virtual/mixed reality, and software development. His email address is brandon.scott@jhuapl.edu.

Mary E. Luongo, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Mary E. Luongo is a biomedical engineer in APL's Research and Exploratory Development Department. She has a BS in biomedical engineering from the Catholic University of America and an MSE in systems engineering from Johns Hopkins University. Working in APL's Intelligent Systems Center, Mary's focus is on human subjects research in human performance, neuroscience, biometrics, human–machine teaming, and biomechanics. Her work includes project planning, experimental design, institutional review board protocols, test development, tissue preparation and instrumentation, human subjects testing, intake, analysis, data collection, and test execution. Her email address is mary.luongo@jhuapl.edu.