Check for updates

HUMAN, AI, ROBOT TEAMING AND THE FUTURE OF WORK: BARRIERS AND OPPORTUNITIES FOR ADVANCEMENT

Co-Moderators:
Erin K. Chiou, Human Systems Engineering, Arizona State University
Eric Holder, U.S. Army Combat Capabilities Development Command Army Research Laboratory Human Research and Engineering Directorate (CCDC ARL-HRED)

Panelists:
Igor Dolgov, Elevate, Uber Technologies, Inc.
Kaleb McDowell, CCDC ARL-HRED
Lance Menthe, RAND Corporation
Rod D. Roscoe, Human Systems Engineering, Arizona State University
Shivam Zaveri, Human and Social Dimensions of Science and Technology, Arizona State University

Copyright 2020 by Human Factors and Ergonomics Society. All rights reserved. 10.1177/1071181320641018

Global investments in artificial intelligence (AI) and robotics are on the rise, with the results to impact global economies, security, safety, and human well-being. The most heralded advances in this space are more often about the technologies that are capable of disrupting business-as-usual than they are about innovation that advances or supports a global workforce. The *Future of Work at the Human-Technology Frontier* is one of NSF's 10 Big Ideas for research advancement. This panel discussion focuses on the barriers and opportunities for a future of human and AI/robot teaming, with people at the center of complex systems that provide social, ethical, and economic value.

INTRODUCTION

A common misunderstanding in the design of new technologies is to view technology as simply a *tool* to fulfill a function, when they are more akin to *processes* that impact workers and other important outcomes of an organization. A systems-level approach to technology design must include the people, process workflows, and work environment factors to realize both human and technology potential and to avoid known mistakes of the past. Accelerating this effort in the face of current technology-centered investments in AI and robotics will take immediate and meaningful engagement across groups, and a unified effort that spans disciplines and work domains.

Applying the concept of teaming to human-technology systems is not just a "magic concept" (Pollitt & Hupe, 2011) that would be difficult to realize (Groom & Nass, 2007), but rather it is a way to build upon past work on human-automation interaction and system design through the lens of teaming and teamwork. This framing allows human factors scientists and practitioners to work in parallel with ongoing advances in AI and robotics, to move beyond the levels of automation/autonomy categories (Parasuraman, Sheridan, & Wickens, 2000), toward addressing more complex, networked, and adaptive relationships with automation, as automation becomes increasingly distributed, agentic, and autonomous.

Much is known about the effectiveness of human teams, as well as the control of synthetic teams of robotic or software agents. For example, we generally know what a good human teammate looks like, and we also know what good team problem-solving looks like (Paris, Salas, & Cannon-Bowers, 2000; Salas, Burke, & Janis, 2000; Salas, Dickinson, Converse, & Tannenbaum, 1992; Salas, Wilson, Murphy, King, & Salisbury, 2008). However, this understanding of human teams is almost always descriptive, with a focus on the characteristics that contribute to good team performance (e.g., Duhigg, 2016). While useful, such description can be difficult to translate into control system design for robots or AI, which depend on quantifiable inputs and outputs.

There are also well-established formal controls or optimization frameworks for designing and engineering robotic or software agent teams that can negotiate and achieve shared goals (Sugihara & Suzuki, 1994; Sycara & Zeng, 1996). These formalized frameworks are suited to developing machines yet they tend to be for limited, well-defined applications that are likely to break when introduced to the reality of human systems. Even if such frameworks acknowledge human interactions, the human in the system is often treated as a black box endpoint, with its internal processes excluded (Fong, 2001; Michaud et al., 2010). Such approaches prevalent in certain engineering disciplines, often ignore the coordination and negotiation of personal, cultural, political, technological, and organizational processes that occur in the dynamic and complex world of humans – and the role that technology can play in facilitating or impeding those processes. There is a need to invest more in efforts that move beyond testing applications in rigidly-controlled environments, in which mathematical and computational modeling approaches thrive while obfuscating key human elements of the larger system, and instead bring focus to understanding people, AI, and robots working together in naturalistic settings.

OBJECTIVE

In line with the Organisation for Economic Co-operation and Development's Principles on Artificial Intelligence, adopted by 42 countries (OECD, 2019), a newly formed HFES Technical Group focused on Human, AI, Robot

Copyright 2020 by Human Factors and Ergonomics Society. All rights reserved. 10.1177/1071181320641018

Teaming envisions a future in which people and AI-imbued agents or robots are thoughtfully integrated to improve care, quality, flexibility, productivity, safety, and security. The assembled panelists will provide an overview of recent work related to human, AI, robot teaming for the future of work in transportation, military, education, and manufacturing. Key issues related to opportunities or barriers for advancing the science and engineering of human and AI/robot teams will be addressed.

PANELIST STATEMENTS

Human-Machine Teaming in Urban Air Mobility

Igor Dolgov, Ph.D.; Lead Human Factors Engineer at Uber Elevate. The term Urban Air Mobility (UAM) describes a vision for a new era of metropolitan transportation systems that will move people and goods through the air. A key advantage of this type of transport is sustainability. UAM aircraft will be electric and will not pollute the air. Additionally, advances in batteries also enable these types of aircraft to be more cost-effective than their fuel-powered counterparts. While many technologies have matured in recent decades, making UAM a reality is not without its difficulties (Thipphavong et al., 2018; Vascik, Hansman, & Dunn, 2018).

The challenges facing UAM proponents are manifold and include the development of new vehicles, batteries, infrastructure, communication technologies and networks, airspace services and networks, multi-modal transportation services, and regulations. Moreover, these technologies will need to be integrated, interoperable, and certified to be safe (Holden & Goel, 2016).

One avenue for handling this complexity is through automated, autonomous, robotic, and artificially intelligent tools and agents. These technologies can help reduce human workload, stress, and fatigue while improving situation awareness and decision making (e.g., Dolgov et al., 2017). While advantageous, the use of these technologies can be particularly challenging in life-critical applications (Brynjolfsson & Mitchell, 2017). Such contexts require human operators to maintain meaningful control of autonomous, robotic, or artificially intelligent systems, which has a number of design implications and challenges for human-machine teaming (Santoni de Sio & Van den Hoven, 2018).

Namely, these types of systems will need to be human-interpretable, entailing that they are transparent, explainable, and accountable (Roundtree, Goodrich, & Adams, 2019; Wachter, Mittelstadt, & Floridi, 2017). Meeting these criteria requires nuanced solutions. One major constraint is that many types of artificially intelligent and autonomous systems are neither readily explainable nor accountable (Samek, Wiegand, & Muller, 2017). For instance, while deep learning models exhibit excellent learning performance they suffer from a lack of transparency and explainability. This makes them challenging to implement in life-critical applications and requires the use of additional tools that can help alleviate these issues (Gunning, 2017).

Furthermore, while transparency is traditionally seen as unequivocally positive for system safety and performance, it can sometimes be a double-edged sword. When transparency reveals accuracy and reliability on the part of an automated aid, humans tend to trust it more and performance improves. But, trust wanes and performance suffers when transparency reveals errors or inadequacies (e.g., Kaltenbach & Dolgov, 2017).

In sum, artificial intelligence, automation, autonomy, and robotics can help enable the UAM vision but great care must be taken in designing systems and processes that support complex human-machine teaming and human-system integration.

Igor Dolgov is currently the Lead HF Engineer for Uber Elevate and also Chair of the Aerospace Systems Technical Group for the Human Factors and Ergonomics Society. He was previously a tenured associate professor of human factors/engineering psychology at New Mexico State University, where he led the Perception, Action, and Cognition, in Mediated, Artificial, and Natural Environments (PACMANe/“pacman”) laboratory. He earned a B.S.E. in Computer Science from Princeton University and a Ph.D. in Psychology-Arts, Media, & Engineering from Arizona State University via NSF’s Integrative Graduate Education and Research Traineeship Program.

A Framework for Integrating AI into Human Society: An Army Perspective.

Kaleb McDowell, Ph.D.; Chief Scientist at U.S. Army Combat Capabilities Development Command Army Research Laboratory, Human Research and Engineering Directorate (CCDC ARL HRED)

We present a framework meant to facilitate ideation about strategies to address the challenges of integrating AI into human society. Our ultimate aim is to expand discussion around concepts that we see as essential for the formation of human-AI partnerships, moving beyond the typical oversimplifications made when viewing this problem space as a monolithic entity, wherein a single, generalizable type of interaction is of concern. Example oversimplifications include: designating areas where AI replacing humans will either be inevitable or impossible, assuming that changing an AI's behavior can only be accomplished by scientists and engineers, and approaching human-AI integration as a simple task allocation problem.

We argue that the fundamental nature of future human AI partnership is task relative, and it thus depends critically on the complexity of the problem to be solved, the certainty of information about the problem, and the time available to enact a solution. Our framework proposes that viewing interactions between AI and humans in the context of complexity, certainty, and time clearly emphasizes the reality that there is not one sort of interaction that must be considered and supported. Rather, there are myriad ways that humans and AIs may cooperate and interact. Critically, the framework provides a mechanism for researchers and engineers to consider their problem space and indicates how different types of human-AI interactions can be employed together for increased robustness and reliability over the lifespan of the technology. This is important as we consider that continually advancing AIs may be expected to evolve away from being mere tools for human application and will instead advance towards integration that is better described as novel, team-like partnerships (DeCostanza et al., 2018)

Kaleb McDowell is currently the Chief Scientist of the CCDC ARL HRED. Since joining ARL, Dr. McDowell has developed a strong record of publication and impact within government, industry, and academic research and development communities and he has led several major research and development programs focused on neuroscience and neuroengineering, indirect vision systems, vehicle mobility, and human-agent teaming; receiving Army Research and Development Achievement awards in 2007 and 2009; and ARL Awards for Leadership and Engineering in 2011 and 2013.

Integrating AI into Military Intelligence Analysis

Lance Menthe, Ph.D., Senior Physical Scientist at RAND Corporation. Military intelligence analysis is a massive team effort. AI can improve the processing, exploitation, and dissemination of intelligence in many ways, but integrating AI into the analysis workflow requires careful consideration of the larger collaborative process to realize the promised benefits and avoid creating new bottlenecks.

We present a framework for categorizing levels of automation and synthesis in the intelligence analysis process to highlight how AI can be helpful. We also present a set of evaluation criteria to indicate what kind of improvements should be sought. Finally, we outline a method of mapping the data flow in team-based analysis processes to identify where these capabilities will fit, and also where less sophisticated techniques—e.g. small scripts, software linkages, and workflow improvements—are required. Without this kind of consideration, the most expensive investments in AI may ultimately end up gathering dust in the corner.

We find that many intelligence analysis tasks can be fully or partially automated, but human involvement will continue to be necessary in more complex tasks. AI can also free analysts to address new intelligence problems and develop supporting technologies to enable more complex analysis. However, analysts will need new skills, both to facilitate use of AI and to take advantage of opportunities to conduct more-advanced analysis. We have used this framework and data flow mapping process to assess how AI investments can support team-based intelligence analysis processes within the U.S. Air Force and U.S. Army.

Lance Menthe is a senior physical scientist at RAND and a member of the Pardee RAND Graduate School faculty. He works primarily on intelligence, surveillance, and reconnaissance issues, including employment of remotely piloted aircraft and machine learning technologies for processing, exploitation, and dissemination. Other recent work includes analyzing the potential for light attack systems to provide close air support in counterinsurgency and counterterrorism operations. Menthe is the lead developer of RAND's Systems and CONOPs Operational Effectiveness Model (SCOPEM), an agent-based model of air, ground, and space domains. Prior to joining RAND, Menthe received a Ph.D. in physics from the University of California, Los Angeles, with a thesis on the physics of twisting conformations of DNA.

Learning with Robots

Rod D. Roscoe, PhD, Associate Professor of Human Systems Engineering. Human team members learn from each other by sharing knowledge, demonstrating skills, giving and receiving feedback, and more. Research on educational technologies suggests that robot teammates could contribute to team learning in several ways.

First, researchers have explored how AI-based systems can teach learners in complex domains. Intelligent tutoring systems use AI-based algorithms to mimic expert human tutors, such as detecting learners' inputs, actions, and solutions and then responding with appropriate feedback (Kulik & Fletcher, 2016; VanLehn, 2011). Similarly, automated writing evaluation tools use natural language processing and AI-based algorithms to assess students' writing and provide feedback to aid revision (Shermis & Burstein, 2013; Wilson & Roscoe, 2019). In these cases, automated systems "instruct" the learner, which can be delivered via verbal messages, animated characters, or physical robots (Belpaeme et al., 2018).

Second, researchers have used AI techniques to simulate learners—creating teachable agents that "learn" based on user inputs. Such systems leverage "learning by teaching" such that students learn by explaining, demonstrating skills, and testing the performance of the simulated learner (Matsuda et al., 2020; Roscoe et al., 2013). These "learners" can be tangible robots, which introduces embodied learning opportunities (Belpaeme et al., 2018; Thomaz, & Breazeal, 2008).

Despite this potential, automated educational technologies also suggest several cautions. Developers often offload too much of the instructional process onto the software, inadvertently neglecting (a) the expertise of human teachers and (b) the importance of interpersonal relationships (e.g., rapport). This is exacerbated by the fact that automated detection (i.e., of language, behavior, or affect) remains imperfect, which undermines performance and trust (Belpaeme et al., 2018 Roscoe et al., 2017). Rather than automating or "replacing" teachers, learners, and team members, perhaps more attention should be paid to augmenting and facilitating human capabilities via adaptive support (e.g., Matarić, 2017; Rummel et al., 2016).

Rod Roscoe is an associate professor of human systems engineering in the Polytechnic School of the Ira A. Fulton Schools of Engineering, and a Diane and Gary Tooker Professor of Effective Education in STEM. He is affiliate faculty of the Mary Lou Fulton Teachers College and the Center for Human, Artificial Intelligence, and Robot Teaming (CHART). His research investigates how the intersection of learning science, computer science, and user science can inform effective and innovative uses of educational technologies.

Enhancing Workforce Outcomes with AI-Based Training Systems

Copyright 2020 by Human Factors and Ergonomics Society. All rights reserved. 10.1177/1071181320641018

Copyright 2020 by Human Factors and Ergonomics Society. All rights reserved. 10.1177/1071181320641018

Shivam Zaveri, M.S., Graduate Research Assistant at the School for the Future of Innovation in Society. Successfully integrating advanced technologies in an existing work environment is a pervasive challenge for the future of work (Anithes, 2017). Complexities arise from both workers and work environments; workers are an integral part of the work system with different roles and measures of success, and work environments require advanced technologies to consider multiple factors for effective integration and for maintaining a level of work system success.

The recently awarded NSF C-Accel project, Safe Skill-Aligned On-The-Job Training with Autonomous Systems (PI: Srivastava) recognizes the potential and challenges in human-robot teaming technologies. This project is exploring different avenues for enhancing workforce environments and creating successful human-robot teaming technologies. Currently, a multi-disciplinary team of researchers representing five disciplines at Arizona State University are interviewing workers in the logistics, finance, retail, and healthcare industries who interact with advanced technologies and automated systems on a daily basis. These workers are being interviewed about their role requirements, technology interactions, training methods and preferences.

Preliminary findings indicate key insights on how workers view technology systems. Workers are open to having repetitive tasks automated when it makes them more successful in their daily roles. These advanced technologies allow workers to be more effective and efficient. Certain industries, such as healthcare, require that workers adhere to strict policies, like keeping meticulous patient notes while maintaining a level of compassion for the patient. These workers are open to new technologies to aid in their tasks but prioritize patient wellbeing. Human-robot teaming systems have the potential to lessen the burden on healthcare workers but must be seamlessly integrated into the workflow. A primary challenge for future technologies will require worker buy-in and safeguarding work environment standards.

Human-robot teaming systems can potentially provide many benefits to many workers and work environments. Building adaptability into these systems requires a thorough understanding of workplace challenges, and actively incorporating feedback from the users. Each workplace has a unique pace to accept change. This project considers the qualitative and technical aspects of the future of work as it builds human-robot technologies that can enhance work environments evolving over time.

Shivam Zaveri is a graduate student at Arizona State University focused on industrial workforces and entrepreneurship. Shivam received his B.E. in Industrial Engineering in 2014 from the University of Tennessee - Knoxville, where he focused on entrepreneurship, globalization and reliability in a range of systems. He also received his Masters in Science and Technology Policy from Arizona State University in 2017.

ACKNOWLEDGEMENTS

EC and SZ were partially supported by the National Science Foundation [OIA-1936997] during the formation of this panel.

REFERENCES

Anithes, E. (2017). The Shape of Work to Come. *Nature*, 316-319.

Belpaeme, T., Kennedy, Ramachandran, Scassellati, B., & Tanaka, F. (2018). Social robots for education: A review. *Science Robotics*, *3*(21), eaat5954.

Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, *358*(6370), 1530-1534.

DeCostanza, A. H., Marathe, A. R., Bohannon, A., Evans, A. W., Palazzolo, E. T., Metcalfe, J. S., & McDowell, K. (2018). Enhancing Human–Agent Teaming with Individualized, Adaptive Technologies: A Discussion of Critical Scientific Questions. Retreived February 24, 2020, from https://brain.ieee.org/brain-storm/enhancing-human-agent-teaming/

Dolgov, I., Kaltenbach, E.K., Khalaf, A. S., & Toups, Z. O. (2017). Measuring human performance in the field: Concepts and applications for unmanned aircraft systems. *Human Factors in Practice: Concepts and Applications*. CRC Press.

Duhigg, C. (2016, February 25). What Google learned from its quest to build the perfect team. *The New York Times Magazine*. https://doi.org/https://doi.org/10.3929/ethz-b-000238666

Fong, T. (2001). *Collaborative Control : A Robot-Centric Model for Vehicle Teleoperation*. Carnegie Mellon University.

Groom, V., & Nass, C. (2007). Can robots be teammates? *Interaction Studies*, *8*(3), 285–301. https://doi.org/10.1075/gest.8.3.02str

Gunning, D. (2017). *Explainable artificial intelligence (xai) program update November 2017*. Defense Advanced Research Projects Agency, DARPA/I20.

Holden, J., & Goel, N. (2016). *Fast-forwarding to a future of on-demand urban air transportation* [White paper]. Retrieved February 6, 2020 from Uber Technologies, Inc.: https://www.uber.com/elevate.pdf.

Kaltenbach, E.K., & Dolgov, I. (2017, September). On the dual nature of transparency and reliability: Rethinking factors that shape trust in automation. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (Vol. 61, No. 1, pp. 308-312). Los Angeles, CA: SAGE Publications.

Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*, *86*(1), 42-78).

Matarić, M. J. (2017). Socially assistive robotics: Human augmentation versus automation. *Science Robotics*, *2*(4), eaam5410.

Matsuda, N., Weng, W., & Wall, N. (2020. The effect of metacognitive scaffolding for learning by teaching a teachable agent. *International Journal of Artificial Intelligence in Education*. DOI: 10.1007/s40593-019-00190-2

Michaud, F., Boissy, P., Labonté, D., Brire, S., Perreault, K., Corriveau, H., … Létourneau, D. (2010). Exploratory

Copyright 2020 by Human Factors and Ergonomics Society. All rights reserved. 10.1177/1071181320641018

design and evaluation of a homecare teleassistive mobile robotic system. *Mechatronics*, *20*(7), 751–766. https://doi.org/10.1016/j.mechatronics.2010.01.010

OECD. (2019). Forty-two countries adopt new OECD Principles on Artificial Intelligence. Retrieved December 5, 2019, from https://www.oecd.org/science/forty-two-countries-adopt-new-oecd-principles-on-artificial-intelligence.htm

Parasuraman, R., Sheridan, T., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, *30*(3), 286–297. https://doi.org/10.1109/3468.844354

Paris, C., Salas, E., & Cannon-Bowers, J. (2000). Teamwork in multi-person systems: a review and analysis. *Ergonomics*, *43*(8), 1052–1075. https://doi.org/10.1080/00140130050084879

Pollitt, C., & Hupe, P. (2011). The role of magic concepts. *Public Management Review*, *13*(5), 641–658. https://doi.org/10.1080/14719037.2010.532963

Roscoe, R. D., Segedy, J. R., Sulcer, B, Jeong, H., & Biswas, G. (2013). Shallow strategy development in a teachable agent environment designed to support self-regulated learning. *Computers and Education*, *62*, 286-297.

Roscoe, R. D., Wilson, J., Johnson, A. C., & Mayra, C. R. (2017). Presentation, expectations, and experience: Sources of student perceptions of automated writing evaluation. *Computers in Human Behavior*, *70*, 207-221.

Roundtree, K. A., Goodrich, M. A., & Adams, J. A. (2019). Transparency: Transitioning from human–machine systems to human-swarm systems. *Journal of Cognitive Engineering and Decision Making*, *13*(3), 171-195.

Rummel, N., Walker, E., & Aleven, V. (2016). Different futures of adaptive collaborative learning support. *International Journal of Artificial Intelligence in Education*, *26*(2), 784-795.

Salas, E., Burke, C., & Janis, A. (2000). Teamwork: Emerging principles. *International Journal of Management Reviews*, *2*(4), 339–356.

Salas, E., Dickinson, T. L., Converse, S. A., & Tannenbaum, S. I. (1992). Toward an understanding of team performance and training. In R. Swezey & E. Salas (Eds.), *Teams: Their Training and Performance* (pp. 3–29). Norwood, NJ.

Salas, E., Wilson, K., Murphy, C., King, H., & Salisbury, M. (2008). Communicating, coordinating, and cooperating when lives depend on it: Tips for teamwork. *Joint Commission Journal on Quality and Patient Safety / Joint Commission Resources*, *34*(6), 333–341. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/18595379

Santoni de Sio, F., & Van den Hoven, J. (2018). Meaningful human control over autonomous systems: A philosophical account. *Frontiers in Robotics and AI*, *5,*15. https://doi.org/10.3389/frobt.2018.00015

Shermis, M. D., & Burstein, J. (Eds.) (2013). *Handbook of automated essay evaluation*. New York, NY: Routledge.

Sugihara, K., & Suzuki, I. (1994). Distributed motion coordination of multiple mobile robots. In *Proceedings. 5th IEEE International Symposium on Intelligent Control 1990* (Vol. 3, pp. 138–143). IEEE Comput. Soc. Press. https://doi.org/10.1109/ISIC.1990.128452

Sycara, K., & Zeng, D. (1996). Coordination of Multiple Intelligent Software Agents. *International Journal of Cooperative Information Systems*, *05*(02n03), 181–211. https://doi.org/10.1142/S0218843096000087

Thipphavong, D. P., Apaza, R., Barmore, B., Battiste, V., Burian, B., Dao, Q., ... & Idris, H. R. (2018). Urban air mobility airspace integration concepts and considerations. In *2018 Aviation Technology, Integration, and Operations Conference* (p. 3676-3694). https://doi.org/10.2514/6.2018-3676

Thomaz, A. L., & Breazeal, C. (2008). Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence*, *172*, 716-737.

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, *46*(4), 197-221.

Vascik, P. D., Hansman, R. J., & Dunn, N. S. (2018). Analysis of urban air mobility operational constraints. *Journal of Air Transportation*, *26*(4), 133-146.

Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Transparent, explainable, and accountable AI for robotics. *Science Robotics, 2*(6), eaan6080.

Wilson, J., & Roscoe, R. D. (2019). Automated writing evaluation and feedback: Multiple metrics of efficacy. *Journal of Educational Computing Research*, *58*(1), 87-125.