Taking a Constructivist Approach to Human-AI Co-learning Design

Julian Abich IV & Eric Sikorski
Quantum Improvements Consulting
Orlando, FL
jabich@quantumimprovements.net
esikorski@quantumimprovements.net

ABSTRACT

Artificial intelligence (AI) can facilitate personalized experiences shown to impact training outcomes. Trainees and instructors can benefit from AI-enabled adaptive learning, task support, assessment, and learning analytics. Impacting the learning and training benefits of AI are the instructional strategies implemented. The co-learning strategy is the process of learning how to learn with another entity. AI co-learning techniques can encourage social, active, and engaging learning behaviors consistent with constructivist learning theory. While the research on co-learning among humans is extensive, human-AI co-learning needs to be better understood. In a team context, co-learning is intended to support team members by facilitating knowledge sharing and awareness in accomplishing a shared goal. Co-learning can also be considered when humans and AI partner to accomplish related tasks with different end goals. This paper will discuss the design of a human-agent co-learning tool for the United States Air Force (USAF) through the lens of constructivism. It will delineate the contributing factors for effective human-AI co-learning interaction design. A USAF maintenance training use case provides a context for applying the factors. The use case will highlight the initiative of leveraging AI to help close an experience gap in maintenance personnel through more efficient, personalized, and engaging support.

ABOUT THE AUTHORS

Julian Abich IV is a Senior Human Factors Engineer at Quantum Improvements Consulting. He holds a Ph.D. in Modeling and Simulation with a specialization in Human Factors, a B.S. in Psychology, and two certificates in Design for Usability and Instructional Design for Simulations from the University of Central Florida. He has over 16 years of experience applying human factors & ergonomics (HF/E) principles, modeling & simulation (M&S) approaches, and instructional design strategies to assess, predict, and improve human performance. He has extensive experience conducting usability & user experience evaluations across various human-computer interaction and training applications. Continuing to take a user-centered approach, his work focuses on designing, applying, and evaluating the use of immersive technologies for training.

Eric Sikorski is Director of Programs and Research at Quantum Improvements Consulting, where he leads a team of Research Psychologists and Human Factors Engineers. His work at QIC includes staff development, strategic planning for the division, and managing various field studies, such as evaluating cutting-edge immersive training technology and assessing team communication, for multiple DoD customers. Eric holds a Ph.D. and M.S. in Instructional Systems from Florida State University and an M.A. in Human Factors and Applied Cognition from George Mason University.

Taking a Constructivist Approach to Human-AI Co-learning Design

Julian Abich IV & Eric Sikorski
Quantum Improvements Consulting
Orlando, FL
jabich@quantumimprovements.net
esikorski@quantumimprovements.net

INTRODUCTION

Artificial intelligence (AI) can enhance learning and training solutions to provide an adaptive and personalized experience (Kaplan, 2021). It is important to consider AI's instructional and interaction strategies when designing AI solutions for personalized learning and training applications. Human-AI interaction guidelines have been researched to generally inform products infused with AI (Amershi et al., 2018), but further consideration is needed for interactions with learning support solutions. The purpose of learning solutions is for humans to achieve their learning goals, requiring AI to learn how to support that process. Hence, research is needed to investigate strategies that facilitate co-learning among human and AI partners to achieve their goals. Co-learning in the current context refers to entities learning how to learn with each other to achieve related goals. Constructivist learning theory suggests that co-learning partners should actively develop and enhance their knowledge, skills, and abilities (KSAs) through social interactions. However, commercial AI solutions typically assume humans will be passive receivers rather than actively influencing their experience by managing what the AI knows and how it interacts (Kuijer & Giaccardi, 2018). Effective co-learning partnerships will benefit from both entities being socially and actively engaged (Endsley & Kiris, 1995; Krishna, Lee, Fei-Fei, & Bernstein, 2022). Therefore, factors contributing to effective human-AI co-learning partnerships must be identified to inform research guidelines. This paper presents factors to consider when designing an AI-based co-learning tool for a real-world USAF training challenge. The factors align with the constructivist learning theory when considering humans' socially active role in personalizing AI co-learning experiences.

REAL-WORLD TRAINING CHALLENGE

U.S. Air Force (USAF) maintenance training personnel continuously seek innovative methods to increase training efficiency. USAF tactical aircraft maintainers are responsible for ensuring aircraft are safe, airworthy, and operationally fit for missions. Aircraft maintenance is the USAF's largest enlisted career detail, accounting for about 25% of their active duty enlisted personnel (GAO, 2019). In a recent report, two key USAF-tracked maintenance metrics revealed the ongoing need to meet mission-capable rate goals and that the aircraft fix rate has declined since 2015 (GAO, 2022). A 2022 Congressional Budget Office report found that the percentage of time an aircraft was fit for training or missions has declined since 2001. Further, a National Commission on Military Aviation Safety (MAS, 2020) report found that 80% of aviation mishaps were due to human factors, specifically aircrew and maintenance errors. These results indicate that at the organizational level, the USAF needs to make critical changes to achieve fleet readiness.

Several interrelated factors may contribute to maintenance delays, but addressing existing training challenges can help improve aircraft mission and training availability. The USAF has recently addressed maintenance staffing shortages, but the challenge has shifted to a need for qualified and experienced maintainers (GAO, 2022). Staffing these less experienced maintainers has led to aircraft being held longer and reduced the units' ability to troubleshoot because the staff must be trained and certified on their specific aircraft (GAO, 2022). Experienced technicians are needed for longer maintenance periods, which takes away from their ability to train others or upgrade themselves. A major aviation safety concern is that a lack of competency can have a compounding effect, where less experienced instructors and trainers teach the next generation of maintainers, continually reducing the transfer of tacit and organizational knowledge and impacting fleet readiness levels (NCMAS, 2020). Additionally, aircraft maintainers may be reassigned to address staffing shortages for other aircraft, adding demand to provide rapid reskilling. Working to meet mission-capable maintenance rates while short-staffed adds pressure on the maintenance crews (NCMAS, 2020), resulting in fatigue and stress, misperception of hazards, time mismanagement, lack of motivation, and inadequate skill development. These pressures drive experienced maintainers to leave the Services, being replaced by personnel with no expertise (NCMAS, 2020). Without immediate intervention, the USAF will continue to face maintenance readiness issues.

Proposed Application of AI in USAF Maintenance

The U.S. Department of Defense (DoD) has long recognized autonomous agents as a potential force multiplier to help Airmen operate successfully in complex environments (Lin, Bekey, & Abney, 2008). Developing synthetically intelligent agents has been a priority for the U.S. military (Department of Defense, 2013), and the USAF is sponsoring efforts to leverage AI for aircraft maintenance staff training. The sponsored work's focus was twofold: 1) study the dynamic factors contributing to co-learning with synthetic agents and 2) develop AI utilities for human-agent co-learning to help Airmen achieve aircraft maintenance proficiency. The goal was to design and develop persistent AI agents within DoD cloud infrastructure to personalize assistance for Airmen in training and transition as a job aid to the flight line. An environment analysis was conducted to inform where and when an AI-enabled support tool could be utilized. The environment analysis focused on understanding the types of software and hardware available for AI integration during tactical aircraft maintenance training. Integration can begin in the technical schools with AI-enabled interactive courseware accessible on tablets and desktop computers to assist trainees with foundational knowledge development. Both trainees and instructors can help curate the data AI leverages to establish course learning criteria, build trainee models, and optimize instructional approaches for learning. Data analytics informed by trainee behavior and performance, such as content accessed and course assessments, can be accessed through the learning management system and record store. AI can facilitate unique presentation of data analytics that offer enhanced insight into trainees' current and future success. When the trainee graduates from technical school and transitions into a maintenance unit, the AI will also transition from primary learning support to a personalized job aid. All Airmen are assigned tablets through which training doctrine and maintenance checklist can be accessed, providing AI assistance and tracking opportunities. The Airman can leverage the job aid to support task execution, such as procedural knowledge on a task that the Airman has little experience performing. Instructors can input observational assessment data regarding the proficiency level of the Airman that further informs the training model. Instructor and trainer benefits can also emerge from AI integration by generating updated courseware, assessments, and trainee models capable of real-time and predictive analytics. The trainers' assessment data can inform the quality of the technical school curriculum and strategies. AI can learn the most effective strategies through direct and indirect feedback to facilitate the human's goals. Leveraging AI to rapidly and effectively advance human skill development is a crucial area of research, particularly in areas experiencing staffing shortages.

CO-LEARNING AND CONSTRUCTIVISM

Co-learning is the process of learning how to learn with another entity. It comprises two sub-processes: learning together and mutual adaptation (van Zoelen, van den Bosch, & Neerincx, 2021). Co-learning is an iteratively emergent process in which group members adapt their behavior and provide feedback to each other. Behavioral adaptations are established and maintained over time and across contexts through bidirectional communication, providing the basis for mutual understanding. Inherent within the co-learning process are the activities that necessitate interaction between humans and AI. These activities drive the generation and transfer of information to group members. Group members communicate information through socially learned behaviors. When humans adapt their performance and learning strategies to prevent the AI from making mistakes, the AI can learn an effective strategy and choose the process that aligns with the human's strategy (van Zoelen, van den Bosch, & Neerincx, 2021).

Constructivism provides a theoretical basis for designing human-AI co-learning. The concepts of constructivism transcended through cybernetics studies and eventually made their way into AI (Hof, 2021). The constructivist paradigm has influenced AI researchers notably with the idea that learning is a multistage process building from simple to complex structures supported by active social interaction. Today, different AI approaches rely on human input to help build the AI's knowledge base and follow a similar development pattern. Applying this to the context of human-AI co-learning, strategies emerging from the constructivist paradigm can guide the design of interactions that foster learning for humans and AI partners. Constructivism is built on the concept that learners actively develop their knowledge, build on their prior experience, and construct meaning through socially supported exploration in real-world contexts (Hodson & Hodson, 1998). The notion that social interactions make exchanging ideas possible and make new knowledge available has led to work on socially situated AI (Krishna, Lee, Fei-Fei, & Bernstein, 2022). Constructivism states that through assimilation and accommodation, humans gather and interpret new knowledge or alter their pre-existing mental models or schemas based on newly gained information. Learners are encouraged to challenge their existing knowledge or pre-conceived notions to expand their knowledge base through the guidance of an instructor or skilled peer. Research into agents capable of counternarrative intelligence has looked at the benefits of AI that oppose assigned goals or plans to purposefully encourage humans to take broader and diverse perspectives on a topic (Coman & Aha, 2017). Active real-world engagement leads to higher learning gains through collaboration and creativity (Behling & Hart, 2008). Trainees should participate in dynamic, motivational, practical activities and engaging course material that stimulates novel problem-solving approaches. Proposed AI support for complex problem-solving has been recommended based on the increased digitization of the workplace (de Laat, Joksimovic, & Ifenthaler, 2020). As identified within the human-computer interaction literature, these core attributes are already in practice, but the explicit connection to constructivism and the practical implications have not been thoroughly examined, including for human-agent co-learning design.

FACTORS FOR EFFECTIVE HUMAN-AI CO-LEARNING

Investigation of the theoretical and empirical literature on human-agent teaming, synthetic agents as learning companions, team-based learning, and trust in automation revealed many factors and subfactors that contribute to effective collaboration. Key factors were identified that are proposed to contribute to an effective human-AI co-learning partnership. To form effective co-learning partnerships, humans and AI must engage in extended bidirectional communication, establish trust, gain an implicit understanding of each other's knowledge and skills, and be flexible in their roles. These factors align with the tenets of constructivism and are proposed to benefit human-AI co-learning.

Trust

Trust is defined as "the willingness of a party to be vulnerable to the actions of another party based on the expectation that the other will perform a particular action important to the trustor, irrespective of the ability to monitor or control that other party (Mayer, Davis, & Schoorman, 1995)." To benefit from constructivist social activities, co-learning partners must develop trust in members' competence (Scheer, Noweski, & Meinel (2012). Developed over time, trust is dynamic and fluctuates as perceptions of other parties are developed. Performance accuracy, reliability, and transparency affect trust, which plays a moderating role in the human-AI relationship (Chancey, Bliss, Yamani, & Handley, 2017; Wright, Chen, & Lakhmani, 2020). Transparent intent and reasoning processes aid in establishing trust, especially during initial interactions, and reliability maintains trust over time. AI intent is associated with the state of commitment to achieve a goal through some action (Miller, 2019). While AI can be programmed to achieve a goal and offer transparency, there are often unintentional outcomes. In this case, AI should be programmed to track and present the causal history or enabling factors to support the determination of its cause (Malle, 2004). Tracking its actions will present the steps that led to the outcome without relying on AI to reason its actions.

Perceived reliability is affected by the human's knowledge of the agent's actions, reasoning, and projected future status (Daronnat, Azzopardi, Halvey, & Dubiel, 2020). One study found significant associations between perceived system accuracy, reliance, and trust, with an accuracy level of 70% as the threshold between increasing and decreasing trust (Yu, Berkovsky, Taib, Zhou, & Chen, 2019). Errors in performance erode humans' trust in autonomous teammates, and a lengthy period of reliable performance is required to repair feelings of trust (Daronnat, Azzopardi, & Halvey, 2020). Trust can be restored if an autonomous agent promptly corrects itself or apologizes clearly and decisively. Additionally, the agent can ask the human for help understanding the situation and then display competence by applying the new information to arrive at the correct solution or perform the correct action (Akgun, Cagiltay, & Zeyrek, 2010; de Visser, Pak, & Shaw, 2018; Inbar & Meyer, 2015).

Common Ground

Common ground needs to be established to develop trust (Glikson & Woolley, 2020). Common ground refers to the knowledge, facts, and beliefs shared between participants in some joint activity (Clark, 1996). Establishing and maintaining common ground among group members is an important and ongoing process for co-learning partnerships. Common ground (i.e., mutual understanding) is the basis by which two entities can work together effectively and assists with establishing and maintaining a shared mental model (Woods, Bradshaw, Hoffman, & Feltovich, 2004). Constructing a shared mental model and developing common ground involves actively engaging co-learners to develop a shared understanding of each partner. For a human to develop and maintain a mental model of the AI, its explanations of actions should be sound (i.e., correct), complete (i.e., all underlying causes identified) but not overwhelming (Kulesza, Burnett, Wong, & Stumpf, 2015). By balancing those three explanation criteria, humans can learn to adapt their behavior to achieve the desired AI outcomes. Since mental models are considered context-specific or -sensitive (Scheutz, DeLoach, and Adams 2017) and AI does not have the range of cognitive abilities of a human, it becomes challenging for AI to develop and share their models of a human. At a minimum, the AI needs a task and a human model. A task model implies a limit for the type of tasks the AI could perform and does not capture inferences made by humans (Andrews, Lilly, Srivastava, & Feigh, 2023). Attempts to develop AI systems capable of developing human models have been successful by making restrictive assumptions about humans. Recently, inverse reinforcement learning, which uses a reward function employed by humans, has shown success in the AI's human model development (Walsh & Feigh, 2021).

To support the development of common ground and shared mental models, AI needs the ability to explain or interpret its decision path. Decision paths are the steps the AI takes to achieve a specific output. The steps can showcase what the AI has learned to generate and present the output in many formats, such as decision trees. These decision paths aim to facilitate transparency of the AI and support explanations of its outputs. Overall, attempts to generate common ground through explainable AI seem to uncover more challenges than results (Gunning, Stefik, Choi, Miller, Stumpf, & Yang, 2019), with progress partially hindered by a lack of consistent terms and effective assessment measures (Andrews, Lilly, Srivastava, & Feigh, 2023). While more research is needed to understand how common ground develops in human-AI co-learning partnerships, there is evidence that common ground can be maintained through awareness of the group's contributions to the learning partnership (Janssen & Bodemer, 2013).

Group Awareness

Group awareness is the group members' knowledge of how the group functions and how expertise is divided among group members (Janssen & Bodemer, 2013; Schnaubert & Bodemer, 2019). Human and AI agents can work together to achieve more than either could independently by developing roles based on each other's strengths and an awareness of their weaknesses. Consistent with constructivism, co-learning partnerships allow learners to rely on members for support in areas where their knowledge or skills are lacking. Managing actions through group awareness can reduce coordination effort, increase efficiency, and reduce the chance of errors (Gutwin & Greenberg, 2004). The benefits of developing group awareness, specifically cognitive and social group awareness, found in human-human teams also impacts human-AI co-learning performance (Klein, Woods, Bradshaw, Hoffman, & Feltovich, 2004).

Cognitive group awareness (CGA) is the awareness that results from information about group members' knowledge, the information they possess, or the opinions they hold. CGA can be used to coordinate activities in the content space of co-learning. When CGA is high, the effort to coordinate activities in the co-learning content space is reduced. CGA can trigger beneficial learning behavior by providing a basis for adapting learners' contributions to their partners' knowledge (Zufferey, Bodemer, Buder, & Hesse, 2011). In human-synthetic agent teams, these adaptations are beneficial at the individual and group levels (van Zoelen, van den Bosch, Neerincx, 2021; Xu et al., 2012). CGA allows learners to identify gaps and deficiencies in their knowledge, avoid illusions of comprehension, and initiate learning activities to gain missing knowledge (Sangin, Molinari, Nüssli, & Dillenbourg, 2011), all of which impact AI transparency. Perceived differences between learners regarding knowledge, assumptions, or opinions can trigger collaborative elaboration processes such as thoroughly discussing controversial perspectives (Bodemer & Dehler, 2011), leading to meaningful learning. Further, CGA can initiate information-sharing processes involving social group awareness.

Social group awareness (SGA) refers to group members' awareness of the social situation of the rest of the group (Gross, Stary, & Totter, 2005). It refers to awareness about what group members are doing, with whom they are communicating, how they contribute to the common group goal and their roles. Developing social awareness in AI requires the agent to pursue social goals and be aware of the social effects of its actions (Dignum, Prada, & Hofstede, 2014). SGA is generated by information about group members' collaborative behavior, such as the quality of participation and the number of contributions to online discussions. Social group awareness tools provide members with quantitative information about everyone's role in the collaborative learning processes. SGA is achieved by informing learners about group members' participation rates, such as how actively learners are involved in group processes, emphasizing the need for effective communication.

Communication

Effective human and AI co-learning partnerships will rely heavily on communication. At a minimum, bidirectional communication must occur between humans and AI to establish successful transmission and reception of the conveyed information. As with most learning theories, communication is a fundamental aspect of constructivism and essential for learners to share knowledge and experiences to construct personal meaning. Communication patterns and transparency of the agent and the agent's understanding of the human affect trust and human performance (Lakhmani, Wright, Schwartz, & Barber, 2019). Specifically, a bidirectional communication pattern combined with a transparent agent provides a context for humans to understand the agent's process, purpose, and performance. Communication can be either explicit or implicit. Explicit forms of communication entail directly sending meaning through an overt modality, such as human speech. Implicit communication refers to non-verbal behavior that requires the receiver to interpret the signaled information. Non-verbal communication can be used if there is an established meaning between the sender and receiver (Breazeal, Kidd, Thomaz, Hoffman, & Berlin, 2005). Humans communicate implicitly through body language, facial expressions, or voice tone. Humans rely on both explicit and implicit forms of communication. While both may be preferred, explicit forms of communication aim to ensure the message is clearly understood without needing interpretation. Understanding each member's communication capabilities and nuances early in the co-learning partnership will help alleviate errors and frustrations as the interaction progresses. Although AI partners may not always be capable of implicit communication, they can utilize multimodal communication to ensure the human receives and understands the message.

Multimodal communication provides co-learning partners with multiple methods to convey information. There are many benefits of multimodal communication. Communication modalities can be chosen to support dual-task efficiency if the information is sent and received through modalities that pose little cognitive demand during a task. Further, multimodal communication allows information to be sent in multiple ways, leveraging redundancy gains by allowing the same message to be received through different modes. By supporting multimodal communication in human-AI co-learning, both partners can adapt their communication modes for the success of the co-learning partnership.

Mutual adaptation

Constructivism contends that communication is the process of mutual adaptation among learning partners. Mutual adaptation involves the dynamic adaptive process and behaviors exhibited by humans and AI in achieving agreed-upon goals (Nikolaidis, Hsu, & Srinivasa, 2017). Behavioral adaptations are established and maintained over time and across contexts through bidirectional communication, providing the basis for mutual understanding. Humans and synthetic agents that adapt to each other during the co-learning process improve learning at the individual and team levels (van Zoelen, van den Bosch, Neerincx, 2021; Xu et al., 2012). Longitudinal studies that have examined mutual adaptation during human-agent teaming have found that synthetic agents can adapt to users' preferences and needs. Furthermore, users' perceptions of the agent can evolve, and people adjust their behaviors in response to agent-initiated social behaviors (Serholt & Barendregt, 2016). When humans adapt their performance and learning strategies to prevent AI from making mistakes, AI can learn and select an effective strategy aligning with the human's strategy (van Zoelen, van den Bosch, & Neerincx, 2021). Identifying interaction patterns, such as curating or exploring (Grabe, González-Duque, Risi & Zhu, 2022), can support shared awareness and provide a foundation for team communication about mutual adaptation. The interactions between human-AI co-learning partners could be used to develop a shared vocabulary to identify and discuss the adaptations each engages in during the co-learning process. A shared vocabulary can foster shared situational awareness and encourage adaptation to support each human and AI partner. An agent that informs the human about the level of uncertainty underlying the action it is performing can help humans appropriately calibrate expectations of agent performance and adapt as needed for successful goal achievement (Chen, Lakhmani, Stowers, Selkowitz, Wright, & Barnes, 2018; Johnson, Demir, McNeese, Gorman, Wolff, & Cooke, 2021).

CONCLUSION

Exploration of human-AI co-learning adds a new perspective to today's AI revolution. In contrast to typical teaming applications, co-learning partnerships can work together to achieve different but related goals. Co-learning aims to improve KSAs through interdependence and systematic interaction rather than counterbalancing team member deficits. The success of a human-AI co-learning partnership is proposed through developing and maintaining trust, common ground, group awareness communication, and mutual adaptation. Adhering to constructivist learning principles can facilitate the effectiveness of the co-learning process and support the contributing factors for successful partnership. The partnership should focus on helping the human develop the necessary KSAs for task execution. The AI should focus on learning about humans, developing a learner model, and improving their path to KSA development. Enabling AI to enhance personalized training experiences requires reciprocal patterns of interactions between the human and AI that results in a mutual understanding of each other's strengths and weaknesses (van Zoelen, van den Bosch, & Neerincx, 2021) Consistent with constructivist theory, human-AI co-learning should encourage group interaction to develop the personal level of proficiency needed for successful job performance.

Co-learning may require multiple interaction strategies for a successful partnership. The human and AI will evolve individually and together as they learn about the content, themselves, and their partner. The application context will shift during co-learning if integrated into career-supported tools. The dynamic relationship will require balancing the factors contributing to human-AI co-learning. The factors presented are based on theory and informed by related research literature. Therefore, there is a need for empirical support focused specifically on understanding the role each factor plays and the effectiveness of constructivist strategies to enhance co-learning.

ACKNOWLEDGMENTS

The research reported in this document was performed in support of CHI Systems for the USAF. 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 CHI Systems, the USAF, or the U.S. Government.

REFERENCES

Akgun, M., Cagiltay, K., & Zeyrek, D. (2010). The effect of apologetic error messages and mood states on computer users' self-appraisal of performance. *Journal of Pragmatics*, *42*(9), 2430-2448.

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P.N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019, May). Guidelines for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1-13).

Andrews, R. W., Lilly, J. M., Srivastava, D., & Feigh, K. M. (2023). The role of shared mental models in human-AI teams: a theoretical review. *Theoretical Issues in Ergonomics Science*, *24*(2), 129-175.

Behling, K., & Hart, D. (2008). Universal design: A model for professional development. In S.E. Burgstahler & R.C. Cory (Eds.), *Universal design in higher education: From principles to practice* (pp. 109–125). Harvard Education Press.

Bodemer, D. & Dehler, J. (2011). Group awareness in CSCL environments. *Computers in Human Behavior*, *27*(3), 1043-1045.

Breazeal, C., Kidd, C. D., Thomaz, A. L., Hoffman, G., & Berlin, M. (2005). Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. *Proceedings of The 2005 Ieee/Rsj International Conference on Intelligent Robots and Systems* (pp. 708-713). IEEE.

Chen, J. Y., Lakhmani, S. G., Stowers, K., Selkowitz, A. R., Wright, J. L., & Barnes, M. (2018). Situation awareness-based agent transparency and human-autonomy teaming effectiveness. *Theoretical Issues in Ergonomics Science, 19*(3), 259-282.

Chancey, E. T., Bliss, J. P., Yamani, Y., & Handley, H. A. (2017). Trust and the compliance–reliance paradigm: The effects of risk, error bias, and reliability on trust and dependence. *Human Factors*, *59*(3), 333-345.

Clark, H.H. (1996). Using Language. Cambridge University Press.

Coman, A., & Aha, D. W. (2017). Cognitive support for rebel agents: Social awareness and counternarrative intelligence. *Proceedings of the Fifth Conference on Advances in Cognitive Systems*, 1-6.

Congressional Budget Office. (2022). Availability and use of aircraft in the Air Force and Navy (Publication No. 57433). https://www.cbo.gov/system/files/2022-01/57433-aircraft.pdf

Daronnat, S., Azzopardi, L., & Halvey, M. (2020, December). Impact of agents' errors on performance, reliance, and trust in human-agent collaboration. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *64*(1), 405-409.

Daronnat, S., Azzopardi, L., Halvey, M., & Dubiel, M. (2020, November). Impact of agent reliability and predictability on trust in real time human-agent collaboration. *Proceedings of the 8th International Conference on Human-Agent Interaction*, 131-139. https://doi.org/10.1145/3406499.3415063 de Laat, M., Joksimovic, S. & Ifenthaler, D. (2020). Artificial intelligence, real-time feedback and workplace learning analytics to support in situ complex problem-solving: A commentary. *International Journal of Information and Learning Technology, 37*(5), 267-277. https://doi.org/10.1108/IJILT-03-2020-0026

de Visser, E. J., Pak, R., & Shaw, T. H. (2018). From 'automation' to 'autonomy': the importance of trust repair in human–machine interaction. *Ergonomics*, *61*(10), 1409-1427.

Endsley, M.R. & Kiris, E.O. (1995). The out-of-the-loop performance problem and level of control in automation. *Human Factors*, *37*(2), 381-394. https://doi.org/10.1518/001872095779064555

Glikson, E., Woolley, A.W. (2020). Human Trust in Artificial Intelligence: Review of Empirical Research. *Academy of Management Annals. 14*(2): 627–660. https://www.10.5465/annals.2018.0057

Government Accountability Office. (2019). Actions needed to improve poor conditions of facilities and equipment that affect maintenance timeliness and efficiency (Publication No. GAO-19-242). https://www.gao.gov/assets/gao-19-242.pdf

Government Accountability Office. (2022). Air Force and Navy aviation: Actions needed to address persistent sustainment risks (Publication No. GAO-22-104533). https://www.gao.gov/assets/gao-22-104533.pdf

Grabe, I., González-Duque, M., Risi, S., & Zhu, J. (2022). Towards a framework for human-AI interaction patterns in co-creative GAN applications. *Proceeding of the 3rd Workshop on Human-AI Co-Creation with Generative Models (HAI-GEN '22) at ACM IUI Workshops*.

Gross, T., Stary, C., & Totter, A. (2005). User-centered awareness in computer-supported cooperative work-systems: Structured embedding of findings from social sciences. *International Journal of Human-Computer Interaction*, *18*(3), 323-360.

Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., & Yang, G. Z. (2019). XAI—Explainable artificial intelligence. *Science Robotics*, *4*(37). https://doi.org/10.1126/scirobotics.aay7120

Gutwin, C., & Greenberg, S. (2004). The importance of awareness for team cognition in distributed collaboration. In E. Salas & S. M. Fiore (Eds.), *Team cognition: Understanding the factors that drive process and performance* (pp. 177–201). American Psychological Association. https://doi.org/10.1037/10690-009

Hodson, D. & Hodson, J. (1998). From constructivism to social constructivism: A Vygotskian perspective on teaching and learning science. *School Science Review*, *79*(289), 33-41.

Hof, B. (2021). The turtle and the mouse: How constructivist learning theory shaped artificial intelligence and educational technology in the 1960s. *History of Education*, *50*(1), 93-111.

Inbar, O., & Meyer, J. (2015, September). Manners matter: trust in robotic peacekeepers. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *59*(1), 185-189. https://doi.org/10.1177/1541931215591038

Janssen, J. & Bodemer, D. (2013). Coordinated computer-supported collaborative learning: Awareness and awareness tools. *Educational Psychologist*, *48*(1), 40-55.

Johnson, C. J., Demir, M., McNeese, N. J., Gorman, J. C., Wolff, A. T., & Cooke, N. J. (2021). The impact of training on human–autonomy team communications and trust calibration. *Human Factors.* https://doi.org/10.1177_00187208211047323

Kaplan, A. (2021). *Higher education at the crossroads of disruption*: *The University of the 21st century*. Emerald Publishing.

Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. J. (2004). Ten challenges for making automation a" team player" in joint human-agent activity. *IEEE Intelligent Systems, 19*(6), 91-95.

Krishna, R., Lee, D., Fei-Fei, L., & Bernstein, M. S. (2022). Socially situated artificial intelligence enables learning from human interaction. *Proceedings of the National Academy of Sciences*, *119*(39). https://doi.org/10.1073/pnas.211573011

Kuijer, L., & Giaccardi, E. (2018, April). Co-performance: Conceptualizing the role of artificial agency in the design of everyday life. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, *125*(1), 1-13.

Kulesza, T., Burnett, M., Wong, W.-K., & Stumpf, S. (2015). Principles of explanatory debugging to personalize interactive machine learning. *Proceedings of the 20th International Conference on Intelligent User Interfaces, ACM*, pp.126–137.

Lakhmani, S. G., Wright, J. L., Schwartz, M. R., & Barber, D. (2019, November). Exploring the effect of communication patterns and transparency on performance in a human-robot team. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *63*(1), 160-164. https://doi.org/10.1145/3173574.3173699

Lin, P., Bekey, G., & Abney, K. (2008). *Autonomous military robotics: Risk, ethics, and design*. California Polytechnic State Univ San Luis Obispo.

Malle, B.F. (2004). *How the Mind Explains Behavior: Folk Explanations, Meaning, and Social Interaction*. MIT Press.

Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review, 20*(3), 709-734.
National Commission on Military Aviation Safety. (2020). Military aviation losses FUY 2013- 2020: Report to the President and the Congress of the United States. https://www.globalsecurity.org/military/library/report/2020/ncmas_final_report_20201201.pdf
Nikolaidis, S., Hsu, D., & Srinivasa, S. (2017). Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research*, *36*(5-7), 618-634.
Sangin, M., Molinari, G., Nüssli, M. A., & Dillenbourg, P. (2011). Facilitating peer knowledge modeling: Effects of a knowledge awareness tool on collaborative learning outcomes and processes. *Computers in Human Behavior, 27(*3), 1059-1067. https://doi.org/10.1016/j.chb.2010.05.032
Serholt, S. & Barendregt, W. (2016). Robots tutoring children: Longitudinal evaluation of social engagement in child-robot interaction. *Proceedings of the 9th Nordic Conference on Human-Computer Interaction, 6*(1), 1-10. https://doi.org/10.1145/2971485.2971536
Scheer, A., Noweski, C., & Meinel, C. (2012). Transforming constructivist learning into action: Design thinking in education. *Design and Technology Education: An International Journal, 17*(3), 8-19.
Scheutz, M., DeLoach, S.A. & Adams, J.A. (2017). A Framework for Developing and Using Shared Mental Models in Human-Agent Teams. *Journal of Cognitive Engineering and Decision Making, 11*(3), 203–224. https://doi.org/10.1177/1555343416682891
Schnaubert, L., & Bodemer, D. (2019). Providing different types of group awareness information to guide collaborative learning. *International Journal of Computer-Supported Collaborative Learning, 14*(1), 7-51.
Serholt, S. & Barendregt, W. (2016). Robots tutoring children: Longitudinal evaluation of social engagement in child-robot interaction. *Proceedings of the 9th Nordic Conference on Human-Computer Interaction, 6*(1), 1-10. https://doi.org/10.1145/2971485.2971536
van Zoelen, E. M., van den Bosch, K., & Neerincx, M. (2021). Becoming team members: Identifying interaction patterns of mutual adaptation for human-robot co-learning. *Frontiers in Robotics and AI, 8*(1). https://doi.org/10.3389/frobt.2021.692811
Walsh, S.E. & Feigh, K.M. (2021). Differentiating 'human in the loop' decision strategies. In *IEEE Conference on Systems Man and Cybernetics*. Melbourn, Australia: IEEE.
Wright, J. L., Chen, J. Y., & Lakhmani, S. G. (2019). Agent transparency and reliability in human–robot interaction: The influence on user confidence and perceived reliability. *IEEE Transactions on Human-Machine Systems, 50*(3), 254-263.
Xu, Y., Ohmoto, Y., Okada, S., Ueda, K., Komatsu, T., Okadome, T., Kamei, K., Sumi, Y., & Nishida, T. (2012). Formation conditions of mutual adaptation in human-agent collaborative interaction. *Applied Intelligence, 36*(1), 208-228.
Yu, K., Berkovsky, S., Taib, R., Zhou, J., & Chen, F. (2019, March). Do I trust my machine teammate? An investigation from perception to decision. *Proceedings of the 24th International Conference on Intelligent User Interfaces*, *24*(1), 460-468. https://doi.org/10.1145/3301275.3302277
Zufferey, J.D., Bodemer, D., Buder, J., & Hesse, F. W. (2010). Partner knowledge awareness in knowledge communication: Learning by adapting to the partner. *The Journal of Experimental Education, 79*(1), 102–125. https://doi.org/10.1080/00220970903292991.