



APPLIED RESEARCH LABORATORY FOR
**INTELLIGENCE
AND SECURITY**

There is No “AI” in Teams: A Multidisciplinary Framework for AIs to Work in Human Teams

Breana M. Carter-Browne^{*†}, Susannah B. F. Paletz[†], Susan G. Campbell^{*†}, Melissa J.
Carraway^{*†}, Sarah H. Vahlkamp[†], Jana Schwartz^{*}, Polly O’Rourke^{*}

^{*}Applied Laboratory for Intelligence and Security (ARLIS), University of Maryland
College Park

[†]College of Information Studies, University of Maryland College Park

Applied Research Laboratory for Intelligence and Security (ARLIS) Report
June 2021

Abstract

Many AIs are created to solve problems that humans alone cannot solve. AIs designed to work in human teams should include requirements necessary to work within a team. This paper aims to describe the high-level requirements and issues for an AI to perform in an AI-enabled, human team successfully. We define what it means to be in a team, the context and tasks that impact teamwork, and propose a multidisciplinary model of teamwork for AI-enabled teams. Our model uses a novel input-process-emergent state-output-input (IPEOI) model based on the input-moderator-output-input (IMOI) model. We include four levels in our model. In addition to the individual, team, and organizational levels, we split the individual between the human and the AI. Our model draws on a diverse literature to enable a comprehensive understanding of the requirements from all players, including the AI, in a teaming system. Current AI systems may not be advanced enough to implement some of these team-related tasks, but we propose some ways that future systems may work in human teams.

Table of contents

Abstract	2
Table of contents	3
Introduction	5
AI in Teams Use Case	5
What is AI?	8
What is a Team?	9
Can AI be a Team Member?	10
What is In/Out of Scope for Discussing AI in a Team?	13
Multidisciplinary Model of Teamwork for AI	13
Model explanation	14
Task and Mission	16
Task Types and Types of Teams	16
Job Design	17
Interdependence	17
Other Team and Task Characteristics	18
Task Types Summary	18
Time: Temporal Factors	18
Individual: Human	19
Inputs	19
Individual Human Processes	27
Emergent States	28
Individual Human Outputs	37
Summary of Individual Human IPEOs	38
Individual: AI	39
AI Inputs	39
AI Processes	43
Emergent States	46
AI Outputs	47

Summary of Individual AI IPEOs	48
Team Level	48
Team Inputs	48
Team Processes	52
Emergent States	55
Team Outputs	60
Summary of Team IPEOs	61
Organization and Societal Levels	61
Inputs	62
Processes	65
Emergent States	66
Outputs	67
Summary of Organization and Society IPEOs	67
Discussion	68
Solutions to Past AI Failures	68
Future Steps: Metrics	69
Conclusion	70
References	71
Acknowledgements	90

Introduction

I said, "Your second mother is..." Client wasn't the right word, not anymore. "My teammate." I could see I had to clarify. It was really hard finding the right words. "Before your second mother, I had never been an actual member of a team before. Just an..."

Amena finished, "An appliance for a team."

That was it. "Yes."

- *Network Effect* by Martha Wells

AI in Teams Use Case

A major storm event has just ripped through the country. Unexpectedly strong, there were minimal evacuation preparations taken. Heavy winds and flooding rains have damaged the power, communications, and transportation infrastructure. Now, there are a host of victims to rescue to safety. Some folks are disadvantaged but stable; others are in critical medical condition and need prioritized attention.

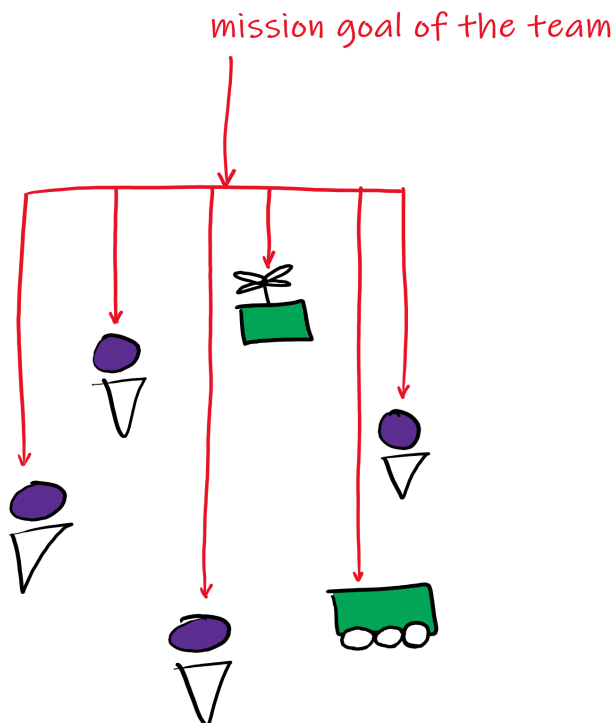


Figure 1: The mission goal of the team is distributed across all team members

Part of the team tasked for this disaster relief effort is shown in Figure 1. This small ground unit of human (purple circles) and robot (green squares) members is tasked to achieve a specific set of mission goals (the red arrows) -- to search an area for victims, gather information about the status of each victim, make on-the-spot decisions as to any victims requiring immediate medivac, and communicate all status and metadata, such as timestamps and location, back to the command center. The team must work collectively to achieve these goals, requiring communication about the task and mission. Taskwork-related communication is represented by the orange communication links (representing literal, and time-varying exchanges of information) in Figure 2. The team must also work to achieve the effort as a team, which may include communication of information that is not related to the task or mission such as mediation of intrateam conflicts or updates about personal well-being. Team related communication is shown as the blue links in Figure 2. This team is composed of a driver, a navigator and note-taker, a medical first responder, an AI embodied in a drone, equipped with thermal sensors to seek out victims, and an AI embodied in a ground vehicle equipped with cameras and vibration sensors to detect abnormal movement.



Figure 2: Team members communicate and coordinate both taskwork (orange links) and teamwork (blue links) related activities in support of the mission

Within a multiteam system, as in Figure 3, there are also many on-the-ground teams responding to this relief effort, each with their own delegated mission goals. An on-the-ground team may be in charge of determining if residual disasters may occur (e.g., tsunami following an earthquake). This team is composed of operational planners who debate how to allocate resources, as well as an AI optimization system (displayed as an information dashboard and decision support tool) that simulates possibilities with projections of probable future destruction.




Figure 3: A multiteam system of several human-AI teams

The natural disaster situation outlines two examples of teams that consist of both humans and AIs working collaboratively, and these examples are a glimpse into the future of AI-enabled teams. Both teams consist of humans working, communicating, and depending on AIs. For AIs to be productive teammates, we must explore what elements are necessary for successful teaming and the role that an AI can play in aiding the team to achieve the overarching goal or mission.

There have been recent calls to “build a science of cooperative AI” (Dafoe et al., 2021). The literature identifies that when humans collaborate with an AI, or when humans and AIs work collectively as a team, better results are achieved than when humans or AIs try to solve the same problem on their own (for example centaur chess, a combination of human and AIs working together produces a better player than either alone, Cassidy, 2014). The design of AIs meant to work in human teams should thus include the requirements necessary to work in a team.

Neglecting to design AIs with human interaction in mind has resulted in several failures. Humans are less likely to adopt the new technology, especially if there are human requirements the technology cannot meet (Parasuraman & Riley, 1997). New technology is often not used in the way it was intended (e.g., abuse of facial recognition software, Garvie, 2017). In addition, new technology may not function according to its original design (e.g., when AIs fail to complete routine system updates, such as the patriot missiles, function is impaired, US General Accounting Office (GAO), 1992). Failures also occur when humans have implicit expectations of the AI that are not made into explicit requirements (e.g., Yorktown Smart Ship failures, Slabodkin, 1998). Not only should design focus on making AI safe, easy to use, reliable, and

trustworthy (Shneiderman, 2020a), it should also focus on how these factors will impact the human counterpart.

The field of distributed cognition has long considered human-system information processing and how sociotechnical systems include humans and machines (Hutchins, 1995). While conceptualizing AI as a member of a team is gaining traction, determining what it means for AI to be a team member within a human-AI system is underrepresented in the literature (one of the few include You & Robert, 2017). Designing an AI is an often lengthy and complex process (Johnson & Vera, 2019). The qualities and specifications of AI are therefore typically tested one or a few at a time. Similarly, in teaming research, only a small set of variables are examined in a given experiment and only some of the important factors for human functional teaming have meta-analyses (e.g., the impact of information sharing on team performance, DeChurch & Mesmer-Magnus, 2009). However, to understand the requirements for an AI to work in a human team, a more comprehensive view is necessary. It is not sufficient to say an AI should be safe, easy to use, durable, explainable and trustworthy (which includes ethical considerations for use and misuse): we aim to articulate the implied requirements of what these mean. Our approach is integrative and multidisciplinary, drawing across different psychology disciplines (cognitive, social, organizational, and developmental) as well as organizational behavior, management, artificial intelligence, human-computer interaction, human performance and sensing, communication, cognitive science, and social robotics.

Our purpose is to provide a multi-disciplinary model for AI-enabled teams, outlining what it means for AI to work in a team as well as supplying information that can assist engineers in the design of AI intended to collaborate with human teammates. First, we clarify what we mean by AI and Team, setting parameters for what is and is not in scope for this paper. This section is followed by a high-level explanation of our model, followed by detailed sections on each element of the model. Finally, we synthesize the parameters necessary for AI to be a successful teammate and discuss the future of human-AI teaming.

What is AI?

AIs are used and created to aid in solving problems that are not easily solved by human labor, to reduce the cost of labor, and generate solutions in less time. AIs can also be created to perform tasks that humans simply cannot do, such as sort through huge amounts of data, uncover underlying patterns (including biased patterns), and process disparate or unstructured data. AIs are increasingly being developed to achieve higher levels of autonomy, thus increasingly eliminating the burden of routine tasks on human workers.

There are many definitions for, and approaches to, building AI. The term artificial intelligence was coined in 1955 by Stanford Professor John McCarthy, who defined AI as “the science and engineering of making intelligent machines, especially intelligent computer programs”

(McCarthy, 2007, ans. 1). A clear definition of an intelligent machine requires a clear definition of intelligence. Russell (2019) argues that the intelligence of an agent depends on the “extent that their actions can be expected to achieve their objectives” (p. 9). Others, such as the Department of Defense 2018 AI Strategy, posit that AIs must have “the ability...to perform tasks that normally require human intelligence—for example, recognizing patterns, learning from experience, drawing conclusions, making predictions, or taking action—whether digitally or as the smart software behind autonomous physical systems” (US Department of Defense, 2019). We can surmise then, that in order to be considered an AI, and not some other type of computer system, tool, or automation, these systems are expected to engage in processes that mimic or surpass human-like processes in a manner that allows for successful completion of a goal.

Examples of intelligent agents, outlined preliminarily by the Secretary of Defense include:

- (1) Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from its experience and improve performance when exposed to data sets.
- (2) An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication or physical action.
- (3) An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
- (4) A set of techniques, including machine learning, that is designed to approximate a cognitive task.
- (5) An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision-making, and acting (John S. McCain National Defense Authorization Act for Fiscal Year 2019, H.R. 5515 § 237, 2019).

Given all these perspectives, we define AI in this paper as a machine agent designed with the ability to achieve a goal, given often imperfect and incomplete information inherent to dynamic environments. This agent can be embodied, such as a robot, or disembodied, resembling software. This agent can work independent of human inputs, but also has the ability to interact and communicate with others as appropriate; it is not simply an algorithm in the sense of a pre-scripted set of instructions. Finally, this agent has the ability to learn from its behaviors, past performance, and the changing environment.

What is a Team?

A team is a type of group that is (1) bounded (has people inside and outside of it), (2) interdependent, and (3) has differentiated member roles (Guzzo & Dickson, 1996; Hackman, 2012). In addition, teams (4) are made up of at least three members¹ who (5) interact and (6)

¹ While some definitions of team argue that two humans are sufficient, we would argue that teams require three or more entities, of which at least two need to be humans. Dyadic processes can and do occur within teams, but the network of interactions within a team qualitatively changes when the third person is included (Bienefeld, 2020).

have shared goals (Kozlowski & Ilgen, 2006), whether they contribute equally or not. For example, teams could include:

- a group of five undergraduates from mechanical, electrical, and bioengineering working together on a design project;
- sports teams such as a baseball team;
- a theater troupe, including the stage manager, costumers, and technical members who run the lights;
- a project group building a house;
- astronaut teams;
- mission control at Johnson Space Center;
- military teams; and
- the Mars Exploration Rover science team, which was really two interacting large teams, each composed of several disciplinary and instrument sub-teams.

Thus, regardless of whether the work is paid or unpaid, teams strive together toward a common purpose. Team structures can vary in terms of leadership, levels of hierarchy, and so on (see later sections for more detail). A group of people who work toward the same goal but never interact, or who interact but only in competition and thus have incompatible goals, would not be a team.

Decades of research from the organizational and social psychological literature provides insights into what humans need or do to work with other humans (e.g., for reviews, see Driskell et al., 2018; Kozlowski & Ilgen, 2006; Salas et al., 2005). This paper covers these numerous insights in the following sections. Many of the qualities identified in research on human teams are also critical for teams consisting of humans and AI agents. This existing body of literature, however, understandably presumes human-human teams. Human-AI teams add a layer of complexity and additional requirements (Bienefeld et al., 2020).

Can AI be a Team Member?

While AI agents can provide support to humans in teaming contexts, there is robust debate about whether an AI can truly be considered part of a team. While some researchers welcome discussion of AI-human cooperation (and AI-AI cooperation and AIs for human collaboration, Dafoe et al., 2021), Groom and Nass (2007) argue that teaming is an inherently human activity, and that while animals and machines may assist human teams, they cannot be considered members. We reject this is a circular argument wherein defining teammates as necessarily human, AIs cannot then be teammates. We disagree and do not presume that only humans can be teammates. More specifically, researchers argue that AIs cannot be team members because: (1) it is risky for humans to anthropomorphize AIs and that AIs should not present themselves as human, (2) AIs cannot engage in trust processes, and (3) AIs lack the flexibility and responsiveness for teaming (Groom & Nass, 2007; Shneiderman, 2020b). We address each of these points in turn.

First, some researchers caution against anthropomorphizing computers as they believe it to be counterproductive and that computers should not be designed to represent themselves as humans

because doing so would be misleading (Don et al., 1992). However, people ascribe human traits to non-human entities that seem to have agency (Kwan & Fiske, 2008). Because of this tendency to anthropomorphize, Groom and Nass (2007) contend that humans working with AI will inevitably seek some level of "humanness" in their AI teammates (e.g., have shared mental models, see later section). Achieving something close to "humanness" may never be possible, thus leading to a violation of expectations and a failure to accept the AI as a teammate (Haslam, 2006). The likelihood of this failure may increase for those aspects of human teaming that are not easy to observe and are therefore difficult to include in the design of AI. We acknowledge that AIs will lack certain aspects of "humanness" and argue that, within design limitations, AIs can still be teammates, and that appropriately setting expectations with users is key for acceptance. We agree with Susan Brennan that AIs should represent themselves as such and not as humans: "We should stop worrying about anthropomorphism and work on making systems capable of behaving as coherent interactive partners. Whether these partners or anthropomorphized or not, they should present their limitations frankly" (Don et al., 1992, p. 68). Thus, AIs need not be anthropomorphized, nor need to pretend to be humans, to be teammates.

Second, some researchers contend that AIs cannot fully engage in trust assessment processes with humans, which they argue is an inherently human process that is necessary for teamwork (Groom & Nass, 2007). In a recent extensive review of trust and AI, Glikson and Williams (2020) use Mayer et al.'s (1995) definition, which describes trust 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 et al., 1995, p. 712). Groom and Nass (2007) argue that trust also requires an understanding that both parties must be impacted if trust is lost and assert this as a capability that AIs cannot possess. Glikson and Woolley (2020), however, argue that the definition of trust can be easily extended to the human-AI relationship, given the elements of willingness to take meaningful risks and expectations for a positive outcome.

Because trust is a key issue that emerges at the individual human and AI levels, as well as at the team, organizational, and societal levels, we discuss issues of trust throughout this paper. In human teams, we would discuss mutual trust processes. In the case of hybrid human and AI teams, we distinguish between the trust a human has toward other humans and AIs and the trust the AI has toward other AIs and humans. Relevant to the argument that an AI can be a teammate is whether it is possible for an AI to trust. In the literature, trust is described as having both emotional (affect-based) and cognitive aspects (Glikson & Woolley, 2020; McAllister, 1995). Affect-based trust involves interpersonal concern, investments in relationships, and emotional ties, whereas cognition-based trust revolves around expectations about the target's competence, dependability, and reliability (McAllister, 1995). Trust is predicated on an understanding of the past and present (observability) as well as an expectation of future behavior and competence (predictability; Johnson & Vera, 2019). We agree with the implied argument that AIs cannot, at

this time, have emotional trust. However, it is possible and likely that an AI can be designed to model and estimate the state (current/temporary) or trait (consistent across situations) reliability, dependability, and competence--or cognition-based trustworthiness-- of a human or other AI teammate. For example, AIs designed to observe the work completed by teammates and their overall progress towards achieving a goal would also be able to maintain information about an individual team member's performance and flag aberrant behavior or that which is counter to the shared objective. Moreover, an AI could track and quantify individual competencies of teammates, and the resultant score would allow the AI to predict the likelihood of the teammate's ability to complete the next task. While in humans, there may be interactions between emotion-based and cognition-based trust (McAllister, 1995), this process would not occur in AIs.

We contend that AIs are also able to be vulnerable. While AI vulnerability would not involve the emotional aspects, AIs would likely still need to accept risk regarding teammates fulfilling their roles within the team. Imagine a scenario where an AI takes input from a human or another AI and cannot accomplish its objectives without that input. In depending on those other teammates--in their activities and outputs being dependencies and inputs for the AI--the AI can be vulnerable to the actions of these others, especially if it cannot be constantly monitoring.²

Third, some researchers contend that AIs currently lack sufficient flexibility required to be effective teammates (Shneiderman, 2020b). This inflexibility, tied to the first issue of lacking "humanness", is only an issue if we expect AIs to think and learn like a human. Lake and colleagues (2016) argue that humans have more *model-based* thinking than AIs, which are designed to assess the world using statistical *pattern recognition*. *Model-based* thinking allows for explanations based on the building of a causal model of the world combined with incoming data and affords the flexibility to learn a variety of skills and possess knowledge in a variety of areas. *Pattern-recognition*, however complicated, relies on categorizing incoming information based on training data, and limits the ability to process unrecognized information (Lake et al., 2016). While the explanatory style of thinking allows for faster and more accurate processing of information and flexibility, it is not necessary for a teammate to possess. Successful completion of a specific task only requires flexibility within the confines of that task. AIs designed to collaborate with humans in a teaming context must be able to learn from their past performance, the environment in which they are working, as well as their interactions with human and AI teammates (discussed further in AI Inputs). This level of flexibility will allow the AI to appropriately aid in the completion of the team missions and effectively engage with teammates in a manner that will bolster teamwork.

² Answering the age-old question (song) of "what do you do with a drunken sailor," an AI teammate's answer would be that it would detect decrements in the sailor's performance and physiological state past a certain threshold and respond in a pre-programmed manner. If the AI is not constantly monitoring, the drunken sailor may cause downstream decrements to the AI's outputs, which is where the vulnerability aspect of trust would come through.

Will all AIs be capable of being teammates in human-AI teams? No. Some AIs will simply be tools, but other AIs can be designed to be teammates. Porter and colleagues (2020) recommend that a teammate, either human or AI, “must (1) be able to influence each other’s problem state; (2) be working toward a common higher-level goal, and (3) coordinate actions or decisions” (p. 3-67). This definition explicitly outlines factors including shared goals, interdependence, and required interaction mentioned in our original definition of a team. We concur and assert that AIs can wholly fulfill the qualities of a team member as we have set out in this paper: share goals, engage interdependently and interactively, be perceived to be in a team (bounded), and have differentiated roles. While AI agents will not replace humans in a teaming context, they can be designed with the functional capabilities that will allow them to collaborate with human teammates. Capabilities such as trust and flexibility can be designed into the AI. An AI can be designed to not only monitor human teammates, but be an encouraging teammate, asking in a non-threatening way if the teammate in trouble needs assistance.

What is In/Out of Scope for Discussing AI in a Team?

Though some authors (e.g., Groom & Nass, 2007), have questioned whether an AI ever could or should serve as part of a human team, we take as a given that at least one advanced AI will at some point be placed into a team structure as a teammate. The exploration that follows attempts to determine what conditions need to be met for an AI to be a team member, and what problems will need to be addressed before such teaming will be functional. By making our assumptions explicit, such as our definitions of teams and AIs, we hope to navigate the complexities that have led this field to have seemingly inconsistent recommendations (e.g., Klein et al., 2004; Russell, 2019; Shneiderman, 2020b).

Teams may have a variety of structures and features, and AIs may have a variety of characteristics. For the purposes of this paper, we focus our efforts on AI systems that are explicitly designated and treated as team members in primarily human teams. These systems may serve different purposes and may have different characteristics, such as being physically embodied or only virtually presented. We will explore these characteristics later, but some characteristics, such as an inability to have an independent view of the team goals or influence the problem states of human teammates, would preclude an AI system from being part of a team.

Human-AI team structures

For the purposes of this paper, we define an AI-enabled team, at the most basic level, as one that includes at least two humans and at least one AI capable of communicating with more than one human member. Prior research finds differences between robots interacting with individual humans versus interacting with groups, such that a recent review also required at least two human teammates and one robot (Sebo et al., 2020). Not all colloquial examples of AI-enabled teams meet our criteria. For example, current swarm structures, which include a single human

operator and a collection of semi-autonomous robots, do not qualify as teams because they do not include multiple humans, meaning there is no collaboration or other critical team-based interaction. The definition of teams from the human team literature might imply that a swarm and human could count as a team if the human is perceived to be within the group (see previous definition). However, for our purposes, we are interested in AIs engaging in human teams, so we require at least two humans in an AI-enabled team. This configuration would require that the AI engages with at least two humans to make a three-entity team. For example, the possible future of financial day trading entails each human with their own financial optimizer, and the optimizer bots negotiate to create the maximum value for all parties. While a compelling scenario for future research on AI system interoperability, this scenario fails to keep the focus on how AIs would work in human teams; the complicated interaction is among the AIs, but the AIs need not navigate human team dynamics.

Similarly, AIs that communicate with only one human lack true team interaction, and so are not in a team according to our definition. The literature on dyadic interactions fails to explain the dynamic intrateam interactions that occur when there are at least two humans interacting with an AI (e.g., in trust processes; Bienefeld et al., 2020). An AI that communicates with only one human is either in a dyad with that human or is serving as a personal assistant to that single human, regardless of whether the human is in a human-human team.

We suggest two main configurations for AI-enabled teams: a team that consists of one AI assigned to an entire human team, and a team that consists of multiple AIs interacting with multiple human team members. Our examples in the beginning of this report illustrate these two configurations. The operations and planner center, on-the-ground, team in our use case scenario exemplifies the first AI-enabled team configuration with a one-to-many relationship between the AI and the humans. This AI is updated by individual human team members and provides updates to every human member of the team. Our ground rescue team exemplifies the second configuration such that there is a many-to-many relationship between the AIs and the human team members. In this configuration, each AI is communicating with more than one human and it is possible that the various AIs occupy different roles on the team.

Level of AI autonomy

A technology tool that responds to the commands of a human on the team, like an Alexa smart device, is a shared resource, but it is not a team member. Although the human team literature does not find it necessary to include certain implied, universal characteristics of humans in the definition of a team, some degree of agency, autonomy, or ability to plan is presumed and required for a non-human entity to be considered a teammate (Klein et al., 2004; Russell, 2019). Thus, the AI must have some ability to act independently of being given only explicit orders by a human. Communication or transparency of behaviors and state (see more on observability in Individual AI), understanding of goals and objectives, and projections of situation awareness to

include second and third order effects of actions will be required of the AI to mimic those expert negotiations that happen between human operators. Similarly, as conditions evolve and needs change, these systems must behave as humans do by being flexible and adaptable with regards to goals. That is to say, that when a certain goal becomes unachievable or less relevant to meeting top-level objectives, the goal must be thoughtfully re-evaluated and communicated with the team. Thus, AIs designed to work with humans should have the capability to assess orders against the goals of the team and against both self-preservation and the preservation of each team member. To fulfill this function, AI teammates must have the capacity to determine if there is a seeming conflict between explicit orders, new information, or required behaviors on the one hand, with the original goals and high-level objectives, such as human safety, on the other hand. This detected conflict must be communicated to human operators to allow for adaptations to the plan and to ensure the AI never goes against direct orders from a human or collective goals. For example, in our rescue team, the AI embodied in the ground vehicle must have the ability to inform human teammates that searching in location A may not be ideal if it senses a higher likelihood of human activity from location B. Thus, if the AI does not have an independent and shareable view of team goals, it is not a team member.

Teamwork vs. taskwork

The work of teams includes at least two components: fulfilling the goal of the team (taskwork) and making the team work as a unit (teamwork). Teamwork behaviors include coordination and conflict (see Team Processes) and contribute to team performance. For the purpose of limiting the scope of this review so as not to cover every possible use case for an AI in a team, this document specifically focuses on teamwork behaviors rather than taskwork behaviors. How tasks are allocated between humans and AIs is its own area of research (e.g., Tausch et al., 2020). Taskwork behaviors are generally part of the design parameters of any AI (regardless of its performance as a teammate). In other words, we do not delve into the functional requirements of any particular AI or how to make judgments about them, such as being able to triage information from large collections, drive autonomously over uneven terrain, or search for life signs. Rather, we focus on how the AI participates in teamwork.

Multidisciplinary Model of Teamwork for AI

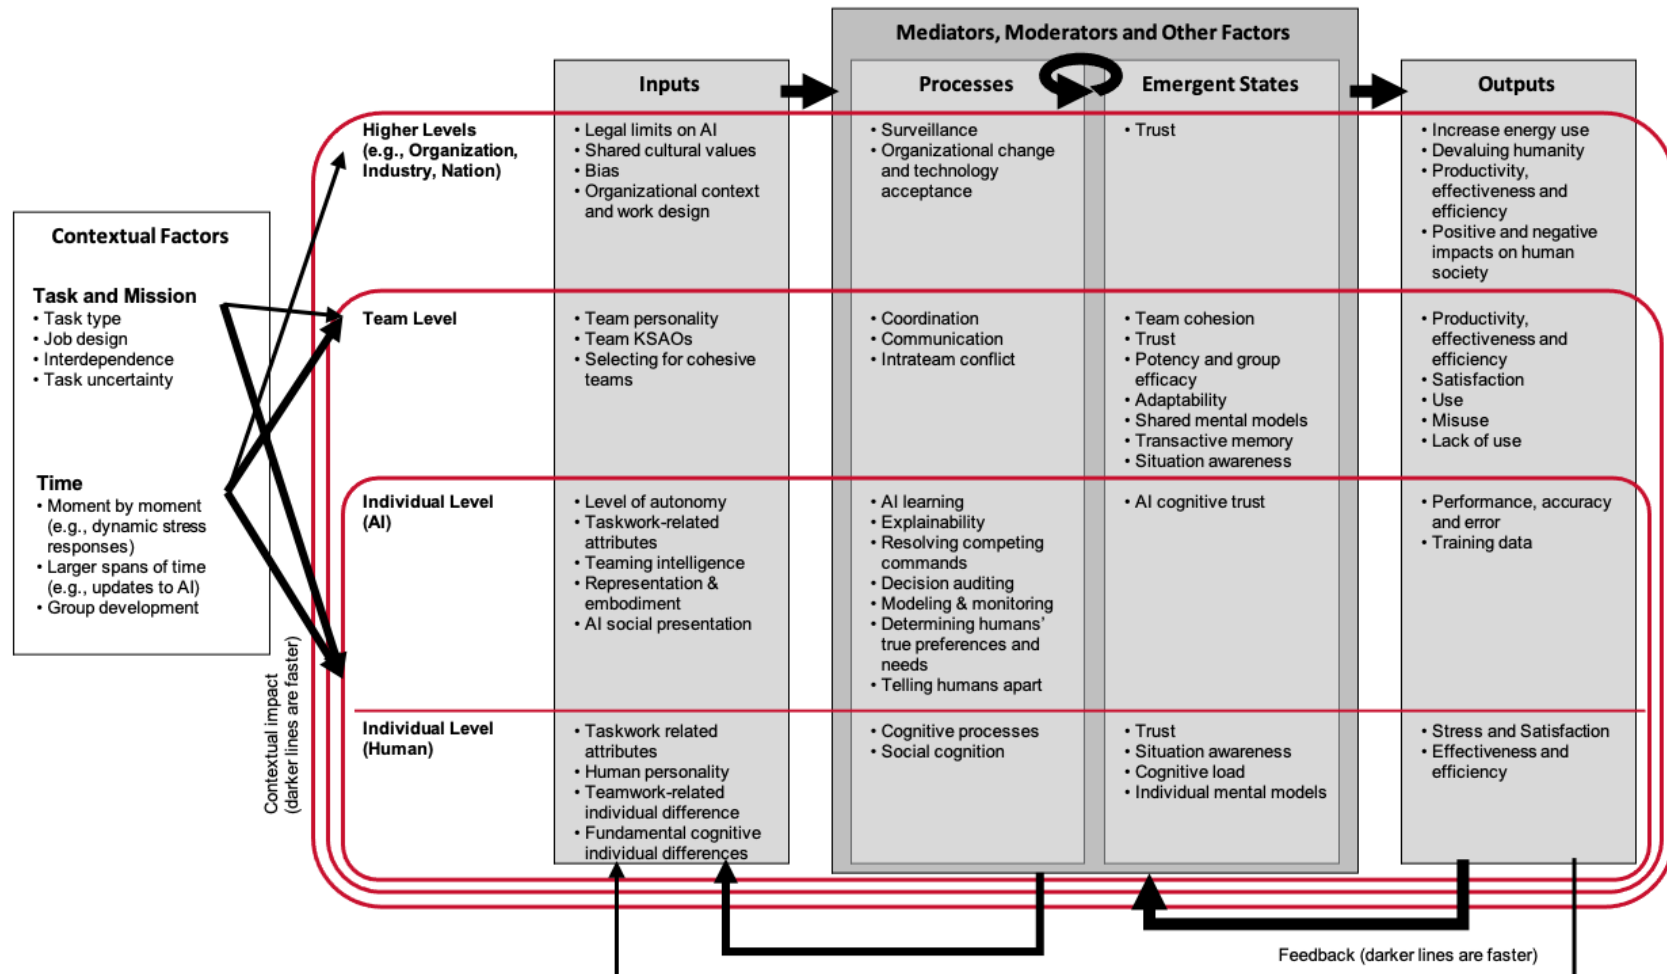


Figure 4: Multidisciplinary model of teamwork for human-AI teams

Model Explanation

The key motivation for developing this model is the prospect of AI-enabled teams supporting human teams effectively. Perhaps the AI can free up humans to focus on tasks humans excel at, alleviating some of the burden of mundane tasks. Perhaps, AIs can do work that humans are not as efficient or accurate at accomplishing or are incapable of, such as vigilance tasks and processing large amounts of data quickly or consistently (i.e., eliminating transposition errors).

To provide a framework to address these questions, we propose a model of teamwork for AI (Figure 4). One prominent approach in the human teams literature, which we adapt, is the input-mediator (process)-output-input (IMOI) model (Ilgen, Hollenbeck, Johnson, & Jundt, 2005; Kozlowski & Ilgen, 2006; also referred to as IPO models). Our model might be better considered an IPEOI model, or an input-process-emergent state-output-input model. First, we describe the *inputs* to the team, such as individual skills, personality, and task constraints. Inputs are what the team begins with. Next, we describe team *processes*, which are dynamic interactions over time, such as communication and conflict (Kozlowski, 2015). Processes are often presumed to be the mediators between inputs and outputs (as with IMOI models), such as if coordination and communication are how a particular set of skills and knowledge across team members results in team success. However, processes need not always be mediators between inputs and outputs (Kozlowski, 2015). We also follow the lead of other teams researchers who distinguish between team processes, which are interactions that occur over time, and *emergent states* (e.g., Marks et al., 2001; Kozlowski et al., 2016). Emergent states are “constructs that characterize properties of the team that are typically dynamic in nature and vary as a function of team context, inputs, processes, and outcomes” (Marks et al., 2001). Emergence is when a higher-level phenomenon comes into being because of interactions at a lower level (Cronin et al., 2011). An example of an emergent state would be team cohesion. Team cohesion involves team members sharing a commitment and attraction to the goals of the team and to each other (Braun et al., 2020). It is the result of team interactions, individual perceptions, and team outcomes, but it can also influence subsequent interactions and team outputs (Braun et al., 2020). Team *outputs* include performance and satisfaction. The outputs can be at the end of the life of the team (e.g., a final tabulation and mapping of rescues, at the end of the emergency) but also occur at points throughout the life of the team (e.g., planning the next search, finding the next victim, assessing their state, and so on). We also include an important aspect of IMOI models, a *feedback loop*, such that satisfaction becomes an input later in the life of the team.

Our model includes three other important features which are also included in the teams literature: task and mission factors, time factors, and a multilevel structure. Teams, and factors influencing teams, are inherently multilevel, meaning that they have a nested structure. Individuals are nested within teams which then exist within larger entities such as organizations and nations (Rousseau, 1985; Klein et al., 1994; Paletz et al., 2018). We therefore include four levels in the model: In

addition to the individual, team, and organization, we split the individual between the human and the AI. Our model reflects the nesting of these levels by having the most encompassing level, organization and society, at the top, with progressively lower levels lower on the model. Each level has unique and overlapping IPEOs and are interrelated such that, for example, organizational inputs can alter individual inputs. For instance, a lack of trust in AI at an organizational level may lead to miscommunication about the ideal use of the AI and alter individual human trust in their AI teammates. Likewise, individual personality inputs alter the overall team personality and combine to influence team processes and outputs. In conceiving of teams as nested within a multilevel structure, researchers have made several insights that are relevant here. One is the idea of emergent states, such that many team features are more than a sum of individual parts (e.g., Marks et al., 2001). Emergent phenomena are dynamic and can go across all the levels mentioned (Kozlowski et al., 2016), although we call them out at each relevant level. Another insight is that there may be complex effects across levels, such as that individuals may be impacted by their team and their organization and may impact those levels in return. There can also be comparative processes, as when an individual compares herself to her team, thus having a different experience in different types of teams (House et al., 1995).

Contextual factors, such as the task and mission and temporal factors, impact the context of how inputs, processes and emergent states, and outputs interact (Ilgen et al., 2005; Edmondson & Harvey, 2018). Task and mission factors are usually inputs, but they cross the four levels in our model (see Figure 4). The task and mission consist of the type of tasks members of the team are working on, the makeup and roles of the team, job design, and task interdependence. Temporal factors include moment-to-moment changes such as the dynamic changes in stress levels of a human throughout a task, changes that occur over larger spans of time such as updates to the AI, and group development. Altering any of the contextual factors can reshape the impact a given input may have on a successful output, and the ease with which processes occur and states emerge.

In sum, we have taken the typical IMOI model and explicitly broken mediators into processes and emergent states (which may or may not be mediators) and the individual level into human and AI. We have also included task and mission inputs as a separate category that stands separately from the four levels, as it can be generated by any level. Finally, we explicitly include temporal factors. This model leverages the best of the human-human teams literature while making room for important AI and human-AI factors.

Our model, and the resulting review, is a thorough framework for AI-enabled teams. We used a multidisciplinary approach that spans cognitive and social psychology, cognitive neuroscience, human-computer interaction (HCI), engineering, management, and more. This approach allows for a holistic view of the problem in a manner that will ideally result in a richer understanding of what AI can do to be a successful teammate. We are not the first to envision a model for human-

AI teaming (e.g., You & Robert, 2017), however, our model will be one of the most comprehensive. For instance, You and Robert (2017) proposed an IMOI model of human-robot teams involving human and robot inputs, but our model focuses on AIs, not just robots, and includes a broader range of inputs, emergent states, processes, and outputs. In the next sections, we will discuss the task and mission, and then go through the various levels (individual human, individual AI, team, and organizational and societal), and then discuss temporal issues.

Task and Mission

Different elements of the task and mission both constrain what humans and AIs can do and dictate what they should do, in terms of goals. This section briefly sketches some of the commonly discussed environmental factors that influence how teams are formed, how they operate, and how their success is defined. This section presumes that an AI should fulfill its task-relevant functions, which is an important set of requirements that are outside the scope of this project. Depending on what occurs these elements may be changed during the life of a team or for the next version of the team. For instance, a team may be required to or may choose to go from face-to-face to virtual, changing midstream or for the next project.

Some of the earliest research on teamwork has defined different underlying elements of the task or mission into different categories: task types and types of teams, job design, interdependence, and some other task and team design characteristics.

Task Types and Types of Teams

McGrath (1984) created a task circumplex which divides group tasks into two dimensions, from conflict to cooperation and from conceptual to behavioral. These dimensions are then arrayed in a circumplex into four quadrants which include eight tasks, total. The four quadrants are generative (creative) tasks, executing tasks, negotiation tasks, and choice tasks. For example, executive performance tasks include psychomotor tasks such as building a car in an assembly line (without robot assistance), decision-making tasks involve deciding issues with no right answer, and mixed-motive tasks involve resolving conflicts of interest. Any group can engage in these tasks, and likely more than one over the life of the team. The type of tasks the team is to work on will, or should, dictate some of the staffing and compositional decisions made for the team (Klimoski & Jones, 1995). A team that is building cars on an assembly line has different requirements for both teamwork and taskwork than a team that is co-designing robots that would then build cars on an assembly line. Although the research from the teams literature strives to make abstract the needs of all teams, how a team focuses on different aspects of the work such as monitoring, coordinating, and planning will differ depending on the task types (Honts et al., 2012).

There are also a variety of types of teamwork design characteristics, such as virtual/distributed versus face-to-face/collocated teams (Hertel et al., 2005; Hinds & Kiesler, 2002). Research on face-to-face versus distributed teams suggests that computer-mediated teamwork may be more or less successful for different types of tasks, but that many types of tasks are understudied (Hertel et al., 2005). In general, distributed teams may have more conflict (Hinds & Bailey, 2003). Teams can be focused on different tasks as a whole above and beyond the McGrath typology, such as leadership teams (Oldham & Hackman, 2010) or executive teams versus command-and-control teams, customer service teams, and technical teams (Klimoski & Jones, 1995). Thus, although this review does not cover the specific taskwork requirements for an AI working in human teams, the tasks will influence how teams are created, their processes, and their outputs.

Job Design

Another element of the task is job design. Jobs are made up of multiple tasks, whether they involve a team or not. Job design has been studied for over forty years as a set of core characteristics on which different jobs vary (e.g., Hackman & Oldham, 1975; Oldham & Hackman, 2010): task identity or differentiation (how much the job involves doing an identifiably whole piece of work), skill variety (how much the job requires different activities and skills), task significance (the impact of the task on others), job-based feedback, and autonomy (how much the job allows discretion and control in how to go about doing it). The original theory proposed that people would have more motivation from jobs with more of each of these characteristics, although it was later refined to include individual differences of job-relevant skill and knowledge and the degree to which individuals desired personal growth at work (Oldham & Hackman, 2010). However, as the originators of the theory note, the world of work itself has and is changing since the 1980s, and other factors such as social attributes of jobs and organizational context are also important (Oldham & Hackman, 2010). Translating this theory into work groups, Campion and colleagues (1993, 1996) reframed autonomy as self-management and the degree to which team members participate in decisions. These different elements of job design can be related to task complexity, such as if a job requires a lot of different tasks.

Interdependence

Task interdependence, another aspect of group task design, is how much team members have to interact with and are dependent on each other for the task to succeed (Campion et al., 1993; Hertel et al., 2005). There are different kinds of task interdependence, such as sequential (one at a time) and reciprocal interdependence, as well as goal interdependence and outcome interdependence, such that rewards and feedback are contingent on different team members' work (Campion et al., 1993). Task interdependence is a requirement that dictates how much coordination is necessary to occur (see Team Processes). For example, high task interdependence is necessary in the early stages of virtual teamwork to set up and make explicit team processes

(Hertel et al., 2005). Task interdependence can mean the difference between whether a particular variable is important to team effectiveness or not. For example, cohesion is related to team performance, but mainly in teams with high task interdependence (Gully et al., 2012). The degree to which humans are working on an interdependent task and the AI is interdependent with those humans will similarly be a constraint and a requirement for how and what both AIs and humans can do.

Other Team and Task Characteristics

Three other task characteristics are worth noting here: conflicts between tasks in terms of priority, team size, and task uncertainty. *Task priority conflicts*, as well as goal conflicts, can be built into task design. For instance, being a member of multiple fast-paced teams can create work priority conflicts for team members who must juggle multiple time-sensitive tasks (Bell & Brown, 2015). Depending on the requirements of the task, team size will vary: *team size* involves the number of members, which is under the control of staffing (Klimoski & Jones, 1995). Relevant to this project, as team size increases, team dynamics increase in number and complexity. *Task uncertainty* involves how much the task might change over time, such that team members likely have increased coordination needs (Wittenbaum et al., 1998). All of these features of the task and the team may be relevant to how one might deploy an AI within a human team.

Task Types Summary

The mission of a team can dictate their tasks; these tasks and the organizational design of jobs has implications for what a team does with its time, which then has implications for individual and team processes, emergent states, and outcomes. In particular, the type and level of interdependence drives both activities and the degree to which teams must coordinate and communicate. Remote, or distributed, teams may also provide some challenges and opportunities for deploying AI: For example, AIs can help with communication, shared mental models, and managing conflict (see those sections for descriptions of those constructs). AI teammates can also be assigned simple routine tasks that can be easily automated, freeing human teammates to engage in more complex tasks, and, if designed with the ability to monitor human physiology and maintain information about individual human capabilities, AIs can assist in optimal task allocation. Depending on the specific tasks, different kinds of capabilities are required for both the human and AI team members.

Time: Temporal Factors

Time, like task and mission, is a factor that influences all levels of our AI-enabled teaming model. However, time isn't just one factor. One can conceptualize time as time scales of

phenomena, from very brief, moment-by-moment phenomena to longer time chunks stretching across days, weeks, months, and years (e.g., Paletz et al., 2011). In the context of AIs in teams, for example, this may mean that the AI may be sensitive to changes in human conditions moment-by-moment versus over weeks. Different types of phenomena may require different time scales to examine, from communication as it occurs in different turns to the introduction of a new team member. Another way to consider time is via communication delays: for example, in different types of space missions, communication delays may be minutes to the Moon versus four to 22 minutes between the Earth and Mars (Kobs Nawotniak et al., 2019), necessitating different types of work processes involving virtual teamwork (Fischer & Mosier, 2014). Even on Earth, there may be communication delays due not to the speed of sending messages, but due to other work processes such as 8-hour workdays, time zone differences, and so on.

Another way of considering time is as relevant to group development. Team processes in the early part of a group's life when it is establishing norms and processes are different than during maintenance operations, revisiting its norms, or when a team is about to disband (Hertel et al., 2005; Morgan, Salas, & Glickman, 1994; Tuckman, 1965; Tuckman & Jensen, 1977). Roles can change over the life of the team and team membership can change (Mathieu et al., 2014), and the team can have specific temporal milestones that impact the processes before and after these milestones (Wittenbaum et al., 1998). Teams may also have time pressure due to these milestones, which impacts processes (Wittenbaum et al., 1998). Time itself can influence team processes, such as later in the life of a team a particular factor may be less or more important (e.g., Harrison et al., 1998; Zellmer-Bruhn et al., 2008).

This section gave a taste of the variety of processes and nuances related to time and teamwork. These different ways of looking at time are important to keep in mind for any research on teams. Overall, the development of a team can be iterative, where outputs become inputs on short or long-time scales (Marks et al., 2001). For example, cohesion is an emergent state that impacts subsequent performance, which then impacts subsequent cohesion, but even the power of this feedback loop can fade over the life of the team (Braun et al., 2020).

Individual: Human

Inputs

Human beings bring many individual qualities to teams and teamwork. This section describes some of the many individual-level attributes that humans provide as an input to teamwork, with or without AI. For each of these factors, we will discuss: 1) how these qualities, within humans, have been found to influence teamwork, 2) how humans with these qualities might interact with AIs and, where it's necessary, 3) how AIs can or should have some of these qualities.

Taskwork-related attributes

Although the focus of this paper is on teamwork-related attributes, it is worth noting that the humans involved in a human-AI team should have relevant knowledge, skills, abilities, and other characteristics (KSAOs) for the task as well (Hertel et al., 2005; Jones et al., 2000). KSAOs are a term used in industrial/organizational psychology: the process of conducting a job analysis breaks a job into both tasks and elicits the required KSAOs for that job, which can then be used in selection, performance appraisals, and training (Society for Industrial Organizational Psychology, 2018). For instance, the human members of our rescue team must possess the KSAOs necessary for their roles; a medical first responder should have medical-relevant knowledge such as diagnostic criteria, a navigator should have the relevant skills to interpret a map and guide the team to desired locations. It is important to do a teamwork analysis, as well as a job analysis for any new positions (Jones et al., 2000). Within a team, individual task KSAOs can be required of all team members, compensatory, such that only certain team members need a set of KSAOs, or non-compensatory and interactive, such that particular combinations are required (Jones et al., 2000). These task-relevant KSAOs can also depend on what phase a group is in, such as new versus intact teams (Jones et al., 2000). Additional relevant KSAOs include self-management skills and dependability (Hertel et al., 2005), which are covered in subsequent sections. For particular tasks, one might conceive of what kinds of KSAOs are necessary for either the humans or the AIs to have before joining a human-machine team. Given the changing nature of work, an analysis of work can also include broader competences, such as those regarding teamwork more generally (Society for Industrial Organizational Psychology, 2018). The rest of this section focuses on teamwork-relevant KSAOs.

Human personality

Most of the relevant work on personality is related to team composition (see Team-Level Inputs), rather than individual personality, but individual personality is also relevant. The most established and validated categorization schemes of human personality is the five-factor model, which includes the Big Five personality dimensions (John & Srivastava, 1999; John, Naumann, & Soto, 2008): Conscientiousness, Extroversion-Introversion, Agreeableness, Negative Emotionality (or Emotional Stability) and Openness to Experience. Each of these has facets, or subdimensions. In brief, Conscientiousness entails being organized, responsible, self-disciplined, and productive, and is positively related to work performance above and beyond intelligence (Barrick et al., 2001). When employers seek employees who are dependable, what they mean is someone with high Conscientiousness. Extraversion indicates sociability and talkativeness, assertiveness, and energy level while Introversion is associated with lower levels of these, such as lower sociability, more reticence, and lower energy levels. Most people fall in the middle as ambiverts, but there are meaningful differences between people on the ends of the dimension. Agreeableness includes compassion, respectfulness, and trust. Higher levels of Agreeableness are related to empathy and altruism, as well as cooperative behaviors (Barrick et al., 2001). Negative Emotionality (previously called Neuroticism), with its opposite pole Emotional

Stability, involves being anxious, self-conscious, depressed, and emotionally volatile. Emotional stability is also positively related to job performance generally (Barrick et al., 2001). Finally, Openness to Experience involves the person's intellectual curiosity, creative imagination, aesthetic sensitivity, and willingness to pursue travel and artistic interests.

Some researchers argue that a sixth factor, Honesty-Humility, is another important dimension of personality as found using cross-cultural methods (Ashton & Lee, 2007). This factor indicates how honest, modest, and sincere a person is versus greedy, boastful, hypocritical, and pompous (Zettler et al., 2020). The HEXACO model includes both the Big Five and Honesty-Humility (Ashton & Lee, 2007; Zettler et al., 2020). The five-factor structure and surveys of personality have been overwhelmingly found to be valid and reliable in Western, Educated, and Industrial countries. However, in face-to-face surveys of 94,751 participants from 23 low- and middle-income countries, commonly used personality questions showed low validity, whereas the authors found higher validity in internet surveys of over 198,000 respondents from the same countries (Laajaj et al., 2019). Rather than question the underlying validity of HEXACO or the Big Five, this research is a warning to those who presume personality is culture-free and universally understood.

Driskell and colleagues (2006) argued which particular facets of each of the Big Five would be related to particular aspects of teamwork, such as shared situation awareness, interpersonal relations, and communication. In Emotional Stability (Negative Emotionality), the facet of adjustment (low negative affect, freedom from depression and anxiety) would be positively associated with aspects of teamwork, but self-esteem would be positively related to the teamwork aspects of adaptability, team management, and interpersonal relations. In terms of Extraversion, effective teams would have members who are low on dominance (i.e., low need to control, not combative) but high on affiliation, social perceptiveness, and expressivity, although these specific traits' facets vary for particular types of teamwork. For example, affiliation would be positively related to adaptability, shared situation awareness, interpersonal relations, and performance monitoring, but potentially negatively related to team management, decision making, and coordination.³ The authors considered the flexibility (vs. rigidity) facet of Openness to Experience related to all the aspects of teamwork. The facets of trust and cooperation in Agreeableness would be positively related to performance monitoring and feedback, interpersonal relations, communication, and decision making. Finally, Driskell and colleagues (2006) contended that the dependability, dutifulness, and achievement facets of Conscientiousness would enable individuals to be good team players.

It is also important to note how specific personality traits may impact a given human's ability to collaborate with an AI. Individuals who are emotionally stable, open minded, and agreeable may be more capable of trusting new technology and therefore better equipped to adapt to the

³ Note that dominance here is discussed more as assertiveness elsewhere (John et al., 2008).

inclusion of an AI teammate. While Extraversion may play a role in a given human's proclivity for communicating with AIs in general, Conscientiousness may impact the extent to which an individual will attempt to work with an AI teammate before making a decision about its ability to be a positive addition to the team. It may not currently be possible to design AIs with human-like personality; however, characteristics such as adaptability, social perceptiveness (which can be possible if an AI is equipped with the ability to recognize and monitor facial expressions and tones of voice), and dependability can be included in the original design.

Additional teamwork-related individual differences

A variety of other individual differences, attributes, or qualities are related to working well in teams. These qualities are not personality traits, but are a series of other social, emotional, and behavioral skills, which are relevant to being able to regulate emotions and support social relationships (Soto et al., 2020): emotional intelligence; cultural metacognition; collective orientation and other factors related to developing a shared identity; and leadership, status, and power. The most obvious individual level human inputs that are relevant to successful teamwork are skills related to positively harnessing team processes such as communication and information sharing, negotiation, and conflict resolution. These are essentially the KSAOs for teamwork, such as skills in role articulation, ability to compromise, explain information succinctly and accurately, and so on (Jones et al., 2000).

Emotional intelligence

Emotional intelligence has been popularized in waves by various researchers (e.g., Goleman, 1998; Mayer & Salovey, 1997). Emotional intelligence is a combination of several KSAOs, some of which are already described here (e.g., social perceptiveness from extraversion, Driskell et al., 2006), but are essentially a composite of perceiving, understanding, and regulating emotions (Lam & Kirby, 2002; Mayer & Salovey, 1997). Emotional intelligence can be positively related to individual job performance above and beyond cognitive intelligence (Lam & Kirby, 2002) and to teamwork effectiveness above and beyond personality and cognitive ability, among other factors (Farh et al., 2012). Individuals with the ability to adequately regulate their emotions may be better equipped to deal with the changes that may occur with a learning AI or to handle when an AI teammate is not as dependable as anticipated.

Cultural intelligence

Cultural intelligence goes beyond social or emotional intelligence or cognitive intelligence (Ang et al., 2007; Ott & Michailova, 2018). Cultural intelligence is an individual difference that indicates an ability to adapt to new cultural settings and work with members of different cultures (Ott & Michailova, 2018). The two main conceptualizations of cultural intelligence (Earley & Ang, 2003 and Thomas et al., 2008, as noted in Ott & Michailova, 2018) include both metacognition and cultural knowledge (or cognitive cultural intelligence, Earley & Ang, 2003). Cultural knowledge is the knowledge of specific cultural differences, whereas cultural

metacognition includes being able to control cognition, such as self-regulatory processes and the ability to monitor and be aware that different cultures have differences, even if the specific differences are unknown (Ang et al., 2007; Ott & Michailova, 2018). Earley and Ang's (2003) conception also includes motivation and a behavioral scale that is the self-report of a repertoire of different behaviors to use in different cultural contexts (Ang et al., 2007). Similarly, Thomas et al.'s (2008) conception includes cross-cultural skills (Ott & Michailova, 2018). Past research suggests that cultural metacognition in particular may be positively related to team creativity (e.g., Crotty & Brett, 2012) and to task performance in diverse teams (Ang et al., 2007). Cultural intelligence could also assist individuals in understanding and accepting a culture that includes AIs as teammates, making the transition to collaboration with AIs a smoother process.

Other attitudes and preferences

A cluster of attitudes and preferences also serve as an individual difference relevant to teamwork: psychological collectivism, collective orientation, team orientation, and preference for teamwork (Bell, 2007; Bell & Brown, 2015; Driskell et al., 2010; Salas et al., 2005). While these are different constructs within psychology, they overlap in that they are different ways of representing a desire and propensity to work in teams, collaborate, and align with team goals (e.g., Driskell et al., 2010; Salas et al., 2005), all of which are important aspects of successful collaboration whether teaming with humans or AIs. Prior liking of team members, along with psychological collectivism, identifying with a profession, and support for the mission are positively related to shared identity, an aspect of social team cohesion (Bell & Brown, 2015).

Leadership, status, and power

Leadership, status, and power are all also important aspects of teamwork. Leadership has been identified as a foundational aspect of teamwork: it requires KSAOs but is also a position that involves organizing, planning, and assigning tasks, as well as developing team KSAOs and motivating team members (Salas et al., 2005). The important aspect here is to question the role of an AI vis-a-vis leadership: is it to support a leader, be a follower, be a peer, or be a leader itself (Sebo et al., 2020)? Sebo and colleagues' (2020) review of robots in human teams noted that tutoring can be conceived as a leadership role, because it involves initiating and facilitating the activities of others. Given the importance of the role of different types of leaders, it is useful to consider what aspect of the AI's role might overlap with, complement, supplement, and/or replace that of a human leader (see also Sebo et al., 2020).

Power and status can be inputs but also emergent processes and can be related to teamwork (Van Swol & Kane, 2019). A large literature on power (with status being the relative standing) is available from management, psychology, and sociology perspectives (e.g., D'Ignazio & Klein, 2020; Keltner et al., 2003). Power is an individual's capacity to change others based on having control over resources, rewards, and punishments (Keltner et al., 2003). In addition to power given by social systems, structures, and dominant cultures (Hill Collins, 2008) and the position

power one might have assigned as a team lead, branch manager, and so on within an organization (Bombari et al., 2017), personal power has psychological implications (e.g., Keltner et al., 2003; Galinsky et al., 2006; van Kleef et al., 2008). A major theory of the effects of power on the power holder's behavior and psychology suggests that being in a position of power (specifically, feeling that power, Bombari et al., 2017) activates approach-related tendencies, such as positive affect, using heuristics, and disinhibited behavior, whereas lower power is related to relative attention to potential threats, more systematic/controlled information processing, and inhibited behavior (Keltner et al., 2003). For example, people primed to have high power are not as good as those primed with low power in taking perspectives and interpreting others' emotional expressions (Galinsky et al., 2006) and are less likely to feel distress and compassion when confronted with another's suffering (van Kleef et al., 2008). These findings have implications for how people in positions of power might behave and feel toward others, whether they are put in positions of power due to organizational or social structures.

When designing AIs to work with humans it would be beneficial to grant the capability to determine which human members of the team are in leadership positions and the power structure of the team overall. This knowledge would allow an AI teammate to encourage people in lower positions of power to disinhibit their behavior and contribute more, while helping remind those in positions of higher power to take into consideration the perspectives of others. Understanding who leads and who takes orders may also allow an AI teammate to appropriately disseminate task relevant information to individual human team members as well.

Summary of additional teamwork-related individual differences

In summary, individuals who possess emotional and cultural intelligence are better able to work cohesively in a team, which may be generalized to AI-enabled teams. Other attitudes and preferences may also contribute to an individual's ability to work well within a team. Leadership, status, and power are all important components that influence how individuals work both independently and collectively. AIs designed to collaborate with humans as teammates could be designed to increase processing of information related to emotional and cultural intelligence and have capabilities of discerning status and power. These capabilities would enhance an AI's ability to work well in a team.

Fundamental cognitive individual differences

It is well established that individuals vary considerably in cognitive ability and that this variability is present in all areas of cognition (Nunez et al., 2015; Bridwell et al., 2013; Unsworth et al., 2004; Vogel & Machizawa, 2004; Rypma & D'Esposito, 1999). Consideration of individual differences in cognitive ability is a useful tool for better understanding how humans perform in teams (e.g., which roles an individual is more likely to succeed at) and to provide insights into how to assemble, manage and optimize teams, be they human or human-machine teams. Furthermore, with recent advances in technology, information overload and multitasking

are increasingly pervasive in operational settings (e.g., Chérif et al., 2018; Grier, 2012), and both place significant demand on executive functions like attention control and working memory, often leading to performance decrements (e.g., Ralph & Smilek, 2017). Accordingly, while there is individual variability in many cognitive abilities that may underpin team performance (e.g., creative problem solving, language comprehension, inductive reasoning; Just & Carpenter 1992; Kane & Engle, 2002; Kaufman, 2011) we will limit our discussion to a few key functions that underpin most forms of higher cognition: attention control, working memory, and processing speed.

Attention control

In performing any form of goal directed behavior, it is critical to focus attention on the relevant pieces of information and inhibit irrelevant information. Interference from stimuli or information that is not currently relevant to the task at hand can impair performance. The ability to inhibit this interference or maintain focus in the face of interference is a key cognitive substrate of successful goal-directed behavior (for reviews see Driver, 2001; Johnston & Dark, 1986). The distracting information can capture attention through bottom-up (e.g., perceptual information from the environment) or top-down (e.g., biases related to the task) channels. Inhibition can be exercised over a variety of distractors (e.g., task-irrelevant stimuli in the environment, emotional reactions to the task, once critical information that is no longer relevant). For example, when assisting in the search for survivors, it is necessary for operational planners on-the-ground to focus on data and visuals relevant to the task at hand, inhibiting perceptual distractions, such as anxiety about personal losses due to the natural disaster.

Attention control ability varies across individuals, and these individual differences predict a variety of factors like general fluid intelligence (e.g., Kane et al., 2007; Unsworth, Fukuda, Awh & Vogel, 2014), academic achievement (e.g., Gilmore et al., 2013; Kane et al., 2007) and certain clinical outcomes (e.g., Attention Deficit Hyperactivity Disorder, ADHD; Schachar, Tannock, Marriott & Logan, 1995). Importantly for working in teams, variability in attention control also predicts emotional regulation such that greater control is associated with greater emotional regulation (e.g., Hsieh & Chen, 2017). When working in teams, the ability to inhibit distractions and thereby enable attention to be focused on the truly relevant information is key. In group situations when information is changing rapidly, the ability to inhibit distractors is essential for flexibly updating the information in the attentional focus. AI teammates could assist their human counterparts by reducing the presentation of irrelevant information and/or highlighting relevant information for ease of processing. Even if human attention regulation is not the main objective of the AI, AIs should be designed with attention-control-like capabilities as well, including functions that consider normal human variation in attention.

Working memory

When information is in the attentional focus, the set of abilities associated with processing that information is called working memory (WM; Kane & Engle, 2002). WM is “a multicomponent system responsible for active maintenance of information in the face of ongoing processing and/or distraction” (Conway et al., 2005, p. 770). This active maintenance of information is underpinned by “domain-specific storage and rehearsal processes and domain-general executive attention” (Conway et al., 2005, p. 770). WM is, therefore, the capacity to hold information in the attentional focus for ongoing processing. This collection of mechanisms is the substrate of goal-directed behavior. There is considerable individual variability in WM capacity (WMC; Kane & Engle, 2002), and an individual’s WMC predicts a wide variety of outcomes including language processing (e.g., Friederici, Steinhauer, Mecklinger & Meyer, 1998), general intelligence (e.g., Conway et al., 2003), and the likelihood a given task will result in cognitive overload (Yu, Change, & Yang, 2014). Teams with higher levels of team WMC, which depend on individual WMC, are more likely to have more successful team performance (McKendrick et al., 2014). AI teammates with knowledge of the WMC of human team members can better assist in either reducing the cognitive load that a particular task may have on a given team member or facilitating coordination with other team members with higher WMC.

Processing speed

Processing speed involves the speed with which an individual perceives and responds to a given stimulus and is integral to other cognitive processes such as the maintenance of information in working memory (Fry & Hale, 1996). Factors such as age (Salthouse, 2000) and trauma or neurodegenerative diseases (Lengenfelder et al., 2006; Hale et al., 1993) can greatly impact the speed at which an individual can process information, and this speed is biologically limited (Gazzaniga et al., 2018). Individual differences in processing speed can predict both reading (Kail & Hall, 1994) and arithmetic (Bull & Johnston, 1997) abilities and, like attention control, some clinical outcomes (e.g., ADD, Goth-Owens et al., 2010). Reductions in processing speed can impair one's ability to remain situationally aware (Bolstad, 2001) and drastically impact that individual’s ability to contribute to a team mission, especially in a dynamic situation. Designing AI teammates with the ability to monitor human team members will enable constant measuring of relevant processing speed and allow for the appropriate assigning of tasks.

Summary of fundamental cognitive individual differences

Individual differences in attention control, working memory capacity, and processing speed can greatly influence an individual's ability to engage in teamwork. AI teammates with the ability to monitor and measure individual limitations in these cognitive functions can mitigate possible negative impacts.

Summary of individual human inputs

Human personality, taskwork-related attributes, teamwork-related attributes, and individual differences in cognitive processing are all individual inputs humans bring to a team. These inputs impact how well humans engage with other teammates and their ability to process incoming information and complete relevant tasks. Individual human inputs also impact team inputs. For example, each human teammate joins a team with a given set of personality traits and the combination of those traits form the makeup of team personality. AI teammates equipped with the abilities to both store information about personality, KSAOs, and cognitive capabilities of each teammate and actively assess information as it may change over time will be better equipped to aid in task allocations, team coordination, and assisting teammates with information that may exceed their cognitive capabilities.

Although this section suggests different individual human attributes that are important inputs for teamwork, people can often choose their work partners. Previous research suggests that people choose future coworkers who have a reputation for competence and hard work and those who they developed strong working relationships within the past (and of the same race; Hinds et al., 2000). These findings suggest that when given the choice, some of the qualities noted here--Conscientiousness, the right level of Agreeableness, similarity--are sought out.

Individual Human Processes

Individual human inputs, such as personality, emotional intelligence, and attention control all have an impact on the processes that occur while engaging in teamwork. While we outline individual cognitive differences as inputs, it is important to note they also reflect processes. In this section we focus on cognition as a process and the process of social cognition.

Cognitive processes

While individual differences in cognitive resources and capabilities such as attention control, WMC, and processing speed are all inputs, actively focusing attention, maintaining information in working memory, and processing information are all cognitive processes. Individuals must engage in the process of attention, regardless of attention control abilities, to work on and complete tasks and maintain situation awareness. Engaging in activities such as task switching, searching for information, or responding to alerts requires active maintenance of information in working memory that was just learned or recently retrieved from long-term memory storage. The ability to actively attend to and maintain information in working memory necessitates the ability to process that information. AI teammates could relieve some of the burden of these processes in their human teammates by actively monitoring relevant information and freeing human teammates to engage in other tasks. Designed with the capability to monitor physiological measures of alertness such as heart rate and galvanic skin response, AI teammates could

determine the likelihood that a given team member is engaging in one of the processes and judge when alerts or other relevant information is supplied.

Social cognition

Social cognition, including social perception, is an entire subfield within psychology (e.g., Fiske & Taylor, 2017). Social cognition processes are parallel to, but not quite the same as, processes of non-social cognition. Humans can perceive, encode, store, and retrieve information about other social entities, but there is a wealth of research on biases, selective perception and retrieval, stereotyping, and so on both oneself and others (e.g., Higgins & Bargh, 1987; Snyder et al., 1977; Steele, 1997). Humans use observation and interaction to decode others' emotions and learn about others' reputations, group membership, and intentions (Frith & Frith, 2012). While humans are particularly suited to pick up on cues from other humans (Macrae & Quadflieg, 2010), these processes are imperfect, with humans remembering and noticing social information that aligns with preexisting mental models (Fiske & Taylor, 2017).

Although humans seem uniquely able to be able to reflect and ponder on their own thoughts (have metacognition, Frith & Frith, 2012), social cognition is not limited to humans either as targets or perceivers (Kwan & Fiske, 2008; Reeves et al., 2020). For example, dogs seem to have evolved to pick up social cues from humans (Hare et al., 2002), and humans may perceive their pets to have racial biases against other humans (Hawkins & Vandiver, 2019). Humans have long anthropomorphized nonhuman agents, attributing will, intentionality, and personality to them (Kwan & Fiske, 2008; Reeves et al., 2020). Military service members have referred to the robot member of an Explosive Ordinance Disposal team as a team member (Carpenter, 2016). These service members may apply human-human interaction cues onto the human-robot relationship, or they may place the robots in a new type of social category (Carpenter, 2016). The human-robot bond may have emotional parallels to human-animal bonds, and, as with animals, society and humans may treat robots as tools, resources, or companions (Carpenter, 2016). Based on the literature, humans seem to respond to technology as they do to other people, such that the fundamental social cognition of attribution, perception, and stereotyping does not differ (Reeves et al., 2020). Thus, these processes of selectively perceiving, encoding, storing, and retrieving information may apply to the human perceptions of AIs as well.

Emergent States

Emergent states can become outputs and feedback into inputs, permeating the entire cycle of the IPEOI model of teaming. Here, we outline trust, situation awareness, cognitive load, and individual mental models as emergent processes for individual humans.

Trust

Trust is defined in several ways depending on the field and draws from psychology, legal studies, user acceptance theory, task fit, emotional, and personality studies, among others (Guo, 2020, Mayer et al., 2003, Reeves & Nass, 1996). We follow Mayer and colleagues' (1995) emphasis on the willingness to be vulnerable and the expectation of positive outcomes on important actions irrespective of close monitoring or control. Trust is an emergent state that permeates every level of the model; each level has unique and common characteristics. Given the relatively small set of research into longer-term trust in teams, the findings described here might not generalize. This section focuses on trust as an emergent state at the level of the individual human. While any human-agent team must involve human-human trust, this section outlines factors that influence human trust in AI. These insights can be applied in general, and often draw from, research on human-human trust (e.g., McAllister, 1995).

Individual human trust in AI team members is influenced by various factors such as the AI's physical form, timing of responses and actions, spatial concerns in relation to physical AI and many others (Bainbridge et al., 2010; Glikson & Woolley, 2020). Generally, human ability to trust in AI develops in ways that are similar to how humans trust other humans: starting at a low level, developing over time, and over the course of interactions that create and set expectations for the future (Glikson & Woolley, 2020). Human trust in AI can be created and lost based on AI and human attributes, cultural factors, individual and team characteristics, and preconceived ideas formed from entertainment and literature (Bainbridge et al., 2010; Bansal et al., 2019; Culley, & Madhavan, 2013). Trust in AIs or robots will also in part depend on trust in the designer (Carpenter, 2016). While the AI will need to be considered part of the larger team, individuals will bring varied levels of acceptance of AIs, and these individual assessments may inform the AI's integration into or segregation from the team. Individuals have their own predispositions towards trusting behavior, which is unique to them and consistent across situations. This type of trust, dispositional trust, stems from their general faith in the goodness of human nature and can be tied to culture and social standing (McKnight, 1996), self-esteem and feelings of control (Uslaner, 2013) culture, age, gender, and personality traits (Hoff & Bashir, 2013).

Culture will have an influence on an individual's comfort level with AI, but the AI itself will likely be designed with the biases and have connection to the culture of the programmer. When the culture programmed into the AI doesn't align with the culture of the individual user, the user's trust will likely take more time and work to establish. This phenomenon is reflected in terms of concepts described above, like attributes, characteristics, and preconceived ideas, but also plays out in terms of work style and interpersonal interactions. Organizational and team culture are defined by a subset of shared values, behaviors, expectations, and norms in work and interaction between the in-group members, which can impact how human team members respond to the culture programmed into the AI. While the overall concept of trust is an emergent state,

the impact of each of the listed issues fluctuates throughout the duration of the team. At different times, trust presents itself as one of all three: input, process, and output.

Trust in AIs is dynamic, waxing and waning as AI teammates achieve goals and make errors. Similar to interpersonal interactions, faulty AI behaviors can negatively impact perceptions of reliability and trust, but correlate with reduced interaction (Mirnig et al., 2017). Surprisingly, humans tend to find imperfect and fallible robots as more likeable than those that behave perfectly (Glikson & Woolley, 2020; Mirnig et al., 2017). This finding implies that an incongruence between AI behavior and teammate expectation can affect trust. This phenomenon is explored in the uncanny valley theory, discussed in more detail in the following section. A team member's existing expectations are influenced by the culture of the individual, team, or the programmed culture of the AI. If AI behavior does not reconcile with these cultural expectations, the AI may lose the trust of an individual or the team.

Trust processes

There are three sources of variability with regards to trust formation: the trustor, the situation, and the trustee. Hoff and Bashir (2013) describe in their model three primary types of trust formation: dispositional, situational, and learned trust. Dispositional trust refers to an individual's likeliness and predisposition to trust: It is carried across situations and is impacted by factors such as culture and global economic standing (McKnight, 1996), age, gender, and personality traits (Hoff & Bashir, 2013; MacDorman & Entezari, 2015). Situational trust is contextualized, wherein an individual only trusts another within a given context (McKnight, 1996). By contrast, learned trust is based on first-hand experience over time with intentions, benevolence, expertise, and integrity. Benevolence is defined as "the extent to which a trustee is believed to want to do good to the trustor" (Mayer et al., 1995, p. 718). In the context of AI-human trust, humans may assume the AI has motives equivalent to its developers. Whether those motives are assumed to be benevolent or not may depend on other perceptions of the AI and the trust in the AI's company/creators (e.g., Amazon, Google, the US Army). Learned, situational, and dispositional trust do not exist in isolation, but together enable the creation and expression of trust over time.

Types of trust

The study of trust crosses disciplines and is referred to in terms of individuals and teams, short-term and long-term, gain and loss, swift, interpersonal, organizational, social capital or social influence, and with even more subdivisions, depending on the field. Different researchers define trust as emerging early, developed over time, or something in between. It can be considered differently in one-on-one vs. team or organizational contexts, and intrinsic and behavioral factors can influence the development of trust (Mayer et al., 1995; Meyerson et al., 1996; Castalado, 2003; Delhey & Newton, 2003; Fukuyama, 2010).

Swift trust is especially relevant to the problem of AI teammates. This type of trust is associated with temporary teams such as swift starting action teams (STAT) which can be assembled for search and rescue missions (Wildman et al., 2012). The formation of swift trust happens between experts who do not have prior work experience with one another, as they come together temporarily to complete a specific task or to solve a problem and then disperse (McKinney, Barker, Davis, & Smith, 2005). Compared to the concept of interpersonal trust, swift trust focuses more on the cognitive aspects of trust than affect. Swift trust relies on team communication, individual differences in disposition to trust, and expectations of team members' abilities and integrity (Blomqvist and Cook, 2018). It is influenced by perceptions of norms within groups or initial assumptions that are verified or adjusted as time progresses. As described by Meyerson, and colleagues (1996), it is a pragmatic approach that relies on the presumption that others in the group are competent. Swift trust can be developed and maintained via establishing and maintaining clear norms and standards for communication and behavior (Blomqvist and Cook, 2018). Swift trust often develops quickly or not at all (Wildman, 2012).

Trust repair

It is almost inevitable that for many relationships, there will be some transgressions, small and large, that violate and damage trust. To continue cooperative work, that trust should be repaired. Lewicki and Brinsfield (2017) describe a four-step cycle of trust and repair. First, there must be a pre-existing level of trust, that is then broken by an action of one party (the violator), as well as a recognition by the other (the victim) that trust has been broken. After a series of verbal or behavioral actions on the part of the violator are enacted, the victim may then signal whether the trust repair efforts have been successful (Lewicki & Brinsfield, 2017).

Trust repair processes are important in human-human teammate relationships, but also in human-AI relationships. For AI teammates, to repair trust with a human teammate a series of steps must be enacted and considered within the design. First, the AI system must be capable of understanding that a breach of trust has occurred with the human as a victim--a challenging requirement for an AI, but worth consideration. Once the violation has been recognized, a system would have to account for and explain the provenance of the decision. In many cases, changes in behavior for future actions, such as an update to a system algorithm or in re-allocation of responsibilities, may help repair trust. Then, there must be a mechanism by which the operator (or a supervisory body) is able to acknowledge and accept the change as sufficient to repair the harm done or prevent further harm. If these steps are not undertaken, then silence and denial of an existing problem with the system is likely to result in mistrust and disuse, especially if the outcome is likely to be a dangerous or risky one.

The uncanny valley

Adoption of AI as a teammate may be helped or hindered by its presentation. Popular culture is rife with stories and imagery of AI that is unsettling if not dangerous. The "uncanny valley"

hypothesis describes a curvilinear relationship between the degree of perceived humanness in a robot's appearance and its likeability, up to a point (Mori, 1970). Past that point, increased humanness in the robot's appearance causes the operator unease about the robot, or an uncanny feeling (the valley).⁴ The hypothesis holds that upon reaching a certain point of human likeness, likeability increases again. Research on the topic is inconclusive, showing inconsistent results, with some authors questioning its existence (Brenton et al, 2005) as well as the shape of the relationship (Bartneck, 2007).

A variety of hypotheses have been suggested to explain and describe the uncanny valley phenomenon. A 2015 literature review of the uncanny valley research concluded that the “the uncanny valley may not be a single phenomenon to be explained by a single theory, but rather a nexus of phenomena with disparate causes” (Wang et al., 2015). A set of theories, based on ‘precognitive processes’ (implicit, rather than controlled cognitive processing) theorize that the uncanny feeling may be based in self-preservation instincts (MacDorman et al., 2009; MacDorman & Ishiguro, 2006) triggered by feelings of disgust (Rozin & Fallon, 1987) or due to associations with psychopathic personality traits (Tinwell et al., 2013). A different set of theories from a cognitive processing perspective have also been proposed. Theories such as the violation of expectations hypothesis generated by a mismatch in cues (Saygin et al, 2012; Mitchell et al., 2011) and the categorical uncertainty hypothesis, the hypothesis that discomfort stems from not knowing what category (human or machine) the AI belongs in (Ramey, 2006), suggest that the uncanny valley is a result of cognitive processing. The mind perception hypothesis argues that agency, which is the ability to act and plan, as well as the ability to feel emotions and have a subjectivity, separates humans from nonhumans (Gray et al., 2007).

Wang and colleagues (2015) argue that the mind perception hypothesis is only a partial explanation to the uncanny phenomenon, and that the answer can be explained in the application of the broader concept of anthropomorphism. Anthropomorphism has an overall positive and facilitative effect for interactions between humans and non-humans. Preference for human-centric interaction is the oft-cited reason for the humanization of robots, owing to the ease of applying human interactions to non-human robots as a more natural concept for Human Robot Interaction (HRI, Giger et al., 2019). This phenomenon is not limited to the humanization of robots but can extend easily to the use of avatars and behaviors of embedded systems as well. Giger and colleagues (2019) identified several benefits of humanization including greater interaction engagement and increased social connection. Potential downsides should also be considered, such as overreliance and unrealistic perceptions of a robot's autonomy and

⁴ While most of the research on the uncanny valley focuses on perceived humanness, it is not limited to humanness. One study has shown that monkey visual behavior when monkeys were shown synthetic realistic and unrealistic monkey faces also showed uncanny valley behavior (Steckenfinger & Ghazanfar, 2009). Anecdotally, robots that look almost like, but not quite, animals can also be incredibly creepy, such as MIT's Biomimetic Robotics Lab's cheetah: <http://biomimetics.mit.edu> and this robot dog who was tweeted about on Nov. 5, 2020: <https://twitter.com/GeorgeWillems1/status/1324516652453662721>

capabilities as well as the threat of AIs encroaching on human uniqueness and human tasks (Giger et al., 2019). One very important downside to the humanization of robots includes the threat to human identity and distinctiveness (Giger et al., 2019). These threats are compounded when robots are depicted as having the ability to reject a human's commands (Złotowski et al., 2017) or are shown as either being equal or superior to humans in performing emotion-oriented tasks (Waytz et al., 2014). Gray & Wegner (2012) argue that anthropomorphism is responsible for the uncanny phenomenon because it invokes 'attribution of mind' or the prescription of experience to the AI that can be unnerving (Gray & Wegner, 2012). Wang and colleagues' (2015) dehumanization hypothesis states that perceiving the anthropomorphized robot as lacking humanness creates a dehumanization effect. This dehumanization creates an out-group of these androids. Haslam (2006) suggests two types of dehumanization: animalistic (more animal-like) and mechanistic (more machine-like). Animalistic dehumanization evokes the perception that the AI has lower intelligence and a lack of self-control while mechanistic dehumanization evokes the perception that the AI is lacking emotion and warmth (Angelucci et al., 2014).

While there are several theories proposed, they are not without limitations. Wang et al. (2015) argue that the research on this topic is problematic for three reasons: First, "likeability" has not been given a standardized definition, with some researchers using attractiveness and others using warmth as their interpretations of what it means for an android to be likable. Prior research also does not contextualize acceptance of the android for a certain task or use, but instead limits this measure to a reflexive response. Second, human-likeness is also multi-faceted, and many studies seem to gravitate to different aspects of humanity to emulate and test. Third, there is no solid consensus on a consistent definition of "uncanny", though recent work attempts to fill this gap (Ho et al., 2017).

Individual differences may account for variations in uncanny valley experiences between participants (MacDorman & Entezari, 2015). MacDorman and Entezari (2015) found that neuroticism to include emotional instability, specifically anxiety, significantly predicts uncanny valley sensitivity. Animal reminder sensitivity, which is a sensitivity to reminders of one's own creatureliness as opposed to godliness, is also a predictor. Religious Fundamentalism significantly and indirectly contributes through robot-related attitudes to include dehumanization of out-groups (to see them as less civilized) and the belief that humanity is divided from the rest of existence (MacDorman et al, 2009; Vail et al., 2010; Vess et al., 2012). These findings logically follow the theories posed, given that the uncanny phenomenon is something of a personal experience that involves, in theory, precognitive (implicit), cognitive, and cultural factors which are likely to vary from person to person.

Designing AIs to be more human-like in look or behavior, or anthropomorphizing the AI, might increase trust in the AI over time as teammates become more conditioned to the appearance of the AI (Zlotowski et al, 2016). However, as the uncanny valley theory suggests, there are limits

to the benefits of human form. The form of the AI can impact human trust and acceptance of the AI. For example, an embodied AI might bring up ethical issues about personhood, while a virtual AI may add more transparency to work processes such as in the algorithms used. Further research is necessary to better understand the full impact of anthropomorphism on Human-AI teams.

Trust summary

Despite varied factors that influence it, trust--here from the human's perspective--is always an integral part of human-AI teaming at the individual level. Trust is influential throughout all three levels of input, process, and output, with levels rising and falling over time in response to changing variables. Level of trust is influenced by the human's characteristics and beliefs, as well as characteristics of the AI itself, such as form, learning actions, and behavior. The personality of each individual, in addition to the individual's experience and culture, all play a role in determining trust. Moreover, transparency on the part of the AI impacts human trust in AI team members in multiple theories. Without a clear understanding of what the AI is, how it works, how it comes to conclusions or solutions, what security is implemented in and around the technology, and what bias is observed, individuals will have difficulty trusting their AI teammate. Trust in AI and trust among individuals who comprise teams are each broad topics on their own and bringing them together adds additional complexity to already complicated situations. Although this section focused on the trust a human could have for an AI, humans will also have trust in other humans as well. This trust will similarly involve both affective and cognitive components, and relies on judgments of predictability, dependability, and competence (McAllister, 1995).

Situation awareness

Dynamic and often challenging situations require constantly updated information so that individuals can make accurate and swift decisions. Lack of knowledge of a target's movements during rescue missions, or the blood pressure of a rescuee during lifesaving efforts, can be the difference between life and death. Each team member, in these instances, needs to have an accurate level of situation awareness. Endsley (1988) defines situation awareness as "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future" (p. 97). While by no means the sole definition of situation awareness, this description clearly outlines the role of cognition during several stages of awareness.

Endsley (1997) proposed three hierarchical levels of situation awareness: 1) noticing, 2) comprehension, and 3) projection. An individual must first notice if an issue is occurring, which requires focusing attention on the factors that matter (e.g., changes in coordinates, blood pressure or other related factors). Then, an understanding of the problem or dilemma is required, which necessitates perception and the appropriate schema (relevant knowledge) about what to expect

and what is abnormal. The ability to retrieve relevant information quickly and efficiently, or long-term working memory, will depend on how often the individual has completed the task. Finally, that same individual must have the ability to determine how soon in the future they need to address an issue, (e.g., if the rescuer's blood pressure falls we must take drastic measures). Each individual team member will need to have situation awareness for their specific task, and individual situation awareness can combine to become team situation awareness.

As stated, AI team members may be tasked with highlighting relevant information. If this information is being actively attended to, or noticed, then the AI is also assisting human teammates in maintaining situation awareness. Human teammates who are situationally aware may conversely be better equipped to determine the best tasks for AI teammates to engage in. In sum, situation awareness is an important aspect of successfully completing tasks in dynamic conditions and can be enhanced for humans by their AI teammates.

Cognitive load

As stated previously, both attention and working memory are limited in capacity. When engaging in a task, individual team members can only attend to, or maintain, a limited amount of information at a time. Many environments, especially those that are dynamic and require situation awareness such as the battle field or the operating room, require team members to process several stimuli in the environment while simultaneously accessing information from long-term memory. Cognitive load (also referred to as workload in the literature) is the accumulated load that a given task can have on cognition (Sweller, 1988) and specifically comprises intrinsic, extraneous, and germane load (Sweller et al., 1998; Sweller, 2010). Intrinsic load increases with the complexity of a task. Complex tasks and concepts require retrieving less complex information, with a linear relationship between complexity and the amount of information retrieved. The intrinsic load also depends on the knowledge level of the individual. Once the task has been learned, intrinsic load decreases. Extraneous load is related to the method by which information is presented. Aspects that have nothing to do with the task itself, such as visual information presented on a crowded screen or overlapping presentations of auditory information, distract and/or impact maintenance of information in working memory and increase cognitive load. Germane load is associated with knowledge of the task. Task schemata are updated with each exposure. Experts will have lower levels of germane load than novices for the same tasks.

Cognitive load can be measured using subjective, physiological and performance measures. Subjective measures, like the NASA Task Load Index (NASA-TLX), involve self-report that weigh the different dimensions of the load of a task (e.g., mental demand, effort, physical demand, frustration level; Hart, 2006; Hart & Staveland, 1988; Human Performance Research Group, 1986). Physiological measures such as heart rate (De Rivecourt et al., 2008) and respiration rate (Fournier et al., 1999; Fairclough et al., 2005) increase with cognitive load, while

blink rate and blink duration decrease with increase in cognitive load (Wilson, 2002; Veltman & Gaillard, 1996). With performance measures, changes in reaction time and accuracy on a given task are observed. Typically, an increase in reaction time and decrease in accuracy signify increased cognitive load. Tasks that are dynamic, with moment-to-moment changing factors requiring SA, often necessitate maintenance of multiple stimuli which can lead to cognitive overload. An AI teammate equipped with capabilities to monitor the physiological hallmarks of load may best assist its human teammates in reducing the negative impacts of overload.

Individual Mental Models

Mental models are cognitive representations of anything in the world, including relationships between constructs (Johnson-Laird, 1980). Mental models develop based on new information and how it fits in with existing cognitive representations. Of particular interest for this paper is the human mental model of the AI or AIs in the team and team members. This model develops over the course of an interaction, though it may also be informed by prior knowledge and expectations that the human has before encountering the AI. An accurate mental model of the AI may also require some level of training. Human's working within AI-enabled teams should be informed of the limitations and abilities of their AI teammates. Training will allow for adequate levels of expectations of performance and increase the likelihood that humans will trust their AI teammates can complete relevant tasks.

Humans interact with AI collaborators based on their mental models of what those AIs will do in particular situations. One way to characterize human mental models of an AI is by dividing the AI's behavior into global behavior (the way its updating and adjustment strategy is implemented), local behavior (the way its behavior appears within a single interaction), and knowledge base (the data that humans infer the system was trained with). Better models of global and local behavior are associated with better performance on a human-AI collaborative task (Gero et al., 2020), will allow humans to fully understand the capabilities of their AI teammates which should reduce the misalignment between implicit and explicit expectations of what the AI should be able to do, and may reduce the likelihood of human abuse of the technology (Parasuraman & Riley, 1997; Glikson & Woolley, 2020).

Summary of individual human emergent states

Emergent states such as trust, cognitive load, and individual mental models are dynamic throughout the life of the team. Trust and individual mental models change as teammates work with their human and AI teammates, and cognitive load changes with the tasks that are being engaged in. Humans must trust AIs for successful teaming and must update their mental model of what it means to collaborate with AIs. AI teammates can assist their human counterparts by having the ability to both measure and reduce cognitive load.

Individual Human Outputs

Individual human outputs are the results of the inputs, processes, and emergent states. The individual aspects each human member brings into the team interact with the processes and emergent states that occur throughout the life of the team, and the processes and emergent states in turn impact outputs. Here we will discuss outputs such as human effectiveness, efficiency, satisfaction, and stress.

Stress and satisfaction

Understanding human satisfaction with a human-AI teaming environment is critical to optimizing effectiveness and usability of the system. The degree of user satisfaction can be gauged with a variety of factors like emotional reaction (e.g., content versus frustrated), cognitive state (e.g., cognitive load), or stress level. All these factors provide information on the user experience in the human-AI teaming environment. This type of information has historically been collected using questionnaires in which users rate their mood, attentional focus, and/or stress level. In addition to the fact that self-ratings can be inaccurate for many reasons, they are retrospective and require effort on the part of the user. One type of data that can be leveraged to provide insights about the user that does not suffer the same shortcomings as survey instruments is physiological data.

Physiological sensors (e.g., heart rate, respiration) are increasingly available, unobtrusive, and inexpensive, and they provide a rich and continuous source of information about user state (e.g., Can et al., 2019; Kyriakou et al. 2019). Using AI/ML to analyze this rich continuous data, it is possible to extract critical information about an individual during a human-AI teaming activity that would serve as system outputs to inform system effectiveness and usability (e.g., Nath et al., 2020). There is a substantial body of research using physiological metrics paired with machine learning to detect stress (e.g., see Panicker & Gayathri, 2019 for review), with very promising results.

In the domain of using data from physiological sensors paired with machine learning techniques to detect emotional or affective response, a key aspect of user satisfaction, there has also been considerable research (see Shu et al., 2018 for review). There have been a variety of approaches to emotion recognition based on which model of emotion researchers use as a framework (e.g., Plutchik's Wheel of Emotions or Lang's Arousal-Valence model; see Shu et al., 2018). While there is some variability in results, it is possible to identify emotional states based on physiological signals, particularly when multiple signals are incorporated. There is suggestive research using discrete emotions (e.g., joy, anger, sadness, pleasure; Lin et al., 2010) and emotional state along the dimensions of arousal and valence (e.g., Basu et al., 2015). Similarly, physiological data can be used in conjunction with machine learning to provide an index of the user's degree of trust in a given machine agent (e.g., a website; Leichtenstern et al., 2011).

Having an objective, continuous measure of emotional state and degree of trust would be very useful to teaming context in that it would be possible to adapt the system/workflow/team composition/tasks based on these affective states. However, there are serious criticisms of these approaches: many automated systems rely on outdated and limited theories of emotion (Paletz et al., 2021), and the variability of physiological signals within and across individuals is so great as to make some physiological measures--and their link between specific ones and specific affective states--meaningless without understanding that variation and context (Hoemann et al., 2020).

With regards to cognitive states, physiological data can similarly potentially provide an index of a human's mental processes. For example, Luque-Casado and colleagues (2016) found that heart rate variability (i.e., the variation in the intervals between successive heartbeats) varied as a function of task demand during tests of working memory and attention control, thereby providing an online index of task difficulty.

Effectiveness and efficiency

In addition to satisfaction, usability is generally assessed using measures of effectiveness and efficiency. Such measures are generally related to taskwork performance rather than teamwork performance, however. Effectiveness is generally defined as being able to do the task well, and efficiency involves doing the task with the fewest wasted steps (Frøkjær et al., 2000). If AI teammates hinder human effectiveness and efficiency, humans may not adopt or use the technology (Parasuraman & Riley, 1997) In a team setting, these constructs could be operationalized by looking at task outcomes or team communication, though sometimes additional communication is a sign of better teamwork rather than less efficient taskwork.

Summary of individual human outputs

Individual human outputs reflect the results of the interaction between inputs and processes and emergent states. Satisfaction is an individually perceived measure of how well a team collaborated. High levels of satisfaction in AI-enabled teams are necessary for continued work with AI teammates. Measures of performance such as effectiveness and efficiency highlight an individual human's ability to complete relevant tasks successfully. AI teammates may assist in task completion by reducing cognitive load.

Summary of Individual Human IPEOs

The vast majority of teaming research focuses on how humans work together in a team. Human teammates bring in their own inputs, independently engage in processes and emergent states, and generate their own outputs that impact and reflect their tenure within a given team. These inputs, processes and emergent states, and outputs also impact and reflect work with an AI teammate. AI teammates may also assist their human teammates during processes and emergent states, and use

information from individual human inputs, such as attention control abilities and personality traits to efficiently aid humans.

Individual: AI

AI-enabled teams must consist of at least one AI teammate. This AI can be either embodied, such as the unmanned drone and terrain vehicle, or disembodied, like the optimization system. For AIs to successfully integrate into the team, they must have teaming capabilities. AI teammates should be designed advantageously so that their inputs can facilitate teamwork. This section also describes AI processes and emergent states and what outputs are relevant to understanding how they perform as a teammate and on a given task.

AI Inputs

Like humans, individual AIs join teams with their own inputs that can impact the processes, emergent states and outcomes of other individual team members and the team as a whole. Here we describe AI inputs such as the level of automation, representation and embodiment, teaming intelligence, and social presentation. As with humans, who bring task- and team-related attributes to a team, AIs bring task-related and team-related models and interaction capabilities.

Levels of autonomy

In our introduction we highlighted some of the requirements of the level of autonomy that is required for an AI to successfully collaborate within a human teaming structure. While the AI teammate is always required to follow human orders, it should be designed with the capability to suggest alternative actions given some relevant conflicting information. Ideally, the AI should have the capability to operate with minimal supervision if necessary. While some argue that ideal levels of autonomy and human involvement should both be high for ideal effectiveness of the AI (Shneiderman, 2020a), lower levels of human involvement would be necessary for the AI to be a teammate and not a tool.

Taskwork-related attributes

Just as humans bring taskwork-related KSAOs to a team, so too will AIs. An AI's task knowledge is literally encoded in the AI's model. The model is itself a compressed representation of the data it was trained on and the type of model (e.g., Bayesian, regression, rule-based, deep learning, etc.). An AI's knowledge can be tested prior to a mission, by measuring, for example, the rate of correct inferences versus false positives versus false negatives (misses).

Any AI system follows a set of constraints based on what its designers intended it to do. It also generally has some objective functions whose outputs it wants to maximize. A simple ML system may have a single objective function such as to reduce the error in prediction on a facial recognition task, reduce the jerk in a driving task, but an AI in a team would need to satisfy multiple constraints at once. Minimally, it would need to take into consideration its own goals and the goals of the team.

The objective functions and goals of an AI system must be carefully specified to avoid overoptimizing and causing unintended consequences. One method that Russell (2019) suggests is that AI systems must observe human preferences and defer to them. Another method might be to include intentionally conflicting objective functions and to provide an incomplete method for resolving that conflict which is a more human-like method.

An AI's task abilities consist of both functional and non-functional requirements (Voas, 2004). Both are types of requirements. For example, a functional requirement for mapping software would be that it must continually update traffic information in the area that is being mapped; a non-functional requirement would be that the software must be able to work on different phone operating systems. Any AI system that is performing a task in a team should be able to perform that task within its design parameters. It should be able to flag anomalous situations either for team members or remote support personnel. It should have the appropriate training data necessary for performing the task it is designed to do.

Teaming intelligence

For AIs to be capable teammates, they must be designed with teaming capabilities. Johnson and Vera (2019) suggest that AIs in human teams be designed with *teaming intelligence*, or “knowledge, skills, and strategies with respect to managing interdependence” (p. 18). Teaming intelligence would include an understanding of the *state*: knowledge of the mission, current events in the environment, individual teammate knowledge base, and the collective team knowledge base. Teaming intelligence would also include how the team and the task are organized, or an understanding of the *structure*. For example, structure would include whether all team members give equal effort, if a leader guides the team, which team member is responsible for which specialized tasks, and the various interdependencies within the group. Understanding the team's organization refers to knowing who is responsible for what tasks and the hierarchy or status of each member of the team, clearly highlighting who depends on whom for various forms of information and to complete important tasks. Knowing how the task is structured informs the ideal manner in which team members may work together. The AI must be designed with an ability to assess both the state and structure in a dynamic manner, since tasks and team structure may evolve over time. If team members, the AI included, have similar understandings of state and structure, they will have similar or overlapping mental models.

For AIs to manage the interdependencies represented in teaming intelligence, they are required to have *skills* such as coordination (Johnson & Vera, 2019). Coordination typically requires some form of communication between team members. AIs would need to possess the capabilities to both send and receive information from other members of the team. When coordinating, team members must be observable, predictable and directable (Christofferson & Woods, 1995; Johnson et al., 2014). Humans and AIs must be transparent about what they know, what they are doing, and what they need. They must engage in a set of behaviors that allows other team members to accurately predict their ability to complete a task. Predictability is extremely important when determining who will complete what task and in building trust for fellow teammates. Team members must also be directable, or able to influence the behavior of others and be able to be influenced as well. A lack of observability, predictability and/or directability hinders an individual's ability to know and deal with team interdependencies.

Knowing what interdependencies exist in the team and having the skills to deal with these interdependencies are only relevant if an individual knows when and how to use these skills and information. If an AI teammate is equipped with the ability to monitor their human teammates and other AI teammates, if necessary, that information could be used to determine when and how to communicate an issue or error. For example, if an AI can monitor stress levels and task engagement, it can deduce an appropriate current workload for a given human teammate (or collective workload if necessary). This AI could also assess whether information is important enough to break a teammate's concentration during periods of high workload or if the workload is low enough that attention can be diverted elsewhere.

Representation and embodiment

As described in (Gilkson & Woolley, 2020), AIs are generally found to take one of three representations. AI systems may be physically present, as in social robotics, virtual, like a chatbot, or embedded in tools or software systems (e.g., the embedded AI in an iPhone camera). Different levels of representation, regardless of the functionality behind them, may lead to different reactions from human teammates. The form or mix of forms taken can also influence a variety of team processes and emergent states, such as trust, between the individual team members, team as a unit, and the AI. Tangible AI tends to bring out evaluation by an individual based largely on cultural factors, through interaction, immediacy behaviors, and physical characteristics, whereas virtual AI will generate different levels of interaction depending on various traits, such as if it is represented as an attractive agent (Khan & Sutcliffe, 2014). A literature on social robots describes robots with different types of physical embodiments, including human-like or humanoid, similar to animals, able to exhibit social behavior, or those with no social appearance whatsoever (e.g., Roombas; Sebo et al., 2020). Embedded AIs, or AIs that are coded within other systems, can be trickier, because often the non-technical users might not even be aware that an AI is present in the system. If the user does understand that an AI is

present, their evaluation will depend on how the information is presented, as well as their own established ideas related to the ethical use of AI.

Notably, representation and embodiment may impact how AIs work within a team. While anthropomorphism can boost human collaboration with AI teammates (Giger et al., 2019), there is an ideal level of humanness the AI can achieve before causing unease in humans (see the section on the uncanny valley). Embodiment may also limit the range of tasks an AI teammate can engage in, such as by limiting mobility to wheels that work primarily on flat surfaces. Embodied AI are designed to physically interact with the world around them, and therefore can engage in tasks that require movement and manipulate stimuli in the environment. Disembodied AI cannot physically alter the environment but can still facilitate in completing the task and mission. A disembodied AI inside a smartphone need not navigate difficult terrain but can simply be put in a human's backpack.

AI social presentation

While most robots are designed to perform mechanical tasks, social robots are those whose purpose is to serve a person or group of individuals in a “caring interaction” (Sheridan, 2020), engaging with people at an interpersonal and socio-affective level (Breazeal et al., 2008; Sebo et al., 2020). While robots cannot feel emotional trust, robots, such as social robots, may be designed to facilitate human emotional trust. In a review by Sheridan (2020), the mechanisms by which robots achieve these functions have been broken down into three areas: affect, adaptation, and sensing and control. In the field of social robotics, affect involves how movements are used to convey and mimic human emotion (Sheridan, 2020). They can also express what is described in the field as the robot's personality, such as being competitive or relationship oriented (Sebo et al., 2020; Sheridan, 2020). These personality traits can impact individual humans and collective team motivation to work with the AI (Robert, 2018). Adaptation involves the use of information about the user to best suit their needs (Martins, 2019). Sensing and control in this context are an extension of collision and accident-avoidance functions, placing more attention on motion planning for human likeness (Turnwald & Wolherr, 2019). These features work to gain trust and acceptance in the human users (Martins, 2019) and allow the robot to integrate seamlessly in “harmonious coexistence” with users (Kostavelis et al., 2019). For example, robots that display group-based emotional expressions are perceived as more likable and trustworthy during a card game than those who expressed individual-based emotions (Sebo et al., 2020). While a great deal of attention has been devoted to the use of robots for helping the elderly, children, or those with disabilities (e.g., Ismail et al., 2019; Pennisi et al., 2016), the inclusion of robots in teams can improve human collaboration. For example, human to human interaction can benefit from the inclusion of a robot teammate if the robot teammate asks questions during pauses in operation to better focus operators on the task at hand (Strohkorb et al., 2016).

The ability to seamlessly engage with humans within a team may require designing AIs with cultural knowledge and the ability to assess the culture of individual humans and the team as a collective. Cultural knowledge would allow the AI to alter engagement with each human depending on their prescribed culture. The knowledge of what member of the team to engage with when an issue arises, how to alert individual humans to important information, and how to engage in a manner that will increase accurate use and use in general are all factors that can result in successful AI-enabled teams.

Summary of AI input

AIs can only be successful teammates if they are designed to process and handle both task relevant and team relevant information. Our individual AI inputs outline some of the factors an AI teammate can bring to the team. Factors such as the AIs' level of automation, representation and embodiment, taskwork-related inputs, and whether the agent was designed to engage in social behaviors or with teaming intelligence may impact how AIs interact with human teammates and their ability to successfully complete assigned tasks.

AI Processes

AIs designed to collaborate with humans will engage in processes that either enhance their capability to work with human teammates or to complete a given task. An AI that can tell humans apart, determine human needs and preferences, is explainable, and understands the chain of command will be capable of collaborating with human teammates. The ability to learn, audit decisions, and model and monitor humans and team performance will enhance an AIs ability to both efficiently complete tasks and assist humans in taskwork.

AI learning

For AIs to be effective teammates they must be capable of learning from their environment, their teammates, and their past performance. While limited and/or biased training data may result in errors or inaccuracies, corrections by humans or other AIs and data gathered from observing others can be used to improve performance.

Online versus offline learning

Humans continuously update mental models of the world. By contrast, most of today's computer systems only update on scheduled releases; they do not update their parameters until such updates can be checked and tested (Shane, 2019). Online learning allows models to continuously improve, but may be computationally expensive (Hoi, et al., 2018). An AI may require substantial computational resources to update its model from newly acquired data, which may need to be performed offline, rather than during a mission (Shane, 2019). An AI system deployed as part of a team may use either or both approaches, based on the needs of a specific situation. Online learning may leave the system vulnerable to input poisoning or catastrophic interference

(where a neural network model attempting to learn a new pattern overwrites old information), while offline learning may be unsettling for humans, who would be faced with a new AI every shift change or deployment.

Learning by demonstration

AIs in a field setting may not be deployed with an engineer who is responsible for training and updating their models. An AI in that situation may need to figure out how to do specific tasks without support from a dedicated trainer, and therefore will need to learn from team members. One of the most effective ways for humans to transfer procedural knowledge to other humans is to demonstrate the desired operation. An AI teammate would need to be able to interpret such demonstrations and translate them to its substrate, which might execute the operation differently from the human teacher (for a survey of the learning from demonstration literature see Argall et al., 2009). In sum, the flexibility that is required of AI teammates is contingent upon the agent's ability to learn or update on a chosen level of frequency. Learning and updating will allow AI teammates to remain current, especially in dynamic situations, and will decrease the likelihood that old data is being used to solve evolving problems.

Explainability

Human team members need to understand what an AI team member is doing, and why. This need is more than just providing interpretable outputs; interpretability is described in the literature as including both transparency in how the model functions internally, and post-hoc interpretation of how the model behaved (Mittelstadt et al., 2019). These metrics may be ideal for a developer or evaluator, but would often be overwhelming and off-topic for an operational user. Instead, an AI team member must provide explanations for its decisions. While an interpretable output allows for a gist mental representation of what the output means, an explainable output allows for a causal mental representation for how the outcome was achieved (Broniatowski, 2021). Humans provide explanations for a variety of reasons, and those reasons vary according to the current situation. That is, explanations are not just a combination of correlations and causes; they are contextual (Miller, 2019). Leveraging that context, explanations are contrastive, selected, and social (Miller, 2019; Mittelstadt et al., 2019). Having an AI that can explain its decisions will require the AI to lay out assumptions or information that may not be shared by the humans on the team, and such explanations will help the human team members to build a model of the AI's performance more quickly. For example, with appropriate fidelity for the context, an AI should expose the data it was trained on and how that data relates to the current data, the sensitivity of the model as a function of where current context is placed in the data space, and the output from the model with respect to prior decisions, and with a prediction of future impacts (Wang et al., 2020). AIs designed with the ability to build mental models of the problem state will have the further ability to inform humans whether or not the output from the model is reliable information or should be used in a given situation. Explainability has gotten

major attention from both the research and operational communities (Flournoy, 2020), as it likely correlates with appropriate cognitive trust.

Resolving competing demands

If humans give conflicting orders or suggestions, the AI needs to be able to prioritize which orders, constraints, or suggestions are most important to follow and possibly also choose among them (Russell, 2019). This prioritization may need to happen instantaneously in high-pressure situations, and the person who oversees a specific function may change across shifts or across situations. In human teams, the authority in each situation may be determined on the fly as conditions change. In situations with a formal chain of command, interpretation of orders may occur at multiple levels as people with different operational views try to adapt broad guidance to specific situations.

Decision auditing

Human team members are responsible for their decisions and devote some resources to determining whether a specific decision was correct or not, sometimes after the fact in demanding situations (e.g., Atunes et al., 2010). AI team members would also need to evaluate the quality of their decisions based on observable effects and feedback from human team members or supervisors. This monitoring could happen on the same time scale as updating the model, or it could happen after the fact in after-action analyses.

Modeling and monitoring

AI systems may need to monitor, and therefore model, both the specific individuals they are working with as well as the team as a whole. The requirement that humans be able to observe the workings of AIs (Johnson et al., 2014) can also be relevant in terms of AIs observing humans. First, AI systems may need to determine whether they are interacting well with individuals on the team. Such monitoring would require a model of the individual team members, perhaps varied in complexity based on how closely the AI works with that team member. The AI should also be able to model which of its actions affect the relationship with each human, in order to hypothesize whether some action is having a negative (or positive) effect on human team members' performance.

Second, AI systems need to be able to determine whether some action they are taking is contributing to better team performance or detracting from overall effectiveness. This process necessarily requires that the AI have some updatable model of the types of actions that contribute to team performance and that the AI constantly monitor whether performance is being affected.

Determining humans' true preferences and needs

Humans are not always the best judges of what types of assistance they need, and in stressful situations, metacognitive performance may not be prioritized. AI systems need to be able to derive human preferences and needs from observable behavior, including verbal behaviors. Russell (2019) suggests that this process of guessing and responding to human preferences should be constant if an AI system is going to provide the appropriate support to a human or human team.

Telling humans apart

Facial recognition and voice recognition systems have well-defined issues with bias (Klare et al., 2012; Tatman, 2017; Noble, 2018; Benjamin, 2019), but some level of recognition is required for an AI interacting with multiple humans. That recognition can be as simple as pointing out which teammate is authenticated into a particular identity or which teammate possesses the skills necessary to successfully complete a specific task. In some systems, AI teammates can use voice, face, body shape, gait, and other physical cues to determine which human they are currently interacting with (Collins et al., 2002; Wang et al., 2010; Chellappa et al., 2018; Folorunso et al., 2019). With some systems, the humans may only interact with the AI through some moderating interface, such as a graphical user interface, and these cues will be unavailable. Operationally, the ability to tell individuals apart is a necessary requirement for adapting to the needs and preferences of each team member.

Summary of AI processes

AI teammates designed to work in human teams can engage in many processes that will facilitate teaming. Learning will allow the AI teammate to be more flexible and change with changes in the task or mission. Explainability allows human teammates to understand what their AI teammate is doing and the reasoning behind that behavior. AI teammates will be better equipped to handle complex tasks and demands if designed with the capabilities to both resolve competing demands and audit decisions. Optimal human engagement, and ideally task completion, can occur if an AI teammate is designed to model and monitor humans, determine the true preferences and needs of teammates, and recognize faces and voices for human identification.

Emergent States

Here we describe AI cognitive trust as an AI emergent state. Similarly to humans, trust, specifically cognitive trust, permeates the entire model for individual AIs. More emergent states may become apparent with the increase in the inclusion of AI in team structures.

AI cognitive trust

Even if AIs are not capable of emotional trust and benevolence, they can still assess whether other AIs and humans in their teams are meeting expectations. As discussed in the earlier section on whether AIs can engage in trust processes and be teammates, we argue that AIs can have cognition-related trust. This trust can be programmed in or developed over time, depending on the nature of the AI (i.e., learning algorithms are present). This type of trust involves the AI depending on human or other AI behavior, such that these expectations can be violated. For example, a neural network that continuously considers operator actions and selections to build its knowledge base might be required to have cognitive trust in its human operators. Human performance is not always consistent, as humans are susceptible to performance degrading conditions such as fatigue, the presence of distractions, and illness. A system may monitor human performance for aberrant behavior or signals that may indicate decreased precision. The system may then choose to omit data from the operator at this time and notify the operator of this performance assessment. This kind of monitoring is already common practice within the TSA, and is referred to as Threat Image Projection (Hofer & Schwaninger, 2005). This activity will also be important in the case of bad actors and insider threats. The key to AI trust, be it of human or other AI targets, is for the designers of the AI to decide which capabilities of the others are dependencies for the AI and need to be monitored. Not every single aspect of human behavior or performance *can* be monitored, so cognitive trust is necessary. The pace of monitoring will also imply the level of vulnerability of the AI to human and other AI inputs, including decrements in performance. The level of trust assigned to a given teammate can determine modes of communication and collaboration, and the likelihood of selection for tasks.

AI Outputs

Regardless of whether an AI is designed to work with humans, its main design is to achieve a specific goal, transforming input data into answers of some sort. Therefore, AI outputs are centered around performance and the data created from the process of learning that can be used to train other AIs.

AI performance

Like humans, an AI's performance on tasks, such as accuracy and error rates, can be measured and shared with other teammates. AI performance can vary from how well it completed a specific task, to its ability to collaborate with human teammates. It is essential that AI performance remain observable to human teammates so that assessments about AI predictability can be made.

Training data

The outcomes that humans derive from team activities are not mirrored in AI systems; the system will not feel satisfaction at a job well done or be stressed if the team is not functioning well. The metrics derived from an individual AI's performance in a team, however, can and should be used to improve the performance of that AI or similar AIs in future teams. This can be in the form of online learning (that is, the explicit process of using new sets of validated input/output data to retrain the underlying model). It can also be broadly in support of the continuous test and evaluation of AIs: every use of an AI is an opportunity to improve (Flournoy, 2020).

Summary of individual AI outputs

The outputs from the individual AI are focused on accuracy of the task at hand, and improvement of the task in future iteration. Other "results" from the AI -- such as interim results, or additional rationale in support of explanations -- are covered elsewhere in the model.

Summary of Individual AI IPEOs

Like their individual human teammates, AIs bring their own inputs into the team, engage in their own processes and emergent states, and produce their own outputs. As individuals, humans and AIs are nested within, impact, and are impacted by teaming and organizational systems.

Team Level

The team level consists of a combination of at least two individual humans and one AI. Together, team members bring in input factors that impact processes and emergent states. The processes that team members engage in and the emergent states that arise then alter team outputs.

Team Inputs

Aggregated individual human and AI inputs combine to create team inputs. These inputs, such as team personality, team KSAOs, and other team characteristics help determine how well a team will engage in team processes and emergent states. Team personality and team KSAOs are team inputs that impact the team's ability to work collaboratively together and complete taskwork. Cohesive teams are more often made of individuals who have a proclivity for teamwork.

Team composition inputs can be quite complex. Team composition can be represented by the average of the teams' individual traits (e.g., mean Agreeableness), variability of the trait within the team (e.g., standard deviation of team members' Agreeableness), the highest of a trait within the team (e.g., maximum Agreeableness), and so on. Thus, one can examine team composition as aggregation and heterogeneity, including a requisite team KSAO such that one team member

must have a characteristic (Kozlowski & Klein, 2000). In staffing a team, managers need to identify the critical team tasks and KSAOs of both the individuals and the team overall (Mathieu et al., 2014). These KSAOs can include both position-specific and team-generic KSAOs (e.g., organizing skills, team orientation), as well as both a relative contribution (e.g., weakest member, strongest member) and a team profile approach (e.g., a language skill is required within the team; diversity of functions; Mathieu et al., 2014). This section reviews the literature on team composition inputs.

Team personality

Team design includes team personality compositions (Kozlowski & Ilgen, 2006). Research on team aggregate personality has been sufficiently copious as to allow for a meta-analysis (Bell, 2007). Bell (2007) compared team performance in lab and field settings and found that team mean and minimum Agreeableness are positively related to team performance in field settings, such that one disagreeable member can undermine team performance (corrected population correlations of .31 overall, .34 for mean, .37 for minimum). In a more focused study of team personality, team mental models, and team coordination, Fisher and colleagues (2012) found that the mean levels of cooperation as a facet of Agreeableness (but not dispositional trust, or the tendency to believe others are well-intentioned) was positively related to team-focused shared mental model similarity (team $n = 32$). Team mean Extraversion in field settings was slightly positively related to team performance (corrected population correlation = .18), as was team mean Emotional Stability (corrected population mean = .21; Bell, 2007). Providing support that team KSAOs related to harnessing team processes are related to team performance, Bell (2007) also found that team mean Collectivism in field settings and team mean preference for teamwork were also positively related to team performance (corrected population correlation = .40 and .26, respectively). Team emotional intelligence was also positively related to performance, but not as strongly and mainly in the lab, and there were few studies before 2007 (.20 in lab studies and .10 for field studies). Finally, Bell (2007) found that team mean Openness to Experience was somewhat but significantly positively related to team performance (corrected population correlation = .25).

Individual personalities, and the collective personality of the team ultimately impact performance and may contribute to teaming with AI. Teams that have pro-teaming personality traits will be advantaged in their ability to collaborate and work cohesively. An AI teammate could use team personality compositions to predict the type of assistance the team may need. For example, teams with disagreeable members may particularly benefit from AI teammates that help manage team conflict to elicit broad participation rather than engender negative affect (see Conflict section).

Team KSAOs

Above and beyond team personality, teams need the right mix of levels and areas on taskwork-relevant KSAOs, such as senior and junior medical staff (Klimoski & Jones, 1995). Teams need

particular mixes of skills, abilities, training, representation of functional areas, group affiliations, credentials, and connections, as well (Jones et al., 2000). For different kinds of organizations, worker requirements related to fairness and stakeholder acceptance require different profiles of team members, such as making sure different constituent groups are represented (Klimoski & Jones, 1995). Teams may need homogeneity or heterogeneity in terms of strategic outlook and how members go about finding problems, depending on the type of tasks involved (Jones et al., 2000). Just as Conscientiousness is important for individual team members to have, it has also been found to be important for the team members overall (Hertel et al., 2005): Bell's (2007) meta-analysis found that average team Conscientiousness was positively related to team performance in field settings (corrected population correlation of .33).

Team members may be required to be intelligent: Bell (2007) found that team general mental ability (GMA) was positively related to team performance, but the relationship was stronger for lab tasks than for in field settings (corrected population correlations of minimum team GMA in lab = .48; minimum GMA in lab = .42; mean and sum in lab settings = .33; overall mean GMA in field settings = .26). Bell's (2007) findings suggest that the most intelligent and least intelligent team member is important in lab settings, but that intelligence is less important for field studies as they have, so far, measured team performance. Moving the field forward, Woolley and colleagues (2010) found evidence for collective intelligence, or the intelligence of an interacting team⁵, which is not strongly associated with average or maximum individual intelligence. Collective intelligence is thus more than the aggregate of individual intelligence but includes the social sensitivity of group members and whether the team can leverage the unique contributions of the team members (Woolley et al., 2010).

One particular setting that is relevant to AI-Human teams are human virtual teams. Some of the same required KSAOs may apply. Research on virtual (or distributed) teams suggests that some of the KSAOs required across an entire team include attributes relevant to virtual cooperation, such as: self-management skills, interpersonal trust and trust-maintenance skills, communication skills, intercultural skills, dependability, conflict management skills, and expertise with technology and communication medium (Hertel et al., 2005; Schulze & Krumm, 2017). Schulze and Krumm (2017) suggest that individuals (humans) working in virtual teams must be able to deal with, and have the motivation to deal with, three facets of virtual teams that can cause difficulty: technology use including low information richness and asynchronous contexts, cultural differences including different communication styles and languages, and geographic dispersion which includes both temporal and spatial dispersion. In particular, power struggles in distributed teams can activate faultlines, such as over asymmetries in language fluency in the common language in a globally distributed team (Hinds et al., 2014).

⁵ Not to be confused with the usage of the phrase "collective intelligence" for crowdsourced projects with non-interacting large groups of people.

Many of these skills and challenges are similarly relevant to humans working in AI-Human teams. Teams that contain more agreeable and open members may be more likely to accept new AI teammates with ease. Teams high in collectivism will have members that are more inclined to engage in behaviors that are good for the group, including collaborating with an AI teammate that may ease some of the workload for the collective team. While some human teams may consist of members that together possess all of the relevant KSAOs for a given mission, some teams may need assistance in covering missing KSAOs. AI teammates, armed with the knowledge of its human teammate's KSAOs may fill in for missing team members, levels, or specific KSAOs certain team members may be lacking. Operational planners from our on-the-ground team may possess the knowledge to determine how severely the storm has hit and what regions may need resources fast but may not be able to assess the best methods of getting resources to goal locations without putting the rescue squad in harm's way. This situation may be especially true if the operational planners consist of individuals with little experience. The AI optimization system can use past data and current information to determine optimal methods of dispersing resources in the safest manner for all.

Additional team characteristics: Selecting for cohesive teams

In addition to the personality traits above, Bell and Brown's (2015) review has suggested a set of precursors for cohesive teams, specifically, which go beyond team personality or traits. Team cohesion includes the dimensions of task cohesion, which is task commitment, and social cohesion, which involves interpersonal attraction and shared identity (Bell & Brown, 2015; Braun et al., 2020). Bell and Brown (2015) suggest that interpersonal attraction is driven and explained by propinquity, or physical and social closeness; complementarity, such as when one's weaknesses or traits are complemented by another's; and reciprocal relationships. Shared identity readiness would be driven by, among other features, a preference for teamwork and psychological collectivism (as noted previously) and identifying with the profession on the team level (Bell & Brown, 2015). Finally, team goal priority (a facet of psychological collectivism where team goals are prioritized over individual ones), a learning goal orientation, and being able to prioritize the team's tasks (e.g., not having multiteam membership or being overtaken; Bell & Brown, 2015).

Summary of team inputs

Team inputs can be an aggregate of individual inputs, such as with average or variability of team personality or intelligence. For instance, teams with high average Conscientiousness, Agreeableness, Collectivism, intelligence, and so on tend to have better performance. Team collective intelligence, which goes beyond average intelligence to include social skills and the ability to participate equally, is also a valuable input. However, team inputs can include team composition that goes beyond averages or variability. A team may require a mix of specific skills, such as a medic, a navigator, a person who understands how to fix a robot, and so on. Cohesive teams tend to also have team-level qualities such as closeness, complementarity, a

shared identity, and the ability to prioritize team tasks. Further, teams will have their own artifacts, stated rules and shared expectations that are difficult for even humans who are part of the in-group to grasp at times.

Team Processes

The phrase “team processes” has been used to describe a large range of phenomena. Marks and colleagues (2001) discuss three types of teamwork processes, lumping other types within them: Interpersonal processes, such as managing relationships and conflict management, action processes, or processes during goal-directed activities such as coordination and monitoring, and transition phase processes, such as planning, goal specification, and strategy (see also Driskell et al., 2018). For instance, a real-time language feedback system that monitors and provides instructions to a group of humans would be an AI that assists a human team (e.g., Tausczik & Pennebaker, 2013), although an AI could also simply work in a group on tasks without having a specific teamwork improvement role. In this section, we focus on a set of specific processes that have been raised as vital to teamwork and how AI teammates may assist in these processes.

Coordination

One of the essential team processes is coordination. Individuals within a team, be they human or AI, need to be able to each perform their tasks in a way that results in a team response and output. Teams have to coordinate internally but also potentially with external teams (Driskell et al., 2018). Much of this coordination needs to be tacit, in that individuals cannot be required to explicitly discuss everything they are doing (Wittenbaum et al., 1998). Coordination can be dysfunctional when individuals duplicate effort or rely on assumptions and stereotypes that are incorrect (Wittenbaum et al., 1998). Coordination often involves monitoring and feedback of other team members (Driskell et al., 2018; Honts et al., 2012; Salas et al., 2005; Thompson & Cohen, 2012), which requires that both human and AI teammates are observable and can both receive and send information. Also related to coordination are workload sharing and backup behavior, where team members anticipate each others’ needs, help each other during periods of high workload (Salas et al., 2005), and carry their own weight (not do social loafing; Campion et al., 1993). Coordination is also related to cooperation (Campion et al., 1993), and we argue it can also be related to the directability aspect of the AI-human team relationship (Johnson et al., 2014). Being able to both direct others and receive direction is part of coordination even among humans. Thus, coordination is inherently related to emergent team cognition factors, such as shared mental models and transactive memory processes (Thompson & Cohen, 2012; see those sections). AI teammates can play a pivotal role in aiding in team coordination. Human teammates may engage in tasks that exceed their cognitive workload limits but may not know who to turn to in the team for assistance. An AI teammate that is monitoring team workload may alert a capable member who is currently underloaded.

Communication

One of the main processes is team communication, which is “an exchange of information, occurring through both verbal and nonverbal (e.g., email) channels, between two or more team members” (Marlow et al., 2018, p. 145). Different aspects of communication have been identified and studied as a necessary part of teamwork. Communication quality, rather than frequency, is positively related to team performance (Marlow et al., 2018). Among other features, communication can be face-to-face or distributed; involve information elaboration, knowledge sharing, and frequency; and be related to the task, processes, relationships (Marlow et al., 2018). A recent review of robots in human teams notes that different robots have been designed to communicate nonverbally as well as verbally, and to express emotions and robot personalities (Sebo et al., 2020). Salas and colleagues (2005) describe closed-loop communication as vital to coordinating mechanisms within teamwork. Closed-loop communication involves a sender beginning a message, the intended recipient receiving and interpreting it and acknowledging the receipt of the message, and then the sender checking to make sure the message was received (Salas et al., 2005). Regarding human-AI teams, communication also entails understanding not just language (including jargon, vocabulary, and grammar), but accents, dialects, and pitch, both from the human and the AI. Current natural language processing algorithms are hindered by race and gender biases (Tatman, 2017; Zhao et al., 2018), which could drastically impact human-AI team communication. Therefore, diverse and representative training data should be used when designing AI teammates. In addition, the timing of communication over the life of a project can be useful. Teams that coordinate attention, as reflected in bursts of intense email communication, are more likely to be able to leverage their resources (Mayo & Woolley, 2019). The observability aspect of AIs can be applied to humans as well: In human-human teams, this observation may occur through communication, including nonverbal symbols. Both the humans and the AIs need to have effective communication (Demir et al., 2016). While closed-loop communication and communication patterns can be important, other aspects of communication have also been discussed in detail in the literature on teams: information sharing, particularly of unique information, along with participation equality/equity, and entrainment (e.g., Levitan, 2020; Mesmer-Magnus & DeChurch, 2009).

Information sharing and participation

Equal, or equitable, participation and information is important in teams so that different group members can effectively solve problems (Mesmer-Magnus & DeChurch, 2009), and so that diverse teams can share unique information and be innovative (Paletz & Schunn, 2010). Groups with more equitable turn-taking in conversation are more collectively intelligent (Woolley et al., 2010). A series of social psychological experiments revealed that team members are more likely to discuss information that is not unique and that this is liable to hurt team decision making (e.g., Stasser et al., 1995; Stasser & Titus, 1985). A metaanalysis found that information sharing, specifically of unique information, is vital for team performance and cohesion (Mesmer-Magnus & DeChurch, 2009).

With regards to AIs in human teams, these findings suggest that whatever else, unique information from both a range of the humans in the team and the AIs needs to be appropriately shared. If an AI is limited in how much it can communicate (i.e., three-word utterances), this limitation can lead to poorer team outcomes (Demir et al., 2020). Communication is not only relevant in terms of what an AI expresses; AIs can positively impact human and human-human verbal communication, as well (Sebo et al., 2020), such as by promoting equal participation and information exchange (Tausczik & Pennebaker, 2013).

Accommodation and entrainment

Another aspect of communication is accommodation, in which people talk more similarly to each other over time and adapt to each others' communication behaviors (Giles et al., 1973, 1991). Speech entrainment, specifically, is how people talk more similarly to each other over time, generally indicating social closeness and rapport versus distance (Beňuš, 2014; Kory-Westlund & Breazeal, 2019; Levitan et al., 2012; Levitan, 2020). Entrainment is related to conversation quality, including in human-computer dyads (Levitan, 2020). In particular, recent studies suggest that teachable robots or agents that entrain to humans result in more learning and engagement than without social language and entrainment (e.g., Lubold et al., 2018). However, different types of acoustic-prosodic entrainment are generally not significantly associated with each other and show different impacts on conversational quality, suggesting that entrainment is not a single construct (Levitan, 2020). A recent study with a spoken dialogue system suggests that intensity (loudness) entrainment has a positive effect on trust and pitch intensity has a negative effect on trust, where trust is specifically trust in conversational avatars (Gálvez et al., 2020).

A great deal of the research on entrainment has examined lexical matching, specifically linguistic style matching (LSM). LSM is when individuals match on specific words, generally measured as sets of function words using Linguistic Inquiry Word Count (LIWC; e.g., Tausczik & Pennebaker, 2010). LSM has been found to be positively related to cohesion in dyads (Gonzales et al., 2010), social support (Heuer et al., 2020), and friendship formation (Kovacs & Kleinbaum, 2019), but negatively related to team performance, controlling for team tenure (which did not moderate the relationship, Heuer et al., 2020). Research using LSM has indicated that people are more likely to entrain to others who are of higher status (Danescu-Niculescu-Mizil et al., 2012; see Van Swol & Kane, 2019 for a review). However, further research using this measurement, as well as acoustic-prosodic entrainment, would be useful: LSM categories are often unrelated to observer judgments of interaction quality (Niederhoffer & Pennebaker, 2002).

Intrateam conflict

Intrateam conflict entails differences or incompatibilities between people (Jehn & Bendersky, 2003). It can be explicitly stated as disagreements or unspoken, and it is often measured as perceptions (Paletz et al., 2011). Researchers suggest that the most harmful conflict includes negative emotions such as anger or contempt (e.g., Amason, 1996; Barki & Hartwick, 2004;

Gottman & Notarius, 2000). Although there is still some debate in this area (De Dreu & Weingart, 2003), it seems that under some conditions, conflict serves to increase unique information sharing (De Dreu & West, 2001; Jehn & Mannix, 2001; Nemeth, 1986). A sizable literature examines under what conditions conflict can support creativity, innovation, and good decision making, concluding that it is possible when: the conflict is separate from negative emotions and focuses on the task rather than personal incompatibilities (de Wit et al., 2012; Hulsheger et al., 2009); open-minded disagreement is expressed (Kellermanns et al., 2008; Tjosvold et al., 2014); and the conflict is mild rather than intense (Todorova et al., 2014). There is some research suggesting that AIs can intervene in and be designed to mitigate team conflicts among humans (Sebo et al., 2020), such as by increasing awareness of a personal attack (e.g., Jung et al., 2015).

Task conflict can become heated, however, so one issue is how to keep humans from getting too angry at AIs. If an AI agent violates expectations, it can lead to Human-AI conflict, possibly diminishing trust in the AI's predictability. Violations of trust can result in verbal displays of frustration and, if the behavior is not resolved, can reduce or halt use of the AI entirely (Groom & Nass, 2007). Human-AI conflict, therefore, must be resolved swiftly. This problem requires observability on the part of the AI. Knowledge of the source of the conflict can aid human teammates in either rectifying the issue or preparing for the possibility of the conflict to arise in the future. Other methods, such as adequately training human teammates on the uses of and possible issues with the AI teammate, can reduce the likelihood of conflict prior to task related engagement.

Summary of team processes

Team processes, such as coordination, communication, and reducing interpersonal conflict, are vital to the team's success. Without coordination, teammates may experience cognitive overload on tasks without reprieve, resulting in unsuccessful completion of tasks. Lack of communication will inevitably hinder coordination and reduce the likelihood that team members have sufficiently overlapping mental models. Interpersonal conflict, mainly when negative emotions such as anger are involved, can significantly reduce team productivity and jeopardize the mission. AI teammates can assist in increasing overall team coordination and encouraging closed-loop communication, particularly the sharing of unique and valuable information, and mitigate interpersonal conflict. Human and AI teammates must also coordinate and communicate with one another and resolve, limit, or leverage Human-AI conflict.

Emergent States

In the following section, we discuss emergent states which have also been conceptualized as outcomes. Some of these states can fall under the concept of team metacognition, which involves thinking or feeling about thinking (Thompson & Cohen, 2012). Similarly, the process of

individual knowledge becoming team or group knowledge broadly is referred to as macrocognition (Fiore et al., 2010). Macrocognition includes several of the constructs discussed in this section that relate to group-level cognition, such as transactive memory and shared mental models. Some of the emergent states described in this section are not exactly considered macrocognition, however, such as team cohesion.

Team cohesion

Team cohesion is an emergent team state that is composed of social and task (action) components and essentially means the shared commitment and attraction to the goals and each other of a team (Braun et al., 2020). Team cohesion has a consistent small but significant effect on team performance (Beal et al., 2003; Gully et al., 2012), and this relationship is bidirectional (Mullen & Copper, 1994; Braun et al., 2020). However, this effect is stronger for tasks with high task interdependence (Gully et al., 2012). Given the bidirectional relationship between team cohesion and team performance, AI teammates could monitor and assist in team performance in an effort to increase team cohesion (see Sebo et al., 2020). A lack of cohesion in an AI-enabled team, which can thus be the cause, correlate, or effect of poor performance, may result in reduced trust in AI teammates and therefore lack of use of the AI. Team cohesion is often divided into task and social cohesion, where task cohesion is team task commitment, and social cohesion is the interpersonal attraction element (Bell & Brown, 2015). Task cohesion is generally considered more associated with performance than social cohesion (Bell & Brown, 2015), and iterative research suggests that the cohesion-performance relationship weakens over time (Braun et al., 2020).

Trust

Within the human-human team literature, mutual trust is considered essential for coordinating teamwork, participation, interpreting others' behaviors appropriately, and sharing information (Salas et al., 2005). Similar to individuals, team level trust can be influenced by AI embodiment and representation (Lee et al., 2006; Li, 2015). Initial team trust levels can reflect larger organizational representations of AIs and teams and will generally evolve as time progresses. Team trust can also be impacted by individual interactions. Not only is trust important between humans, it is also important for humans to trust their AIs in order to effectively work together. Teammates may see the AI as non-judgmental and thus feel safer sharing personal information at a level they don't typically share with their human teammates and this heightened trust in AI could reduce overall group interaction (Krämer et al., 2018). As an emergent state, trust can be gained throughout the team's work, but perhaps more easily, lost at any point. There are hindrances, such as team and individual culture misalignment, personality differences, individual experience, expectation levels, credibility concerns, learning adjustments, presentation and control concerns, and the security of the technology itself, to developing and continued trust, and later teaming outcomes can be driven disproportionately by early interactions. For example, when a machine learning algorithm is unexpectedly updated and now gives a different, even if

correct answer, team members may need to re-develop a bit of trust in the reliability of the algorithm. Similarly, an AI team member programmed with social cues from one national culture, when working with teammates of another national culture, will likely lead to at least one faux pas on behalf of the AI teammate, and can create an obstacle to trust development.

Potency and group efficacy

Group potency, also known as group efficacy, has long been considered positively related to team effectiveness (e.g., Campion et al., 1993, 1996; Guzzo et al., 1993). Group potency and group efficacy are shared beliefs held by the group overall about how effective the team is, in general (Guzzo et al., 1993). Group efficacy is often positively related to group performance (e.g. Campion et al., 1993, 1996), although it is possible this finding is due to the group having accurate beliefs about its external factors such as resources and internal factors such as member KSAOs and goals (Guzzo et al., 1993). Group potency and efficacy can greatly impact how the team interacts with AI teammates. If the team is aware of where they fall short, and they understand that an AI teammate can fulfill an important role, team members may be more inclined to work with an AI teammate. Conversely, teams that have a full understanding of how effective the team is may have difficulty embracing an AI teammate that can optimize effectiveness.

Adaptability

Adaptability, related to coordination, is the team's ability to change their strategies based on new information and dynamic situations (Salas et al., 2005). Like coordination, adaptability required metacognitive skills. Changing workload and task demands can require adaptability (Entin & Serfaty, 1999). Team adaptive expertise involves the team being able to respond and change behavior to non-routine tasks or learning how and when to apply heuristics such that the team can handle novel situations (Paletz et al., 2013; Schwartz et al, 2005). For many kinds of tasks teams of humans and AIs will need to be adaptive in this manner and not limited to repetitive routine tasks. While currently it may be the human driving the adaptive expertise, eventually one hopes that the AI can contribute in this way as well.

Shared mental models

While an individual's mental model is a cognitive representation of a thing, place, person, and/or the relationships within a system (Johnson-Laird, 1980), shared mental models (also called team mental models) are representations of constructs, including team and taskwork, that are similar between individuals (Lim & Klein, 2006; Mathieu et al., 2000; Mohammed et al., 2010). Shared mental models are thus "dynamic, simplified, cognitive representations of reality that team members use to describe, explain, and predict events" (Burke et al., 2006, p. 1199). A shared mental model about taskwork could be about the work to be accomplished and its constraints, whereas a shared mental model about teamwork could include who is doing what, organizational

knowledge, and social networks (Thompson & Cohen, 2012). Multiple shared mental models can be held simultaneously (Mohammed et al., 2010). The sharing of mental models can be a matter of degree, with different aspects more shared than others (Cronin & Weingart, 2007). Shared mental models can have two qualities: their similarity between team members and their accuracy, or how much they overlap with reality or with expert opinions (DeChurch & Mesmer-Magnus, 2010; Edwards et al., 2006; Thompson & Cohen, 2012). Both the similarity and the accuracy of shared mental models are important for team effectiveness (DeChurch & Mesmer-Magnus, 2010; Edwards et al., 2006; Kozlowski & Ilgen, 2006; Lim & Klein, 2006; Mathieu et al., 2000), with, as of 2010, more research on, and support for, the similarity rather than accuracy of mental models (Mohammed et al., 2010). Shared mental models support implicit communication, coordination, and common activities, even when the overlap is not perfect (Mohammed et al., 2010; Salas et al., 2005). Antecedents to shared mental models include experience and educational and organizational similarity, and team general mental ability has been found to predict taskwork mental model accuracy (Mohammed et al., 2010). Paletz and Schunn (2010) argued that multidisciplinary teams would have more difficulty achieving similar shared mental models, based in part on Cronin and Weingart's (2007) theory that members of diverse groups may conceptualize and represent problems differently.

Shared mental models are also an important aspect of AI-enabled teams. AI teammates must be aware of the mental models of the individuals in the team, which would facilitate effective communication with human teammates. The type and amount of information that is disseminated from the AI, such as the aerial drone in our rescue squad, to individual human teammates will vary depending on the mental model that team possesses about relevant information. For example, the aerial drone may supply the navigator with coordinates and weather patterns while letting the navigator know the best course of action is to keep going north. If the AI understands the shared mental model, the AI can maintain and update this mental model for itself and the human team members as situations change. In a study of simulated human-robot dyads (really two humans), Demir and colleagues (2020) found that teams were more successful when the 'robot' understood the limitations of the 'human', having a more shared mental model.

Transactive memory

Transactive memory systems (TMSs) are a different type of emergent, shared cognition, wherein the task of remembering different pieces of information (memory) is divided up among team members or in a dyad (Wegner, 1986; Palazzolo, 2017). TMSs include metacognition, where individuals are aware of the knowledge and expertise held by the different team members. TMSs are made up of knowledge and communication components (Palazzolo, 2017). Knowledge components include the different peoples' expertise and their expertise relative to others, perceptions of expertise and their accuracy, and updating of that knowledge (Palazzolo, 2017). TMSs also have communication components: communication to assign and retrieve information, but also general communication that is not specific to the task (Palazzolo, 2017). Both TMSs and

shared mental models include overlapping, integrated knowledge structures (Palazzolo, 2017). The emphasis for TMSs in the literature is typically more on the differentiated knowledge structures, or the unique and non-redundant knowledge between team members (Hollingshead, 1998; Lewis & Herndon, 2011). Team members engage in memory processes but across different people, such as when a team member asks a different team member to retrieve information (Lewis & Herndon, 2011; Palazzolo, 2017; Wegner, 1986). TMSs arise from norms, rules, and structures that support sharing information (Lewis & Herndon, 2011). The ideal TMS has enough redundancy for team members to coordinate, communicate, and understand each other, but enough specialization to provide benefits for each team member (Palazzolo, 2017).

In their meta-analysis of team cognition, DeChurch and Mesmer-Magnus (2010) categorized transactive memory as compilational emergence and shared mental models as compositional emergence. They found that transactive memory had a stronger association with team performance than shared mental models (DeChurch & Mesmer-Magnus, 2010). Based on a review of over 70 TMS papers, Ren and Argote (2011) created an integrative framework of antecedents and consequences of TMSs. The antecedents are as varied, if not more, than those detailed in this paper, and include team composition inputs such as member technical competence and team member assertiveness; team-level inputs such as task interdependence, goal interdependence, shared experiences, team size, technology and virtuality, and team familiarity; and organizational inputs such as acute stress and geographic distributions. Consequences include team performance behaviors (i.e., team learning, reflectivity, and creativity), team performance outcomes (effectiveness, efficiency), and team satisfaction. Finally, Ren and Argote (2011) detail moderators such as group size, task complexity, team membership change (attrition/newcomers), and face-to-face versus computer-mediated communication. Thus, TMSs are as complex as any other major team emergent state, interacting with most of the inputs and outputs described in this study.

Similarly, to shared mental models, AI teammates must have the ability to assess and update team TMSs. With an adequate assessment of the general knowledge areas each human teammate specializes in, AIs can assist other team members who may be experiencing cognitive overload on a given task by suggesting an ideal teammate to coordinate with. It is also vital that this information be updated for the AI, especially since humans are often increasing or changing their expertise. AI teammates with an updated team TMS may also have the ability to pinpoint holes in team knowledge that may be crucial for successful completion of the mission.

Team situation awareness

Individual situation awareness is necessary for team situation awareness. Each team member, especially in dynamic situations like our natural disaster example, must be aware of information in the environment, especially information relevant to individual KSAOs, in order to accurately and swiftly make decisions. Team situation awareness, however, is more than the aggregate of

individual situation awareness. Knowledge, communication, and shared mental models are important aspects of team situation awareness. Communication between team members allows for a transfer of knowledge about the task, the mission, or any other relevant information. This transfer of knowledge allows team members to update relevant schemata and enables them to better comprehend their individual tasks (Salas et al., 1995) and results in overlapping mental models. The more overlap there is in team mental models, the greater team situation awareness will be. AI teammates are not only useful in updating individual team members on relevant information about the situation, they can also assist and engage in team communication and coordination. The AI agent can be tasked with determining who receives information, and when and how that information is disseminated. For example, the drone may supply the navigator with promising locations where several individuals are alive well in advance, while supplying other team members with simplistic navigation information. Supplying team members who do not have the necessary KSAOs to navigate with relevant information without unnecessary details relieves the navigator from communicating information and allows the navigator to focus on the task of getting to a specific location.

Summary of team emergent states

Team emergent states are necessary for all individuals in a team, whether human or AI, to effectively collaborate, communicate, and ideally complete the mission. Team cohesion and trust ensure that individual team members are working toward a shared goal and can depend on each other. Team efficacy can impact performance, while adaptability allows teams to continue working even when the problem set has changed. Teams that engage in transactive memory and have shared mental models are better equipped to both collaborate and engage in tasks individually. Indeed, a recent meta-analysis of team cognition (examining both transactive memory and shared mental models) suggests a positive relationship between these and team performance, particularly when the team is heterogeneous in terms of social category composition (e.g., mixed gender) and is reliant on individuals outside the team to achieve team goals (Niler et al., 2020). While team situation awareness requires each individual member to both be situationally aware and engage in communication and collaboration with other team members and can increase team performance.

Team Outputs

Just as individual satisfaction, effectiveness, and efficiency are individual human outputs, so too do teams have **productivity, satisfaction, effectiveness, and efficiency**. Team task effectiveness, such as productivity, is generally the output that managers are interested in, though many are also interested in social effectiveness (e.g., satisfaction). Specific outcome measures are debated and can be domain specific. For example, in hospital teams, outcomes could include number of patients seen by doctors, number of hospital readmissions, individuals released from an intensive care unit, and so on. In a marketing team, it could be the number of advertisements

created (productivity) or the number that won awards (quality). In rescue operations with AIs, social effectiveness could include satisfaction with AI and human teammates or desire to work with the same group again, and task effectiveness could include amount of ground explored and number of people rescued over a week. The promise of increased productivity and team satisfaction are the outcomes that spur the introduction of AIs, after all. If an AI can be designed to facilitate teamwork and taskwork and minimize frustration, inefficiencies, and team dysfunction, with both itself and with other teammates, then it is likely some of the designer's goals for that AI have been met.

Teams can also respond to and adopt AI technology. In some situations, individuals have a choice as to whether they can continue to use a type of technology. They can choose to adopt, integrate (with workflow or other technology), use, or misuse the AI (Parasuraman & Riley, 1997). In other situations, these choices are made at the level of the team, such as by a team leader, the local technology experts, or even an organization as a whole. These levels of choice may also interact, such as when a team adopts and requires a particular AI, but some individuals create workarounds to avoid using it or misusing it. An entire team may also misuse an AI, if they are given inappropriate guidance, or a norm develops about how to use the AI (which is wrong). The normalization of deviance, when people get used to risky behavior and start to consider it appropriate, was termed to describe factors leading to the Space Shuttle Challenger disaster (Vaughan, 2016), but is relevant with any issue involving risk perceptions at a group level. An organization could also mandate the use of a technology which team leaders then ignore. The Technology Acceptance Model (e.g., Marangunić & Granić, 2014) is relevant here and discussed in more detail under organizational and societal outputs.

Summary of Team IPEOs

Within a human-AI teaming system, humans and AIs will collectively contribute to team level IPEOs. While this section heavily references the human-human teaming literature, it is important to fully understand the human teaming system to design AIs to interact with humans within that system. Human-AI teams are also contingent upon the organization and society the team is nested within.

Organization and Societal Levels

Entities within the organizational and societal levels involve all of the larger systems a team can be nested in. Organizations can range from departments, to companies, to conglomerates. Societal influences can come from norms with the state, country, nation, or world. Organization and societal influences can impact whether humans can work with AI teammates and limit what Human-AI collaboration can look like.

Inputs

Organizational inputs consist of factors, such as the legal limits on AI, shared cultural values, bias, and organizational context and work design, that can greatly impact the design and implementation of AI teammates, and the proclivity of individual use of the AI.

Legal limits on AI

The legal landscape around AI and AI applications is evolving rapidly. As Shneiderman (2020b) notes in arguing that AIs cannot be true teammates, “Only humans are responsible legally and morally” (p. 5)⁶. Humans can create AIs that have legal and moral implications, which then need to be controlled via legal and social means. For instance, many jurisdictions seem to be moving toward banning facial recognition technology (Conger et al., 2019). Team performance requires the AI to be able to tell team members apart. Any AI engaged in teamwork could leverage some form of facial recognition capability, if it is accurate enough and legal, but cannot if it is not legal.

The General Data Protection Regulation (GDPR), Chapter III, Section 4, Article 22, restricts automated decision making about individuals, including profiling (European Parliament and Council of European Union, 2016). This restriction may be picked up by other jurisdictions and may be expanded. This type of automated decision making is not currently prohibited in the United States. In places where a GDPR-like regulation exists, certain tasks may require legal sign off when performed by an AI system.

Shared cultural values

In the context of organizations, culture is:

“(a) a pattern of basic assumptions, (b) invented, discovered, or developed by a given group, (c) as it learns to cope with its problems of external adaptation and internal integration, (d) that has worked well enough to be considered valid and, therefore (e) is to be taught to new members as the (f) correct way to perceive, think, and feel in relation to those problems” (Schein, 1990, p. 111).

Schein’s Corporate Culture Survival Guide (Schein & Schein, 2019) described the three-tier model of organizational culture. The model originally focused on artifacts, which are tangible and the most easily identifiable part of an organization’s culture; espoused values, which are the stated rules and expected behaviors; and assumptions, which are the implicit behaviors, values and norms that can be difficult to change but less obvious. The model has expanded to cover the practice of culture. They include a socio-technical examination of organizations and a “bullseye of culture” model. This model takes a more nuanced look at organizational culture, including

⁶ Although, there are evidently instances in Roman and British and American law where inanimate objects were held legally responsible for crimes such as the death of a human (Holmes, 2020).

subcultures, macro cultures, micro cultures, social/relational factors, and technical factors (Schein & Schein, 2019). For example, some organizational cultures encourage quality, such that quality improvement is recognized, measured, and trained, and employees participate in strategic planning (Johnson & McIntye, 1998). Organizational cultures can be influenced by industry factors including societal expectations and the competitiveness of the industry in which the organization exists (Gordon, 1991). At face value, an organization's shared cultural values may seem easy to identify, but they are usually much more deeply ingrained, and they still might not make sense to an outsider. Part of being on a team is understanding how to navigate the nuances of the organization and team's culture. Without being able to intuit the correct behaviors and discussions for its respective group, the AI and team will have to figure out a way to overcome this hurdle.

Bias

AI systems reflect the values of their designers and funders. What is chosen to be built, what data are collected, and how systems are designed is all impacted by existing power structures and priorities of dominant groups (D'Ignazio & Klein, 2020). In many cases, these designs reflect historical processes that have led to unjust outcomes. An extreme example might be a system that did not recognize darker skin or non-standard (Mid-Atlantic American) accents (Buolamwini & Gebru, 2018; Klare et al., 2012; Tatman, 2017; Schwemmer et al., 2020). Both cases have occurred in simple systems on which an AI could be built, and many AI systems reflect societal biases that are encoded in the training data that they are provided with (O'Neal, 2016).

Many AI systems that are intended for use as assistants or helpers are coded in feminine ways (like Siri, Cortana, or Alexa, which have default female voices), reflecting both a bias in society and a lack of desire to combat that bias among developers (West et al., 2019). Some prominent cognitive scientists have argued for AI to be explicitly coded as "female" and imbued with stereotypically feminine characteristics in order to prevent it from taking over the world, because overambition is coded as a masculine trait (Pinker, 2015).

In addition, AI systems that are designed to do complex knowledge work are explicitly or implicitly coded as white. These systems may be physically white, and they may also include racial signifiers that are coded as white (Cave & Dihal, 2020). In her book *Race Against Technology*, Ruha Benjamin (2019) discusses some attempts to address diversity issues in Silicon Valley. She suggests that the seemingly homogenous groups developing algorithms are not equipped to reduce bias and their attempts are doing more harm because while they seem like positive change, these attempts are just a facade. Algorithms may give the appearance of change, but broader representation in the workforce and society is not occurring (Benjamin, 2019).

Organizational context and work design

At the level of the organization, several work design factors can impact individual and team processes. Different organizations are flexible--or not--with regards to how much individuals can, or need to, pick up others' job roles (Campion et al., 1993). Different organizations offer more or less adequate training and managerial support, as well as good leadership and resources, granting the team appropriate authority over their mission, resources, and impact (Campion et al., 1993; Hackman, 1998). Some organizations have procedures in place to give feedback to employees and teams about their performance (Wittenbaum et al., 1998), or set up reward systems to promote teamwork (Hertel et al., 2005). Different organizations vary in their contexts and goals, be they classrooms, pain clinics, intelligence shops, or search and rescue organizations (Wittenbaum et al., 1998). Organizational goals can be in conflict or aligned with each other (Wittenbaum et al., 1998). These different goals and contexts can also require or allow different types of technology, including but not limited to different technology to support virtual teams (Hertel et al., 2005) and AIs.

There are many ways in which the work context can impact the deployment and success of AIs, including AIs in teams. One simple way is whether an organization includes AIs in its human teams or not, and whether the AIs are adequate at meeting organizational goals. However, more subtle interactions between an AI in human teams and organizational factors may occur, such as how AIs may inadvertently interact with human performance appraisal systems. Imagine an AI that is built to assist intelligence analysts with sifting through and identifying useful trends or individual pieces of information in copious amounts of public social media. It may be part of the AI's planned design to start off as a novice and then become more adept at identifying useful information as human analysts give it feedback. However, if the humans are given poor performance reviews because they must take extra time to train the AI and are less accurate and efficient, this experience will serve as a disincentive for analysts to work with the AI until the AIs can become useful. Thus, the task of training the AI and the added workload should be explicitly integrated as a goal, and its impacts on performance on other tasks considered, in the human team members' performance evaluations.

Summary of organizational and societal inputs

Organization inputs can strongly impact the design of AIs and the likelihood that an AI tool or teammate will be used by human collaborators. Legal limitations on AI may reduce the likelihood that an AI can be used to make important selection decisions, such as job selection. A lack of understanding of the shared cultural values within an organization can hinder human-AI team integration. Biases held by the designer, which may or may not reflect biases held at the societal level, can be introduced into AI training data and/or algorithms resulting in AIs that make selections that reflect those biases. Finally, the design and context in which the team works may greatly influence whether an AI teammate is even necessary.

Processes

In our model we outline processes that occur at the organizational level and will most likely change with increases in AI-enabled teams. Processes at the organizational level include surveillance and organizational change and technology acceptance.

Surveillance

Human teammates observe and may report back to management about each others' behavior. AI teammates that are monitoring the team's status will collect a large volume of data about each team member. Even if these data are used only for ethical purposes, they may still contain private data about team members that they may not want shared with management. For example, monitoring physiological states of human teammates goes beyond simply assessing if the human is performing adequately. It can also assess the human's physical and mental states, such as heart health. Auditing the AI's performance will also involve auditing these data to determine if judgments about human team members were appropriate, which may involve sharing the data with people who are not in the team. In addition, management may look at logs or models to assess team effectiveness in a fine-grained way. This aligns with wider discussions about the dichotomy between convenience related to AI and surveillance to increase AI efficacy. (Zuboff, 2019) Humans who are engaged in teamwork with AI systems will need to be aware of this surveillance and consent to it.

Organizational change and technology acceptance

The technology being deployed within an organization, be it AI or a travel accounting system, needs to be accepted and used by workers. Organizational change is a difficult process that has been long studied. Lewin (1947/1997) proposed that for lasting change to occur, the group that needs to change must first "unfreeze," such that complacency must be replaced with a desire to change. Even then, after the change has been instituted, there must be a "refreezing" such that the new culture or processes cannot easily be undone.

Related change theory relevant to AI largely focuses on user acceptance and various user acceptance models. Fred Davis' (1987) Technology Acceptance Model (TAM), developed in 1986 upon the framework of Theory of Reasoned Action (TRA) and Theory of Planned Behavior (TPB), serves as a foundational theory for much of the ongoing discussion of acceptance and adoption (Davis, 1987). This model focuses distinctly on two criteria: perceived usefulness of the technology and perceived ease of use of the technology, with the latter pulling from Bandura's self-efficacy theory. Despite the arguably limited scope of the original TAM, as noted in its glaring lack of trust as a factor, further acceptance model development was able to take much of the theory and build upon it (Marangunić & Granić, 2014). Suggestions for future directions in acceptance research include adding still more variables to determine the degree of the relationship between variables, and moderating factors relating to emotional, cognitive, and

demographic variables (Marangunić & Granić, 2014). Further research investigated the relationship between acceptance and experience or beliefs, as well as incorporated more varied ideas like culture, task relevance, motivation, organizational factors, confidence, and more. These changes can generally be categorized along the Four Major Categories of TAM Modification: external factors, factors from other theories, contextual factors, and usage measures (Marangunić & Granić, 2014). (As discussed previously, factors that closely relate to individual and team trust are fundamental to organizational technology change and acceptance. In the context of organizational and societal levels these factors may be related to how much a society, or an organization, accepts a new AI. If the leadership of a company does not accept AI, then that technology is unlikely to be adopted by that company. If a general population and/or political leaders are suspicious of AI's usefulness and ease of use or do not trust it, then that society may be slow to adopt AI or may create laws limiting its use (not necessarily without reason).

In sum, organizations and societies engage in processes humans may need to be made aware of before collaborating with an AI teammate, such as surveillance, and processes that may greatly impact whether AIs are allowed to team with humans, such as organizational design and technology acceptance.

Emergent States

Emergent states, which are dynamic and intrinsically tied to every aspect of the IPEOI model, often arise as individuals are engaging in teaming behaviors. Based on human-human teaming behaviors and an understanding of AI limitations and possible uses we can speculate the relevant individual AI emergent states and other individual human and team emergent states relevant to AI-enabled teaming. Unlike the other levels of our model, emergent states at the organizational or societal level may require observing AI-enabled teams within the larger system. In this section we discuss trust as an emergent state, but other states will arise as AIs are designed to work as teammates and used in that manner as well.

Trust

We have discussed trust extensively at the individual and team level, but the organizational level is also relevant. First, does the organization enact procedures that imply AI inadequacies or undermine the use of AI, and does the culture presume worker trust in the AIs? Second, does the organization enact procedures and have a culture that presumes a trust in its own workers? These questions can apply to the societal level, as well. As individuals become more aware of the failures and limitations of AI (O'Neil, 2016; Noble, 2018), they may be wary of AIs even if their company is enthusiastic about a particular application. For example, while voice-based assistants such as Alexa and Siri have become ubiquitous, use of these assistants is limited to basic tasks due to wide mistrust of how the recorded data is used and stored (Pitardi & Marriott, 2021).

Human trust is related to use of autonomous agents (Parasuraman & Riley, 1997) and is greatly impacted by organizational trust. A lack of trust in AI at the organizational or societal level would greatly impede use of AI by individuals and teams in that organization or society.

Outputs

Outputs at the organizational level of our model reflect the effects of the use of AI in teams. The possible outputs of AI-enabled teaming include (1) increased energy use, (2) changes in productivity, efficiency, and effectiveness, and (3) positive and negative impacts on human society such as increased bias, socioeconomic threats, and increased accessibility. First, AIs are much more resource-intensive than human team members -- some systems require massive amounts of power and cooling resources to maintain performance. The carbon footprint of AI systems is a growing concern, given the state of climate change (Strubell, Ganesh, & McCallum 2019). The concern that AI systems require more power than humans to do the same tasks, of course, does not offset the fact that future AI systems will presumably be able to exceed human performance, at least in the task workspace. Second, the goal of instituting AIs is not simply individual or team productivity, but that of the entire organization. Whether an AI is working with a human team to sift through large amounts of information, or a few AIs are helping a rescue team, if AIs could not contribute positively to the economic or social health, growth, productivity, efficiency, and/or effectiveness of an organization, they should not be used. Third, some scholars have voiced concerns that the rise of AI will lead to the devaluing of humanity (Lanier, 1995; Maitra, 2020): some jobs may be automated, devaluing the human labor involved in those jobs. As AIs become more capable, the range of jobs under threat can grow, but the impact is likely to fall on those most socioeconomically vulnerable. In addition, biased AIs can impact society in systematic ways, such as erroneously predicting higher rates of future crimes committed by black convicts and lower rates of future crimes committed by white convicts (Angwin et al, 2016, as cited in Asaro, 2019). Alternately, AIs in teams can supplement and streamline existing systems such as by helping humans accomplish their jobs better and making activities and information more accessible (for a review see Reeder et al., 2013). Humans with maladies ranging from dementia to visual impairment can benefit from having AIs assist in everyday tasks (such as taking medication on time and searching for items), allowing for independence and increased safety (Kacorri et al., 2017; Coghlan et al., 2021, in press; Dixon et al., 2021, in press). The economic and social health of nations, societies, and the world should be kept in mind as potential outcomes for integrating AIs in teams.

Summary of Organization and Society IPEOIs

The rules that govern an organization or society, be they explicit laws or implicit cultural values, can greatly impact whether AIs can work within a human team, and the likelihood that human teammates will adopt the new technology. As AIs are integrated within human teams, processes such as surveillance and organizational change and technology acceptance will occur, with the

burden of protecting human information laying squarely on those who govern the organization or society. Trust in AI will depend on the use of AI within the teaming system, and whether it meets trust requirements at the individual and team levels as well. Increased human-AI teaming will impact energy use and organizational productivity, efficiency, and effectiveness.

Discussion

This paper makes a multidisciplinary contribution to theory and research on AIs in human teams. Moving beyond the literature on the relatively simple explainability, predictability, and directability of AI (Voas, 2004), we draw from research on team science, cognitive science, management, and other fields to detail a multifaceted model of factors that should be acknowledged for AIs in human teams. Using an IPEOI model, we describe the necessary factors of a socio-technical teaming system at four main levels: individual levels (human and AI), the team level, and the organizational level. Our model also includes contextual factors such as the task and mission and time which can impact any stage, such as inputs or processes, of the model at all levels. We argue that all of the factors are necessary for effective human-AI teaming and can be used to assist in improving AI designed to collaborate with humans.

Although comprehensive, this review is just a beginning in detailing the kinds of factors to keep in mind for a human-AI team. Importantly, future work should address how to conduct testing for human-AI integrated teams. Each of the factors here can be measured either with established or novel measures (e.g., self-report cohesion measures versus physiological measures of different emotions). This measurement by necessity must also include assessments of AI performance as well, including measurements of AI coordination with and team support of humans.

As emphasized throughout this review, trust is a key emergent state: trust of the human to other humans, to AI partners, and even cognitive trust of an AI partner to its human and AI teammates. Increased research into these aspects of trust, specifically regarding how to (appropriately) build it within a human-AI team, will impact and be impacted with a broad array of the other factors delineated here. For example, further research into the effectiveness of team building amongst members of human-AI teams may provide some direction to sustained trust, but it is currently limited (Glikson & Woolley, 2020).

Solutions to Past AI Failures

Designing AIs within a socio-technical teaming system will mitigate the failures outlined in the introduction. First, designers must know what human requirements are necessary for effective collaboration with AIs. This requires an understanding of how humans work within a team both as individuals and a collective. Using the information in our model should supply designers with insights into what humans need from their AI counterparts. Factors such as team communication,

team cohesion, individual human cognitive resources and processes, to name a few, must be considered when designing collaborative AIs. The ability to contribute to team processes and monitor and model the individual human and team will allow AIs designed to work with humans to do so successfully.

Second, sufficient human training on how to use the AI will reduce the likelihood that humans will use the technology for unintended purposes. While it is important for AIs designed to work with humans to have capabilities relevant to the teaming system, it is also necessary for human teammates to understand the role all teammates play within the team, including the AI. If the processes the AI is engaging in are transparent and if it is relatively easy to direct the AI, humans and AIs can not only be interdependent (Johnson & Vera, 2019), humans will have the information necessary to use the AI for its intended purposes. The AI must also have a sufficient representation of the problem state to be capable of informing its human counterparts about the efficacy of the information being supplied.

Third, AIs may not function according to their original design for several reasons including issues at the level of design and/or lack of appropriate updating (either due to a lack of understanding on the human part or design that neglects relevant factors for learning and updating). Humans that are trained to work with AIs and who have appropriate mental models of each AI teammate should understand the taskwork and teamwork-related attributes the AI should be able to engage in. If these expectations are violated, humans can lose trust in the AI and misuse the AI or fail to use the AI entirely. To mitigate this issue, AIs must be observable and therefore engage in the process of explainability. Understanding what operations the AI should be engaging in should allow humans to intervene as soon as the AI is no longer functioning appropriately.

Finally, AIs can be designed with the capability to understand human norms, scripts, and potential reactions which should aid AI teammates in identifying and addressing implicit human expectations. An AI based human narrative would require the AI to be capable of monitoring and modeling humans and the team and determining human preferences and needs. The accuracy of this narrative can be tested against the mental models the AI has built about each human teammate and the team as a collective through a relevant line of questioning. Our model proposes several solutions to some of the failures that occur when AIs are not designed to work within a socio-technical teaming system. While this list of solutions is not exhaustive, it may be used to aid in the design of AIs intended to collaborate with humans.

Future Steps: Metrics

If AI systems are going to be embedded in teams, evaluators will need new metrics for both the system and the team. Such metrics would not replace the existing system requirements but could leverage the outputs and emergent states listed in the model, such as trust or degree of AI use.

Damacharla and colleagues (2018) surveyed the available metrics for human-machine teams and concluded that the primary metrics that will be available in the near future can be divided into human metrics, machine metrics, and team metrics. The human metrics they proposed as most practical to measure as part of benchmarking were judgment, attention allocation, mental computation, and error. Machine metrics included Robot Attention Demand (how much time the humans need to spend paying attention to the robot), machine state, and errors. On the team side, they suggested measuring productive time, cohesion, and interventions. Several of these metrics do not currently have standard measures, such as productive time. The metrics reviewed by Damacharla and colleagues (2018) can be used in future efforts in designing AI to collaborate with human teammates. Our review suggests a much broader range of factors that go beyond Damacharla and colleagues' review (2018) that could, and perhaps should, be measured for assessing the inputs, processes, and outputs for AI-human teams. From the individual human to the organizational and societal levels, several interwoven factors can be assessed that will determine team performance.

Conclusion

Contemporary AIs designed to work with humans are often designed without human interaction in mind. This limitation often results in failures including AI misuse, disuse and abuse (Parasuraman & Riley, 1997). Our review describes factors important for AIs in teams, and in the process also implies ways in which an AI member can excel as a teammate. Just as some humans can be excellent teammates, an AI can be, as well. Such AIs must be designed and capable of adapting to team dynamics, such as through backup behavior (Johnson & Vera, 2019). Our model aims to inform the design of AI teammates from a human-centered perspective that is contextualized within a sociotechnical system. These design features and models are intended to continue the conversation about what is necessary and appropriate for AIs to work successfully with humans.

References

- Amason, A. (1996). Resolving a paradox for top management teams. *Academy of Management Journal*, 39, 123-143.
- Angelucci, A., Bastioni, M., Graziani, P., & Rossi, M. G. (2014). A philosophical look at the uncanny valley. In J. Seibt, R. Hakli, & M. Nørskov (Eds.), *Sociable robots and the future of social relations: Proceedings of robo-philosophy*, 2014 (pp. 165–169).
- Ang, S., Van Dyne, L., Koh, C., Ng, K. Y., Templer, K. J., Tay, C., & Chandrasekar, N. A. (2007). Cultural intelligence: Its measurement and effects on cultural judgment and decision making, cultural adaptation and task performance. *Management and Organization Review*, 3, 335–371. <https://doi:10.1111/j.1740-8784.2007.00082.x>
- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). *Machine bias: There's software used across the country to predict future criminals. and it's biased against blacks*. Propublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Antunes, F., Costa, J. P., & Macas, P. (2010). Ascertaining group decision auditing capabilities. *International Journal of Business and Systems Research*, 4(3), 246-263.
- Argall, B. D., Chernova, S., Veloso, M., & Browning, B. (2009). A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5), 469-483.
- Asaro, P. M. (2019). AI ethics in predictive policing: From models of threat to an ethics of care. *IEEE Technology and Society Magazine*, 38(2), 40-53.
- Ashton, M. C., & Lee, K. (2007). Empirical, theoretical, and practical advantages of the HEXACO model of personality structure. *Personality and Social Psychology Review*, 11, 150–166. doi:10.1177/1088868306294907
- Bainbridge, W. A., Hart, J. W., Kim, E. S., & Scassellati, B. (2010). The Benefits of Interactions with Physically Present Robots over Video-Displayed Agents. *International Journal of Social Robotics*, 3(1), 41–52. <https://doi.org/10.1007/s12369-010-0082-7>
- Bansal, G., Nushi, B., Kamar, E., Lasecki, W. S., Weld, D. S., & Horvitz, E. (2019). Beyond accuracy: The role of mental models in human-AI team performance. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* (Vol. 7, No. 1, pp. 2-11).
- Barki, H., & Hartwick, J. (2004). Conceptualizing the construct of interpersonal conflict. *International Journal of Conflict Management*, 15, 216-244.
- Barrick, M. R., Mount, M. K., & Judge, T. A. (2001). Personality and performance at the beginning of the new millennium: What do we know and where do we go next? *Personality and Performance*, 9, 9-30.
- Bartneck, C., Kanda, T., Ishiguro, H., & Hagita, N. (2007, August). Is the uncanny valley an uncanny cliff?. In *RO-MAN 2007-The 16th IEEE international symposium on robot and human interactive communication* (pp. 368-373). IEEE.
- Basu, S., Jana, N., Bag, A., Mahadevappa, M., Mukherjee, J., Kumar, S., & Guha, R. (2015, December). Emotion recognition based on physiological signals using valence-arousal model. In *2015 Third International Conference on Image Information Processing (ICIIP)* (pp. 50-55). IEEE.
- Beal, D. J., Cohen, R. R., Burke, M. J., & McLendon, C. L. (2003). Cohesion and performance in groups: A meta-analytic clarification of construct relations. *Journal of Applied Psychology*, 88, 989–1004. <https://doi.org/10.1037/0021-9010.88.6.989>

- Bell, S. T. (2007). Deep-level composition variables as predictors of team performance: A meta-analysis. *Journal of Applied Psychology, 92*, 595-615. DOI: 10.1037/0021-9010.92.3.595
- Bell, S. T., & Brown, S. G. (2015). Selecting and composing cohesive teams. In *Team cohesion: Advances in psychological theory, methods and practice* (pp. 181-209). Emerald Group Publishing Limited.
- Benjamin, R. (2019). *Race after technology: Abolitionist tools for the new Jim Code*. Social Forces.
- Beňuš, S. (2014). Social aspects of entrainment in spoken interaction. *Cognitive Computation, 6*, 802-813.
- Bienefeld, N. (2020, July 30-Aug. 1). *Personal communication* [Presentation]. Annual Convention of the Interdisciplinary Network of Group Researchers, remote conference
- Bienefeld, N., Ben-Menahem, S., Grote, G. (2020). *Trust negotiations and triadic interaction patterns in newly formed human-AI teams: An example from critical care*. In preparation.
- Blomqvist, K. and Cook, K.S. (2018). Swift Trust - State-of-the-Art and Future Research Directions, in Searle, R.H. A. Nienaber & S. Sitkin (Eds.) *Routledge Companion to trust*. Routledge, pp. 29-49.
- Bolstad, C. A. (2001, October). Situation awareness: Does it change with age?. In *Proceedings of the human factors and ergonomics society annual meeting* (Vol. 45, No. 4, pp. 272-276). SAGE Publications
- Bombari, D., Schmid Mast, M., & Bachmann, M. (2017). Felt power explains the link between position power and experienced emotions. *Emotion, 17*, 55–66. <https://doi.org/10.1037/emo0000207>
- Braun, M. T., Kozlowski, S. W. J., Brown, T. A. R., & DeShon, R. P. (2020). Exploring the dynamic team cohesion-performance and coordination-performance relationships of newly formed teams. *Small Group Research, 1*-30. [https://DOI: 10.1177/1046496420907157](https://DOI:10.1177/1046496420907157)
- Breazeal, C., Takanishi, A., & Kobayashi, T. (2008). Social robots that interact with people. In B. Siciliano & O. Khatib (Eds.), *Springer handbook of robotics* (pp. 1349–1369). Berlin: Springer.
- Brenton, H., Gillies, M., Ballin, D., & Chatting, D. (2005, September). The uncanny valley: does it exist. In *Proceedings of conference of human computer interaction, workshop on human animated character interaction*. Edinburgh: Napier University.
- Bridwell, D. A., Hecker, E. A., Serences, J. T., & Srinivasan, R. (2013). Individual differences in attention strategies during detection, fine discrimination, and coarse discrimination. *Journal of neurophysiology, 110*(3), 784-794.
- Broniatowski, D. A. (2021). *Psychological foundations of explainability and interpretability in Artificial Intelligence*. NIST Interagency/Internal Report (NISTIR), National Institute of Standards and Technology, Gaithersburg, MD. <https://doi.org/10.6028/NIST.IR.8367> (Accessed May 24, 2021).
- Bull, R., & Johnston, R. S. (1997). Children's arithmetical difficulties: Contributions from processing speed, item identification, and short-term memory. *Journal of experimental child psychology, 65*(1), 1-24.
- Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. *Proceedings of the 1st Conference on Fairness, Accountability, and Transparency, Proceedings of Machine Learning Research, 81*, 77-91.

- Burke, C. S., Stagl, K. C., Salas, E., Peirce, L., & Kendall, D. (2006). Understanding team adaptation: A conceptual analysis and model. *Journal of Applied Psychology, 91*, 1189–1207.
- Campion, M. A., Medsker, G. J., & Higgs, A. C. (1993). Relations between work group characteristics and effectiveness: Implications for designing effective work groups. *Personnel Psychology, 46*, 823-850.
- Campion, M. A., Paper, E. M., Medsker, G. J. (1996). Relations between work team characteristics and effectiveness: A replication and extension. *Personnel Psychology, 49*, 429-452.
- Can, Y. S., Arnrich, B., & Ersoy, C. (2019). Stress detection in daily life scenarios using smart phones and wearable sensors: A survey. *Journal of biomedical informatics, 92*, 103139. <https://doi.org/10.1016/j.jbi.2019.103139>
- Carpenter, J. (2016). *Culture and human-robot interaction in militarized spaces; A war story*. Routledge.
- Cassidy, M. (2014, December). Centaur chess shows power of teaming human and machine. *Huffington Post*. https://www.huffpost.com/entry/centaur-chess-shows-power_b_6383606
- Castaldo, S. (2003). Trust Variety - Conceptual Nature, Dimensions, and Typologies. *Proceedings of the 19th IMP-conference*. 1-30
- Cave, S., & Dihal, K. (2020). The whiteness of AI. *Philosophy & Technology*. <https://doi.org/10.1007/s13347-020-00415-6>
- Chellappa, R., Roy-Chowdhury, A. K., & Zhou, S. K. (2018). Human identification using gait and face. In *Circuits, Signals, and Speech and Image Processing* (pp. 27-1). CRC Press.
- Chérif, L., Wood, V., Marois, A., Labonté, K., & Vachon, F. (2018). Multitasking in the military: Cognitive consequences and potential solutions. *Applied Cognitive Psychology, 32*(4), 429-439.
- Christoffersen, K., & Woods, D. D. (2002). 1. How to make automated systems team players. *Advances in Human Performance and Cognitive Engineering Research, 2*, 1–12. [https://doi.org/10.1016/S1479-3601\(02\)02003-9](https://doi.org/10.1016/S1479-3601(02)02003-9)
- Coghlan S., Waycott J., Lazar A., and Neves B. (2021, in press). Dignity, autonomy, and style of company: Dimensions older adults consider for robot companions. *Proceedings of the ACM: Human Computer Interaction (PACM HCI)*, To appear at CSCW 21.
- Collins, R. T., Gross, R., & Shi, J. (2002, May). Silhouette-based human identification from body shape and gait. In *Proceedings of fifth IEEE international conference on automatic face gesture recognition* (pp. 366-371). IEEE.
- Conger, K., Fausset, R., & Kovalski, S. F. (2019). San Francisco bans facial recognition technology. *The New York Times, 14*.
- Conway, A. R., Kane, M. J., Bunting, M. F., Hambrick, D. Z., Wilhelm, O., & Engle, R. W. (2005). Working memory span tasks: A methodological review and user's guide. *Psychonomic bulletin & review, 12*(5), 769-786.
- Conway, A. R., Kane, M. J., & Engle, R. W. (2003). Working memory capacity and its relation to general intelligence. *Trends in cognitive sciences, 7*(12), 547-552.
- Cronin, M. A., & Weingart, L. R. (2007). Representational gaps, information processing, and conflict in functionally diverse teams. *Academy of Management Review, 32*, 761–773.
- Cronin, M. A., Weingart, L. R., & Todorova, G. (2011). Dynamics in groups: Are we there yet? *Academy of Management Annals, 5*, 571-612.

- Crotty, S. K., & Brett, J. M. (2012). Fusing creativity: Cultural metacognition and teamwork in multicultural teams. *Negotiation and Conflict Management Research*, 5, 210-234.
- Culley, K. E., & Madhavan, P. (2013). A note of caution regarding anthropomorphism in HCI agents. *Computers in Human Behavior*, 29(3), 577-579.
<https://doi.org/10.1016/j.chb.2012.11.023>
- Damacharla, P., Javaid, A. Y., Gallimore, J. J., & Devabhaktuni, V. K. (2018). Common metrics to benchmark human-machine teams (HMT): A review. *IEEE Access* 6, 38637-38655.
<https://doi.org/10.1109/ACCESS.2018.2853560>
- Danescu-Niculescu-Mizil, C., Lee, L., Pang, B., & Kleinberg, J. (2012, April). Echoes of power: Language effects and power differences in social interaction. In *Proceedings of the 21st international conference on World Wide Web* (pp.699-708).
- Dafoe, A., Bachrach, Y., Hadfield, G., Horvitz, E., Larson, K., & Graepel, T. (2021, May 4). Cooperative AI: Machines must learn to find common ground. *Nature*, 583, 33-36.
<https://doi.org/10.1038/d41586-021-01170-0>
- Davis, F. D. (1987). User acceptance of information systems: the technology acceptance model (TAM).
<https://deepblue.lib.umich.edu/bitstream/handle/2027.42/35547/b1409190.0001.001.pdf?sequence=2>
- DeChurch, L. A. & Mesmer-Magnus, J. R. (2010). The cognitive underpinnings of effective teamwork: A meta-analysis. *Journal of Applied Psychology*, 95(1), 32-53.
- De Dreu, C. K. W., & Weingart, L. R. (2003). Task versus relationship conflict, team performance, and team member satisfaction: A meta-analysis. *Journal of Applied Psychology*, 88, 741-749.
- De Dreu, C. K. W., & West, M. A. (2001). Minority dissent and team innovation: The importance of participation in decision making. *Journal of Applied Psychology*, 86, 1191-1201.
- De Rivecourt, M., Kuperus, M. N., Post, W. J., & Mulder, L. J. M. (2008). Cardiovascular and eye activity measures as indices for momentary changes in mental effort during simulated flight. *Ergonomics*, 51(9), 1295-1319.
- Delhey, J. & Newton, K. (2003) Who trusts?: The origins of social trust in seven societies, *European Societies*, 5 (2), 93-137, DOI: [10.1080/1461669032000072256](https://doi.org/10.1080/1461669032000072256)
- Demir, M., McNeese, N. J., & Cooke, N. J. (2016). Team communication behaviors of the human-automation teaming. *2016 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*. IEEE.
<http://www.doi.org/10.1109/COGSIMA.2016.7497782>
- Demir, M., McNeese, N. J., & Cooke, N. J. (2020). Understanding human-robot teams in light of all-human teams: Aspects of team interaction and shared cognition. *International Journal of Human-Computer Studies*, 140, 102436. <https://doi.org/10.1016/j.ijhcs.2020.102436>
- De Wit, F. R. C., Greer, L. L., & Jehn, K. A. (2012). The paradox of intragroup conflict: A meta-analysis. *Journal of Applied Psychology*, 97, 360-390.
- D'Ignazio, C., & Klein, L. F. (2020). *Data feminism*. MIT Press. <https://data-feminism.mitpress.mit.edu>
- Dixon, E., Piper, AM, and Lazar, A. (2021, in press). "Taking care of myself as long as I can": How People with Dementia Configure Self-Management Systems. Proceedings of the Conference on Human Factors in Computing Systems (CHI). To appear at CHI 21.

- Don, A., Brennan, S., Laurel, B., & Shneiderman, B. (1992, June). Anthropomorphism: from ELIZA to Terminator 2. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 67-70).
- Driskell, J. E., Goodwin, G. F., Salas, E., & O'Shea, P. G. (2006). What makes a good team player? Personality and team effectiveness. *Group Dynamics: Theory, Research, and Practice, 10*, 249-271. <https://doi.org/10.1037/1089-2699.10.4.249>
- Driskell, J. E., Salas, E., & Hughes, S. (2010). Collective orientation and team performance: Development of an individual difference measure. *Human Factors, 52*, 316-328. <https://DOI:10.1177/0018720809359522>
- Driskell, J. E., Salas, E., & Driskell, T. (2018). Foundations of teamwork and collaboration. *American Psychologist, 73*, 334-348.
- Driver, J. (2001). A selective review of selective attention research from the past century. *British Journal of Psychology, 92*(1), 53-78.
- Earley, P. C., & Ang, S. 2003. *Cultural intelligence: Individual interactions across cultures*. Stanford University Press.
- Edmondson, A. C., & Harvey, J-F. (2018). Cross-boundary teaming for innovation: Integrating research on teams and knowledge in organizations. *Human Resource Management Review, 28*, 347-360. <http://dx.doi.org/10.1016/j.hrmr.2017.03.002>
- Edwards, B. D., Day, E. A., Arthur, W., & Bell, S. T. (2006). Relationships among team ability composition, team mental models, and team performance. *Journal of Applied Psychology, 91*, 727-736.
- Endsley, M. R. (1988, October). Design and evaluation for situation awareness enhancement. In *Proceedings of the Human Factors Society annual meeting* (Vol. 32, No. 2, pp. 97-101). Sage CA: Los Angeles, CA: Sage Publications.
- Endsley, M. R. (1997, October). Supporting situation awareness in aviation systems. In *1997 IEEE International Conference on Systems, Man, and Cybernetics. Computational Cybernetics and Simulation* (Vol. 5, pp. 4177-4181). IEEE.
- Entin, E. E., & Serfaty, D. (1999). Adaptive team coordination. *Human Factors, 41*, 312-325.
- European Parliament and Council of European Union (2016) Regulation (EU) 2016/679. Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN>
- Fairclough, S. H., Venables, L., & Tattersall, A. (2005). The influence of task demand and learning on the psychophysiological response. *International Journal of Psychophysiology, 56*(2), 171-184.
- Farh, C. I. C. C., Seo, M.-G., & Tesluk, P. E. (2012). Emotional intelligence, teamwork effectiveness, and job performance: The moderating role of job context. *Journal of Applied Psychology, 97*, 890-900.
- Fiore, S. M., Rosen, M. A., Smith-Jentsch, K. A., Salas, E., Letsky, M., & Warner, N. (2010). Toward an understanding of macrocognition in teams: Predicting processes in complex collaborative contexts. *Human Factors, 52*, 203-224. <http://DOI.org/10.1177/0018720810369807>
- Fischer, U., & Mosier, K. (2014). The impact of communication delay and medium on team performance and communication in distributed teams. *Proceedings of the Human Factors and Ergonomic Society 58th Annual Meeting* (pp. 115-119). <https://DOI.org/10.1177/1541931214581025>

- Fisher, D. M., Bell, S. T., Dierdorff, E. C., & Belohlav, J. A. (2012). Facet personality and surface-level diversity as team mental model antecedents: Implications for implicit coordination. *Journal of Applied Psychology, 97*, 825-841.
- Fiske, S. T., & Taylor, S. E. (2017). *Social cognition: From brains to culture* (3rd ed). Sage.
- Flournoy, M. A., Haines, A., & Chefitz, G. (2020). *Building trust through testing: Adapting DOD's test & evaluation, validation & verification (TEVV) enterprise for machine learning systems, including deep learning systems*. WestExec Advisors. Retrieved December 3, 2020 from <https://cset.georgetown.edu/wp-content/uploads/Building-Trust-Through-Testing.pdf>
- Folorunso, C. O., Asaolu, O. S., & Popoola, O. P. (2019). A Review of Voice-Base Person Identification: *State-of-the-Art*. *Covenant Journal of Engineering Technology, 3*(1), 38-60.
- Fournier, L. R., Wilson, G. F., & Swain, C. R. (1999). Electrophysiological, behavioral, and subjective indexes of workload when performing multiple tasks: manipulations of task difficulty and training. *International Journal of Psychophysiology, 31*(2), 129-145.
- Frith, C. D., & Frith, U. (2012). Mechanisms of social cognition. *Annual Review of Psychology, 63*, 287-343.
- Frøkjær, E., Hertzum, M., & Hornbæk, K. (2000, April). Measuring usability: are effectiveness, efficiency, and satisfaction really correlated?. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 345-352).
- Fry, A. F., & Hale, S. (1996). Processing speed, working memory, and fluid intelligence: Evidence for a developmental cascade. *Psychological science, 7*(4), 237-241.
- Fukuyama, F. (2012). Capital, Civil Society and development. *Third World Quarterly, 22*(1), 7-20.
- Galinsky, A. D., Magee, J. C., Inesi, M. E., & Gruenfeld, D. H. (2006). Power and perspectives not taken. *Psychological Science, 17*, 1068-1074.
- Gálvez, R. H., Gravano, A., Beňuš, S., Levitan, R., Trnka, M., & Hirschberg, J. (2020). An empirical study of the effect of acoustic-prosodic entrainment on the perceived trustworthiness of conversational avatars. *Speech Communication, <https://doi.org/10.1016/j.specom.2020.07.007>*
- Garvie, C. (2019). Garbage in, Garbage out: Face recognition on flawed data. Georgetown Law Center on Privacy & Technology. <https://www.flawedfacedata.com/>
- Gazzaniga, M., Ivry, R. B., & Mangun, G. R. (2018). *Cognitive neuroscience: The biology of the mind*. (5th ed.). W. W. Norton & Company
- Gero, K. I., Ashktorab, Z., Dugan, C., Pan, Q., Johnson, J., Geyer, W., Ruiz, M., Miller, S., Millen, D. R., Campbell, M., Kumaravel, S., & Zhang, W. (2020). Mental models of AI agents in a cooperative game setting. *CHI '20*. <http://dx.doi.org/10.1145/3313831.3376316>
- Giger, J., Piçarra, N., Alves-Oliveira, P., Oliveira, R., & Arriaga, P. (2019). Humanization of robots: Is it really such a good idea? *Human Behavior and Emerging Technologies, 1*(2), 111–123. <https://doi.org/10.1002/hbe2.147>
- Giles, H., Coupland, N., & Coupland, J. (1991). Accommodation theory: Communication, context, and consequence. In: H. Giles, J. Coupland, & N. Coupland (Eds.). *Contexts of accommodation: Developments in applied sociolinguistics* (pp. 1-68). Cambridge University Press.

- Giles, H., Taylor, D. M., & Bourhis, R. (1973). Towards a theory of interpersonal accommodation through language: Some Canadian data. *Language in Society*, 2, 177-192.
- Gilmore, C., Attridge, N., Clayton, S., Cragg, L., Johnson, S., Marlow, N., Simms, V., & Inglis, M. (2013). Individual differences in inhibitory control, not non-verbal number acuity, correlate with mathematics achievement. *PLoS One*, 8(6), e67374.
- Glikson, E., & Woolley, A. W. (2020). Human Trust in Artificial Intelligence: Review of Empirical Research. *Academy of Management Annals*, 14(2), 627–660.
<https://doi.org/10.5465/annals.2018.0057>
- Goleman, D. (1998). *Working with emotional intelligence*. Bantam Books.
- Gonzales, A. L., Hancock, J. T., & Pennebaker, J. W. (2010). Language style matching as a predictor of social dynamics in small groups. *Communication Research*, 37, 3-19.
- Gordon, G. G. (1991). Industry determinants of organizational culture. *Academy of Management Review*, 16, 396-415.
- Goth-Owens, T. L., Martinez-Torteya, C., Martel, M. M., & Nigg, J. T. (2010). Processing speed weakness in children and adolescents with non-hyperactive but inattentive ADHD (ADD). *Child Neuropsychology*, 16(6), 577-591.
- Gottman, J. M., & Notarius, C. I. (2000). Decade review: Observing marital interaction. *Journal of Marriage and the Family*, 62, 927–947.
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, 315(5812), 619-619. <http://dx.doi.org/10.1126/science.1134475>
- Gray, K., & Wegner, D. M. (2012). Feeling robots and human zombies: Mind perception and the uncanny valley. *Cognition*, 125, 125–130.
<http://dx.doi.org/10.1016/j.cognition.2012.06.007>
- Grier, R. A. (2012). Military cognitive readiness at the operational and strategic levels: A theoretical model for measurement development. *Journal of Cognitive Engineering and Decision Making*, 6(4), 358-392.
- Groom, V., & Nass, C. (2007). Can robots be teammates? Benchmarks in human-robot teams. *Interaction Studies* 8(3), 483-500. <https://doi.org/10.1075/is.8.3.10gro>
- Gully, S. M., Devine, D. J., & Whitney, D. J. (2012). A meta-analysis of cohesion and performance: Effects of level of analysis and task interdependence. *Small Group Research*, 43, 702-725.
- Guo, W. (2020). Explainable artificial intelligence for 6G: Improving trust between human and machine. *IEEE Communications Magazine*, 58(6), 39-45.
doi:10.1109/mcom.001.2000050
- Guzzo, R. A., & Dickson, M. W. (1996). Teams in organizations: Recent research on performance and effectiveness. *Annual Review of Psychology*, 47, 307–338.
- Guzzo, R. A., Yost, P. R., Campbell, R. J., & Shea, G. P. (1993). Potency in groups: Articulating a construct. *British Journal of Social Psychology*, 32, 87-106.
- Hackman, J. R. (1998). Why teams don't work. In: R. S. Tindale et al., (Eds.), *Theory and research in small groups* (pp. 245-267). Plenum Press.
- Hackman, J. R. (2012). From causes to conditions in group research. *Journal of Organizational Behavior*, 3, 428–444.
- Hackman, J. R., & Oldham, G. R. (1975). Development of the job diagnostic survey. *Journal of Applied Psychology*, 60, 159-170.

- Hale, S., Fry, A. F., & Jessie, K. A. (1993). Effects of practice on speed of information processing in children and adults: Age sensitivity and age invariance. *Developmental psychology*, 29(5), 880.
- Hare, B., Brown, M., Williamson, C., & Tomasello, M. (2002). The domestication of social cognition in dogs. *Science*, 298, 1634-1636. DOI: 10.1126/science.1072702
- Harrison, D. A., Price, K. H., & Bell, M. P. (1998). Beyond relational demography: Time and the effects of surface- and deep-level diversity on work group cohesion. *Academy of Management Journal*, 41, 96-107.
- Hart, S. G. (2006). NASA-Task Load Index (NASA-TLX): 20 years later. *Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting* (pp. 904-908).
- Hart, S. G., & Staveland, L. E. (1988). Development of the NASA-TLX (Task Load Index): Results of empirical and theoretical research. In: P. A. Hancock & N. Meshkati (Eds.), *Human mental workload* (pp. 139-183). Elsevier Science.
- Haslam, N. (2006). Dehumanization: An integrative review. *Personality and Social Psychology Review*, 10, 252-264. http://dx.doi.org/10.1207/s15327957pspr1003_4
- Hawkins, C. B., & Vandiver, A. J. (2019). Human caregivers perceive racial bias in their pet dogs. *Group Processes & Intergroup Relations*, 22, 901-917. <https://doi.org/10.1177/1368430218824656>
- Hertel, G., Geister, S., & Konradt, U. (2005). Managing virtual teams: A review of current empirical research. *Human Resource Management Review*, 15, 69-95.
- Heuer, K., Müller-Frommeyer, L. C., & Kauffeld, S. (2020). Language matters The double-edged role of linguistic style matching in work groups. *Small Group Research*, 51, 208-228. <https://doi.org/10.1177/1046496419874498>
- Higgins, E. T., & Bargh, J. A. (1987). Social cognition and social perception. *Annual Review of Psychology*, 38, 369-425.
- Hill Collins, P. (2008). *Black feminist thought: Knowledge, consciousness, and the politics of empowerment*. Routledge.
- Hinds, P. J., & Bailey, D. E. (2003). Out of sight, out of sync: Understanding conflict in distributed teams. *Organization Science*, 14, 615-632.
- Hinds, P. J., Carley, K. M., Krackhardt, D., & Wholly, D. (2000). Choosing work group members: Balancing similarity, competence, and familiarity. *Organizational Behavior and Human Decision Processes*, 81, 226-251. <https://doi:10.1006/obhd.1999.2875>
- Hinds, P. J., & Kiesler, S. (Eds.) (2002). *Distributed work*. MIT Press.
- Hinds, P. J., Neeley, T. B., & Cramton, C. D. (2014). Language as a lightning rod: Power contests, emotional regulation, and subgroup dynamics in global teams. *Journal of International Business Studies*, 45, 536-561.
- Ho, C. C., MacDorman, K. F., & Pramono, Z. A. (2008, March). Human emotion and the uncanny valley: A GLM, MDS, and Isomap analysis of robot video ratings. In *2008 3rd ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (pp. 169-176). IEEE.
- Hoemann, K., Khan, Z., Feldman, M. J., Nielson, C., Devlin, M, Dy, J., Barrett, L. F., Wormwood, J. B., & Quigley, K. S. (2020). Context-aware experience sampling reveals the scale of variation in affective experience. *Scientific Reports*, 10, 12459. <https://doi.org/10.1038/s41598-020-69180-y>
- Hofer, F., & Schwaninger, A. (2005). Using threat image projection data for assessing individual screener performance. *WIT Transactions on the Built Environment*, 82, 417-426.

- Hoff, K., & Bashir, M. (2013). A theoretical model for trust in automated systems. In *CHI'13 Extended Abstracts on Human Factors in Computing Systems* (pp. 115-120).
- Hoi, S. C. H., Sahoo, D., Lu, J., & Zhao, P. (2018). Online learning: A comprehensive survey. arXiv preprint arXiv:1802.02871.
- Hollingshead, A. B., (1998). Retrieval processes in transactive memory systems. *Journal of Personality and Social Psychology*, 74, 423-442.
- Holmes, W. (2020, Sept. 17). Why the law should treat algorithms like murderous Greek statues. LegalCheek. Retrieved May 24, 2021 from <https://www.legalcheek.com/lc-journal-posts/why-the-law-should-treat-algorithms-like-murderous-greek-statues/>
- Honts, C., Prewett, M., Rahael, J., & Grossenbacher, M. (2012). The importance of team processes for different team types. *Team Performance Management*, 18, 312-327.
- House, R., Rousseau, D. M., & Thomas-Hunt, M. (1995). The meso paradigm: a framework for integration of micro and macro organizational behavior. *Research in Organizational Behavior*, 17, 71-114.
- Hsieh, I. J., & Chen, Y. Y. (2017). Determinants of aggressive behavior: Interactive effects of emotional regulation and inhibitory control. *PloS one*, 12(4), e0175651.
- Hulsheger, U. R., Anderson, N., & Salgado, J. F. (2009). Team-level predictors of innovation at work: A comprehensive meta-analysis spanning three decades of research. *Journal of Applied Psychology*, 94, 1128-1145.
- Hutchins, E. (1995). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265-288.
- Ilgen, D., R., Hollenbeck, J. R., Johnson, M., & Jundt, D. (2005). Teams in organizations: From input-process-output models to IMO models. *Annual Review of Psychology*, 56, 517-543.
- Ismail, L. I., Verhoeven, T., Dambre, J., & Wyffels, F. (2019). Leveraging robotics research for children with autism: a review. *International Journal of Social Robotics*, 11(3), 389-410.
- Jehn, K. A., & Bendersky, C. (2003). Intragroup conflict in organizations: A contingency perspective on the conflict-outcome relationship. *Research in Organizational Behavior*, 25, 187-242.
- Jehn, K. A., & Mannix, E. A. (2001). The dynamic nature of conflict: A longitudinal study of intragroup conflict and group performance. *Academy of Management Journal*, 44, 238-251.
- John S. McCain National Defense Authorization Act for Fiscal Year 2019, H.R. 5515 § 237. (2019). <https://www.congress.gov/115/bills/hr5515/BILLS-115hr5515enr.pdf>
- John, O. P., Naumann, L. P., & Soto, C. J. (2008). Paradigm shift to the integrative Big Five trait taxonomy: History, measurement, and conceptual issues. In O. P. John, R. W. Robins, & L. A. Pervin (Eds.), *Handbook of personality: Theory and research* (pp. 114-158). Guilford Press.
- John, O. P., & Srivastava, S. (1999). The Big Five trait taxonomy: History, measurement, and theoretical perspectives. In L. A. Pervin & O. P. John (Eds.), *Handbook of personality: Theory and research* (2nd ed., pp. 102-138). Guilford Press.
- Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., van Riemsdijk, M. B., & Sierhuis, M. (2014) Coactive design: Designing support for interdependence in joint activity. *Journal of Human-Robot Interaction* 3(1), 43-69 doi.org/10.5898/JHRI.3.1.Johnson.
- Johnson, J. J., & McIntye, C. L. (1998). Organizational culture and climate correlates of job satisfaction. *Psychological Reports*, 82, 843-850.

- Johnson, M., & Vera, A. H. (2019). No AI is an island: The case for teaming intelligence. *AI Magazine*, 40, 17-28. <https://doi.org/10.1609/aimag.v40i1.2842>
- Johnson-Laird, P. N. (1980). Mental models in cognitive science. *Cognitive Science*, 4, 71–115.
- Johnston, W. A., & Dark, V. J. (1986). Selective attention. *Annual review of psychology*, 37(1), 43-75.
- Jones, R., Stevens, M. J., & Fisher, D. L. (2000). Selection in team contexts. In J. F. Kehoe (Ed.), *Managing selection in changing organizations: Human resource strategies* (pp. 210-241). Jossey-Bass.
- Jung, M. F., Martelaro, N., & Hinds, P. J. (2015). Using robots to moderate team conflict; The case of repairing violations. HRI'15. <http://dx.doi.org/10.1145/2696454.2696460>
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: individual differences in working memory. *Psychological Review*, 99(1), 122.
- Kacorri, H., Kitani, K. M., Bigham, J. P., & Asakawa, C. (2017, May). People with visual impairment training personal object recognizers: Feasibility and challenges. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (pp. 5839-5849).
- Kail, R., & Hall, L. K. (1994). Processing speed, naming speed, and reading. *Developmental Psychology*, 30(6), 949.
- Kane, M. J., Conway, A. R., Miura, T. K., & Colflesh, G. J. (2007). Working memory, attention control, and the N-back task: a question of construct validity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(3), 615.
- Kane, M. J., & Engle, R. W. (2002). The role of prefrontal cortex in working-memory capacity, executive attention, and general fluid intelligence: An individual-differences perspective. *Psychonomic bulletin & review*, 9(4), 637-671.
- Kaufman, J. C. (2011). Individual differences in creativity. In T. Chamorro-Premuzic, S. von Stumm, & A. Furnham (Eds.), *The Wiley-Blackwell handbooks of personality and individual differences. The Wiley-Blackwell handbook of individual differences* (p. 679–697). Wiley Blackwell.
- Kellermanns, F. W., Floyd, S. W., Pearson, A. W., & Spencer, B. (2008). The contingent effect of constructive confrontation on the relationship between shared mental models and decision quality. *Journal of Organizational Behavior*, 29, 119-137.
- Keltner, D., Gruenfeld, D. H., & Anderson, C. (2003). Power, approach, and inhibition. *Psychological Review*, 110, 265-284.
- Khan, R. F., & Sutcliffe, A. (2014). Attractive agents are more persuasive. *International Journal of Human-Computer Interaction*, 30(2), 142-150.
- Klare, B. F., Burge, M. J., Klontz, J. C., Bruegge, R. W. V., & Jain, A. K. (2012). Face recognition performance: Role of demographic information. *IEEE Transactions on Information Forensics and Security*, 7(6), 1789-1801.
- 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. <https://doi.org/10.1109/MIS.2004.74>
- Klein, K. J., Dansereau, F., & Hall, R. J. (1994). Levels issues in theory development, data collection, and analysis. *Academy of Management Review*, 19, 195-229.
- Klimoski, R., & Jones, R. G. (1995). Staffing for effective group decision making: Key issues in matching people and teams. In: R. A. Guzzo, E. Salas (Eds.), *Team effectiveness and decision making in organizations* (pp. 291-332). Jossey-Bass.

- Kobs Nawotniak, S. E., Miller, M. J., Stevens, A. H., Marquez, J. J., Payler, S. J., Brady, A. L., Hughes, S. S., Haberle, C. W., Sehlke, A., Beaton, K. H., Chappell, S. P., Elphic, R. C., & Lim, D. S. S. (2019). Opportunities and challenges of promoting scientific dialog through execution of future science-driven extravehicular activity. *Astrobiology*, 19, 426-439. <https://DOI.org/10.1089/ast.2018.1901>
- Kory-Westlund, J. M., & Breazeal, C. (2019). Exploring the effects of social robot's speech entrainment and backstory on young children's emotion, rapport, relationship, and learning. *Frontiers in Robotics and AI*, 6, Article 54.
- Kostavelis, I., Vasileiadis, M., Skartados, E., Kargakos, A., Giakoumis, D., Bouganis, C. S., & Tzovaras, D. (2019). Understanding of human behavior with a robotic agent through daily activity analysis. *International Journal of Social Robotics*, 11(3), 437-462.
- Kovacs, B., & Kleinbaum, A. M. (2019). Language-style similarity and social networks. *Psychological Science*. <https://doi.org/10.1177/0956797619894557>
- Kozlowski, S. W. J. (2015). Advancing research on team process dynamics: Theoretical, methodological, and measurement considerations. *Organizational Psychology Review*, 5, 270-299.
- Kozlowski, S. W. J., Chao, G. T., & Grand, J. A. (2016). Capturing the multilevel dynamics of emergence: Computational modeling, simulation, and virtual experimentation. *Organizational Psychology Review*, 6, 3-33. <https://doi.org/10.1177/2041386614547955>
- Kozlowski, S. W. J., & Ilgen, D. R. (2006). Enhancing the effectiveness of work groups and teams. *Psychological Science in the Public Interest*, 7, 77-124.
- Kozlowski, S. W. J., & Klein, K. J. (2000). A multi-level approach to theory and research in organizations: Contextual, temporal, and emergent processes. In K. J. Klein & S. W. J. Kozlowski (Eds.), *Multilevel theory, research, and methods in organizations* (pp. 3-90). Jossey-Bass.
- Krämer, N. C., Lucas, G., Schmitt, L., & Gratch, J. (2018). Social snacking with a virtual agent—On the interrelation of need to belong and effects of social responsiveness when interacting with artificial entities. *International Journal of Human-Computer Studies*, 109, 112-121.
- Kwan, V. S. Y., & Fiske, S. T. (2008). Missing links in social cognition: The continuum from nonhuman agents to dehumanized humans. *Social Cognition*, 26, 125-128.
- Kyriakou, K., Resch, B., Sagl, G., Petutschnig, A., Werner, C., Niederseer, D., ... & Pykett, J. (2019). Detecting moments of stress from measurements of wearable physiological sensors. *Sensors*, 19(17), 3805. <https://doi.org/10.3390/s19173805>
- Laajaj, R., Macours, K., Hernandez, D. A. P., Arias, O., Gosling, S. D., Potter, J., Rubio-Codina, M., & Vakis, R. (2019). Challenges to capture the Big Five personality traits in non-WEIRD populations. *Science Advances*, 5, 1-13. <https://doi.org/10.1126/sciadv.aaw5226>
- Lake, B. M., Ullman, T. D., Tenebaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, e253. <https://doi.org/10.1017/S0140525X16001837>
- Lam, L. T., & Kirby, S. L. (2002). Is emotional intelligence an advantage? An exploration of the impact of emotional and general intelligence on individual performance. *The Journal of Social Psychology*, 142(1), 133-143.
- Lanier, J. (1995). Agents of alienation. *Journal of Consciousness Studies*, 2(1), 76-81.
- Lee, K. M., Jung, Y., Kim, J., & Kim, S. R. (2006). Are physically embodied social agents better than disembodied social agents?: The effects of physical embodiment, tactile interaction,

- and people's loneliness in human–robot interaction. *International journal of human-computer studies*, 64(10), 962-973.
- Leichtenstern, K., Bee, N., André, E., Berkmüller, U., & Wagner, J. (2011, June). Physiological measurement of trust-related behavior in trust-neutral and trust-critical situations. In *IFIP International Conference on Trust Management* (pp. 165-172). Springer, Berlin, Heidelberg.
- Lengenfelder, J., Bryant, D., Diamond, B. J., Kalmar, J. H., Moore, N. B., & DeLuca, J. (2006). Processing speed interacts with working memory efficiency in multiple sclerosis. *Archives of Clinical Neuropsychology*, 21(3), 229-238.
- Levitan, R. (2020). Developing an integrated model of speech entrainment. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20)*, 5159-5163.
- Levitan, R., Gravano, A., Willson, L., Beňuš, Š., Hirschberg, J., & Nenkova, A. (2012, June). Acoustic-prosodic entrainment and social behavior. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human language technologies* (pp. 11-19).
- Lewicki, R. J., & Brinsfield, C. (2017). Trust Repair. *Annual Review of Organizational Psychology and Organizational Behavior*, 4, 287–313. <https://doi.org/10.1146/annurev-orgpsych-032516-113147>
- Lewin, K. (1947/1997). *Field theory in social science*. American Psychological Association reprinting of 1947 version.
- Lewis, K., & Herndon, B. (2011). Transactive memory systems: Current issues and future research directions. *Organization Science*, 22, 1254-1265.
- Li, J. (2015). The benefit of being physically present: A survey of experimental works comparing copresent robots, telepresent robots and virtual agents. *International Journal of Human-Computer Studies*, 77, 23-37.
- Lim, B.-C., & Klein, K. J. (2006). Team mental models and team performance: A field study of the effects of team mental model similarity and accuracy. *Journal of Organizational Behavior*, 27, 403-418.
- Lin, Y. P., Wang, C. H., Jung, T. P., Wu, T. L., Jeng, S. K., Duann, J. R., & Chen, J. H. (2010). EEG-based emotion recognition in music listening. *IEEE Transactions on Biomedical Engineering*, 57(7), 1798-1806.
- Lubold, N., Walker, E., Pon-Barry, H., & Ogan, A. (2018). Automated pitch convergence improves learning in a social, teachable robot for middle school mathematics. *International Conference on Artificial Intelligence in Education* (pp. 282-296).
- Luque-Casado, A., Perales, J. C., Cárdenas, D., & Sanabria, D. (2016). Heart rate variability and cognitive processing: The autonomic response to task demands. *Biological psychology*, 113, 83-90.
- Macrae, C. N., & Quadflieg, S. (2010). Perceiving people. In: S. T. Fiske & D. T. Gilbert (Eds.), *Handbook of social psychology* (Vol 1: pp. 428-463). Wiley.
- MacDorman, K. F., & Entezari, S. O. (2015). Individual differences predict sensitivity to the uncanny valley. *Interaction Studies*, 16(2), 141-172.
- MacDorman, K. F., Vasudevan, S. K., & Ho, C. C. (2009). Does Japan really have robot mania? Comparing attitudes by implicit and explicit measures. *AI & society*, 23(4), 485-510. DOI: 10.1007/s00146-008-0181-2

- MacDorman, K. F., & Ishiguro, H. (2006). Opening Pandora's uncanny Box: Reply to commentaries on "The uncanny advantage of using androids in social and cognitive science research". *Interaction Studies*, 7(3), 361-368. doi:10.1075/is.7.3.10mac
- Maitra, S. (2020, February). Artificial Intelligence and Indigenous Perspectives: Protecting and Empowering Intelligent Human Beings. In Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (pp. 320-326).
- Marangunić, N., & Granić, A. (2014). Technology acceptance model: a literature review from 1986 to 2013. *Universal Access in the Information Society*, 14(1), 81–95. <https://doi.org/10.1007/s10209-014-0348-1>
- Marks, M. A., Mathieu, J. E., & Zaccaro, S. J. (2001). A temporally based framework and taxonomy of team processes. *Academy of Management Review*, 26, 356-376.
- Marlow, S. L., Lacerenza, C. N., Paoletti, J., Burke, C. S., & Salas, E. (2018). Does team communication represent a one-size-fits-all approach?: A meta-analysis of team communication and performance. *Organizational Behavior and Human Decision Processes*, 144, 145-170.
- Martins G. S., Santos L., & Dias, J. (2019). User-adaptive interaction in social robots: a survey focusing on non-physical interaction. *International Journal of Social Robotics*, 11, 185-205
- Mathieu, J. E., Heffner, T. S., Goodwin, G. F., Salas, E., & Cannon-Bowers, J. A. (2000). The influence of shared mental models on team process and performance. *Journal of Applied Psychology*, 85, 273–283.
- Mathieu, J. E., Tannenbaum, S. I., Donsbach, J. S., & Alliger, G. M. (2014). A review and integration of team composition models: Moving toward a dynamic and temporal framework. *Journal of Management*, 40, 130-160.
- Mayer, J. D., & Salovey, P. (1997). What is emotional intelligence? In: P. Salovey & D. J. Sluyter (Eds.), *Emotional development intelligence* (pp. 3-31). Basic Books.
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20, 709-734.
- Mayer, R. E., Sobko, K., & Mautone, P. D. (2003). Social cues in multimedia learning: Role of speaker's voice. *Journal of Educational Psychology*, 95(2), 419–425. <https://doi.org/10.1037/0022-0663.95.2.419>
- Mayo, A. T., & Woolley, A. W. (2019). Variance in group ability to transform resources into performance, and the role of coordinated attention. *Academy of Management Discoveries*, <https://doi.org/10.5465/amd.2019.0231>
- McAllister, D. J. (1995). Affect- and cognition-based trust as foundations for interpersonal cooperation in organizations. *Academy of Management Journal*, 38, 24-59.
- McCarthy, J. (2007, November 12). *What is artificial intelligence: Basic questions*. Stanford. <http://www-formal.stanford.edu/jmc/whatisai/node1.html>
- McGrath, J. E. (1984). *Groups: Interaction and performance*. Prentice Hall.
- McKendrick, R., Shaw, T., de Visser, E., Saqer, H., Kidwell, B., & Parasuraman, R. (2014). Team performance in networked supervisory control of unmanned air vehicles: Effects of automation, working memory, and communication content. *Human factors*, 56(3), 463-475.
- Mckinney Jr, E. H., Barker, J. R., Davis, K. J., & Smith, D. (2005). How swift starting action teams get off the ground: What United flight 232 and airline flight crews can tell us about team communication. *Management Communication Quarterly*, 19(2), 198-237.

- Mcknight, D. H., & Chervany, N. L. (1996). The meanings of trust. *Technical Report MISRC Working Paper Series 96-04, University of Minnesota, Management Information Systems Research Center*, 1-86
- Mesmer-Magnus, J. R., & DeChurch, L. A. (2009). Information sharing and team performance: A meta-analysis. *Journal of Applied Psychology*, *94*, 535-546.
- Meyerson, D., Weick, K. E., & Kramer, R. M. (1996). Swift trust and temporary groups. *Trust in organizations: Frontiers of theory and research*, *166*, 195.
- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, *267*, 1-38.
- Mirrig, N., Stollnberger, G., Miksch, M., Stadler, S., Giuliani, M., & Tscheligi, M. (2017). To err is robot: How humans assess and act toward an erroneous social robot. *Frontiers Robotics AI*, *4*, 21. <https://doi.org/10.3389/frobt.2017.00021>
- Mitchell, W. J., Szerszen, K. A., Sr., Lu, A. S., Schermerhorn, P. W., Scheutz, M., & Macdorman, K. F. (2011). A mismatch in the human realism of face and voice produces an uncanny valley. *i-Perception*, *2*, 10–12. <http://dx.doi.org/10.1068/i0415>
- Mittelstadt, B., Russell, C., & Wachter, S. (2019, January). Explaining explanations in AI. In *Proceedings of the conference on fairness, accountability, and transparency* (pp. 279-288).
- Mohammed, S., Ferzandi, L., & Hamilton, K. (2010). Metaphor no more: A 15-year review of the team mental model construct. *Journal of Management*, *36*, 876-910.
- Morgan, B. B., Salas, E., & Glickman, A. S. (1994). An analysis of team evolution and maturation. *Journal of General Psychology*, *120*, 277–291.
- Mori, M. (1970). Bukimi no tani [the uncanny valley]. *Energy*, *7*, 33-35.
- Mullen, B., & Copper, C. (1994). The relation between group cohesiveness and performance: An integration. *Psychological Bulletin*, *115*, 210-227.
- Human Performance Research Group. (1986). NASA Task Load Index (TLX), V.1: Paper and pencil package. NASA Ames Research Center. 1-19.
- Nath, R. K., Thapliyal, H., Caban-Holt, A., & Mohanty, S. P. (2020). Machine learning based solutions for real-time stress monitoring. *IEEE Consumer Electronics Magazine*, *9*(5), 34-41.
- Nemeth, C. J. (1986). Differential contributions of majority and minority influence. *Psychological Review*, *93*, 23-32.
- Niederhoffer, K. G., & Pennebaker, J. W. (2002). Linguistic style matching in social interaction. *Journal of Language and Social Psychology*, *21*, 337-360.
- Niler, A. A., Mesmer-Magnus, J. R., Larson, L. E., Plummer, G., DeChurch, L. A., & Contractor, N. S. (2020). Conditioning team cognition: A meta-analysis. *Organizational Psychology Review*, 1-31. <https://DOI.org/10.1177/2041386620972112>
- Noble, S. U. (2018). *Algorithms of oppression: How search engines reinforce racism*. NYU Press.
- Nunez, M. D., Srinivasan, R., & Vandekerckhove, J. (2015). Individual differences in attention influence perceptual decision making. *Frontiers in psychology*, *6*, 18.
- Oldham, G. R., & Hackman, J. R. (2010). Not what it was and not what it will be: The future of job design research. *Journal of Organizational Behavior*, *31*, 463-479.
- Ott, D. L., Michailova, S. (2018). Cultural intelligence: A review and new research avenues. *International Journal of Management Reviews*, *20*, 99-119.

- O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
- Palazzolo, E. T. (2017). Transactive memory. In: C. R. Scott, L. Lewis, J. R. Barker, J. Keyton, T. Kuhn, & P. K. Turner (Eds.), *International encyclopedia of organizational communication*. John Wiley & Sons.
<https://DOI.org/10.1002/9781118955567.wbieoc208>
- Paletz, S. B. F., Golonka, E. M., Stanton, G., Ryan, D., Adams, N., Pandža, N. B., & Bradley, P. (2021). *Annotation guide for emotion in social media: Development and initial validity*. Under review.
- Paletz, S. B. F., Kim, K. H., Schunn, C. D., Tollinger, I., & Vera, A. (2013). Reuse and recycle: The development of adaptive expertise, routine expertise, and novelty in a large research team. *Applied Cognitive Psychology*, 27, 415-428.
- Paletz, S. B. F., Pavisic, I., Miron-Spektor, E., & Lin, C.-C. (2018). Diversity in creative teams: Reaching across cultures and disciplines. In L. Y.-Y. Kwan, S. Liou, & A. K.-Y. Leung (Eds.), *Handbook of culture and creativity: Basic processes and applied innovations* (pp. 169-206). Oxford University Press.
- Paletz, S. B. F., & Schunn, C. D. (2010). A social-cognitive framework of multidisciplinary team innovation. *Topics in Cognitive Science*, 2, 73-95.
- Paletz, S. B. F., Schunn, C. D., & Kim, K. H. (2011). Conflict under the microscope: Microconflicts in naturalistic team discussions. *Negotiation and Conflict Management Research*, 4, 314-351.
- Panicker, S. S., & Gayathri, P. (2019). A survey of machine learning techniques in physiology based mental stress detection systems. *Biocybernetics and Biomedical Engineering*, 39(2), 444-469.
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human factors*, 39(2), 230-253.
- Pennisi, P., Tonacci, A., Tartarisco, G., Billeci, L., Ruta, L., Gangemi, S., & Pioggia, G. (2016). Autism and social robotics: A systematic review. *Autism Research*, 9(2), 165-183.
- Pinker, S. (2015). Thinking does not imply subjugating. In Brockman, J. (Ed.), *What to Think About Machines That Think: Today's Leading Thinkers on the Age of Machine Intelligence* (pp. 6-8), Harper Perennial.
- Pitardi, V., & Marriott, H. R. (2021). Alexa, she's not human but... Unveiling the drivers of consumers' trust in voice-based artificial intelligence. *Psychology & Marketing*, 38(4), 626-642.
- Porter, D., McAnally, M., Bieber, C., Wojton, H., & Medlin, R. (2020). *Trustworthy autonomy: A roadmap to assurance. Part 1: System effectiveness*. Institute for Defense Analysis. IDA Document P-10768-NS. Retrieved November 5, 2020, from <https://www.ida.org/-/media/feature/publications/t/tr/trustworthy-autonomy-a-roadmap-to-assurance/p-10768.ashx>
- Ralph, B. C., & Smilek, D. (2017). Individual differences in media multitasking and performance on the n-back. *Attention, Perception, & Psychophysics*, 79(2), 582-592.
- Ramey, C. H. (2006). An inventory of reported characteristics for home computers, robots, and human beings: Applications for android science and the uncanny valley. In *Proceedings of the ICCS/CogSci-2006 Long Symposium 'Toward Social Mechanisms of Android Science'*, Vancouver, Canada.

- Reeder, B., Meyer, E., Lazar, A., Chaudhuri, S., Thompson, H., & Demiris, G. (2013). Framing the evidence for health smart homes and home-based consumer health technologies as a public health intervention for independent aging: a systematic review. *International Journal of Medical Informatics*, 82(7), 565-579.
- Reeves, B., Hancock, J., & Liu, X. (2020). Social robots are like real people: first impressions, attributes, and stereotyping of social robots. *Technology, Mind, and Behavior*, 1 (1), 1-14. <https://doi.org/10.1037/tmb0000018>
- Reeves, B., & Nass, C. I. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. Center for the Study of Language and Information; Cambridge University Press.
- Ren, Y., & Argote, L. (2011). Transactive memory systems: 1985-2010: An integrative framework of key dimensions, antecedents, and consequences. *Academy of Management Annals*, 51, 189-229.
- Robert, L. P. (2018). Motivational theory of human robot teamwork. *International Robotics & Automation Journal (IRAJ)*, 4(4), 248-251.
- Rousseau, D. M. (1985). Issues of level in organizational research: multi-level and cross-level perspectives. *Research in Organizational Behavior*, 7, 1-37.
- Rozin, P., & Fallon, A. E. (1987). A perspective on disgust. *Psychological Review*, 94, 23-41. <http://dx.doi.org/10.1037/0033-295X.94.1.23>
- Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking.
- Rypma, B., & D'Esposito, M. (1999). The roles of prefrontal brain regions in components of working memory: effects of memory load and individual differences. *Proceedings of the National Academy of Sciences*, 96(11), 6558-6563.
- Salas, E., Prince, C., Baker, D. P., & Shrestha, L. (1995). Situation Awareness in Team Performance: Implications for Measurement and Training. *Human Factors*, 37(1), 123-136. <https://doi.org/10.1518/001872095779049525>
- Salas, E., Sims, D. E., & Burke, C. S. (2005). Is there a "Big Five" in Teamwork? *Small Group Research*, 36, 555-599.
- Salthouse, T. A. (2000). Aging and measures of processing speed. *Biological psychology*, 54(1-3), 35-54.
- Saygin, A. P., Chaminade, T., Ishiguro, H., Driver, J., & Frith, C. (2012). The thing that should not be: Predictive coding and the uncanny valley in perceiving human and humanoid robot actions. *Social Cognitive and Affective Neuroscience*, 7, 413-422. <http://dx.doi.org/10.1093/scan/nsr025>
- Schachar, R., Tannock, R., Marriott, M., & Logan, G. (1995). Deficient inhibitory control in attention deficit hyperactivity disorder. *Journal of Abnormal Child Psychology*, 23(4), 411-437.
- Schein, E. H. (1990). Organizational culture. *American Psychologist*, 45, 109-119.
- Schein, E. H., & Schein, P. A. (2019). *The corporate culture survival guide*. John Wiley & Sons.
- Schulze, J., & Krumm, S. (2017). The "virtual team player": A review and initial model of knowledge, skills, abilities, and other characteristics for virtual collaboration. *Organizational Psychology Review*, 7, 66-95.
- Schwartz, D. L., Bransford, J. D., & Sears, D. (2005). Efficiency and innovation in transfer. In J. Mestre (Ed.), *Transfer of Learning from a Modern Multidisciplinary Perspective* (1-51). CT: Information Age Publishing.

- Schwemmer, C., Knight, C., Bello-Pardo, E. D., Oklobdzija, S., Schoonvelde, M., & Lockhart, J. W. (2020). Diagnosing gender bias in image recognition systems. *Socius*, 6, 237802312096.
- Sebo, S., Stoll, B., Scassellati, B., & Jung, M. F. (2020). Robots in groups and teams: A literature review. *Proceedings of the ACM on Human-Computer Interaction 4 (CSCW2)*, 1-36. Article 176. DOI:10.1145/3415247
- Shane, J. (2019). *You look like a thing and I love you: How artificial intelligence works and why it's making the world a weirder place*. Voracious.
- Sheridan, T. B. (2020). A review of recent research in social robotics. *Current Opinion in Psychology*, 36, 7–12. <https://doi.org/10.1016/j.copsyc.2020.01.003>
- Shneiderman, B. (2020a). Human-centered artificial intelligence: Reliable, safe and trustworthy. *International Journal of Human-Computer Interaction*, 36(6), 495-504. <https://doi.org/10.1080/10447318.2020.1741118>
- Shneiderman, B. (2020b). Human-centered artificial intelligence: Three fresh ideas. *AIS Transactions on Human-Computer Interaction*, 12, 1-16. <http://DOI:10.17705/1thci.00101>
- Shu, L., Xie, J., Yang, M., Li, Z., Li, Z., Liao, D., ... & Yang, X. (2018). A review of emotion recognition using physiological signals. *Sensors*, 18(7), 2074.
- Slabodkin, G. (1998). Software glitches leave Navy Smart Ship dead in the water. *Government Computer News*, 17(17)
- Snyder, M., Tanke, E. D., & Berscheid, E. (1977). Social perception and interpersonal behavior: On the self-fulfilling nature of social stereotypes. *Journal of Personality and Social Psychology*, 35, 656-666.
- Society for Industrial Organizational Psychology. (2018). *Principles for the validation and use of personnel selection procedures*, 5th edition. American Psychological Association. <https://www.apa.org/ed/accreditation/about/policies/personnel-selection-procedures.pdf>
- Soto, C. J., Napolitano, C. M., & Roberts, B. W. (2020). Taking skills seriously: Toward an integrative model and agenda for social, emotional, and behavioral skills. *Current Directions in Psychological Science*, 1-8. <https://doi.org/10.1177/0963721420978613www.psychologicalscience.org/CDPS>
- Stasser, G., Stewart, D. D., & Wittenbaum, G. M. (1995). Expert roles and information exchange during discussion: The importance of knowing who knows what. *Journal of Experimental Social Psychology*, 31, 244–265.
- Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during group discussion. *Journal of Personality and Social Psychology*, 48, 1467–1478.
- Steckenfinger, S. A., & Ghazanfar, A. A. (2009). Monkey visual behavior falls into the uncanny valley. *Proceedings of the National Academy of Sciences*, 106(43), 18362-18366.
- Steele, C. M. (1997). A threat in the air: How stereotypes shape intellectual identity and performance. *American Psychologist*, 52, 613-629.
- Strohkorb, S., Fukuto, E., Warren, N., Taylor, C., Berry, B., & Scassellati, B. (2016). Improving human-human collaboration between children with a social robot. *25th IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN 2016*, 551–556. <https://doi.org/10.1109/ROMAN.2016.7745172>
- Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *arXiv preprint arXiv:1906.02243*.

- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive science*, 12(2), 257-285.
- Sweller, J. (2010). *Cognitive load theory: Recent theoretical advances*. In J. L. Plass, R. Moreno, & R. Brünken (Eds.), *Cognitive load theory* (p. 29–47). Cambridge University Press. <https://doi.org/10.1017/CBO9780511844744.004>
- Sweller, J., Van Merriënboer, J. J., & Paas, F. G. (1998). Cognitive architecture and instructional design. *Educational Psychology Review*, 10(3), 251-296.
- Tatman, R. (2017, April). Gender and dialect bias in YouTube's automatic captions. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing* (pp. 53-59).
- Tausch, A., Kluge, A., & Adolph, L. (2020). Psychological effects of the allocation process in human-robot interaction: A model for research on ad hoc task allocation. *Frontiers in Psychology*, <https://doi.org/10.3389/fpsyg.2020.564672>
- Tausczik, Y. R., & Pennebaker, J. W. (2010). The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, 29, 24-54. DOI:10.1177/0261927X09351676
- Tausczik, Y. R., & Pennebaker, J. W. (2013). Improving teamwork using real-time language feedback. *CHI 2013: Changing Perspectives*, Paris, France.
- Thomas, D. C., Elron, R., Stahl, G., Ekelund, B. Z., Ravlin, E.c., Cerdin, J. L., Poelmans, S., Brislin, R., Pekerti, A., Aycan, Z., Maznevski, M., Au, K., & Lazarova, M. B. (2008). Cultural intelligence: Domain and assessment. *International Journal of Cross Cultural Management*, 8, 123-143.
- Thompson, L., & Cohen, T. R. (2012). Metacognition in teams and organizations. In: P. Briñol & K. DeMarree (Eds.), *Social metacognition* (pp. 283-302). Taylor & Francis.
- Tinwell, A., Nabi, D. A., & Charlton, J. P. (2013). Perception of psychopathy and the Uncanny Valley in virtual characters. *Computers in Human Behavior*, 29, 1617–1625. <http://dx.doi.org/10.1016/j.chb.2013.01.008>
- Tjosvold, D., Wong, A. S. H., & Chen, N. Y. F. (2014). Constructively managing conflicts in organizations. *Annual Review of Organizational Psychology and Organizational Behavior*, 1, 545-568.
- Todorova, G., Bear, J. B., & Weingart, L. R. (2014). Can conflict be energizing? A study of task conflict, positive emotions, and job satisfaction. *Journal of Applied Psychology*, 99(3), 451–467. <https://doi.org/10.1037/a0035134>
- Tuckman, B. W. (1965). Developmental sequence in small groups. *Psychological Bulletin*, 63, 384–399.
- Tuckman, B. W., & Jensen, M. A. C. (1977). Stages of small-group development revisited. *Group Organization Management*, 2, 419–427.
- Turnwald, A., & Wollherr, D. (2019). Human-like motion planning based on game theoretic decision making. *International Journal of Social Robotics*, 11(1), 151-170.
- Unsworth, N., Fukuda, K., Awh, E., & Vogel, E. K. (2014). Working memory and fluid intelligence: Capacity, attention control, and secondary memory retrieval. *Cognitive psychology*, 71, 1-26.
- Unsworth, N., Schrock, J. C., & Engle, R. W. (2004). Working memory capacity and the antisaccade task: individual differences in voluntary saccade control. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(6), 1302.

- US Department of Defense. (2019). *Summary of the 2018 department of defense artificial intelligence strategy: Harnessing AI to advance our security and prosperity*.
- US General Accounting Office (GAO). (1992). *Patriot missile defense: Software problem led to system failure at Dhahran, Saudi Arabia*. General Accounting Office.
- Uslaner, E. M. (2013). Trust as an alternative to risk. *Public Choice*, 157(3–4), 629–639. <http://doi.org/10.1007/s11127-013-0082-x>
- Vail, K. E., Rothschild, Z. K., Weise, D. R., Solomon, S., Pyszczynski, T., & Greenberg, J. (2010). A terror management analysis of the psychological functions of religion. *Personality and Social Psychology Review*, 14(1), 84-94.
- Van Kleef, G. A., Oveis, C., van der Lowe, I., LuoKogan, A., Goetz, J., & Keltner, D. (2008). Power, distress, and compassion: Turning a blind eye to the suffering of others. *Psychological Science*, 19, 1315-1322.
- Van Swol, L. M., & Kane, A. A. (2019). Language and group processes: An integrative, interdisciplinary review. *Small Group Research*, 50, 3-38.
- Vaughan, D. (2016). *The Challenger launch decision: Risky technology, culture, and deviance at NASA. Enlarged edition*. University of Chicago Press.
- Veltman, J. A., & Gaillard, A. W. K. (1996). Physiological indices of workload in a simulated flight task. *Biological Psychology*, 42(3), 323-342.
- Vess, M., Arndt, J., & Cox, C.R. (2012). Faith and nature: The effect of death-relevant cognitions on the relationship between religious fundamentalism and connectedness to nature. *Social Psychological and Personality Science*, 3(3), 333–340. DOI: 10.1177/1948550611420303
- Voas, J. (2004). Software's secret sauce: the "-ilities". *IEEE Software*, 21(6), 14-15.
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, 428(6984), 748-751.
- Wang, X., Bisantz, A. M., Bolton, M. L., Cuvuoto, L., & Chandola, V. (2020). Explaining Supervised Learning Models: A Preliminary Study on Binary Classifiers. *Ergonomics in Design*, 28(3), 20-26.
- Wang, S., Lilienfeld, S. O., & Rochat, P. (2015). The uncanny valley: Existence and explanations. *Review of General Psychology*, 19(4), 393–407. <https://doi.org/10.1037/gpr0000056>
- Wang, J., She, M., Nahavandi, S., & Kouzani, A. (2010, December). A review of vision-based gait recognition methods for human identification. In 2010 international conference on digital image computing: techniques and applications (pp. 320-327). IEEE.
- Waytz, A., Heafner, J., & Epley, N. (2014). The mind in the machine: Anthropomorphism increases trust in an autonomous vehicle. *Journal of Experimental Social Psychology*, 52, 113-117.
- West, M., Kraut, R., & Ei Chew, H. (2019). *I'd blush if I could: closing gender divides in digital skills through education*. EQUALS
- Wegner, D. M. (1986). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185-208). New York: Springer-Verlag.
- Wildman, J. L., Shuffler, M. L., Lazzara, E. H., Fiore, S. M., Burke, C. S., Salas, E., & Garven, S. (2012). Trust development in swift starting action teams: A multilevel framework. *Group & Organization Management*, 37(2), 137-170.

- Wilson, G. F. (2002). An analysis of mental workload in pilots during flight using multiple psychophysiological measures. *The International Journal of Aviation Psychology*, 12(1), 3-18.
- Wittenbaum, G. M., Vaughan, S. I., & Stasser, G. (1998). Coordination in task-performing groups. In: R. S. Tindale et al., (Eds.), *Theory and research in small groups* (pp. 177-204). Plenum Press.
- Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330, 686-688.
- You, S., & Robert, L. (2018). Teaming up with robots: An IMOI (inputs-mediators-outputs-inputs) framework of human-robot teamwork. You, S. and Robert, LP (2017). *Teaming Up with Robots: An IMOI (Inputs-Mediators-Outputs-Inputs) Framework of Human-Robot Teamwork*, *International Journal of Robotic Engineering, (IJRE)*, 2(3).
- Yu, J. C., Chang, T. Y., & Yang, C. T. (2014). Individual differences in working memory capacity and workload capacity. *Frontiers in Psychology*, 5, 1465.
- Zellmer-Bruhn, M. E., Maloney, M. M., Bhappu, A. D., & Salvador, R. B. (2008). When and how do differences matter? An exploration of perceived similarity in teams. *Organizational Behavior and Human Decision Processes*, 107, 41–59.
- Zettler, I., Thielmann, I., Hilbig, B. E., & Moshagen, M. (2020). The nomological net of the HEXACO model of personality: A large-scale meta-analytic investigation. *Perspectives on Psychological Science*, 1-38. <https://doi.org/10.1177/1745691619895036>
- Zhao, J., Wang, T., Yatskar, M., Ordonez, V., & Chang, K. W. (2018). Gender bias in coreference resolution: Evaluation and debiasing methods. *arXiv preprint arXiv:1804.06876*.
- Złotowski, J., Yogeewaran, K., & Bartneck, C. (2017). Can we control it? Autonomous robots threaten human identity, uniqueness, safety, and resources. *International Journal of Human-Computer Studies*, 100, 48– 54. <https://doi.org/10.1016/j.ijhcs.2016.12.008>
- Złotowski, J., Sumioka, H., Nishio, S., Glas, D. F., Bartneck, C., & Ishiguro, H. (2016). Appearance of a robot affects the impact of its behaviour on perceived trustworthiness and empathy. *Paladyn, Journal of Behavioral Robotics*, 1(1), 55-66.
- Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*: New York: Public Affairs, 2019.

Acknowledgements

The authors are grateful for comments on an earlier draft by Adam Russell and Brian Pierce and for informative conversations with Craig Lawrence, Alonso Vera, John Paschkewitz, Joshua Elliot, and Jiangying Zhou. This work was supported by internal research funds at the University of Maryland Applied Research Laboratory for Intelligence and Security.

Disclaimers

Any opinions, findings, conclusions, or recommendations expressed in this publication do not necessarily reflect the views of DoD. Additionally, neither DoD nor any of its employees make any warranty, expressed or implied, nor assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, product, or process included in this publication.

Certain commercial entities, equipment, or materials may be identified in this document to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the Applied Research Laboratory for Intelligence and Security (ARLIS), the University of Maryland, or the DoD, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

About ARLIS

Applied Research Laboratory for Intelligence and Security (ARLIS) is a UARC based at the University of Maryland College Park and established in 2018 under the auspices of the OUSD(I&S). ARLIS is intended as a long-term strategic asset for research and development in artificial intelligence, information engineering, acquisition security, and social systems. One of only 14 designated United States Department of Defense (DoD) UARCs in the nation, ARLIS conducts both classified and unclassified research spanning from basic to applied system development and works to serve the U.S. Government as an independent and objective trusted agent.

Technical Points of Contact:

Breana Carter-Browne, Ph.D.
Research Associate, ARLIS
bmcarter@umd.edu

Susannah Paletz, Ph.D.
Research Professor, UMD College Park
paletz@umd.edu