Information Systems Research

Editorial for the Special Section on Humans, Algorithms, and Augmented Intelligence: The Future of Work, Organizations, and Society

Hemant Jain, Balaji Padmanabhan, Paul A. Pavlou, T. S. Raghu

informs
http://pubsonline.informs.org/journal/isre

INFORMATION SYSTEMS RESEARCH
Vol. 32, No. 3, September 2021, pp. 675–687
ISSN 1047-7047 (print), ISSN 1526-5536 (online)

Editorial for the Special Section on Humans, Algorithms, and Augmented Intelligence: The Future of Work, Organizations, and Society

Hemant Jain,[a] Balaji Padmanabhan,[b] Paul A. Pavlou,[c] T. S. Raghu[d]

[a] Gary W. Rollins College of Business, University of Tennessee – Chattanooga, Chattanooga, Tennessee 37403; [b] Muma College of Business, University of South Florida, Tampa, Florida 33620; [c] C.T. Bauer College of Business, University of Houston, Houston, Texas 77204; [d] W.P. Carey School of Business, Arizona State University, Tempe, Arizona 85287

Contact: hemant-jain@itc.edu, https://orcid.org/0000-0001-6184-1565 (HJ); bp@usf.edu, https://orcid.org/0000-0002-3498-0778 (BP); pavlou@bauer.uh.edu, https://orcid.org/0000-0002-8830-5727 (PAP); raghu.santaman@asu.edu, https://orcid.org/0000-0003-1071-6339 (TSR)

https://doi.org/10.1287/isre.2021.1046

Abstract. Recent developments in artificial intelligence (AI) have increased interest in combining AI with human intelligence to develop superior systems that augment human and artificial intelligence. In this paper, augmented intelligence informally means computers and humans working together, by design, to enhance one another, such that the intelligence of the resulting system improves. Intelligence augmentation (IA) can pool the joint intelligence of humans and computers to transform individual work, organizations, and society. Notably, applications of IA are beginning to emerge in several domains, such as cybersecurity, privacy, counterterrorism, and healthcare, among others. We provide a brief summary of papers in this special section that represent early attempts to address some of the rapidly emerging research issues. We also present a framework to guide research on IA and advocate for the important implications of IA for the future of work, organizations, and society. We conclude by outlining promising research directions based on this framework for the information systems and related disciplines.

Keywords: augmented intelligence • artificial intelligence • human-computer symbiosis • collective intelligence • future of work

> The main intellectual advances will be made by men and computers working together in intimate association.
>
> (Licklider 1960, p. 4)

1. Introduction

The desire to overcome the physical and intellectual limitations of human beings has fueled societal progress since prehistoric times. This has often resulted in the development of infrastructures on a massive scale, including transportation, civic, communication, and computing infrastructures. The latest such endeavor to overcome human limitations targets *intelligence* and represents relatively early stages of building out the infrastructure for augmenting human intelligence. The emerging intelligence augmentation (from hereon also referred to as IA) infrastructure has the potential to pool the intelligence of human beings and computers to transform individual work, businesses, institutions, and even society in an unprecedented manner and at scale. Studying the wide variety of issues that will arise in this transformation journey is a multidecade, even possibly multicentury, effort. This commentary and the corresponding special section attempt to showcase some of the early attempts at understanding the societally relevant challenges and opportunities of IA. In addition to a brief review of the papers in this special section, we present a framework to guide research on IA and its impact on the future of work, organizations, and society. Using the framework, we lay out promising research directions for the information systems (IS) and related disciplines.

Although artificial intelligence (AI), machine learning, and other autonomic technologies are usually in the spotlight, many important problems are often solved through humans and computers working alongside each other cooperatively. For the special section, augmented intelligence informally meant computers and humans working together, by design, to enhance one another, such that the intelligence and performance of the resulting system, as a whole, is superior. Our perspective here is intentionally bidirectional, that is, IA happens when AI enhances human intelligence and also when human beings enhance AI. More importantly, rather than focusing on how each side can help the other, we propose that our focus should be on designing IA such that the system, as a whole, operates with superior intelligence than either side alone. Because intelligence itself is hard to measure directly, often the objective in any IA design is to enhance overall performance in a given context, rather than intelligence directly, with the assumption that superior intelligence is likely to enhance performance. Although this approach mirrors today's practical focus when it comes to the widespread interest in IA, we must recognize that *performance* and *intelligence* are distinct concepts, leaving open interesting future possibilities to conceptualizing effective IA designs that directly improve the (augmented) intelligence of the system, irrespective of the overall performance of the system.

In our view, IA therefore requires a focus on design that optimally combines the abilities of human beings with various AI technologies and algorithms. An important open question in any such design is where the ultimate control belongs—with humans or the machine? As such, interesting directions emerge for IS researchers in the design of IA systems that focus on control, interactions between humans and machines, their interface points, and the delegation of work between humans and AI, such that the system as a whole improves (as measured directly by improved intelligence or indirectly by performance) over other alternate designs for the problem at hand. Applications of IA are beginning to emerge in a number of domains, such as cybersecurity, counterterrorism, healthcare, and space exploration, among others. This special section of *Information Systems Research* invited researchers to submit their best work to highlight how they are beginning to seamlessly integrate human and computer intelligence to solve interesting and important problems that impact the future of work, organizations, and society.

IA emerged as an important area of research, even during the earliest days of AI. In the 1960s, Engelbart and Licklider (who both managed research programs at Defense Advanced Research Projects Agency (DARPA)) pioneered the arguments for human-computer symbiosis (Licklider 1960). A fundamental assumption behind the need for human-computer symbiosis is that computers and human brains have different problem-solving abilities. As such, IA research pursues design ideas that are intended to optimize the combined computational potential of humans and computers. One branch of IA very familiar to IS researchers is human computer interaction (HCI). Winograd (2006) describes the tensions between the AI and the HCI camps and the associated "rationalistic" and "design" perspectives that they represent. Some AI researchers attempted to model human beings as cognitive machines and sought to build human-like AI systems. HCI, on the other hand, focused on a design approach that emphasizes interpretation, human behavior, and experimentation. Winograd quotes David Kelley, the renowned design thinker, as saying, "Enlightened trial and error outperforms the planning of flawless intellect," which suggests the importance of iteratively improving by modeling the interaction between humans and AI (Winograd 2006, p. 1257). However, HCI is not the only perspective for human-computer symbiosis. Large-scale computational problems often cannot be solved by either computers or human beings alone and are often termed human computation problems (von Ahn and Dabbish 2008). For instance, crowdsourcing strategies for many messy large-scale image or character recognition problems fall into this domain. Human computation problems rely on harnessing the collective intelligence of human beings to solve problems that computers are not yet good at solving.

Human-computer symbiosis has the potential to address some of the most difficult issues facing society today. Indeed, IS researchers have embraced both AI and IA traditions, highlighting the design and rational schools of thoughts in research papers, notes, and commentaries (see, for example, Dhar et al. 2014, Gregory and Muntermann 2014, Meyer et al. 2014, Clarke et al. 2016, Rai et al. 2019). However, there is still a lack of integrated discussion and a comprehensive body of literature on the direct implications of how IA and AI research can contribute to various applications to individuals, organizations, and society and to their impact on the future of work. This special section of *Information Systems Research* intends to begin a new dialog on the potential synergies between IA and AI within the context of IS research.

In this editorial commentary, we lay out the historical context; we introduce a framework for studying the future of work implications from an individual, organizational, and societal perspective; and we discuss important research directions for IS research and related disciplines in the domain of AI and IA.

2. Historical Context of AI

Alan Turing's question "Can Machines Think?" and the associated Turing Test (Turing 1950) are widely thought of as the beginning of AI over seven decades ago. Turing famously commented that, "at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted (Turing 1950, p. 442)." Indeed, this prophecy was decades ahead of core developments in AI, and much that has played out in AI and machine learning over the last two decades essentially has validated it.

In the early days, AI was often associated with problem solving using basic computational methods (Minsky 1961). Some early examples of problems included playing checkers (the first game-playing example) or chess—activities that humans normally associate with intelligence. This focus on devising computational methods to solve problems that demonstrate intelligence often translated into "search problems," where algorithms had to explore large search spaces of possible solutions to solve problems. In *Steps Toward Artificial Intelligence,* Minsky (1961, p. 8) notes that "when we do not know how to solve a certain problem, we may program a machine (computer) to *search* through some large space of solution attempts"; and he proceeds to lay out pattern recognition, learning, planning, and induction as ways to improve this search process.

Newell and Simon (1976, p. 114) in the tenth Turing Lecture reflected on the initial years of AI, and they argued that the two concepts of (a) storing and manipulating symbols and (b) heuristic search techniques lie at the heart of being able to capture intelligence. They wrote that "both (these) conceptions have deep significance for understanding how information is processed and how intelligence is achieved." Keeping problem-solving as the focus, Simon (1983) noted three typical (but not mutually exclusive) ways of doing this in AI: search, identified reasoning (based on logic, with its axioms and rules of inference), and constraint satisfaction (a precursor to many subsequent optimization problems such as assignment). It was generally understood during this time that *knowledge representation* is itself a critical piece of the puzzle in the quest to build intelligent machines that could reason about complex problems.

Early efforts in this direction were led by researchers in the field of expert systems. These efforts at Stanford University's heuristic programming project (Feigenbaum 1981) produced well-known application systems like MYCIN (Buchanan and Shortliffe 1984). However, limitations of hardware technologies and difficulties in acquiring knowledge limited the practical applications of these types of systems in business and other real-life settings. This resulted in a significant slowdown of business interest and investments in these types of systems throughout the 1980s and 1990s.

Early AI was predominantly symbolic in terms of the how it approached knowledge representation and manipulation. Although there were connectionist ideas (Rosenblatt 1957), these did not develop until much later (Rumelhart et al. 1986, Lecun et al. 1998, Raghu et al. 2001, Hinton et al. 2006) when algorithms that could train reasonably large connectionist models developed. As we now know, much of modern-day AI and deep learning are rooted in the connectionist paradigm, which has led to innovations in a range of problem-solving scenarios, popularly known as "big tech," that demonstrate human or superhuman-like intelligence in areas such as image classification (Krizhevsky et al. 2012), self-driving (Pomerleau 1991), language (Vaswani et al. 2017), playing games (Silver et al. 2016), and protecting individual and organizational privacy (Liu and Pavlou 2021). The widespread availability of voice-activated assistants like Alexa, Siri, and Google assistant; smart cars; and drones have brought AI to the public realm; and this is expected to continue to generate the vast amounts of data needed to train AI models.

Beyond its applications in big tech, AI has the potential to become a fundamental and pervasive technology for all business operations across firms and across industries. It has been used to automate many business tasks, such as selecting applicants, approving loans, identifying fraudulent transactions, and protecting privacy (Liu and Pavlou 2021). However, successful AI applications require large volumes of relatively clean data, which are challenging to get in many practical organizational operations. Repetitive tasks, such as in warehouses, assembly lines, and fast-food restaurants, have been early targets for automation because it is relatively easier to generate and capture quality data in such task scenarios.

With the move into mainstream organizations, other important considerations have started to emerge. For instance, business applications can, in many circumstances, require explanations for recommendations, actions, and decisions. To accept the use of AI applications in such cases, users would have a need to understand the underlying algorithmic process and/or be offered explanations that accompany actions. With increasing AI-powered automation in business, basic questions about how issues like fairness, bias, ethics, and impact are responsibly addressed have led to a growing body of work both in academia and in industry. Operationally, deciding what type of decision making can be delegated to machines remains a challenging question for businesses to resolve today. Management researchers have started to focus on these topics and specifically on how AI relates to and impacts organization design. Krogh et al. (2018, p. 408) argued that "in the long run outsourcing 'intelligence' to machines will neither be useful nor morally right. Although such technologies have many attractive features, they merely emulate human cognitive processes and cannot substitute the great flexibility, adaptability, and generativity we generally associate with human intelligence." Many companies have used AI to automate processes, albeit those that deploy AI mainly to displace employees may see only short-term productivity gains. A research study by Wilson and Daugherty (2018) involving 1,500 companies found that firms achieve the most significant performance improvements when human beings and machines work together, again pointing to the value that a broader intelligence augmentation perspective can bring to organizations.

AI today is moving into high-stakes applications, such as universal self-driving cars, autonomous weapons and warfare, managing entire energy grids, and precision medicine and surgery (Dietterich 2017). In such application contexts, there are now calls for "robust AI" systems that can deal with unknowns in an autonomous manner. This has been motivated in large part from observations by leading researchers (Bengio 2019, Marcus 2020) that a large part of deep learning's success has come from being able to learn from massive data but without necessarily being able to generalize beyond the realms of training data to learn the way humans might learn. The exact way of developing such robust AI is not fully known at this point, and it will likely need another decade or two of research to reach in practice. This call for AI to generalize in a robust manner beyond the bounds of its training data resonates in practical business applications as well. We are seeing AI take over financial advising (Sironi 2016), trading (Dempster and Leemans 2006), supply chains (Min 2010), customer interactions (Adam et al. 2020), and even management (Gladden 2014). Although organizations welcome fully automated solutions with open arms (unlike purely computational applications), there is a greater appreciation for the role of intelligent human beings and domain experts. Hence, one path toward robust AI in organizations will be through the use of human beings and AI systems in newer integrated architectures that can build systems that are better than either human or AI alone, which we turn to next.

3. Intelligence Augmentation and the Future of Work

As a future IA infrastructure develops in academia and practice, one of the pressing questions is what work will be like in the future, a consideration that typically spawns related questions, such as what skills human beings are likely to need, or will require to be trained in; what educational institutions need to teach students for the jobs of tomorrow; and what public policies might be needed to leverage the benefits of IA while properly addressing the human issues that arise in this complex and rapidly evolving process. Our consideration of IA in this section therefore focuses on this multifaceted lens.

To understand how IA is leveraged in future work environments, IS researchers will need to focus on the human and AI abilities that are leveraged in redesigning various work tasks. The labor economics and psychology literatures have developed the core concepts of human abilities (Peterson et al. 2001). The body of knowledge from this literature has been codified in a job classification system by the U.S. Department of Labor known as O*NET (https://www.onetonline.org/). Although this was not the original intent, the human abilities documented in this classification system have become a key basis for predicting future job impacts through the infusion of AI. At the most granular level, there are 52 distinct human abilities broadly grouped under four categories in the taxonomy: cognitive, psychomotor, physical, and sensory abilities. Frey and Osborne (2017) utilized O*NET task classification to identify the key hurdles to cross for task augmentation as perceiving and manipulating objects in unstructured environments, creativity, and emotional and social intelligence.

AI researchers have developed foundational augmentation technologies that are typically implemented as software and/or hardware artifacts to solve basic scientific problems associated with human language (e.g., BERT model), sensing (e.g., computer vision), cognitive inference (e.g., algorithms for learning, problem solving and reasoning), and human movement (e.g., robotics). The research focus at this level is generally independent of specific application areas. A case in point is the research effort around the common objects in context challenge (Lin et al. 2014). The algorithms developed to address the object identification challenge have led to more industrialization of computer vision techniques (Zhang et al. 2021). Such technologies can find applications *across* organizations, industries, sectors, and occupations. To make significant contributions in IA, IS researchers will likely need to keep a close watch, and even contribute to some of the foundational augmentation technologies that offer powerful building blocks for the transformation of future work environments and organizational processes. Although these technologies are usually associated with AI, viewing them explicitly as augmentation technologies can suggest a purpose—beyond many of the common AI use cases—that focuses on augmenting human abilities.

Leveraging the idea of human-computer collaboration, we therefore conceptualize the future of work to emerge from this interplay between basic human abilities critical to work environments and the ability of augmentation technologies to support or perform such tasks. This interplay will define the evolution of task augmentations from both a micro and macro perspective (see Figure 1) while driving societal investments in human capital and technology development to areas where these are the most critical. Fundamental shifts in societal priorities on education are inevitable as augmentation technologies penetrate work, organizational, societal, and even household environments.

Theoretical research in augmenting human abilities at the task level can involve improvements in specific occupations/tasks that are in some ways derived from foundational augmentation technologies and are also independent of any organization-specific nuances. For instance, IA research on augmenting how human beings communicate will leverage foundational language, sensory, and inference technologies. Such augmentation could apply to a variety of tasks in a number of occupations and is therefore not organization specific. Another example would be adaptation of computer vision technologies and robotics for enhancing human ability to

Figure 1. (Color online) Conceptualizing the Future of Work in an Augmented Intelligence World

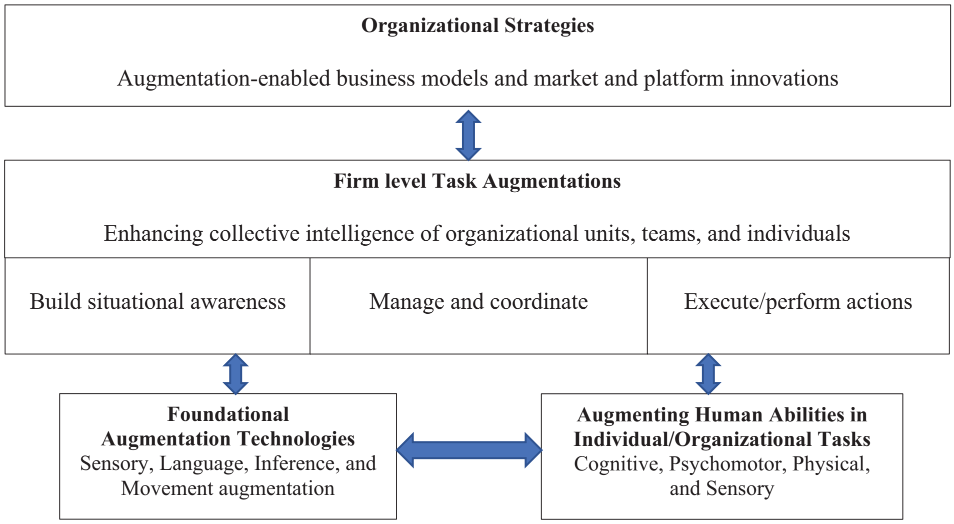

handle hazardous materials and conditions. These types of—typically organization or context independent—augmentation are represented at the lowest level in Figure 1.

At the organization level, based on organizational strategies that considered the business environment and economics, such as labor costs, firms may selectively adapt foundational augmentation technologies. Such *organization-specific adaptations* of occupation/task-specific augmentation technologies support organizational level task augmentation. Thus, we expect organization-specific adaptations to result in an augmented intelligence environment that creatively assembles work environments that leverages organization-specific resources to enhance individual and team abilities in the organization. For instance, a mining company could adapt self-driving technologies (occupation/task augmentation technologies) to mining tasks (firm-specific adaptations). The adaptation process would leverage internal knowledge, industry-specific knowledge, and local economic conditions in augmenting human abilities in the mining operation, such as identifying new areas to mine and drill with intelligent autonomous machines far below the earth's surface. The assembly of such technologies may integrate organization-specific research and development (R&D) in materials and extraction technologies to create unique competitive advantages. Based on the local labor market, certain augmentation technologies or certain human abilities may be more cost-effective. Such organization-level decisions could potentially confer several competitive advantages. As such, IS researchers studying organizational workflows, processes, and operations will need to start incorporating IA as a practical component toward enhancing organizational performance, thereby including another important theoretical construct in research models that seeks to study factors that contribute to organizational performance and competitive advantages.

Whereas the focus of augmentation at the occupation/task level is the human abilities, the organizational level adaptations will additionally need to focus on abilities that augment the collective intelligence of the organization as a whole. The information processing view (Galbraith 1974) of the organization is a useful theoretical framework to parse how firm-specific adaptations of augmented intelligence could occur. Mitigation of task uncertainties is fundamental to how we observe variations in firm performance. The potential for IA to alter the uncertainty calculus of organizations poses interesting questions for organizational and IS researchers. Building upon the information processing view of organizations, Malone (1987) defines multiple coordination structures that have differing performance implications given computational, coordination, and communication costs. The coordination structures can be hierarchical (product or functional) or market based (centralized or decentralized). Performance outcomes of teams and organizations manifest based on how the information processing tasks are organized. The information processing actions can predominantly be placed under three distinct buckets: situational awareness, management and coordination, and execution or action performance. This perspective of organizing is rooted in Galbraith's information processing view of the organization (Galbraith 1974). These three buckets also align well with McGrath's (1984) group task circumplex. Situational awareness aligns with tasks performed in teams to generate ideas and plans based on the available information. Management and coordination map to decision choices performed either through coordination or negotiations. Execution and action performance cover psychomotor and intellective group tasks in the task circumplex.

Augmentation of human intelligence can dramatically alter the structure and magnitude of coordination costs (in addition to the changes to task uncertainties) and thus leads to significantly different organizational structures and performance implications. More broadly, organizational adaptation of IA can lead to new business models, new products/services, and potential entry into new markets. In essence, organization-specific adaptations of IA have the potential to open entirely new theorizations of organizational strategies and create novel sources of competitive advantage.

Framing the IA concept into foundational, task-level, and organization-level developments enables us to assess key focus areas for research that explores the interplay of human and AI capabilities. The three categories also capture the economic significance of the investments in IA technologies to individuals, organizations, and society. Furthermore, the rationale for organizational studies becomes clearer with this perspective as organization-specific and occupation/task-specific intelligence augmentation could be viewed as resource adaptations that generate distinct organizational capabilities that may create competitive advantages.

We also know quite well in IS research the importance of societal and public policy contexts and impacts in technology adaptation (Ganju et al. 2016). We expect IA research to also address how societal, market, and organizational contexts influence—and get influenced by—intelligence augmentation in practice. Many national governments have recognized the potential for AI to revolutionize their societies and have developed national road maps and strategies. For example, Germany's roadmap outlines a number of investment plans and cross-country collaborations to accelerate AI innovations (Germany 2019). The roadmap also reviews and explicitly recognizes the privacy, ethics, and responsible development and oversight of IA technologies. Such policy initiatives will likely have a direct impact on how intelligence augmentation happens in various contexts. A similar report from the United States more directly recognizes the augmentation potential and identifies human-AI collaboration as a strategic national priority (NSTC 2019). The shifts in national strategies and policies are bound to create new dynamics in business environments locally and globally. As such, IS researchers will have new opportunities to refine existing theories and develop new theories on the interactions between technology investments, business environments, and firm strategies that can have broader public policy and societal implications.

4. Strategy and Mechanics of Human-AI Collaboration

Although the framework described in Section 3 offers some ways to think about how organizations can develop IA strategies for competitive advantage, the firm-specific adaptations of IA at the task and strategic levels do require consideration of some more granular architectures of how human beings and AI can collaborate, along with a deeper understanding of many issues related to such collaborative architectures. To highlight some of these architectures, we draw on the most recent United States National AI R&D Strategic Plan developed by the National Science and Technology Council (NSTC) (NSTC 2019), which calls specifically for better research in human-AI collaboration—one of eight strategic areas that are highlighted. For our purpose here, the plan outlines three ways in which AI can work with humans:

- AI works alongside humans to accomplish peripheral tasks and generally looks to support the human expert.
- AI takes over when the human has high cognitive load.
- AI replaces humans in areas where humans have limited strengths or the environment is toxic or when real-time response is key.

An example of a scenario where AI works alongside human beings can be a physical security system, such as the Knightscope security robots (Knightscope 2021). The security robot can monitor and analyze on the ground situations in real time and can report unusual situations to human beings for further action. In hazardous scenarios, such as chemical spills or deep-sea explorations, AI can work alongside human beings for enhancing their situational awareness and support planning. However, in dangerous conditions, AI can completely take over the performance of certain risky actions. The richness of the delegation options available to organizations in designing future work environments creates new opportunities for theorizing optimal AI and human collaborations, delegation patterns, and organizational structures to develop effective IA applications for dangerous conditions and unpredictable environments.

AI may be able to completely replace human beings in limited, structured environments. Given the limitations of current AI capabilities, the need for a structured environment is very critical. One of the reasons for the success of AI research in games like chess, Go, and others is that games often provide an environment where the rules of engagement and boundary conditions are clear. Organizational work environments—even virtual ones—are often not very structured; therefore, even limited replacement of human beings with AI in existing work environments is bound to run into practical challenges, at least for the foreseeable future. This limitation, however, should be considered as an opportunity for organization-level innovation and opportunities for research. For example, a reimagined kitchen configuration is necessary to replace the sous chef in the kitchen with a robot. If the robot and a human being are working together, the rules of engagement become especially important, such that the interactions are intentionally structured, reasonably safe, and even noise free.

The NSTC report (NSTC 2019) also calls for research in better algorithms that are human aware as well as for better systems that take into account human strengths and weaknesses. These are challenging problems that need to take into account complex issues, such as user-oriented design, ease of use, trust, and the building of effective adaptive systems that are also data driven. The kind of work needed to make these happen in a robust and scalable way will have to be socio-technical in nature and will need to model human beings and their behavior as well as model the behavior of AI. These are areas where we have limited understanding today and need significantly more theoretical and empirical research. As a case in point, only recently we are starting to understand effective ways of delegation between human beings and AI (Fügener et al. 2019), where we look at not just how AI can work with human beings but also how human beings can work simultaneously with AI. Yet, there are some exciting new ideas in this context that the papers in this special section offer, which we summarize below. However, these are the proverbial tip of the iceberg, and they are likely to be as useful in exploring the questions they further inspire as the ones that they attempt to solve in their various specific and broader contexts.

Evaluating human beings for the quality of work they provide is at the heart of most human resource management systems. In the context of highly skilled labor provided in online marketplaces, this is arguably even more important because the users of these services rarely have prior first-hand experience in using a certain person for a highly skilled task. Reputation systems used for this purpose are often based on either inputs provided by human users of those systems (i.e., feedback ratings of the workers) or based on machine learning methods that use networks and other attributes. For this problem, the paper by Kokoddis (2021) in this special section presents design principles and a method for combining human inputs with machine learning to build more robust reputation systems, demonstrating the power of IA in the context of customer-driven reputation systems. Interestingly, for the same broader problem is the question of what happens when we use crowdsourcing for low-skilled labor. Different opportunities exist and so do different problems. Low-skilled human labor has been extensively used for labeling work tasks, such as image classification or extracting ground-truth information from data, such as feedback reviews. Often these crowdsourcing tasks are used to build large-scale machine learning systems; in a sense, this use of human labelers along with building AI models is itself a form of intelligence augmentation and one that has been used now for over a decade. The motivations for using crowdsourcing with low-skilled labor range from training AI models to lowering information processing costs. The use of low-skilled labor in crowdsourcing does result in quality problems—an issue tackled by Kokoddis (2021). Focusing on tasks where there are multiple labels that are needed from the workers (e.g., labeling a news article into multiple categories simultaneously or labeling a road image with all the items inside the image), Yin et al. (2021) present a method that exploits label dependency to simultaneously correct for these errors. In this process their method also provides better worker quality estimates, so down the road, these workers can either be removed from those tasks that do not suit their skills or ideally be redirected to other tasks where they do have the necessary skills or expertise.

The above examples do suggest that online labor markets—for low-skilled or high-skilled workers—do present opportunities for clever redesigning from an IA perspective, where AI models and human beings work together in symbiotic and synergistic fashions to jointly bring out augmented intelligence that is greater than each individual component alone. We caution, however, that simply rebranding online labor markets as augmented intelligence systems because humans are automatically assigned tasks by machine learning systems will not do justice to the range of IA opportunities that are present when viewed through a broader lens. The work of Kokkodis (2021) and Yin et al. (2021) in this special section point to deeper and more impactful opportunities for the design of IA applications in the context of online labor markets.

As noted earlier, we know little about the mechanisms by which human beings work with AI systems today. Building a future with augmented intelligence is going to require significant work in this area, and this likely will be an area where our insights and knowledge will evolve over time as humans and AI systems jointly evolve in their mutual understanding of each other's strengths and weaknesses. This process will most likely also be impacted by technology (as AI gets better) and education (as humans learn how to best work with AI). This special section offers three examples in application areas as diverse as customer service, fintech, and healthcare. Schanke et al. (2021) address the issue of what happens when chatbots become more human-like by design. In a field experiment, this paper sheds light on how customers behave as the AI they interact with increasingly has human-like characteristics such as humor. The results suggest that it might help transaction completion, but might simultaneously create a negotiating mindset for consumers. Ge et al. (2021) studied investors interacting with robot advisors in a P2P lending scenario and show that these investors often do not make optimal decisions on when to use these services, suggesting the need for better augmentation capabilities if we are to use AI to help investors make better decisions. Jussupow et al. (2021) highlight a similar scenario in the case of diagnosis decisions in healthcare organizations. This work showed the urgent need for better understanding metacognitive processes around how physicians make decisions when working alongside AI systems. Notably, AI systems do fail, and a proper understanding of how human beings could determine when to override these systems is crucial. Interestingly, Jussupow et al. (2021) also show that there are differences between novice and expert physicians in terms of how they make these decisions, pointing to perhaps the need for more careful education of human beings in every discipline where AI is being used today. Learning to deal with AI errors is crucial for better intelligence augmentation, and there is a significant amount of work needed to understand how best to do so in practice.

5. Potential Research Directions

To envision how IS researchers can navigate the significant research possibilities in this area, we discuss some specific contexts in detail and provide the design implications emerging from IA. The contexts discussed in this section illustrate the intersection of human control and IA systems and the associated research questions that arise. We choose IA in the context of financial markets to illustrate how human cognitive abilities are augmented for better situational awareness and execution. With human beings in the loop, this context throws open a number of research questions pertaining to decision making, human behaviors, delegation of responsibility, and control. We use the gig economy platform context to illustrate how management and coordination augmentation is making it possible to assemble and manage teams of workers to perform tasks. In many gig-work contexts, it is possible for human beings and computers to work alongside each other. The gig context is an example of how all human abilities (cognitive, psychomotor, physical, and sensory) can be leveraged. The third context we illustrate is autonomous driving, given its increasing role in society as an example of how AI might someday fully automate and replace human beings out of many tasks—a scenario we feel needs to be enriched through the discussion of how intelligence augmentation can play out in this context. Augmentation under self-driving primarily targets psychomotor, physical, and sensory abilities; this example shows how it is done at the individual level as well in terms of sensing (situational awareness), processing (management and control), and execution. The organization-level adaptation of self-driving technologies opens up a number of interesting research questions at the organizational as well as the societal level. Thus, the design space for researching the three levels of augmentation we described earlier promises to be a rich problem domain that requires a broader multidisciplinary perspective to design effective human-computer collaboration modes.

5.1. Augmented Intelligence in Financial Markets

The special section submissions included a number of manuscripts that tackled significant issues pertaining to market innovation possibilities with IA. Among the published papers in the section, the article by Ge et al. (2021) investigates an interesting research context pertaining to decision augmentation in financial market interactions. The key takeaway in this paper is the observation that human-in-the-loop can adversely impact performance. There is also evidence for recency bias in the human's use of recommendations made by the AI agent. Both behaviors point to the important need for researchers to examine and understand the human-AI decision interfaces to optimize performance. As Ge et al. (2021) point out, financial markets are ripe for innovation with robo-advising; such advisory services have the potential to democratize investing for large, underserved populations. However, the results in the paper point to the need for a lot more work in understanding the behavioral implications on users of robo-advising features. Machine learning algorithms are often designed without explicit considerations for human rationality assumptions. As outlined in numerous studies, human biases in decision making exhibit heterogenous and distinctive patterns (Tversky and Kahneman 1974, Ariely et al. 2009). Robo advisors in market platforms can be designed with personalization built in. However, inability of learning algorithms to identify and account for individual irrationality in the decisions can affect the quality and acceptability of the recommendations.

Beyond robo-advising, recent fintech innovations have further strengthened the "wisdom of the crowd" in the investment world and demonstrate the operationalization of the collective intelligence concept. For example, a recently introduced exchange-traded fund, SFYF, consists of 50 most widely held stocks of investors on the platform.[1] Variations of this theme can blend algorithmic inferences with human decisions to enrich the variety of financial products available to investors. A number of behavioral and algorithmic research issues emerge in the fintech world in such contexts: How can human intelligence and algorithmic innovations be effectively combined to create safer and better performing investment products? How can investors develop trusting relations with robo-advising and overcome their own decision biases? How can platforms and regulators prevent manipulation of investment products? How would financial industry roles (traders, financial advisors, financial analysts, and others) change in response to the infusion of augmented intelligence in financial products and platforms?

In our framework, the data-rich financial industry is very well poised to exploit developments in foundational augmentation technologies. Given the digital nature of most financial transactions, task-specific innovations are also infusing rapidly into financial markets. Firm-specific adaptation of these technologies can lead to the introduction of new investment and wealth management products. As more investment decisions are delegated to automated tools, investment advisor roles will likely be around "explaining" the decisions to clients as opposed to "executing" the decisions. Similarly, analysts may see their roles changing from "sense makers" to "curators" and "creators" of AI algorithms.

5.2. Augmented Intelligence in Gig Economy Platforms

In the gig economy, automated intermediation in platforms have led to recognition of the "algorithm as the boss" phenomenon. Gig economy platforms, such as Uber, Doordash, Lyft, Airbnb, and others, have reduced information frictions in several service industries. Real-time data on customers and platform workers have led to the creation of innovative digitally mediated services. The gig economy has also transformed white collar work as evidenced on digital platforms, such as Fiverr, freelancer, and Upwork. These platforms have led to the "taylorization of white-collar work" (Taylor and Bain 1999, p. 109). Although gig platforms expand work opportunities, they have also raised concerns on the quality of future work relations. The experience on these platforms for both workers and clients will depend largely on how algorithms and human capabilities combine to create new interaction mechanisms (Wood et al. 2019).

A positive view of algorithmic control would be that data, sensors, and communications enable better situational awareness and rational decision making through algorithmic means. Ride-sharing services can inform, recommend, and dispatch closest drivers to a customer to improve service experience. Algorithms can learn gig worker performance and feedback over time to recommend tasks that match their preferred tasks, times, and geographies. Gig workers who perform specific projects or tasks enjoy choice autonomy along multiple dimensions (customer, task, time, and platform choices). Access to a wider talent base enhances client utility as well. On the other hand, negative perceptions of gig work are also emerging from several perspectives (Kellog et al. 2020). Rational controls enable algorithmic monitoring of workers' behaviors that can sometimes be intrusive and impinge on the workers' privacy. Rating and rankings of workers can introduce bias due to race, religion, and other social identities, raising questions about algorithmic fairness.

Future work environments that integrate gig working models into existing organizational forms and structures open up a wide variety of questions for business and IS researchers. How can projects, tasks, and processes be atomized to enable blended work processes? Under what contexts do specific work modes (fully digital and hybrid) in work environments perform best? What algorithmic control design patterns need to be created to augment the intelligence of workers and customers? How can gig worker and customer reputation ratings be made portable to reduce lock-in and switching costs? How can platforms minimize threats to information validity (hacking, deception, and manipulation) and introduce effective redressal mechanisms? What managerial control mechanisms work best when blended with algorithmic control mechanisms? Several of the tasks performed today by gig workers (ride sharing, grocery delivery, coding, and design) are also tasks that can likely be automated in the next several years. Given this threat, how will gig work platforms evolve to provide new work opportunities in its ecosystem? These are important research and practical questions in the evolving domain of IA in the context of the gig economy.

5.3. Augmented Intelligence in Autonomous Driving Technologies

The automobile industry has greatly enhanced the driver assistance features available in vehicles in recent years. Sensor technologies, such as collision avoidance, blind spot monitoring, driver alert, automatic emergency braking, and others, have been found to greatly reduce the frequency and severity of vehicle crashes (HLDI 2020). The introduction of these technologies in consumer vehicles has enabled large-scale data collection efforts to explore human behavior in the presence of these technologies. A case in point is the Advanced Vehicle Technology Consortium (AVT; https://agelab.mit.edu/avt) consisting of automobile manufacturers, vehicles testers, component manufacturers, and insurance companies. AVT has launched a multiyear study for collecting driver behavior data to both understand human behavior and generate new data for algorithm development to enhance autonomous driving technologies.

Autonomous driving is a context where human psychomotor abilities are specifically targeted—either for replacement or augmentation. Autonomic emergency braking is a driving feature that takes over when humans are either not paying attention or unable to react in time. Blind spot monitoring is a driving feature that works alongside humans and enhances situational awareness. Adaptive cruise control is a vehicle feature that enables humans to delegate driving in a limited way. Although universal 100% autonomous driving is perhaps decades away, the issues that arise with autonomous driving are a microcosm of the issues organizations will face when worker psychomotor abilities are augmented or automated within organizations.

Although a lot of discussion in the public sphere is on automation through self-driving technologies, the augmentation potential (as highlighted by some of the examples above) needs greater focus and attention in research studies. For instance, in high-risk cases if both the human and technologies differ in what they choose to do (as was unfortunately the case in the 2018 Lion Air crash where the pilots and the autopilot were engaged in conflicting maneuvers), how should coordination and control be managed? Is there need for additional sensory data that could be useful (e.g., the human's eye movements or voice that might indicate lack of attention or distraction)? How would ethical issues be integrated? For example, if something does go wrong, does it matter if it was because of the human or the AI? Or, if the AI could have prevented a crash but didn't, was it unethical by design? How would privacy issues be integrated? For example, would humans want to allow in-vehicle technologies to track and use such information in the coordination and control algorithms? Should or how should the human being's own previous driving history be used in such cases to make these ethical or moral decisions? These questions raise many IA ideas for situational awareness, management and control, and execution that would be powerful directions for research.

5.4. Macrolevel Research Issues

Amidst the euphoria of augmented intelligence possibilities, it is also important to have a pragmatic view of the risks and ethical use of AI and IA. It is important that IS, as a discipline, builds a body of knowledge to navigate the ethical challenges of human intelligence augmentation. Privacy continues to be a perennially important topic because of the massive amounts of data collected today from consumers during their interactions with apps and devices and fed into AI algorithms. Data from a recent survey (Pew Research Center 2018) indicate that negative perceptions of AI use in financial services stems from concerns about privacy violations, data inaccuracy, and discriminatory and unfair practices.

Augmentation technologies, by their very nature, will collect massive amounts of data from individuals (e.g., employees and consumers). In some instances, such data may include valuable bio-signals (e.g., heart rate, blood pressure, and body temperature). The intrusive nature of such data can affect employee and customer attitudes toward IA and related technologies. As such, the risks associated with augmentation technologies will go beyond the usual technological and process-related risk factors. As to how data collected from workers get used inside and outside the organization within the regulatory, public policy, and cultural norms remains an open research question. Similarly, collecting consumer-related information signals (e.g., location, biomarkers, and others) will generate new research questions on the ethical, moral, and public policy dimensions. Although much of the IS literature has explored the privacy calculus and various privacy trade-offs (Pavlou 2011), new research models are needed to evaluate the consumer-level trade-offs involved in adopting augmentation technologies that may affect people's privacy. One recent example is the work of Liu and Pavlou (2021) in the context of mobile banking and the privacy trade-offs involved when leveraging an IT solution. More research on the organization-level trade-offs involved in assessing the cultural impacts of data collection balanced against the business value of such data are also needed.

Research in this domain will need to include the social-systems perspective to fully recognize potential risks to human beings and societies when IA systems are implemented within organizations and society (Crawford and Calo 2016). As described in Crawford and Calo (2016), the social-systems perspective can complement existing approaches to monitoring the IA systems usage and impacts by taking an inclusive approach to assessment integrating perspectives of stakeholders at all levels. As there is a growing concern about the quality and comprehensiveness of the data sources used in algorithms, current research models and theories of technological risks need refinement.

Investments in IA will challenge IS researchers to expand the boundaries of technology evaluations to incorporate human-AI collaborations and the associated social benefits and risks. As described in the previous section, IA contexts involve complex human behavioral dimensions and require more elaborate modeling of human behaviors as an integral part of the technology evaluation process. The possibility that improvement options in IA investments can stem from foundational technologies, process innovations, and collective intelligence makes technological investment evaluations research more complex and challenging. There is a rich body of literature in IS that addresses organizational investments in IT innovations from a technological and business perspective (see Benaroch 2018, for example). To address research issues pertaining to the business value of IA, the conceptual framework of our empirical models may need to change. Economists have also begun to notice the limitations of factor-based (i.e., labor and capital) models in explaining productivity and wage growth. Newer frameworks of automation impacts have begun to incorporate automation at the task level in empirical models (Acemoglu and Restrepo 2019). In the Acemoglu and Restrepo model, augmentation is assumed to create a displacement effect (replacement of labor with capital) and a reinstatement effect (introduction of new tasks in the industry). Integrating task-based perspectives can greatly enrich the extant IS business value research literature.

Furthermore, the loss of routine jobs to technology and automation has been extensively documented over the last several decades (Autor and Dorn 2013). Although technology, automation, and AI will continue to affect the future of work, the impact of AI is likely to support, complement, and augment the work and potential of human beings, at least in the foreseeable future. This is consistent with empirical evidence in the literature that technology enhances employment in organizations (Atasoy et al. 2016) and that technology skills enable workers to obtain employment with higher wages (Atasoy et al. 2021).

Finally, the public risk posed by ubiquitous IT (e.g., vulnerabilities in widely adopted operating systems and web browsers) is a precursor to what may happen with ubiquitous IA technologies. Current models of technology valuation, however, rarely factor in the downstream or upstream risks of ubiquitous technologies. Human-computer collaborations will increase the urgency of the need for more comprehensive risk models to help decision makers in making judicious choices on the scope and size of augmentation in products and processes. IS research has also extensively studied cultural issues (both organizational and national) surrounding the adoption, use, and outcomes of IT (Leidner and Kayworth 2006). There is very little research, however, on how technology or AI itself may impact culture. With IA infusion, it is likely that organizational and societal level impacts can themselves have a bearing on cultural values in the future. Grounded theory development on the bidirectional impacts of IA and cultural and societal values is needed to develop new knowledge in this domain.

6. Conclusion

Just as e-commerce and data analytics led to major theoretical and empirical innovations in IS research over the last two decades, intelligence augmentation promises to expand the boundaries of IS research in the coming decades. The rich empirical, analytical, behavioral, and design science research traditions of the IS discipline are likely to prove to be a strong theoretical foundation for the coming demand of cross-disciplinary research efforts in this burgeoning domain.

The IA framework we have described in this commentary highlights the complex interplay between technology, human abilities, and organizational strategies that we will encounter in the future. As such, IS researchers will need to embrace greater methodological and intellectual fluidity in future research studies. Empirical research has shown that organizations achieve significant performance improvements when human beings and machines work together effectively (Wilson and Daugherty 2018), as also shown by the papers in this special section. Often termed “collective intelligence” (Malone 2018, p. 3), the focus in the organizational context is on enhancing each other’s complementary strengths, specifically the leadership, teamwork, creativity, and social skills of human beings and the speed, scalability, and quantitative capabilities of machines. We are optimistic that the IA perspective can enable IS researchers to take the lead on what is likely to be one of the most impactful areas of research over the next few decades.

Finally, given that macrolevel implications, data, and research transparency issues (Burton-Jones et al. 2021) are likely to become more salient in IS research, studies will need to focus on transparency and responsible implementation of intelligence augmentation in organizations and society. As such, some of our related disciplines (economics, computer science, genomics, and others) have recognized the value of creating open data sets as an intellectual pursuit in itself. The IEEE DataPort (https://ieee-dataport.org/datasets) is an example of such an initiative. A similar initiative in the IS community would be needed to support organizational- and societal-level studies of the impacts of intelligence augmentation. Open data sets within specific application domains can also accelerate our research advances such that they match the pace of advances we are witnessing in augmentation technologies.

Acknowledgments

The authors acknowledge the support of the *Information Systems Research* Editor-in-Chief and the Editorial Board for their valuable comments and suggestions for structuring the special section call. The authors also thank the special section editorial team for their hard work in guiding the authors to increase the relevance, impact, and quality of the published manuscripts. Finally, the authors thank all authors who submitted their best works to this special section.

Endnote

[1] See https://www.sofi.com/invest/etfs/sfyf/.

References

Acemoglu D, Restrepo P (2019) Automation and new tasks: How technology displaces and reinstates labor. *J. Econom. Perspect.* 33(2):3–30.

Adam M, Wessel M, Benlian A (2020) AI-based chatbots in customer service and their effects on user compliance. *Electronic Markets*, ePub ahead of print March 17, https://doi.org/10.1007/s12525-020-00414-7.

Ariely D, Bracha A, Meier S (2009) Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *Amer. Econom. Rev.* 99(1):544–555.

Atasoy H, Banker R, Pavlou PA (2016) On the longitudinal effects of IT use on firm-level employment. *Inform. Systems Res.* 27(1):6–26.

Atasoy H, Banker R, Pavlou PA (2021) Information technology skills and labor market outcomes for workers. *Inform. Systems Res.* 32(2):437–461.

Autor H, Dorn D (2013) The growth of low-skill service jobs and the polarization of the US labor market. *Amer. Econom. Rev.* 103(5):1553–1597.

Benaroch M (2018) Real options models for proactive uncertainty-reducing mitigations and applications in cybersecurity investment decision making. *Inform. Systems Res.* 29(2):315–340.

Bengio Y (2019) From system 1 deep learning to system 2 deep learning. *Conf. Neural Inform. Processing Systems*, *Vancouver, BC*.

Buchanan BG, Shortliffe EH (1984) *Rule-Based Expert Systems: the MYCIN Experiment of the Stanford Heuristic Programming Project* (Addison-Wesley, Reading, MA).

Burton-Jones A, Boh WF, Oborn E, Padmanabhan B (2021) Editor's comments: Advancing research transparency at *MIS Quarterly*: A pluralistic approach. *MIS Quart.* 45(2):iii–xviii.

Clarke R, Burton-Jones A, Weber R (2016) On the ontological quality and logical quality of conceptual-modeling grammars: The need for a dual perspective. *Inform. Systems Res.* 27(2):365–382.

Crawford K, Calo R (2016) There is a blind spot in AI research. *NATNEWS* 538(7625):311–313.

Dempster MAH, Leemans V (2006) An automated FX trading system using adaptive reinforcement learning. *Expert Systems Appl.* 30(3):543–552.

Dhar V, Geva T, Oestreicher-Singer G, Sundararajan A (2014) Prediction in economic networks. *Inform. Systems Res.* 25(2):264–284.

Dietterich TG (2017) Steps toward robust artificial intelligence. *AI Magazine* 38(3):3–24.

Feigenbaum EA (1981) Expert systems in the 1980s. Bond A, ed. *State of the Art Report on Machine Intelligence* (Pergamon-Infotech, Maidenhead, England).

Frey CB, Osborne MA (2017) The future of employment: How susceptible are jobs to computerization? *Tech. Forecasting Soc. Change* 114:254–280.

Fügener A, Grahl J, Gupta A, Ketter W (2019) Cognitive changes in human-AI collaboration: Investigating the path towards productive delegation. Technical report, April 8. Available at http://dx.doi.org/10.2139/ssrn.3368813.

Galbraith JR (1974) Organization design: An information processing view. *Interfaces* 4(3):28–36.

Ganju KK, Pavlou PA, Banker RD (2016) Does information and communication technology lead to the well-being of nations? A country-level empirical investigation. *MIS Quart.* 40(2):417–430.

Ge R, Zheng ZE, Tian X, Liao L (2021) Human-robot interaction: When investors adjust the usage of robo-advisors in peer-to-peer lending. *Inform. Systems Res.*, ePub ahead of print July 20, https://doi.org/10.1287/isre.2021.1009.

Germany (2019) Artificial intelligence strategy. Accessed June 21, 2021, https://knowledge4policy.ec.europa.eu/publication/germany-artificial-intelligence-strategy_en.

Gladden ME (2014) The social robot as 'charismatic leader': A phenomenology of human submission to nonhuman power. Seibt J, Hakli R, Norskov M, eds. *Sociable Robots and the Future of Social Relations: Proc. Robo-Philosophy 2014* vol. 273 (Ios Press, Amsterdam), 329–339.

Gregory RW, Muntermann J (2014) Research note—Heuristic theorizing: Proactively generating design theories. *Inform. Systems Res.* 25(3):639–653.

Hinton GE, Simon O, Yee-Whye T (2006) A fast learning algorithm for deep belief nets. *Neural Comput.* 18(7):1527–1554.

HLDI (2020) Compendium of HLDI collision avoidance research. Accessed June 27, 2021, https://www.iihs.org/media/e635cc76-b9bc-4bad-a30a-5d7b78791df2/vxeQ3A/HLDI%20Research/Collisions%20avoidance%20features/37-12-compendium.pdf.

Jussupow E, Spohrer K, Heinzl A, Gawlitza J (2021) Augmenting medical diagnosis decision? An investigation into physicians' decision-making process with artificial intelligence. *Information Systems Research*, ePub ahead of print February 26, https://doi.org/10.1287/isre.2020.0980.

Kellogg KC, Valentine MA, Christin A (2020) Algorithms at work: The new contested terrain of control. *Acad. Management Ann.* 14(1):366–410.

Knightscope (2021) Retrieved June 1, 2021, https://www.knightscope.com/.

Kokkodis M (2021) Dynamic, multidimensional, and skillset-specific reputation systems for online work. *Inform. Sci. Res.*, ePub ahead of print March 31, https://doi.org/10.1287/isre.2020.0972.

Krizhevsky A, Ilya S, Hinton GE (2012) ImageNet classification with deep convolutional neural networks. Pereira F, Burges CJC, Bottou L, Weinberger KQ, eds. *Proc. 25th Internat. Conf. Neural Inform. Processing Systems*, vol. 1 (Curran Associates, Red Hook, NY), 1097–1105.

Krogh VG (2018) Artificial intelligence in organizations: New opportunities for phenomenon-based theorizing. *Acad. Management Discoveries* 4(4):404–409.

LeCun Y, Bottou L, Orr GB, Muller KR (1998) Efficient BackProp. Orr GB, Müller KR, eds. *Neural Networks: Tricks of the Trade*, Lecture Notes in Computer Science (Springer, Berlin), 9–50.

Leidner DE, Kayworth T (2006) A review of culture in information systems research: Toward a theory of information technology culture conflict. *MIS Quart.* 30(2):357–399.

Licklider JC (1960) Man-computer symbiosis. *IRE Trans. Human Factors Electronics* HFE-1(1):4–11.

Lin TY, Maire M, Belongie S, Hays J, Perona P, Ramanan D, Dollár P, Zitnick CL (2014) Microsoft coco: Common objects in context. *Proc. Eur. Conf. Comput. Vision* (Springer, Cham), 740–755.

Liu B, Pavlou PA (2021) Achieving a balance between privacy protection and data collection: A field experimental examination of a theory-driven IT solution. *Inform. Systems Res.* Forthcoming.

Malone TW (1987) Modeling coordination in organizations and markets. *Management Sci.* 33(10):1317–1332.

Malone TJ (2018) *Superminds* (MIT Media Press, Cambridge, MA).

Marcus G (2020) The next decade in AI: Four steps toward robust artificial intelligence. Preprint, submitted February 14, https://arxiv.org/abs/2002.06177.

McGrath JE (1984) *Groups: Interaction and Performance* (Prentice-Hall, Englewood Cliffs, NJ).

Meyer G, Adomavicius G, Johnson PE, Elidrisi M, Rush WA, Sperl-Hillen JM, O'Connor PJ (2014) A machine learning approach to improving dynamic decision making. *Inform. Systems Res.* 25(2):239–263.

Min H (2010) Artificial intelligence in supply chain management: Theory and applications. *Internat. J. Logist.: Res. Appl.* 13(1):13–39.

Minsky M (1961) Steps toward artificial intelligence. *Proc. IRE* 49(1):8–30.

Newell A, Simon HA (1976) Computer science as empirical inquiry: Symbols and search. *Comm. ACM* 19(3):113–126.

NSTC (2019) The national artificial intelligence research and development strategic plan: 2019 update. National Science and Technology Council (NSTC) (US), Select Committee on Artificial Intelligence. Accessed June 26, 2021, https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf.

Pavlou PA (2011) State of the information privacy literature: Where are we now and where should we go? *MIS Quart.* 35(4):977–988.

Peterson NG, Mumford MD, Borman WC, Jeanneret PR, Fleishman EA, Levin KY, Campion MA, et al. (2001) Understanding work using the Occupational Information Network (O* NET): Implications for practice and research. *Personnel Psych.* 54(2):451–492.

Pew Research Center (2018) Public attitudes toward computer algorithms. Accessed June 27, 2021, https://www.pewresearch.org/internet/2018/11/16/attitudes-toward-algorithmic-decision-making/.

Pomerleau DA (1991) Efficient training of artificial neural networks for autonomous navigation. *Neural Comput.* 3(1):88–97.

Raghu TS, Ramesh R, Chang AM, Whinston AB (2001) Collaborative decision making: A connectionist paradigm for dialectical support. *Inform. Systems Res.* 12(4):363–383.

Rai A, Constantinides P, Sarker S (2019) Next generation digital platforms: Toward human-AI hybrids. *MIS Quart.* 43(1):iii–ix.

Rosenblatt F (1957) *The Perceptron, a Perceiving and Recognizing Automaton Project Para* (Cornell Aeronautical Laboratory, Buffalo, NY).

Rumelhart DE, Hinton GE, Williams RJ (1986) Learning representations by back-propagating errors. *Nature* 323(6088):533–536.

Schanke S, Burtch G, Ray G (*2021*) Estimating the impact "humanizing" customer service chatbots. *Inform. Systems Res.*, ePub ahead of print May 24, https://doi.org/10.1287/isre.2021.1015.

Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van Den Driessche G, Schrittwieser J, et al. (2016) Mastering the game of Go with deep neural networks and tree search. *Nature* 529(7587):484–489.

Simon HA (1983) Search and reasoning in problem solving. *Artificial Intelligence* 21(1-2):7–29.

Sironi P (2016) *FinTech Innovation: From Robo-Advisors to Goal Based Investing and Gamification* (John Wiley & Sons, New York).

Taylor P, Bain P (1999) An assembly line in the head: Work and employee relations in the call centre. *Indust. Relations J.* 30(2):101–117.

Turing AM (1950) Computing machinery and intelligence. *Mind* 59(236):433–460.

Tversky AD, Kahneman D (1974) Judgment under uncertainty: Heuristics and biases. *Sci.* 185:1124–1131.

Vaswani A, Noam S, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser K, Polosukhin I (2017) Attention is all you need. Preprint, submitted June 12, https://arxiv.org/abs/1706.03762.

von Ahn L, Dabbish L (2008) Designing games with a purpose. *Comm. ACM* 51(8):58–67.

Wilson HJ, Daugherty PR (2018) Collaborative intelligence: Humans and AI are joining forces. *Harvard Bus. Rev.* 96(4):1–10.

Winograd T (2006) Shifting viewpoints: Artificial intelligence and human–computer interaction. *Artificial Intelligence* 170(18): 1256–1258.

Wood AJ, Graham M, Lehdonvirta V, Hjorth I (2019) Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work Employment Soc.* 33(1):56–75.

Yin J, Luo J, Brown SA (2021) Learning from crowdsourced multilabeling: A variational Bayesian approach. *Inform. Systems Res.* Forthcoming.

Zhang D, Mishra S, Brynjolfsson E, Etchemendy J, Ganguli D, Grosz B, Lyons T, et al. (2021) The AI index 2021 annual report. Preprint, submitted March 9, https://arxiv.org/abs/2103.06312.