

MANAGING ARTIFICIAL INTELLIGENCE¹

Nicholas Berente

University of Notre Dame {nberente@nd.edu}

Bin Gu

Boston University {bgu@bu.edu}

Jan Recker

University of Hamburg {jan.christof.recker@uni-hamburg.de}

Radhika Santhanam

University of Oklahoma {radhika@ou.edu}

Managing artificial intelligence (AI) marks the dawn of a new age of information technology management. Managing AI involves communicating, leading, coordinating, and controlling an ever-evolving frontier of computational advancements that references human intelligence in addressing ever more complex decision-making problems. It means making decisions about three related, interdependent facets of AI—autonomy, learning, and inscrutability—in the ongoing quest to push the frontiers of performance and scope of AI. We demonstrate how the frontiers of AI have shifted with time, and explain how the seven exemplar studies included in this special issue are helping us learn about management at the frontiers of AI. We close by speculating about future frontiers in managing AI and what role information systems scholarship has in exploring and shaping this future.

Keywords: Artificial intelligence, information technology management, autonomy, learning, inscrutability, computing, frontier, ethics

Introduction

Managing artificial intelligence (AI) is unlike information technology (IT) management in the past. AI is not a technology or set of technologies, but a continually evolving frontier of emerging computing capabilities (McCorduck 2004; Meyer 2011). The machine learning technologies that are at the core of contemporary AI have greater autonomy, deeper learning capacity, and are more inscrutable than any of the “intelligent” IT artifacts that have come before (Baird and Maruping 2021). Current AI technologies, which include robots and autonomous vehicles, facial recognition, natural language processing, and virtual agents of all sorts, are being

deployed in an astounding variety of problem domains. By some estimates, more than half of businesses were implementing some form of this new wave of technologies in 2020 (Balakrishnan et al. 2020), and the applications continue to grow at an astounding clip.

These developments are important because AI provides inestimable possibilities for enhancing people’s lives in a variety of areas, including their homes, healthcare, education, employment, entertainment, safety, and transportation (Rahwan et al. 2019; Stone et al. 2016). AI provides businesses with unprecedented opportunities for designing intelligent products, devising novel service offerings, and inventing new business models and organizational forms (Agrawal et al. 2019; Davenport et al. 2020; Townsend and Hunt 2019). But AI is not a technological panacea. Accompanying the horizon

¹This special issue was accepted while Arun Rai was Editor-in-Chief.

of possibilities are a host of emerging thorny and complex challenges around business strategies, human–AI interfaces, data, privacy, security, ethics, labor, human rights, and national security (Faraj et al. 2018; Kellogg et al. 2020; Rahwan et al. 2019; Russell, 2019; Stone et al. 2016). Today’s managers need to deal with both possibilities and challenges that accompany widespread AI.

The role of managers in the burgeoning societal transformation involving AI cannot be overstated. It is the managers that make all key decisions about AI. They oversee the development and implementation of AI-based systems, managers use them in their decision making, leverage them to target customers, and monitor and adjust the decisions, processes, and routines that appropriate AI. Managers allocate resources, oversee AI projects, and govern the organizations that are shaping the future.

For managers, AI introduces a variety of challenges. Many of the challenges are technical, such as those that include finding effective solutions for human interaction; overcoming issues of trust, safety, and security; and being careful to avoid negative consequences (Stone et al. 2016). However, many of the challenges also have a moral and ethical character, including those related to the workforce, labor, and consumers in terms of privacy, fairness, justice, discrimination, bias, deskilling, surveillance, and a host of other thorny issues (Fjeld et al. 2020; Floridi et al. 2018). Some are calling for increasing levels of accountability, that is, for organizations to be responsible for consequences of AI in many circumstances (Martin 2019b).

It falls squarely on the shoulders of managers to communicate, lead, coordinate, and control organizational efforts to navigate the many challenges and realize their goals, while at the same time, avoiding the negative consequences. It is critical for managers to be reflective and carefully shape their organization’s AI-related activity. Thus, it is important for information systems researchers to conduct research into the unique promise and accompanying challenges of AI, and to help managers in their decision making with well-developed, evidence-based practice.

This backdrop served as the motivation to launch this special issue of *MIS Quarterly* that focuses on the management of AI in a variety of forms. This special issue was not intended to start a new conversation, but rather to redirect the existing discourse. There is a great deal of research on AI across disciplines, and this research is growing dramatically (Perrault et al. 2019). However, much of the available research appears siloed. Technical fields, on one hand, focus on the technology and black-box the human and organizational side. Organizational, economic, and behavioral research, on the other hand,

often black-box the technological side of AI. As a key boundary-spanning, interdisciplinary tradition, the information systems field is well poised to contribute to the discourse, both on the technical side, by bringing in the behavioral, organizational, and economic perspective, and the human side, by bringing in the technological perspective. The interface of the two perfectly fits the sociotechnical roots of the field (Sarker et al. 2019; Winter et al. 2014).

The value of such sociotechnical thinking about managing AI is visible in each of the seven papers that are included in this special issue. Each paper conceptualizes AI and distinguishes it from traditional information technologies in different ways. Further, each paper investigates and addresses a different set of critical challenges with managing AI.

In this introduction, we consolidate the insights we gained about managing AI from handling this special issue. We synthesize the insights offered by each of the seven papers with our own views about AI and managing AI. We begin by offering a definition of AI as the frontier of computing and conceptualizing three different, interrelated facets of these frontiers: autonomy, learning, and inscrutability. We review literature on each of these facets and discuss how the research contained in the special issue both marks and departs from the current frontiers across these facets. We then make some suggestions for phenomena that may emerge as the future frontiers alongside these facets. We close by discussing broader implications about sociotechnical scholarship, disciplinary diversity, and scholarly conduct in the age of AI.

AI as the Dynamic Frontier of Computing

Many mark as the origin of serious thinking about AI the famous Dartmouth summer workshop in 1956. John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon defined the project of creating AI in terms of “making a machine behave in ways that would be called intelligent if a human were so behaving” (McCarthy et al. 1955, p. 11). In the subsequent decades, definitions for AI abounded and Russell and Norvig (2009) summed up all of these definitions in an effort to create rational agents. More recent definitions emphasize different aspects of AI, such as its ability to learn (Castelvecchi 2016), or its emulation capability: how it is designed with the intention of mimicking human capabilities and skills (Brynjolfsson and Mitchell 2017). Still, there is no singular, agreed-upon definition for AI.

This definitional ambiguity has been quite generative, leading to all sorts of productive inquiry and technological advance-

ment in a host of different areas over the decades (Stone et al. 2016). Concepts that carry an openness of meaning leave their usage uncertain and create more fertile ground for alternative interpretation and speculation (Kaplan 2017), which means more room to theorize flexibly about AI and how to manage AI.

In this tradition, we conceive of AI as a process, rather than a phenomenon in itself. We define AI as *the frontier of computational advancements that references human intelligence in addressing ever more complex decision-making problems*. In short, AI is whatever we are doing next in computing.

This definition emphasizes several points. First, it emphasizes that AI is not a single, discernable thing—it is not a discrete phenomenon or set of technologies. It is not a device, nor a program, nor an algorithm. It is more of an idea, a concept that represents a moving target of evolving phenomena rather than a phenomenon in itself. AI is a frontier. Viewing AI as a moving frontier is not necessarily a new idea. McCorduck (2004, p. 453) highlighted the paradox in the history of AI that when a problem is solved with AI, it is no longer considered AI:

Practical AI successes, computational programs that actually achieved intelligent behavior, were soon assimilated into whatever application domain they were found to be useful, and became silent partners alongside other problem-solving approaches, which left AI researchers to deal only with the “failures,” the tough nuts that couldn’t yet be cracked. Once in use, successful AI systems were simply considered valuable automatic helpers. MACSYMA, for example, created by MIT’s Joel Moses, building on work by James Slagle, had originally been an AI program designed to solve problems in symbolic algebra. It became an indispensable workhorse for scientists in many disciplines, but few people credited artificial intelligence for having borne and nurtured it. If you could see how it was done, people seemed to think, then it couldn’t be intelligence—a fancy that many people entertain to this day.

As this example illustrates, there is no stable referent for the term AI. Rather, AI is a moving frontier of next-generation advancements in computing. Today, most would not consider early computer programs to be AI, although at that time they were. A good example is Turing’s work on the machine that broke the Enigma code. Although his machine performed calculations that no single human can do, he did not consider this to be “intelligence,” instead preferring to point to a future of advanced digital computing and learning machines in pursuit of this ambiguous target of intelligence (Turing 1950).

Following World War II was the great boom in digital computing rooted in Turing’s work, along with von Neumann’s architecture, cybernetics, and other foundations of artificial intelligence and digital computing in general, all of which culminated in a widespread commercialization of AI in the early 1980s (McCorduck 2004). These technologies built on rule-based algorithms dubbed “expert systems” (e.g., Lamberti and Wallace 1990; Yoon et al. 1995). The goal then was to design IT that could represent domain knowledge and apply specialized reasoning techniques to solve complex problems (Gill 1995). Expert systems were widely considered a type of AI then, but most would not consider them AI today.

Thus, AI is always the frontier of computational advancements that address ever more complex decision-making problems. McCorduck (2004) suggests AI is made up of the “tough nuts,” the next steps in computing that computer scientists are ironing out. Some even consider AI to be simply “what AI researchers do” (Stone et al. 2016) at any point in time, knowing that AI is always necessarily at the edge of our current capabilities and a moving target. Once we are using it in practice, it is no longer AI. At the present time, we often reserve the term AI for machine learning algorithms based on neural networks and data driven prediction. But this view will no longer be satisfactory when the current class of machine learning algorithms gives way to the next generation.

Second, our definition highlights how decision making is core to understanding the role of AI in organizations (Metcalf et al. 2019; Shrestha et al. 2019). Simon (1960) described decision making in terms of choice, or selection, among alternatives. Computing supports decision-making through describing, structuring, and transforming various kinds of information (Denning 1989). How to go about making decisions with computing has been a central, and at times controversial, idea throughout the history of computing (Burton-Jones et al. 2017; Faulkner and Runde 2019; Kent 1978; Stamper 1971; Suchman 1995; Turing 1950). Yet today, we face a turning point. However controversial, AI is fundamentally about making decisions autonomously. AI involves somehow informing or automating some aspect of human decision making, through computer programs that exhibit intelligence of a sort, assuming intelligence involves successful goal-directed action (Russell 2019).

Past generations of AI involved rule-based decision making (Turban and Watkins 1986), whereas current approaches to AI involve predictive models that outperform humans (Agrawal et al. 2019), in some cases even outperforming human crowds (Fu et al. 2021). Although at present, any particular AI model may focus on relatively minor prediction decisions, even such minor decisions can have outsized effects and quickly compound with societal implications (Lindebaum et al. 2020; Vimalkumar et al. 2021).

Third, our definitional focus on decision making also explicitly invokes a relationship between AI and human behavior. Early definitions of AI used humans as a standard for machine intelligence (McCarthy et al. 1955; Turing 1950). Russell and Norvig (2009), however, distinguished between thinking and acting “rationally,” which they thought of as optimally, versus “humanly,” which was not always rational and prone to error. By their own framing of humanly in terms of a less-than-ideal rationality, and also their own goal of creating rational agents, they did not equate artificial intelligence with computer-based approximations of human intelligence. Nevertheless, these two notions of emulating or outperforming humans remain presently at the center of discussions around AI.

AI’s emulation capability (i.e., its ability to think humanly) presents both managerial opportunities and challenges. On the one hand, it has the potential to increase labor productivity whose growth rate has slowed recently, to the worry of many economics and government officials (Byrne et al. 2016). AI’s ability of emulating human decision makers, however, has limitations, especially with regard to innovation. For example, in their contribution to this special issue, Bowen Lou and Lynn Wu, in “AI on Drugs: Can Artificial Intelligence Accelerate Drug Development? Evidence from a Large-Scale Examination of Bio-Pharma Firms,” find that AI is less helpful in generating disruptive drug innovations, but works well to identify drug innovations of a medium-level of novelty.

On the other hand, AI’s emulation capability also raises concerns, such as security and ethics. Many of today’s cyber security measures are built on technologies that distinguish bot behavior from human behavior but AI’s emulation capability often allows hackers to avoid such detection. Through emulation, AI also codifies human biases and errors. These consequences are also the topic of papers in this special issue. Mike Teodorescu, Lily Morse, Yazeed Awwad, and Gerald Kane in their article “Failures of Fairness in Automation Require a Deeper Understanding of Human–ML Augmentation,” propose using human augmentation to address AI biases. Sarah Lebovitz, Natalia Levina, and Hila Lifshitz-Assaf, in “Is AI Ground Truth Really True? The Dangers of Training and Evaluating AI Tools Based on Experts’ Know-What,” point out that the performance of AI is limited by the way in which the underlying ground truth is constructed.

In terms of AI outperforming humans (i.e., its ability to think rationally), doing so implies achieving a certain type high-performance, rational action that humans are not capable of achieving (Russell and Norvig 2009). Rational action involves logically and mathematically engineering optimal decisions for specific objectives (Russell 1997). This sort of rationality is what Weber (1978) termed *zweckrational*, an

instrumental, means–end approach to problem solving, which stands in contrast to *wertrational*, values-based reasoning.² Instrumental rationality prioritizes the codifiable aspects of organizational activity, particularly where there are clear, quantifiably measurable goals.

Unboxing of the form of rationality in AI decision-making is important. Scholars have long understood the difficulties of reasoning in some objectively rational way about problems (Poole et al. 1998), particularly when dealing with incomplete, uncertain, or otherwise deficient information, or issues involving tacit or unqualified knowledge (Luger 2009). Recently, scholars have highlighted how AI is rooted in and reinforces an intensely instrumental rationality, at the expense of more value-oriented rationality (Lindebaum et al. 2020; Raisch and Krakowski 2021). This is consistent with early work on management information systems that pointed out how information systems embody an instrumental rationality that reflects managerial goals of efficiency and control (Argyris 1971; Kling 1980; Zuboff 1998). IT have long been associated with instrumental rationality. For example, expert systems, earlier rule-based algorithms that helped with decision support, were focused on optimization for well-specified efficiency and effectiveness outcomes (e.g., Turban and Watkins 1986; Yoon et al. 1995). But AI involves a sort of instrumental control that is different from traditional direction, evaluation, and discipline because AI can be more comprehensive and interactive than previous generations of IT (Kellogg et al. 2020). Some argue that certain implementations of AI rooted in pattern recognition can get at tacit elements of knowledge (Ilk et al. 2020). But as Lebovitz, Levina, and Lifshitz-Assaf show, the codified “know what” instrumental rationality of AI is not so readily separable from the “know how” critical to knowledge work. In the end, AI inevitably acts with reference to human intelligence—either emulating humans, outperforming them, or underperforming them—given the particular forms of rationality they embody.

Delineating the AI Frontier

Being at the frontier of computing marks AI as both the limit and the most advanced achievement in computing at any point

²Weber’s work was equivocal on the concept of different rationalities. Over the years, our understanding of Weber’s different forms of rationality through which humans pattern action has settled on four types of rationality: practical, theoretical, formal, and substantive (Kalberg 1980). Lindebaum et al. (2020) indicate that AI is a “supercarrier” of formal rationality, but we do not believe it is this simple. Traditional computing algorithms were, perhaps, largely formal, but the current wave of machine learning approaches more resemble practical. Rather than delve into this, we simply distinguish between instrument versus value-based rationality.

in time. By researching and applying AI, we constantly explore what is possible in the future by taking available and plausible actions in the here and now. Thus, AI is always nascent, liminal, and emerging. By departing from the computing practices of the past, AI constantly redefines the frontier of computing. As soon as a new form of AI computing emerges, it no longer marks the frontier; over time, it becomes simply computing. AI is thus a fluid, fleeting, and temporary phenomenon. Managing AI, therefore, means managing ever-changing technology in time and over time. Management of AI can never be finished, instead “managing AI” is a constant process of emergence and performativity.

In that sense, AI as a frontier always reifies its own boundaries. This is most evident in the evolution of the dimensions *performance* and *scope*. The performance dimension of the AI frontier is the ever-improving execution of tasks to which AI is applied. It is tied not only to algorithmic advancement, but also advancement in raw processing performance (National Research Council 2011). This computing performance is directly responsible for the recent ability of machine learning algorithms to tackle complex problems in a variety of domains, including medicine (Topol 2019), engineering (Correa-Baena et al. 2018), or gaming (Schaeffer and van den Herik 2002). Thus, the increased performance of AI technologies is accompanied by an increase in the scope of AI.

The scope dimension of the AI frontier refers to the ever-expanding range of contexts to which AI is applied. Tasks that involve AI are becoming increasingly ubiquitous and with this ubiquity comes complexity (Benbya et al. 2020). Much like the technological inventions that have fostered advances in the symbolic and computational logic underpinning AI, the context of the decision-making problems in which they are developed, applied, and used has advanced in a truly astonishing fashion since the advent of computing in the second half of the 20th century. No longer are AI confined to decision-making problems that reside within organizational containers (Winter et al. 2014); they permeate most if not all aspects of the human experience (Yoo 2010), already well beyond work alone. AI, in different forms and versions, are filling every inhabited corner of the earth, including the polar regions (British Antarctic Survey 2021). We use AI at work just as much as when we select movies to stream, control the temperature of our houses, or search for internet content (Benbya et al. 2020). Complexity is salient to AI because the digital technologies AI involves are communicable and editable (Lyytinen and King 2006; Yoo 2010) and the context of decisions that are solved through or with AI has shifted and expanded dramatically (Avgerou 2019; Benbya et al. 2020).

As the frontier of computational advancements that solve ever more complex decision-making problems, contemporary

forms of AI differ qualitatively from previous generations in three general, interrelated facets that impact managers when they need to deal with the present frontiers of AI: their autonomy, learning, and inscrutability (Baird and Maruping 2021; Glikson and Woolley 2020; Kellogg et al. 2020; Lyytinen et al. 2021; Rahwan et al. 2019). We summarize these in Table 1 and discuss them below.

Autonomy. Contemporary forms of AI have an increasing capacity to act on their own, without human intervention (Baird and Maruping 2021). AI make autonomous decisions and act in the world in a way that has material outcomes—often not only without human intervention, but also without human knowledge (Möhlmann et al. 2021; Murray et al. 2021). Examples have grown at an astonishing rate: software-controlled vehicles that drive autonomously (Frazzoli et al. 2002), robo-advisor software that automatically rebalances investments (Lee and Shin, 2018), and AI underwriters that have sovereignty to process loans (Markus 2017) are increasingly the norm, no longer the exception.

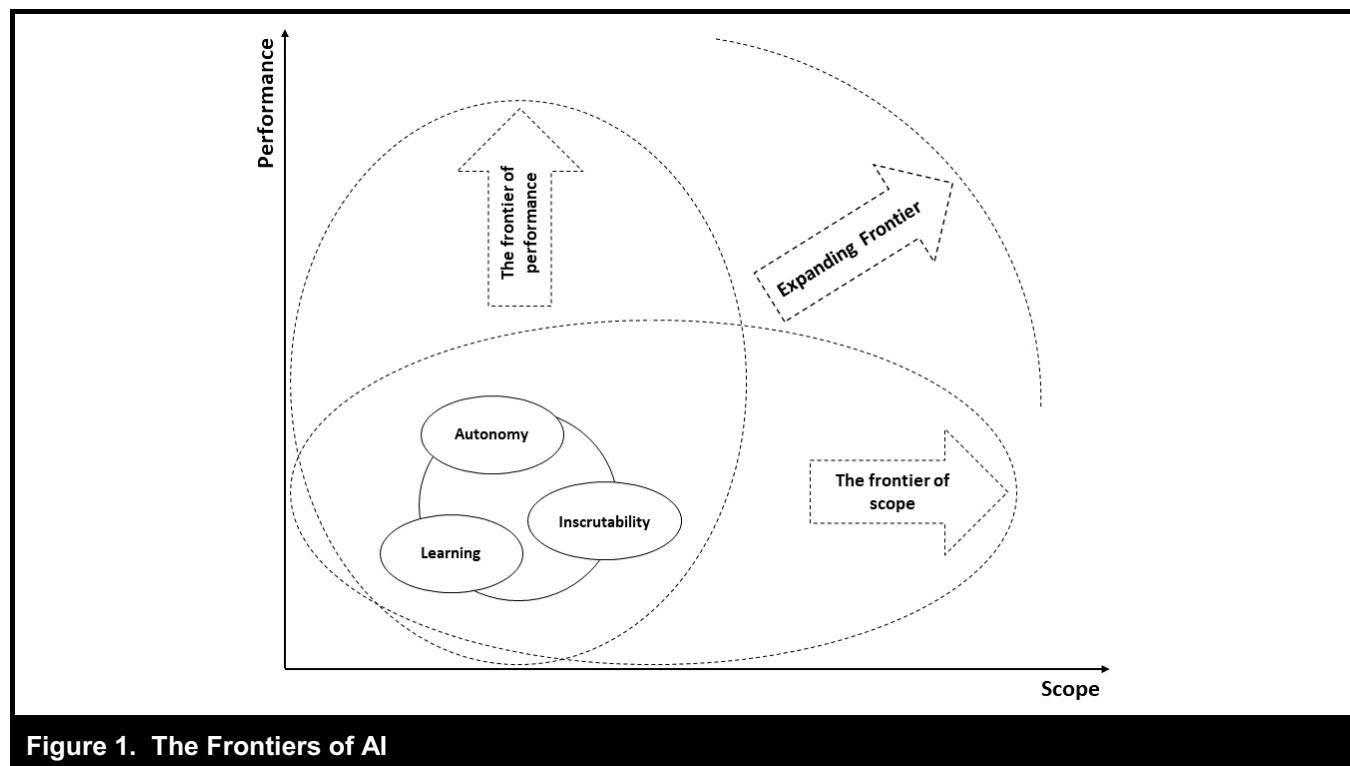
Learning. The ability to inductively improve automatically through data and experience has been a central concept in AI since the beginning (Solomonoff 1964; Turing 1950). While basic problems in supervised and unsupervised learning are now well-understood, other large-scale advances, such as deep or reinforcement learning (LeCun et al. 2015; Sutton and Barto 2018), have only recently been made possible through the availability of big data (Chen et al. 2012; Kitchin 2014). The newfound abilities for learning have enabled AI to make inroads into much more complex decision-making settings, including those that involve audio, speech, and object recognition, or natural language processing.

Inscrutability. With advances in autonomy and learning, contemporary AI also increasingly spawns the ability to generate algorithmic models and outputs that are intelligible only to a select audience while remaining opaque to others, or, in some cases, not intelligible to humans at all. Not only have the algorithms involved in yielding autonomy and learning increased in intricacy, the settings in which AI is being applied also exploded in variety and complexity. Together, these developments have fueled several challenges that are presently being discussed under terms such as the black-box problem (Castelvecchi 2016), explainable AI (Barredo Arrieta et al. 2020), AI accountability (Martin 2019b), or algorithm tractability (Gunning et al. 2019).

In sum, our perspective on AI as the frontier of computing involves considering three interdependent facets of AI, autonomy, learning, and inscrutability, which feed into each other. Learning contributes to, and results from, autonomy. Both autonomy and learning result in inscrutability. Any consequences of AI are shaped by how autonomy, learning, and

Table 1. Key Concepts of AI

Concept	Definition
<i>Artificial Intelligence</i>	The frontier of computational advancements that references human intelligence in addressing ever more complex decision-making problems
Dimensions of the AI Frontier	
<i>Performance frontier</i>	The ever-improving execution of tasks to which AI is applied
<i>Scope frontier</i>	The ever-expanding range of contexts to which AI is applied
Facets of AI	
<i>Autonomy</i>	Acting without human intervention
<i>Learning</i>	Improving through data and experience
<i>Inscrutability</i>	Being unintelligible to specific audiences



inscrutability are managed. The challenge is to manage AI as a moving frontier of both increasing performance and increasing scope, with ever-increasing levels of learning, autonomy, and inscrutability (Figure 1). Next we discuss some issues around managing each of these facets of AI.

Managing the AI Frontier

Management involves communicating, leading, coordinating and controlling the tasks of others in an organization, and decision making is a key activity of managers (Drucker 2008).

From one perspective, management is, at its core, essentially about decision making (Simon 1960). Managers make impactful strategic decisions in organizations, as well as small task-level decisions. These decisions can be fairly routine and predictable, or they can be more complex, value-laden choices among alternatives made under a great deal of uncertainty. Simon (1960) refers to these as programmable and non-programmable decisions, respectively. AI, as a frontier of decision making that is continually evolving in terms of performance and scope, is continually automating or informing the decisions of managers, first, those that are programmable, but increasingly, also those that appear non-programmable. But this situation does not mean that the management profes-

sion is on its way out, merely that the role of managers is changing. Just as AI requires that the role of product designers and engineers must be updated to work in conjunction with the new capabilities brought to bear by AI (Seidel et al. 2019), so, too, must managers adapt their roles with ever new AI technologies. They need to make decisions with the technology and about the technology. Managers need to be informed and understand the relevant facets of AI technologies. Next we reflect on managerial issues around each of the facets of AI that we propose: autonomy, learning, and inscrutability.

Managing Autonomy

Traditionally, managing IT has been driven by the desire to drive positive outcomes through automating and informing (Zuboff 1985). IT is used to automate work by computing codifiable elements of human processes and removing the need for human effort for largely repetitive tasks. At the same time, IT reduces human error and integrates and standardizes the work while driving greater control. IT informs work by helping humans with decisions through all sorts of approaches described in terms of decision support, business intelligence, and big data analytics.

The same pattern can be observed with AI. AI essentially helps automate and informate in new ways (Benbya et al. 2020). One could argue that managers delegate an ever-increasing set of decisions and related tasks to IT. As these technologies become more powerful, managers delegate more to them. What marks a difference in recent waves of AI technologies, however, is that AI technologies increasingly process information that was not necessarily directly delegated by humans. They are truly “autonomous agents” (Baird and Maruping 2021) that do not merely automate and informate what is given to them by humans. In the past, autonomy of AI was human-bracketed, where humans delegated, monitored, and controlled the technologies (Lyytinen et al. 2021). Past decision support systems processed data so that humans did not have to do it, but they did that processing in a way that humans dictated. Similarly, past automation systems automated tasks, often cognitive, formerly accomplished by humans, but they did it in the way that humans explicitly programmed them. But now, newer and newer waves of autonomous agents automate and informate tasks in ways that are not necessarily delegated or dictated by humans, but in new and surprising ways, often driven by their own abilities to control information and make decisions from their own volition, in a sense. Their autonomy is increasingly *generative* (Seidel and Berente 2020). Apple’s Siri, IBM’s Watson, and Google’s Waymo, for example, routinely make autonomous or semi-autonomous decisions, and they do so in

increasing numbers of domains. These and other new autonomous agents are generating and executing decisions, not simply implementing the plans of the humans (Seidel and Berente 2020; Zhang et al. 2021). This generativity must be considered a form of agency on the part of the technology that goes well beyond the generativity of information technologies in the past. In this special issue, Lou and Wu provide an exemplar. They show how AI aids drug development because of the technology’s capacity to facilitate recombination for certain types of drug candidates, a key generative skill (Holmström 2018).

Moreover, as AI technologies become more autonomous, the *interactions* between humans, AI, and with each other, take on a variety of different configurations. On the one hand, people delegate to autonomous agents in a variety of ways including as reflexive, supervisory, anticipatory, and prescriptive agents (Baird and Maruping 2021). On the other hand, autonomous agents also direct, evaluate, and control humans (Kellogg et al. 2020; Möhlmann et al. 2021). Murray et al. (2021) refer to this in terms of different types of “conjoined agency,” which involves a variety of different configurations depending on whether the human or the algorithm selects the actions or develops the protocols for the activity. Different organizational level approaches for dealing with different forms of conjoined agency exist that appreciate the unique contributions of AI and humans side-by-side (Lyytinen et al. 2021; Shrestha et al. 2019).

The key issue associated with thinking through the different interactions between autonomous agents and humans involve understanding their respective strengths. Historically, technologies were often better suited to problem solving, particularly for well-understood problems, whereas humans were necessary for tasks associated with problem formulation, intractability, and complexity (Licklider 1960; Simon 1996; Wiener 1960). However, this balance has started to shift (Russell 2019; von Krogh 2018). Autonomous agents are increasingly acting in ways that resemble those of knowledge workers (Faraj et al. 2018). The stark division between what humans and what machines should do is blurring. For example, autonomous AI tools are being used to generate all sorts of designs in ways that were formerly manually intensive design tasks (Seidel et al. 2020; Zhang et al. 2021). In these cases, AI technologies do not eliminate the need for designers, but instead change the tasks that designers must do to include cycles of monitoring, parameterizing, and adjustment, and of execution (Seidel et al. 2019). As AI tools do more of the creative part of knowledge work, humans do the integrative sensemaking (Verganti et al. 2020). These knowledge workers adapt their identities to their roles in relation to autonomous agents (Strich et al. 2020). Also, knowledge workers may decouple their activity from the interaction with

the algorithms to preserve their unique knowledge in the face of algorithmic expectations (Pachidi et al. 2020).

The exact form of the autonomy and agency can also have implications on how humans interact with each other in a variety of ways. Bots in online communities, for example, can shape the discourse among humans (de Lima Salge et al. 2021) or manage collaborative work practices (Hukal et al. 2019). Autonomous AI agents of all sorts exist that exhibit varying levels of human characteristics and influence all sorts of human reactions and interactions (Glikson and Woolley 2020; Qiu and Benbasat 2009; Traeger et al. 2020).

The interaction between humans and autonomous AI is perhaps the key managerial issue of our time. It has historically been put in terms of “augmentation” (e.g., Lindebaum et al. 2020), but key issues with augmentation remain. In this special issue, Teodorescu and his colleagues argue that the augmentation issue should be understood in terms of the fairness challenges that different augmentation strategies produce. But one could also argue that the computational advantages of AI, which allow them to incorporate more variables relating to a fairness problem than a human ever could consider, should also tilt the responsibility of making “fair” decisions more squarely into the hands of algorithms.

As this example demonstrates, augmentation generates several thorny managerial challenges. For example, firms often invest in AI for codifiable, repeatable tasks (Lacity and Willcocks 2016; van der Aalst et al. 2018). However, most management tasks are not easily codifiable, so automation is inherently limited, which creates the need for augmentation. Augmentation, on the other hand, can be automated so augmentation could lead to automation over time (Raisch and Krakowski 2021).

Another issue of reliance on augmentation could be negative dependency effects. As people increasingly rely on AI to augment their tasks and decisions, they become increasingly dependent on autonomous tools, particularly as tasks become more difficult (Bogert et al. 2021), which can have detrimental consequences. In this special issue, for example, Andreas Fügner, Jörn Grahl, Alok Gupta, and Wolfgang Ketter, in their article “Will Humans-in-the-Loop Become Borgs? Merits and Pitfalls of Working with AI,” demonstrate how unique human knowledge decreases when humans start interacting with AI in group decision environments, potentially undermining the effectiveness of the human–AI collaboration. This is consistent with the learned helplessness expected by some scholars as people grow dependent on AI (Lindebaum et al. 2020). As the frontier of machine autonomy expands, do humans lose their autonomy and their ability to effectively augment those machines? How should managers deal with this tension between automation and augmentation?

Managing Learning

The second facet of AI that pushes the frontier is learning. Learning has been a central issue since the inception of AI (Turing 1950). However, in the past, AI learning was limited by human data analysis, available data, corporate boundaries, and computing performance. Previous generations of AI technologies, such as those used in decision support and expert systems, primarily relied on proprietary corporate data structured, pre-processed, and inputted by human analysts (Gill 1995). AI learning has also always been hampered by the frame problem: managers always needed to understand how AI trained on some proprietary existing data may not generalize to the data they are trying to predict (Salovaara et al. 2019).

At least two key managerial issues are currently associated with AI learning. First, with the widespread availability of digital trace data, AI technologies no longer only learn from proprietary data sets, but instead, feed on all kinds of data, both within and beyond organizational containers (Winter et al. 2014). This brings up managerial issues, such as privacy (Rai 2020) and trust (Glikson and Woolley 2020), as well as legal and data guardianship issues concerning intellectual property rights ownership, security, and governance. As the availability of data burgeons beyond organizational bounds, the information used for AI learning also moves from being explicitly codified by humans to the general patterns associated with tacit knowledge. Some argue that pattern matching from machine learning can capture a particular form of tacit knowledge (Ilk et al. 2020). Whether or not this is the case, understanding the relationship between tacit knowledge and machine learning is certainly an important frontier for managers to consider (Hadjimichael and Tsoukas 2019). By having the ability to design AI that can process large volumes of data and generate its own codified knowledge and rules, we hit upon new limitations. When knowledge is generated by AI, unintended consequences regarding bias, algorithmic fairness, and value-agnostic decisions become more salient and troublesome (Fu et al. 2021). Solutions to these issues could involve approaches, such as oversight and augmentation, as argued by Teodorescu and his colleagues in this special issue. But a more likely scenario is that technical solutions alone are not enough, and neither is augmentation alone (Abbasi et al. 2018). Augmentation cannot be a static process. Those looking to manage AI need to develop and maintain mental models of the AI model, and continually learn about how AI act given different inputs and parameters (Seidel et al. 2019).

Second, technologically, the substantial progress in software and hardware technology for basic operations, such as sensing, recognition, or data storage and processing, have al-

lowed the learning capacity of AI to evolve from basic approaches for inductive learning to large-scale approaches, such as deep, reinforcement, or adversarial learning. As quantum technologies come online, this learning will be accelerated further in an even more dramatic fashion (MacQuarrie et al. 2020). What is common to these new approaches to learning made available by required technologies essentially becoming ubiquitous and cheap is that they involve human oversight less and less, essentially removing human mediation from many domains (Tegmark 2017).

Issues with the management of learning are also central themes in papers included in this special issue. For example, Elmira van den Broek, Anastasia Sergeeva, and Marleen Huysman, in their article “When the Machine Meets the Expert: An Ethnography of Developing AI for Hiring,” examine the tensions that surface when AI learns “truth” both from data and domain experts. Timo Sturm, Jin Gerlach, Luisa Pumplun, Neda Mesbah, Felix Peters, Christoph Tauchert, Ning Nan, and Peter Buxmann, in “Coordinating Human and Machine Learning for Effective Organizational Learning,” explore through a simulation how AI and humans can collectively contribute to organizational learning, and how learning effectiveness can be optimized under certain conditions. Both studies demonstrate how managing AI implies reflexivity of learning in terms of deliberation, correction, and adjustment of both AI and human elements.

Managing Inscrutability

As AI learns more and more and becomes increasingly autonomous, it also grows more inscrutable. By inscrutability, we refer to deficiencies in the intelligibility of AI procedures and outputs in relation to a specific party (Asatiani et al. 2021; Martin 2019b; Samek et al. 2019). The ability to scrutinize, explain, and understand algorithmic activity is critical in establishing trust and accountability in AI (Martin 2019a; Rai et al. 2019). But inscrutability is difficult to assess, because although AI may be intelligible to some, it may not be intelligible to all. Human understanding has different levels of granularity and abstraction and this varies with different purposes. A satisfactory understanding of algorithmic behavior at a general level of abstraction for one purpose may not be adequate for another purpose (Andrulis et al. 2020). This highlights the need to manage human understanding of algorithmic activity with an eye toward the level of that understanding, as well as its purpose.

The original frontiers of AI were determined through deterministic, explicit logic coded into technology. Intelligent expert systems, for example, were rule-based and their use could cut costs, increase financial returns, improve task

performance, and help a firm achieve goals (Sviokla 1990). But data used by AI moved outside the bounds of traditional organizations, and new learning algorithms are applied that rely less on pre-coded rules, human oversight, or supervision. Thus, learning has moved from being deterministic to probabilistic, and inscrutability of AI has taken on new, qualitatively different dimensions of meaning. Inscrutability now carries at least four interdependent emphases (see Vimal-kumar et al. 2021) that move from the algorithm to the human: opacity, transparency, explainability, and interpretability. Structuring them as a continuum highlights different aspects of the inscrutability of algorithmic decisions:

- Opacity refers to the lack of visibility into an algorithm. Opacity is a property of the algorithm and describes the amount of information that can possibly be scrutinized. For example, one opacity issue in machine learning is that the logic of some advanced algorithms is simply not accessible, even to the developers of those algorithms (Knight 2017).
- Transparency refers to the openness or willingness to disclose on the part of the owners of the algorithm, and the amount that those owners wish to disclose or occlude. Transparency implies disclosure and is, thus, a strategic management issue (Granados et al. 2010) predicated on the desire for secrecy (Vimalkumar et al. 2021).
- Explainability refers to an algorithm’s ability to be codified and understood at least by some party (Gregor and Benbasat 1999). Explainability is, thus, a property of an algorithm. Typically, explanation has a purpose and that purpose has a domain that has established language. Adequate explanation in the domain language indicates explainability (Kovalerchuk et al. 2021).
- Interpretability refers to the understandability and sense-making on the part of particular humans. Interpretability highlights the interpreter, not the algorithm. If a particular person can understand what the algorithm is doing, then it is adequately interpretable for that person, but perhaps not for another. Interpretability is dependent on the learning styles and literacy of the human (Vimal-kumar et al. 2021). For example, some people are visual learners, and benefit more from visual explanations, thus, interpretability for them would depend on visualizations (Kovalerchuk et al. 2021).

Regardless of the emphasis, inscrutability is a managerial issue because organizational responsibility is different for different types of AI algorithms and it is critical for managers to understand issues with liability, accountability, culpability, and fiduciary responsibility that their decisions imply (Martin

2019b). When AI is black-boxed, often ethical implications arise (Faraj et al. 2018; Martin 2019a) because making decisions in such a situation requires organizations to rely on approaches such as “envelopment” (Asatiani et al. 2021). Dealing with opacity, transparency, explainability, and interpretability in some form will thus be key to addressing many ethical issues associated with AI (Fjeld et al. 2020; Floridi et al. 2018; Perrault et al. 2019).

Issues with inscrutability also feature in the papers in this special issue. Lebovitz, Levina, and Lifshitz-Assaf, for example, demonstrate the difficulties managers faced when trying to evaluate the performance of machine-learning based AI tools vis-à-vis established performance metrics of human agents until they realized that fundamentally different knowledge (know-what versus know-how) was captured to train the AI tools. A second example is the study by Jingyu Li, Mengxiang Li, Xincheng Wang, and Jason Bennett Thatcher, “Strategic Directions for AI: The Role of CIOs and Boards of Directors,” which asserts the importance of having AI and R&D experience available in member of firms’ upper echelons when making strategic decisions about AI. These authors find that boards must contain at least some ability to scrutinize AI technologies to be able to assess their strategic impact.

To sum up, Table 2 clarifies original and contemporary frontiers that managing AI involves, and shows one way to organize the contributions from this special issue. It also draws attention to issues that cut across the different facets, such as ethics in AI (see Jobin et al. 2019), because they may invoke legal and regulatory issues (see Smuha 2021). We also list one example of a possible future frontier for each dimension that could potentially bear relevance to managing AI in the not-so-distant future. We discuss these next.

Future Frontiers in Managing AI

In this section, we outline four examples, one for each facet of AI and one that cuts across these facets, for research questions that potentially characterize the future frontiers of AI. We have no scientific evidence on their future relevance or importance; they merely characterize directions that we find potentially “interesting” (Davis 1971) when contemplating the future of managing AI.

A Future Frontier in Managing Autonomy: AI and Physicality

One future frontier involves how we deal with physical, tangible material that exists only in relation with AI. We

already put much emphasis on the mutual constitution of our reality through the assemblages of human agents and technologies (Leonardi 2013; Orlikowski and Scott 2008) and how material matters of our lives now originate in the digital realm (i.e., “digital-first,” Baskerville et al. 2020; Recker et al. 2021). Still, one tacit assumption in this line of thinking remains: that physical, tangible constructions remain human-mastered (Zhang et al. 2021) or at least human-imbricated (Leonardi 2011). But AI, in combination with advances in sensor and actuator technology, carries the potential to autonomously construct and enact physical and tangible matters in our lived experiences. We already witness glimpses into what is possible. AI has long begun to take an autonomous, active role in constructing material aspects of reality in digital design environments such as video games (Seidel et al. 2020). But AI-enacted materiality does not stop at the digital-only frontier; AI is increasingly also used for engineering physical matters, from semiconductors (Zhang et al. 2021), to cars (Noor 2017), and pizzas (Wiggers 2018), in increasingly autonomous ways. Manufacturing technologies such as three-dimensional printing and fabrication labs are becoming ever-more powerful and inexpensive, democratizing the manufacturing process (Gershenfeld et al. 2017). Increasingly autonomous physical robots have already started to impact organizational practices involving delicate physical actions, such as surgery, and humans have adapted to this physical automation (Barrett et al. 2012; Sergeeva et al. 2020). Robotics combined with material fabrication and other forms of autonomous physical agency, promise a panoply of future possibilities for how autonomous agents will shape the future (Tegmark 2017). These developments carry several managerial implications, for example, in designing the changing role of human workers as they no longer control but instead collaborate with AI technologies (Zhang et al. 2021). We also need to consider how organizational processes in engineering, construction, and logistics, can best be managed when the routines are not solely carried out by human actors alone.

A Future Frontier in Managing Learning: Adversarial Learning

As more and more autonomous AI technologies come into existence, we will also see a shift toward different forms of AI learning beyond machine and deep learning. For example, several kinds of new algorithms rooted in adversarial approaches to learning exist already today. Generative adversarial networks (Goodfellow et al. 2014), for example, involve two neural networks that make up a generator and a discriminator that function as adversaries. On an intuitive level, the generator constructs some conceptualization, such as a particular distribution, from streams of data, and the discrimin-

Table 2. Original, Contemporary, and Potential Future Frontiers in Managing AI

Facets of AI	Original Frontiers	Contemporary Frontiers	Papers in the Special Issue that Explore the Frontiers	Examples of a Future Frontier	Example of a Cross-Cutting Future Frontier
<i>Autonomy</i>	Human-bracketed AI affordances for automating and informing	Generative agency of AI Conjoined agency between humans and AI	Lou and Wu: The performance of AI for innovation is limited by creativity and novelty. Teodorescu, Mores, Awwad, and Kane: Augmentation strategies bring forth different fairness and bias issues.	AI and physicality	Ethical Issues
<i>Learning</i>	Structured proprietary datasets Human-driven data analysis	Large-scale trace data Human-unmediated analysis	van den Broek, Sergeeva, and Huysman: AI and humans need to engage in mutual learning. Fügener, Grahl, Gupta, and Ketter: In human–AI interaction settings, performance tradeoffs between individual and collective performance must be managed. Sturm, Gerlac, Pumplun, Mesbah, Peters, Tauchert, Nan, and Buxmann: AI learning requires careful human adjustment to be effective under certain conditions.	Adversarial learning	
<i>Inscrutability</i>	Explicit deterministic logic Manually generated reasoning	Opaque and probabilistic algorithmic logic Self-evolving, genetic deep learning algorithms	Li, Li, Wang, and Thatcher: Top management boards require sufficient diversity and experience to be able to scrutinize the strategic potential of AI. Lebovitz, Levina, and Lifshitz-Assaf: Evaluating AI hinges on an understanding of know-what versus know-how.	Social context and interpretability	

ator tests these constructions. Such approaches can identify patterns in large datasets and subject these patterns to evaluation. Currently, these and other kinds of adversarial learning are being used primarily to defend systems from various sorts of attacks. However, applications beyond security are beginning to emerge in domains as diverse as medical imaging (Özdenizci et al. 2020), recommender systems (Ren et al. 2020), and video game design (Seidel et al. 2020). As these examples show, advances to adversarial learning are accelerating the development and testing of AI applications. One can imagine a future, perhaps powered through quantum computing, where even more development work builds on adversarial learning. This would have several managerial implications. For example, constructing systems featuring AI building adversarial learning likely invokes a departure from established approaches to leading, coordinating, and con-

trolling the development of systems. As van den Broek, Sergeeva, and Huysman demonstrate in this special issue, developing systems that feature AI already involves making decisions about a range of complex issues revolving around learning from data and human experts. This situation will likely be even more different with systems involving adversarial learning.

A Future Frontier in Managing Inscrutability: Social Context and Interpretability

Inscrutability is a multidimensional and complex issue. Whether one can explain, interpret, or understand AI decisions depends on the algorithm and its opacity and explainability, as well as the transparency decisions, and

interpretations of humans. But the human side of explanations, the transparency and interpretability of AI decisions, is not only cognitive, but also social (Malle 2004), which has largely been overlooked in much of the literature. Explainable and interpretable AI has largely been treated as a computational problem, a cognitive problem, or both, but not a social problem (Miller 2019). Yet, AI does not come with “universal rules for interpretation” (Suchman 1987, p. 64) and interpretation is always situated within a social context. Interpreters will draw on the social context and generate understanding in a way that maintains coherence with that context. One might consider existing cognitive and computational approaches to inscrutability as consistent with what Boland and Tenkasi (1995) referred to in terms of a “conduit model.” In highly structured communities with well-established interpretive schemes, the conduit model of sense making is useful. However, when diverse knowledge communities need to communicate, establishing some shared understanding is difficult and requires significant reflection on language and assumptions, as well as a great deal of social interaction among communities (Boland and Tenkasi 1995). AI inscrutability is just such an issue that necessarily involves traversing diverse knowledge communities, and the social context must be taken into consideration. Understanding how the shared narratives in different social contexts shape the way AI is interpreted, and what stands for a legitimate, appropriate, and adequate interpretation, in that context is critically important for managers to understand.

A Future Frontier That Cuts Across All Facets of AI: Ethical Issues

Of course, many important managerial issues do not fall neatly into the different facets of AI we described. Autonomy, learning, and inscrutability are interrelated and depend on each other. The frontier of AI moves through the constant interaction and enforcement of each facet through the others. Therefore, much of the future of AI will involve managing issues that span these frontiers. The prime example of such a cross-cutting frontier that is emerging at present are ethical issues surrounding AI. Ethics has now for some time been an important part of management, but the intersection of AI technologies and ethics has taken front-stage only in recent years. This is precisely because of the developments in autonomy, learning, and inscrutability that we describe. A host of thorny ethical issues accompany recent waves of AI technologies.

One example of an ethical issue involves concerns over automation. From one perspective, AI technologies are simply another automation technology impacting the workforce (e.g., Davenport and Kirby 2015). Scholars have long

been apprehensive about the issues around automation and the workforce. Any new frontier of automation has brought forth fears of deskilling and labor substitution (e.g., Brynjolfsson and Mitchell 2017; Chernowicz 1958; Hirschheim 1986), but as Schumpeter (1934) and others have shown, with growth in the economy and aggregate demand, worst case scenarios typically do not play out. This is, in part, because human insatiability reigns supreme and new combinations always require more labor and greater skills in the workforce (Smith 2015). But perhaps AI is different. In the past, machines replaced physical work, but the new set of AI technologies increasingly replace “mental functions,” which makes them markedly different (Leontief 1983). Perhaps it is true that, at some point, intelligent machines will displace humans and be able to handle the related growth without adding to the workforce, particularly given the increased involvement around physical construction that we mentioned above. Such a scenario could have dramatic implications on the workforce, but it is also possible that, while some work is indeed vulnerable to AI automation, many jobs remain resistant (Frey and Osborne 2017). It is not clear yet how managers should address deskilling associated with AI and labor substitution, but it is clear that workforce issues remain a major ethical frontier of AI.

Another major ethical issue is that of data privacy. On the one hand, personal data privacy is important and managers need to navigate the issues around their customer and employee data, as well as institutional requirements. On the other hand, there is greater pressure for widespread surveillance, often for important societal concerns such as public health or violent crime. In the past, there was widespread skepticism of surveillance programs, even for benevolent purposes, such as public health (Bayer and Fairchild 2000). But in recent years the debate has become more complex, with a variety of different perspectives that draw on different ethical standards and address both technological and non-technological aspects of the issue (Angst 2009; Sweeney 2020). Little guidance is provided for managers through existing research beyond simply complying, but compliance is likely not enough. Management research needs to combine knowledge about technical elements with a grounding in ethical philosophy, law, and information systems (Martin 2016), to begin understanding the issue of privacy in the emerging world of AI.

There are a variety of other ethical issues associated with AI, including fairness, justice, discrimination, and bias (Fjeld et al. 2020; Floridi et al. 2018), as well as accompanying legal aspects of culpability, accountability, regulation, and responsibility (de Lima Salge and Berente 2017; Martin 2019b; Smuha 2021). The issue of how to regulate deep fakes alone promises to be an incredibly complex frontier (e.g., Johnson and Diakopoulos 2021).

Just as the industrial age ushered in an era of unprecedented benefits to society, management continues to deal with unintended side effects of industrialization in areas ranging from sustainability and workforce issues to human rights. It is now imperative to proactively address issues around AI and ethics. Given the current pace of investment in AI worldwide, managers will not have decades or centuries to catch up. Managers need to understand effective, ethical, and responsible approaches to the communication, leadership, coordination, and control of AI soon—before it spirals out of their grasp. Managers cannot wait until the future unfolds to understand this emerging and powerful phenomenon. Today's managers need to be actively engaged in shaping the trajectory of AI and its consequences. It is imperative that managers lead the development, application, and governance of artificial intelligence in ways that preserve and generate value.

Conclusion

The objective we pursued with our special issue has been to encourage a shift in the scholarly conversation around AI. The breathtaking array of possibilities and requirements for managing AI are presenting the information systems field with a potential golden age for demonstrating the value of socio-technical thinking. It is hard to imagine how to think about managing AI without taking into account both social and technical components, as well as the interaction between both (Beath et al. 2013; Lee 2001); but with this opportunity comes expectation. The information systems field, with our socio-technical tradition (Sarker et al. 2019; Winter et al. 2014), can proactively inform other fields, including both management and computer science, about the phenomena, problems, and solutions to managing AI that reside at the intersection of social and technological views. Thus, we encourage the pursuit of new ideas around this interesting AI phenomena, rooted in the cumulative knowledge of the human and technical elements in managing information systems (Lee 1999; Sarker et al. 2019). Because of the focus on addressing both the social and the technical, the information systems field is well positioned to become the reference discipline for managing AI.

Our experience in managing this special issue was also that traditional assumptions about disciplinary traditions have begun to fade into the background, in many ways becoming obsolete (Rai 2018). As the exemplars included in this special issue demonstrate, managing AI phenomena lend themselves to the entire variety of disciplinary approaches, from observational and interpretive research, to interventionist and experimental studies, as well as econometric analyses.

In reflecting upon the process of this special issue, much of what we did as editors involved sitting in front of our computers, coffee in hand, and debating the rigor and contribution potential of the papers crafted by the vast array of excellent IS scholars that submitted their work. In this way, our process of handling the special issue was no different than any other review process—but it could have easily been. Reviewing papers for methodological rigor and theoretical contribution may have long been the privilege and responsibility of a select few individuals, but AI technologies are already changing how the review process unfolds (Heaven 2018). Even our decision to drink coffee is already being influenced by AI-powered research on the nutritious effects, good and bad, of coffee consumption (Tingley 2021). And while both of these examples seem funny or perhaps even trivial at first sight, our own research community already is, and will continue to be, impacted by AI as the frontier of computing. It is important that we focus on both the opportunities and challenges that flow from our discipline's placement at the intersection between the ever-evolving frontier of computing and its relationships to business and society, and continuously reflect about how our own norms, processes, outputs, and “ground truth” may be challenged in terms of autonomy, learning, and inscrutability.

Acknowledgments

This special issue would not have been possible without the help of many. We thank Arun Rai, former editor-in-chief for *MIS Quarterly*, for his vision and guidance, and the opportunity to launch the special issue. We received help from many colleagues throughout several rounds of constructive development reviewing. We thank all our peers that contributed anonymous peer reviews for our special issue, and we thank, in no particular order, Michael Barrett, Panos Constantinides, Kevin Crowston, Hailiang Chen, Pei-Yu Chen, Brad Greenwood, Robert Gregory, J. J. Hsieh, Peng Huang, Ke-Wei Huang, Zhengrui Jiang, Karthik Kannan, Deepa Mani, Jeffrey Nickerson, Eivor Oborn, Ilan Oshri, Brian Pentland, Yuqing Ren, Lionel Robert, Rajiv Sabherwal, Nachiketa Sahoo, Stefan Seidel, Mari-Klara Stein, Jingjing Zhang, and Rong Zheng for serving as associate editors. We thank Hong Guo and John Lalor for helping to think through issues of inscrutability. Finally, we thank Andrew Burton-Jones and Arun Rai for reviewing earlier drafts of this editorial. All faults remain ours.

References

- Abbasi, A., Li, J., Clifford, G., and Taylor, H. 2018. “Make ‘Fairness by Design’ Part of Machine Learning, *Harvard Business Review*, August 1 (<https://hbr.org/2018/08/make-fairness-by-design-part-of-machine-learning>).

- Agrawal, A., Gans, J., and Goldfarb, A. (eds.). 2019. *The Economics of Artificial Intelligence: An Agenda*, Chicago: University of Chicago Press.
- Andrulis, J., Meyer, O., Schott, G., Weinbach, S., and Gruhn, V. 2020. "Domain-Level Explainability—A Challenge for Creating Trust in Superhuman AI Strategies," in *Proceedings of the 3rd International Workshop on Explainable AI Planning*, Nancy, France (<http://xaip.mybluemix.net/#/2020>).
- Angst, C. M. 2009. "Protect My Privacy or Support the Common-Good? Ethical Questions About Electronic Health Information Exchanges," *Journal of Business Ethics* (90:2), pp. 169-178.
- Argyris, C. 1971. "Management Information Systems: The Challenge to Rationality and Emotionality," *Management Science* (17:6), pp. B275-B292.
- Asatiani, A., Malo, P., Nagbøl, P. R., Penttinen, E., Rinta-Kahila, T., and Salovaara, A. 2021. "Sociotechnical Envelopment of Artificial Intelligence: An Approach to Organizational Deployment of Inscrutable Artificial Intelligence Systems," *Journal of the Association for Information Systems* (22:2), pp. 325-352.
- Avgerou, C. 2019. "Contextual Explanation: Alternative Approaches and Persistent Challenges," *MIS Quarterly* (43:3), pp. 977-1006.
- Baird, A., and Maruping, L. M. 2021. "The Next Generation of Research on IS Use: A Theoretical Framework of Delegation to and from Agentic IS Artifacts," *MIS Quarterly* (45:1), pp. 315-341.
- Balakrishnan, T., Chui, M., Hall, B., and Henke, N. 2020. *The State of AI in 2020*, McKinsey & Company (<https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/global-survey-the-state-of-ai-in-2020>).
- Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bannetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., and Herrera, F. 2020. "Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI," *Information Fusion* (58), 82-115 (<https://doi.org/10.1016/j.inffus.2019.12.012>).
- Barrett, M., Oborn, E., Orlikowski, W. J., and Yates, J. 2012. "Reconfiguring Boundary Relations: Robotic Innovations in Pharmacy Work," *Organization Science* (23:5), pp. 1448-1466.
- Baskerville, R., Myers, M. D., and Yoo, Y. 2020. "Digital First: The Ontological Reversal and New Challenges for IS Research," *MIS Quarterly* (44:2), pp. 509-523.
- Bayer, R., and Fairchild, A. L. 2000. "Surveillance and Privacy," *Science* (290:5498), pp. 1898-1899.
- Beath, C., Berente, N., Gallivan, M. J., and Lyytinen, K. 2013. "Expanding the Frontiers of Information Systems Research: Introduction to the Special Issue," *Journal of the Association for Information Systems* (14:4), Article 4.
- Benbya, H., Nan, N., Tanriverdi, H., and Yoo, Y. 2020. "Complexity and Information Systems Research in the Emerging Digital World," *MIS Quarterly* (44:1), pp. 1-17.
- Bogert, E., Schechter, A., and Watson, R. T. 2021. "Humans Rely More on Algorithms than Social Influence as a Task Becomes More Difficult," *Nature Scientific Reports* (11:8028) (<https://doi.org/10.1038/s41598-021-87480-9>).
- Boland, R. J., and Tenkasi, R. V. 1995. "Perspective Making and Perspective Taking in Communities of Knowing," *Organization Science* (6:4), pp. 350-372.
- British Antarctic Survey. 2021. *Artificial Intelligence Lab*, Natural Environment Research Council (<https://www.bas.ac.uk/project/ai/>).
- Brynjolfsson, E., and Mitchell, T. 2017. "What Can Machine Learning Do? Workforce Implications," *Science* (358:6370), pp. 1530-1534.
- Burton-Jones, A., Recker, J., Indulska, M., Green, P., and Weber, R. 2017. "Assessing Representation Theory with a Framework for Pursuing Success and Failure," *MIS Quarterly* (41:4), pp. 1307-1333.
- Byrne, D. M., Fernald, J. G., and Reinsdorf, M. B. 2016. "Does the United States Have a Productivity Slowdown or a Measurement Problem?," *Brookings Papers on Economic Activity* (Spring), pp. 109-157.
- Castelvecchi, D. 2016. "Can We Open the Black Box of AI?," *Nature* (538:7623), pp. 20-23.
- Chen, H., Chiang, R., and Storey, V. C. 2012. "Business Intelligence and Analytics: From Big Data to Big Impacts," *MIS Quarterly* (36:4), pp. 1165-1188.
- Chernowitz, G. 1958. "Review: Office Work and Automation by Howard S. Levin; Electronic Data Processing for Business and Industry," *Management Science* (4:4), pp. 475-477.
- Correa-Baena, J.-P., Hippalgaonkar, K., van Duren, J., Jaffer, S., Chandrasekhar, V. R., Stevanovic, V., Wadia, C., Guha, S., and Buonassisi, T. 2018. "Accelerating Materials Development via Automation, Machine Learning, and High-Performance Computing," *Joule* (2:8), pp. 1410-1420.
- Davenport, T. H., Gupta, A., Grewal, D., and Bressgott, T. 2020. "How Artificial Intelligence Will Change the Future of Marketing," *Journal of the Academy of Marketing Science* (48:1), pp. 24-42.
- Davenport, T. H., and Kirby, J. 2015. "Beyond Automation," *Harvard Business Review* (93:6), pp. 58-65.
- Davis, M. S. 1971. "That's Interesting: Towards a Phenomenology of Sociology and a Sociology of Phenomenology," *Philosophy of the Social Sciences* (1:4), pp. 309-344.
- Denning, P. J. 1989. "Computing as a Discipline," *Communications of the ACM* (32:1), pp. 9-23.
- de Lima Salge, C. A., and Berente, N. 2017. "Is That Social Bot Behaving Unethically?," *Communications of the ACM* (60:9), pp. 29-31.
- de Lima Salge, C. A., Karahanna, E., and Thatcher, J. B. 2021. "Algorithmic Processes of Social Alertness and Social Transmission: How Bots Disseminate Information on Twitter," *MIS Quarterly* (forthcoming).
- Drucker, P. F. 2008. *Management* (Rev. Ed.), New York: HarperCollins.

- Faraj, S., Pachidi, S., and Sayegh, K. 2018. "Working and Organizing in the Age of the Learning Algorithm," *Information and Organization* (28:1), pp. 62-70.
- Faulkner, P., and Runde, J. 2019. "Theorizing the Digital Object," *MIS Quarterly* (43:4), pp. 1279-1302.
- Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., and Srikumar, M. 2020. "Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI," Berkman Klein Center for Internet & Society, Harvard University.
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazeland, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., and Vayena, E. 2018. "AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations," *Minds and Machines* (28:4), pp. 689-707.
- Frazzoli, E., Dahleh, M. A., and Feron, E. 2002. "Real-Time Motion Planning for Agile Autonomous Vehicles," *Journal of Guidance, Control, and Dynamics* (25:1), pp. 116-129.
- Frey, C. B., and Osborne, M. A. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?," *Technological Forecasting and Social Change* (114), pp. 254-280.
- Fu, R., Huang, Y., and Singh, P. V. 2021. "Crowds, Lending, Machine, and Bias," *Information Systems Research* (32:1), pp. 72-92.
- Fügener, A., Grahl, J., Gupta, A., and Ketter, W. 2021. "Will Humans-in-the-Loop Become Borgs? Merits and Pitfalls of Working with AI," *MIS Quarterly* (45:3), pp. 1527-1556.
- Gershenfeld, N., Gershenfeld, A., and Cutcher-Gershenfeld, J. 2017. *Designing Reality: How to Survive and Thrive in the Third Digital Revolution*, New York: Basic Books.
- Gill, T. G. 1995. "Early Expert Systems: Where Are They Now?," *MIS Quarterly* (19:1), pp. 51-81.
- Glikson, E., and Woolley, A. W. 2020. "Human Trust in Artificial Intelligence: Review of Empirical Research," *Academy of Management Annals* (14:2), pp. 627-660.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. 2014. "Generative Adversarial Nets," arXiv (1406.2661).
- Granados, N., Gupta, A., and Kauffman, R. J. 2010. "Information Transparency in Business-to-Consumer Markets: Concepts, Framework, and Research Agenda," *Information Systems Research* (21:2), pp. 207-226.
- Gregor, S., and Benbasat, I. 1999. "Explanations from Intelligent Systems: Theoretical Foundations and Implications for Practice," *MIS Quarterly* (23:4), pp. 497-530.
- Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., and Yang, G.-Z. 2019. "XAI—Explainable Artificial Intelligence," *Science Robotics* (4:37).
- Hadjimichael, D., and Tsoukas, H. 2019. "Toward a Better Understanding of Tacit Knowledge in Organizations: Taking Stock and Moving Forward," *Academy of Management Annals* (13:2), pp. 672-703.
- Heaven, D. 2018. "The Age of AI Peer Reviews," *Nature* (563:7733), pp. 609-610.
- Hirschheim, R. 1986. *Office Automation: A Social and Organizational Perspective*, Hoboken, NJ: John Wiley & Sons, Inc.
- Holmström, J. 2018. "Recombination in Digital Innovation: Challenges, Opportunities, and the Importance of a Theoretical Framework," *Information and Organization* (28:2), pp. 107-110.
- Hukal, P., Berente, N., Germonprez, M., and Schecter, A. 2019. "Bots Coordinating Work in Open Source Software Projects," *IEEE Computer* (52:9), pp. 52-60.
- Ilk, N., Shang, G., and Goes, P. B. 2020. "Customer Routing in Contact Centers: An Automated Triage Design Based on Text Analytics," *Journal of Operations Management* (66:5), pp. 553-577.
- Jobin, A., Ienca, M., and Vayena, E. 2019. "The Global Landscape of AI Ethics Guidelines," *Nature Machine Intelligence* (1), pp. 389-399.
- Johnson, D. G., and Diakopoulos, N. 2021. "What to Do About Deepfakes," *Communications of the ACM* (64:3), pp. 33-35.
- Kalberg, S. 1980. "Max Weber's Types of Rationality: Cornerstones for the Analysis of Rationalization Processes in History," *American Journal of Sociology* (85:5), pp. 1145-1179.
- Kaplan, A. 2017. *The Conduct of Inquiry: Methodology for Behavioral Science*, New York: Routledge (originally published in 1964 by Chandler Publishing Company).
- Kellogg, K. C., Valentine, M. A., and Christin, A. 2020. "Algorithms at Work: The New Contested Terrain of Control," *Academy of Management Annals* (14:1), pp. 366-410.
- Kent, W. 1978. *Data and Reality*, Amsterdam: North-Holland.
- Kitchin, R. 2014. *The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences*, Thousand Oaks, CA: SAGE Publications.
- Kling, R. 1980. "Social Analyses of Computing: Theoretical Perspectives in Recent Empirical Research," *ACM Computing Surveys* (12:1), pp. 61-110.
- Knight, W. 2017. "The Dark Secret at the Heart of AI," *MIT Technology Review* (<https://www.technologyreview.com/2017/04/11/51113/the-dark-secret-at-the-heart-of-ai>).
- Kovalerchuk, B., Ahmad, M. A., and Teredesai, A. 2021. "Survey of Explainable Machine Learning with Visual and Granular Methods Beyond Quasi-Explanations," in *Interpretable Artificial Intelligence: A Perspective of Granular Computing* (Vol. 937), W. Pedrycz and S.-M. Chen (eds.), Cham, Switzerland: Springer, pp. 217-267.
- Lacity, M. C., and Willcocks, L. P. 2016. "Robotic Process Automation at Telefonica O2," *MIS Quarterly Executive* (15:1), pp. 21-35.
- Lamberti, D. M., and Wallace, W. A. 1990. "Intelligent Interface Design: An Empirical Assessment of Knowledge Presentation in Expert Systems," *MIS Quarterly* (14:3), pp. 279-311.
- Lebovitz, S., Levina, N., and Lifshitz-Assaf, H. 2021. "Is AI Ground Truth Really True? The Dangers of Training and

- Evaluating AI Tools Based on Experts' Know-What," *MIS Quarterly* (45:3), pp. 1501-1525.
- LeCun, Y., Bengio, Y., and Hinton, G. 2015. "Deep Learning," *Nature* (521:7553), pp. 436-444.
- Lee, A. S. 1999. "Inaugural Editor's Comments," *MIS Quarterly* (23:1), pp. v-xi.
- Lee, A. S. 2001. "Editor's Comments: Research in Information Systems: What We Haven't Learned," *MIS Quarterly* (25:4), pp.
- Lee, I., and Shin, Y. J. 2018. "Fintech: Ecosystem, Business Models, Investment Decisions, and Challenges," *Business Horizons* (61:1), pp. 35-46.
- Leonardi, P. M. 2011. "When Flexible Routines Meet Flexible Technologies: Affordance, Constraint, and the Imbrication of Human and Material Agencies," *MIS Quarterly* (35:1), pp. 147-167.
- Leontief, W. "National Perspective: The Definition of Problems and Opportunities," in *The Long-Term Impact of Technology on Employment and Unemployment*, National Research Council, Washington, DC: The National Academy Press (<https://doi.org/10.17226/19470>).
- Leonardi, P. M. 2013. "Theoretical Foundations for the Study of Sociomateriality," *Information and Organization* (23:2), pp. 59-76.
- Li, J., Li, M., Wang, X., and Thatcher, J. B. 2021. "Strategic Directions for AI: The Role of CIOs and Boards of Directors," *MIS Quarterly* (45:3), pp. 1603-1643.
- Licklider, J. C. R. 1960. "Man-Computer Symbiosis," *IRE Transactions on Human Factors in Electronics* (1:1), pp. 4-11.
- Lindebaum, D., Vesa, M., and den Hond, F. 2020. "Insights from 'The Machine Stops' to Better Understand Rational Assumptions in Algorithmic Decision Making and its Implications for Organizations," *Academy of Management Review* (45:1), pp. 247-463.
- Lou, B., and Wu, L. 2021. "AI on Drugs: Can Artificial Intelligence Accelerate Drug Development? Evidence from a Large-Scale Examination of Bio-Pharma Firms," *MIS Quarterly* (45:3), pp. 1451-1482.
- Luger, G. F. 2009. *Artificial Intelligence: Structures and Strategies for Complex Problem Solving* (6th ed.), Boston: Pearson Education, Inc.
- Lyytinen, K., and King, J. L. 2006. "Standard Making: A Critical Research Frontier for Information Systems Research," *MIS Quarterly* (30:Special Issue), pp. 405-411.
- Lyytinen, K., Nickerson, J. V., and King, J. L. 2021. "Metahuman Systems = Humans + Machines That Learn," *Journal of Information Technology* (forthcoming) (<https://doi.org/10.1177/0268396220915917>).
- MacQuarrie, E. R., Simon, C., Simmons, S., and Maine, E. 2020. "The Emerging Commercial Landscape of Quantum Computing," *Nature Reviews Physics* (11:2), pp. 596-598.
- Malle, B. F. 2004. *How the Mind Explains Behavior: Folk Explanations, Meaning, and Social Interaction*, Cambridge, MA: MIT Press.
- Markus, M. L. 2017. "Datification, Organizational Strategy, and IS Research: What's the Score?," *Journal of Strategic Information Systems* (26:3), pp. 233-241.
- Martin, K. 2016. "Understanding Privacy Online: Development of a Social Contract Approach to Privacy," *Journal of Business Ethics* (137:3), pp. 551-569.
- Martin, K. 2019a. "Designing Ethical Algorithms," *MIS Quarterly Executive* (18:2), pp. 129-142.
- Martin, K. 2019b. "Ethical Implications and Accountability of Algorithms," *Journal of Business Ethics* (160:4), pp. 835-850.
- McCarthy, J. L., Minsky, M. L., Rochester, N., and Shannon, C. E. 1955. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," August 31 (<http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf>).
- McCorduck, P. 2004. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence* (2nd ed.), Natick, MA: A. K. Peters, Ltd.
- Metcalfe, L., Askay, D. A., and Rosenberg, L. B. 2019. "Keeping Humans in the Loop: Pooling Knowledge through Artificial Swarm Intelligence to Improve Business Decision Making," *California Management Review* (61:4), pp. 84-109.
- Meyer, B. 2011. *John McCarthy*, BLOG@CACM, October 28 (<https://cacm.acm.org/blogs/blog-cacm/138907-john-mccarthy/fulltext>).
- Miller, T. 2019. "Explanation in Artificial Intelligence: Insights from the Social Sciences," *Artificial Intelligence* (267), pp. 1-38.
- Möhlmann, M., Zalmanson, L., Henfridsson, O., and Gregory, R. W. 2021. "Algorithmic Management of Work on Online Labor Platforms: When Matching Meets Control," *MIS Quarterly* (forthcoming).
- Murray, A., Rhymer, J., and Sirmon, D. G. 2021. "Humans and Technology: Forms of Conjoined Agency in Organizations," *Academy of Management Review* (46:3), pp. 552-571.
- National Research Council. 2011. *The Future of Computing Performance: Game Over or Next Level?*, Washington, DC: The National Academies Press (<https://doi.org/10.17226/12980>).
- Noor, A. K. 2017. "AI and the Future of the Machine Design," *Mechanical Engineering Magazine* (139:1), pp. 38-43.
- Orlikowski, W. J., and Scott, S. V. 2008. "Sociomateriality: Challenging the Separation of Technology, Work and Organization," *Academy of Management Annals* (2:1), pp. 433-474.
- Özdenizci, O., Wang, Y., Koike-Akino, T., and Erdoğan, D. 2020. "Learning Invariant Representations from EEG Via Adversarial Inference," *IEEE Access* (8), pp. 27074-27085.
- Pachidi, S., Berends, H., Faraj, S., and Huysman, M. 2020. "Make Way for the Algorithms: Symbolic Actions and Change in a Regime of Knowing," *Organization Science* (32:1), pp. 18-41.
- Perrault, R., Shoham, Y., Brynjolfsson, E., Clark, J., Etchemendy, J., Grosz, B., Lyons, T., Manyika, J., Niebles, J. C., and Mishra, S. 2019. *Artificial Intelligence Index 2019 Annual Report*, AI Index Steering Committee, Human-Centered AI Institute, Stanford University, Stanford, CA.

- Poole, D., Mackworth, A., and Goebel, R. 1998. *Computational Intelligence: A Logical Approach*, Oxford, UK: Oxford University Press.
- Qiu, L., and Benbasat, I. 2009. "Evaluating Anthropomorphic Product Recommendation Agents: A Social Relationship Perspective to Designing Information Systems," *Journal of Management Information Systems* (25:4), pp. 145-182.
- Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., Crandall, J. W., Christakis, N. A., Couzin, I. D., Jackson, M. O., Jennings, N. R., Kamar, E., Kloumann, I. M., Larochelle, H., Lazer, D., McElreath, R., Mislove, A., Parkes, D. C., Pentland, A., Roberts, M. E., Shariff, A., Tenenbaum, J. B., and Wellman, M. P. 2019. "Machine Behaviour," *Nature* (568), pp. 477-486.
- Rai, A. 2018. "Editor's Comments: Beyond Outdated Labels: The Blending of IS Research Traditions," *MIS Quarterly* (42:1), pp. iii-vi.
- Rai, A. 2020. "Explainable AI: From Black Box to Glass Box," *Journal of the Academy of Marketing Science* (48:1), pp. 137-141.
- Rai, A., Constantinides, P., and Sarker, S. 2019. "Editor's Comments: Next-Generation Digital Platforms: Toward Human-AI Hybrids," *MIS Quarterly* (43:1), pp. iii-ix.
- Raisch, S., and Krakowski, S. 2021. "Artificial Intelligence and Management: The Automation-Augmentation Paradox," *Academy of Management Review* (46:1), pp. 192-210.
- Recker, J., Lukyanenko, R., Jabbari, M., Samuel, B. M., and Castellanos, A. 2021. "From Representation to Mediation: A New Agenda for Conceptual Modeling Research in a Digital World," *MIS Quarterly* (45:1), pp. 269-300.
- Ren, X., Yin, H., Chen, T., Wang, H., Hung, N. Q. V., Huang, Z., and Zhang, X. 2020. "CRSAL: Conversational Recommender Systems with Adversarial Learning," *ACM Transactions on Information Systems* (38:4).
- Russell, S. J. 1997. "Rationality and Intelligence," *Artificial Intelligence* (94:1-2), pp. 57-77.
- Russell, S. J. 2019. *Human Compatible: Artificial Intelligence and the Problem of Control*, New York: Viking Press.
- Russell, S. J., and Norvig, P. 2009. *Artificial Intelligence: A Modern Approach* (3rd ed.), Upper Saddle River, NJ: Prentice-Hall.
- Salovaara, A., Lyytinen, K., and Penttinen, E. 2019. "High Reliability in Digital Organizing: Mindlessness, the Frame Problem, and Digital Operations," *MIS Quarterly* (43:2), pp. 555-578.
- Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K., and Müller, K.-R. (eds.). 2019. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, Cham, Switzerland: Springer Nature.
- Sarker, S., Chatterjee, S., Xiao, X., and Elbanna, A. R. 2019. "The Sociotechnical Axis of Cohesion for the IS Discipline: Its Historical Legacy and its Continued Relevance," *MIS Quarterly* (43:3), pp. 695-719.
- Schaeffer, J., and van den Herik, H. J. 2002. "Games, Computers, and Artificial Intelligence," *Artificial Intelligence* (132:1), pp. 1-7.
- Schumpeter, J. A. 1934. *The Theory of Economic Development*, Boston: Harvard University Press.
- Seidel, S., and Berente, N. 2020. "Automate, Informate, and Generate: Affordance Primitives of Smart Devices and the Internet of Things," in *Handbook of Digital Innovation*, S. Nambisan, K. Lyytinen, and Y. Yoo (eds.), Cheltenham, UK: Edward Elgar Publishing, pp. 198-210.
- Seidel, S., Berente, N., Lindberg, A., Lyytinen, K., Martinez, B., and Nickerson, J. V. 2020. "Artificial Intelligence and Video Game Creation: A Framework for the New Logic of Autonomous Design," *Journal of Digital Social Research* (2:3), pp. 126-157.
- Seidel, S., Berente, N., Lindberg, A., Lyytinen, K., and Nickerson, J. V. 2019. "Autonomous Tools and Design: A Triple-Loop Approach to Human-Machine Learning," *Communications of the ACM* (62:1), pp. 50-57.
- Sergeeva, A. V., Faraj, S., and Huysman, M. 2020. "Losing Touch: An Embodiment Perspective on Coordination in Robotic Surgery," *Organization Science* (31:5), pp. 1248-1271.
- Shrestha, Y. R., Ben-Menahem, S. M., and von Krogh, G. 2019. "Organizational Decision-Making Structures in the Age of Artificial Intelligence," *California Management Review* (61:4), pp. 66-83.
- Simon, H. A. 1960. *The New Science of Management Decision*, New York: Harper and Brothers.
- Simon, H. A. 1996. *The Sciences of the Artificial* (3rd ed.), Cambridge, MA: MIT Press.
- Smith, A. 2015. *The Wealth of Nations: A Translation into Modern English*, Manchester, UK: Industrial Systems Research.
- Smuha, N. A. 2021. "From a 'Race to AI' to a 'Race to AI Regulation': Regulatory Competition for Artificial Intelligence," *Law, Innovation and Technology* (13:1), pp. 57-84.
- Solomonoff, R. J. 1964. "A Formal Theory of Inductive Inference, Part I," *Information and Control* (7:1), pp. 1-22.
- Stamper, R. K. 1971. "Some Ways of Measuring Information," *Computer Bulletin* (15:12), pp. 432-436.
- Stone, P., Brooks, R., Brynjolfsson, E., Calo, R., Etzioni, O., Hager, G., Hirschberg, J., Kalyanakrishnan, S., Kamar, E., Kraus, S., Leyton-Brown, K., Parkes, D., Press, W., Saxenian, A. L., Shah, J., Tambe, M., and Teller, A. 2016. "Artificial Intelligence and Life in 2030." One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, Stanford, CA (<http://ai100.stanford.edu/2016-report>)
- Strich, F., Mayer, A.-S., and Fiedler, M. 2020. "What Do I Do in a World of Artificial Intelligence? Investigating the Impact of Substitutive Decision-Making AI Systems on Employees' Professional Role Identity," *Journal of the Association for Information Systems* (22:2), pp. 304-324.
- Sturm, T., Gerlach, J. P., Pumplun, L., Mesbah, N., Peters, F., Tauchert, C., Nan, N., and Buxmann, P. 2021. "Coor-

- minating Human and Machine Learning for Effective Organizational Learning," *MIS Quarterly* (45:3), pp. 1581-1602.
- Suchman, L. A. 1987. *Plans and Situated Actions: The Problem of Human Machine Communication*, Cambridge, UK: Cambridge University Press.
- Suchman, L. A. 1995. "Making Work Visible," *Communications of the ACM* (38:9), pp. 56-64.
- Sutton, R. S., and Barto, A. G. 2018. *Reinforcement Learning: An Introduction* (2nd ed.), Cambridge, MA: MIT Press.
- Sviokla, J. J. 1990. "An Examination of the Impact of Expert Systems on the Firm: The Case of XCON," *MIS Quarterly* (14:2), pp. 127-140.
- Sweeney, Y. 2020. "Tracking the Debate on COVID-19 Surveillance Tools," *Nature Machine Intelligence* (2:6), pp. 301-304.
- Tegmark, M. 2017. *Life 3.0: Being Human in the Age of Artificial Intelligence*, New York: Alfred A. Knopf.
- Teodorescu, M., Morse, L., Awwad, Y., and Kane, G. C. 2021. "Failures of Fairness in Automation Require a Deeper Understanding of Human-ML Augmentation," *MIS Quarterly* (45:3), pp. 1483-1499.
- Tingley, K. 2021. "Is Coffee Good for Us? Maybe Machine Learning Can Help Figure it Out," *The New York Times Magazine*, March 24 (<https://www.nytimes.com/2021/03/24/magazine/coffee-heart-machine-learning.html>).
- Topol, E. J. 2019. "High-Performance Medicine: The Convergence of Human and Artificial Intelligence," *Nature Medicine* (25:1), pp. 44-56.
- Townsend, A. M., and Hunt, R. A. 2019. "Entrepreneurial Action, Creativity, and Judgment in the Age of Artificial Intelligence," *Journal of Business Venturing Insights* (11), e00126 (<https://doi.org/10.1016/j.jbvi.2019.e00126>).
- Traeger, M. L., Strohkorb Sebo, S., Jung, M., Scassellati, B., and Christakis, N. A. 2020. "Vulnerable Robots Positively Shape Human Conversational Dynamics in a Human-Robot Team," *Proceedings of the National Academy of Sciences* (117:12), pp. 6370-6375.
- Turban, E., and Watkins, P. R. 1986. "Integrating Expert Systems and Decision Support Systems," *MIS Quarterly* (10:2), pp. 121-136.
- Turing, A. M. 1950. "Computing Machinery and Intelligence," *Mind* (59:236), pp. 433-460.
- van den Broek, E., Sergeeva, A., and Huysman, M. 2021. "When the Machine Meets the Expert: An Ethnography of Developing AI for Hiring," *MIS Quarterly* (45:3), pp. 1557-1580.
- van der Aalst, W. M. P., Bichler, M., and Heinzl, A. 2018. "Robotic Process Automation," *Business and Information Systems Engineering* (60:4), pp. 269-272.
- Verganti, R., Vendraminelli, L., and Iansiti, M. 2020. "Innovation and Design in the Age of Artificial Intelligence," *Journal of Product Innovation Management* (37:3), pp. 212-227.
- Vimalkumar, M., Gupta, A., Sharma, D., and Dwivedi, Y. K. 2021. "Understanding the Effect that Task Complexity has on Automation Potential and Opacity: Implications for Algorithmic Fairness," *AIS Transactions on Human-Computer Interaction* (13:1), pp. 104-129.
- von Krogh, G. 2018. "Artificial Intelligence in Organizations: New Opportunities for Phenomenon-Based Theorizing," *Academy of Management Discoveries* (4:4), pp. 404-409.
- Weber, M. 1978. *Economy and Society: An Outline of Interpretive Sociology*, G. Roth and C. Wittich (eds.), Berkeley, CA: University of California Press.
- Wiener, N. 1960. "Some Moral and Technical Consequences of Automation," *Science* (131:3410), pp. 1355-1358.
- Wiggers, K. 2018. "MIT Students Use AI to Cook up Pizza Recipes," *VentureBeat*, September 10 (<https://venturebeat.com/2018/09/10/mit-students-use-ai-to-cook-up-pizza-recipes/>).
- Winter, S., Berente, N., Howison, J., and Butler, B. S. 2014. "Beyond the Organizational 'Container': Conceptualizing 21st Century Sociotechnical Work," *Information and Organization* (24:4), pp. 250-269.
- Yoo, Y. 2010. "Computing in Everyday Life: A Call for Research on Experiential Computing," *MIS Quarterly* (34:2), pp. 213-231.
- Yoon, Y., Guimaraes, T., and O'Neal, Q. 1995. "Exploring the Factors Associated with Expert Systems Success," *MIS Quarterly* (19:1), pp. 83-106.
- Zhang, Z., Lindberg, A., Lyytinen, K., and Yoo, Y. 2021. "The Unknowability of Autonomous Tools and the Liminal Experience of Their Use," *Information Systems Research* (forthcoming).
- Zuboff, S. 1985. "Automate/Informate: The Two Faces of Intelligent Technology," *Organizational Dynamics* (14:2), pp. 5-18.
- Zuboff, S. 1998. *In The Age Of The Smart Machine: The Future of Work and Power*, New York: Basic Books.