

**ARTIFICIAL INTELLIGENCE AND MANAGEMENT:
THE AUTOMATION-AUGMENTATION PARADOX**

Review Essay

Sebastian Raisch
University of Geneva
sebastian.raisch@unige.ch

Sebastian Krakowski
University of Geneva
and
Stockholm School of Economics
sebastian.krakowski@hhs.se

Pre-print version

Article accepted for publication in the *Academy of Management Review*

Acknowledgements

The authors wish to thank associate editor Subramanian Rangan and two anonymous reviewers for their guidance on developing the paper. Furthermore, they are grateful to Tina Ambos, Jean Bartunek, Jonathan Schad, Katherine Tatarinov, Xena Welch Guerra, Udo Zander, and the participants of research seminars at HEC Lausanne, Nova School of Business and Economics, and the Geneva School of Economics and Management for their valuable comments on earlier versions of the paper. This research received partial support from the Swiss National Science Foundation [SNSF grants 181364 and 185164] and the European Union [EU Grant 856688].

ABSTRACT

Taking three recent business books on artificial intelligence (AI) as a starting point, we explore the automation and augmentation concepts in the management domain. Whereas automation implies that machines take over a human task, augmentation means that humans collaborate closely with machines to perform a task. Taking a normative stance, the three books advise organizations to prioritize augmentation, which they relate to superior performance. Using a more comprehensive paradox theory perspective, we argue that, in the management domain, augmentation cannot be neatly separated from automation. These dual AI applications are interdependent across time and space, creating a paradoxical tension. Over-emphasizing either augmentation or automation fuels reinforcing cycles with negative organizational and societal outcomes. However, if organizations adopt a broader perspective comprising both automation and augmentation, they could deal with the tension and achieve complementarities that benefit business and society. Drawing on our insights, we conclude that management scholars need to be involved in research on the use of AI in organizations. We also argue that a substantial change is required in how AI research is currently conducted in order to develop meaningful theory and to provide practice with sound advice.

Keywords: artificial intelligence, automation, augmentation, paradox theory, business and society

The rise of powerful AI will be either the best or the worst thing ever to happen to humanity. We do not yet know which.

Stephen Hawking, Theoretical physicist

What all of us have to do is to make sure we are using AI in a way that is for the benefit of humanity, not to the detriment of humanity.

Tim Cook, CEO of Apple

Artificial intelligence (AI) refers to machines performing cognitive functions usually associated with human minds, such as learning, interacting, and problem solving (Nilsson, 1971). Organizations have long used AI-based solutions to automate routine tasks in operations and logistics. Recent advances in computational power, the exponential increase in data, and new machine-learning techniques now allow organizations to also use AI-based solutions for managerial tasks (Brynjolfsson & McAfee, 2017). For example, AI-based solutions now play important roles in Unilever's talent acquisition process (Marr, 2018), in Netflix's decisions regarding movie plots, directors, and actors (Westcott Grant, 2018), and in Pfizer's drug discovery and development activities (Fleming, 2018).

In the 1950s, pioneering research predicted that AI would become essential for management (Newell, Shaw, & Simon, 1959; Newell & Simon, 1956). However, initial technological progress was slow and the discussion of AI in management was "effectively liquidated" in the 1960s (Cariani, 2010: 89). Scholars subsequently adopted a contingency view: The routine operational tasks that machines could handle were separated from the complex managerial tasks reserved for humans. Consequently, AI was researched in computer science and operations research, whereas organization and management studies focused on humans (Rahwan et al., 2019; Simon, 1987). Management scholars have therefore provided

very little insight into AI during the last two decades (Kellogg, Valentine, & Christin, in press; Lindebaum, Vesa, & den Hond, in press). Nonetheless, these scholars' understanding will be required, because AI is becoming increasingly pervasive in managerial contexts.

In this *review essay*, we strive to reposition AI at the crux of the management debate. Three highly influential business books on AI (Brynjolfsson & McAfee, 2014; Daugherty & Wilson, 2018; Davenport & Kirby, 2016) serve as a source of inspiration to challenge our thinking and spark new ideas in the management field (Bartunek & Ragins, 2015).

The three books develop a common AI narrative for practicing managers. The authors distinguish two broad AI applications in organizations: automation and augmentation. Whereas automation implies that machines take over a human task, augmentation means that humans collaborate closely with machines to perform a task. Taking a normative stance, the authors accentuate the benefits of augmentation while taking a more negative viewpoint on automation. Their combined advice is that organizations should prioritize augmentation, which they relate to superior performance. In addition, the two more recent books (Daugherty & Wilson, 2018; Davenport & Kirby, 2016) provide managers with ample advice on how to develop and implement such an augmentation strategy.

Assuming a more encompassing paradox theory perspective (Schad, Lewis, Raisch, & Smith, 2016; Smith & Lewis, 2011), we argue that augmentation cannot be neatly separated from automation in the management domain. These dual AI applications are interdependent across time and space, which creates a paradoxical tension. Over-emphasizing either augmentation or automation fuels reinforcing cycles that not only harm an organization's performance, but also have negative societal implications. However, organizations adopting a broader perspective comprising both automation and augmentation are not only able to deal with the tension, but also achieve complementarities that benefit business and society.

We conclude by discussing our insights' implications for organization and management research. The emergence of AI-based solutions and humans' increasing interactions with them creates a new managerial tension that requires research attention. Management scholars should therefore play a more active role in the AI debate by reviewing prescriptions for managerial practice and developing more comprehensive perspectives. They could do so by changing the ways they currently conduct research in order to accurately analyze and describe AI's implications for managerial practice.

REVIEWED MATERIALS

We started with a review of three recent business books on the use of AI in organizations. While there are many other books on this topic, we selected the following three, which have been widely influential in managerial practice, filling the void that the lack of scholarly research caused. The New York Times bestseller *The Second Machine Age* by MIT Professors Erik Brynjolfsson and Andrew McAfee (Brynjolfsson & McAfee, 2014) was called “the most influential recent business book” in a memorandum that Harvard Business School's Dean sent to the senior faculty (The Economist, 2017). The much-debated *Only Humans Need Apply* (Press, 2016) is the latest book on AI by Babson Professor Thomas H. Davenport and Harvard Business Review contributing editor Julia Kirby (Davenport & Kirby, 2016). Finally, the recently published *Human + Machine* by Accenture leaders Paul R. Daugherty and H. James Wilson (Daugherty & Wilson, 2018) has had an immediate impact on both academia and practice (Wladawsky-Berger, 2018).

Collectively, the three books suggest that we are on the cusp of a major transformation in business, comparable to the industrial revolution in scope and impact. During this “first machine age,” which started with the invention of the steam machine in the 18th century, mechanical machines enabled mass production by taking over manual labor tasks at scale.

Today, we face an analogous inflection point of unprecedented progress in digital technology, taking us toward the “second machine age” (Brynjolfsson & McAfee, 2014: 7). Instead of performing mechanical work, machines now take on cognitive work, traditionally an exclusively human domain. However, machines still have many limitations, which means we are entering an era in which the human-machine relationship is no longer dichotomous, but evolving into a machine “augmentation” of human capabilities. Rather than being adversaries, humans and machines should combine their complementary strengths, enabling mutual learning and multiplying their capabilities. Instead of fearing automation and its effects on the labor market, managers should acknowledge that AI has the potential to augment, rather than replace, humans in managerial tasks (Davenport & Kirby, 2016: 30–31).

Building on this analysis, the three books advise organizations to focus on augmentation rather than on automation. The two more recent books explicitly relate such an augmentation strategy to superior firm performance. For example, Daugherty and Wilson (2018: 214) conclude that companies using “AI to augment their human talent (...) achieve step gains in performance, propelling them to the forefront of their industries.” Conversely, companies focusing on automation may “see some performance benefits, but those improvements will eventually stall” (Daugherty & Wilson, 2018: 214). Similarly, Davenport and Kirby (2016: 214) predict that “a company whose strategy all along has emphasized augmentation, not automation (...) will win big.” Consequently, Davenport and Kirby advise companies to prioritize augmentation (“don’t automate, augment”) (2016: 59), which they hail as “the only path to sustainable competitive advantage” (2016: 204).

The two more recent books also provide managers with ample advice on how to develop and implement such an augmentation strategy in their organizations. Davenport and Kirby (2016: 89) describe five strategies for “post-automation” human work, all involving some form of augmentation. In addition, they provide a seven-step process for planning and developing

an augmentation strategy (Davenport & Kirby, 2016: 201). Daugherty and Wilson (2018: 105ff) describe a range of new jobs that organizations could create and in which managers complement machines and machines augment managers. The authors further describe how augmentation could be implemented across domains, ranging from sales, marketing, and customer service to research and development (Daugherty & Wilson, 2018: 67ff).

Consistent with the books' recommendations, companies have started adopting an augmentation strategy. For example, CEO Satya Nadella has announced that Microsoft will "build intelligence that augments human abilities and experiences. Ultimately, it's not going to be about human vs. machine" (Nadella, 2016). Similarly, in the preamble to its AI guidelines, Deutsche Telekom states that "AI is intended to extend and complement human abilities rather than lessen or restrict them" (Deutsche Telekom, 2018). At IBM, the corporate principles declare that "the purpose of AI and cognitive systems developed and applied by the IBM company is to augment human intelligence" (IBM Think Blog, 2017). In her speech at the World Economic Forum, IBM's President and CEO Ginni Rometty suggested replacing the term "artificial intelligence" with "augmented intelligence" (La Roche, 2017).

THE AUTOMATION-AUGMENTATION PARADOX

Taking the three books as a starting point, we use a paradox theory perspective (Schad et al., 2016; Smith & Lewis, 2011) to explore organizations' use of AI further. A paradox lens allows us to elevate the level of analysis to study both automation and augmentation, which reveals a paradoxical tension between these dual AI applications in management. Following the Smith and Lewis (2011) paradox framework, we will analyze the *paradoxical tension*, the *management strategies* to address it, and their *outcomes*.

Paradoxical Tension

The three books describe the relationship between automation and augmentation as a trade-off decision: Organizations attempting to use AI have the choice of either automating the task or using an augmentation approach. If they opt for automation, humans hand over the task to a machine with little or no further involvement. The objective is to keep humans out of the equation to allow more comprehensive, rational, and efficient processing (Davenport & Kirby, 2016: 21). In contrast, augmentation implies continued close interaction between humans and machines. This approach allows for complementing a machine's abilities with humans' unique capabilities, such as their intuition and common-sense reasoning (Daugherty & Wilson, 2018: 191f). The nature of the task determines whether organizations opt for one or the other approach. Relatively routine and well-structured tasks can be automated, while more complex and ambiguous tasks cannot, but can be addressed through augmentation (Brynjolfsson & McAfee, 2014: 138ff; Daugherty & Wilson, 2018: 107ff; Davenport & Kirby, 2016: 34ff).

The arguments that the books provide are hard to refute, but their perspective is largely limited to a given task at a specific point in time. Paradox theory, however, warns that such a narrow trade-off perspective does not represent reality adequately (Smith & Lewis, 2011). A paradox lens can help increase the scale or level of analysis for a more systemic perspective (Schad & Bansal, 2018), which allows organizations to not only perceive the contradictions, but also the interdependencies between automation and augmentation. A more comprehensive paradox perspective (both/and) then replaces the traditional trade-off perspective (either/or). The essence of paradox is that the dual elements are *both* contradictory *and* interdependent – forming a persistent tension (Schad et al., 2016).

Automation and augmentation are contradictory, because organizations choose either one or the other approach to address a given task at a specific point in time. This choice creates a tension, since these AI approaches rely on competing logics with different organizational

demands. For example, Lindebaum et al. (in press) maintain that automation instills a logic of formal rationality in organizations that conflicts with the logic of substantive rationality, or the human capacity for value-rational reflection, whereas augmentation preserves this substantive rationality. The tension is further reinforced, because some organizational actors prefer augmentation (e.g., managers at risk of losing their jobs to automation), while others prioritize automation (e.g., owners interested in efficiencies) (Davenport & Kirby, 2016: 61).

While these contradictions are real, they only reveal a partial picture. If we increase our analysis's temporal scale (from one point in time to its evolution over time) and spatial scale (from one to multiple tasks), we comprehend that, in the management domain, the two AI applications are not only contradictory, but also interdependent.

Increasing the temporal scale. Taking a process view of paradox reveals a cyclical relationship between opposing forces (Putnam, Fairhurst, & Banghart, 2016; Raisch, Hargrave, & van de Ven, 2018). Engagement with one side of the tension may set the stage or even create the conditions necessary for the other's existence; in addition, over time there is often a mutual influence between the opposing forces, with swings from one side to the other (Poole & van de Ven, 1989). Elevating the temporal scale from one point in time to the process over time allows for exploring this cyclical relationship between automation and augmentation.

As the books suggest, the process of using AI for a managerial task starts with a choice between automation and augmentation. Organizations addressing a well-structured routine task, such as completing invoices or expense claims, could opt for automation. They could do so by drawing on codified domain expertise to program rules into the system in the form of algorithms specifying the relationships between the conditions ("if") and the consequences ("then") (Gillespie, 2014).¹ Such rule-based automation requires an explicitly stated domain model, which optimizes the chosen utility function (Russell & Norvig, 2009).² With clear rules in place, managers can relinquish the task to a machine.

However, most managerial tasks are more complex, and the rules and models are therefore not fully known or readily available. In such cases, rule-based automation is impossible, but managers could use an augmentation approach to explore the problem further (Holzinger, 2016). This choice allows managers to remain involved and to collaborate closely with machines on these tasks. It is a common misconception that this augmentation process can be delegated to the IT department or external solution providers. While rule-based automation allows such delegation, because the rules can be explicitly formulated, codified, and passed on to data scientists, complex tasks' augmented learning relies on domain experts' tacit knowledge, which cannot be easily codified (Brynjolfsson & Mitchell, 2017). Data scientists can provide technical support, but domain experts need to stay "in the loop" in augmented learning (Holzinger, 2016: 119).

Augmentation is therefore a co-evolutionary process during which humans learn from machines and machines learn from humans (Amershi, Cakmak, Knox, & Kulesza, 2014; Rahwan et al., 2019). In this iterative process, managers and machines interact to learn new rules or create models and improve them over time. The type and extent of human involvement vary with the specific machine-learning solution (Russell & Norvig, 2009).³ Human domain expertise is the starting point for supervised learning. Managers provide a machine with a set of labeled training data specifying the inputs (or features) and the corresponding outputs. The machine analyzes the training data and generates rules and/or models. In contrast, unsupervised learning allows managers to induce patterns, of which they were not previously aware, directly from the unlabeled data (Jordan & Mitchell, 2015).

In both applications, managers then use their domain expertise to evaluate, select, and complement machine outputs. Spurious correlations or other statistical biases need to be weeded out. For example, machines generally learn from large, noisy datasets containing random errors. Overfitting is a key risk in this context, which means that a machine may learn

a complete model that also explains the errors, consequently failing to generalize appropriately beyond its training data (Fan, Han, & Liu, 2014). The experts' revision of the learned knowledge is therefore an important part of the augmented learning process (Fails & Olsen, 2003).⁴ In each iteration, managers assess the current model's quality, subsequently deciding on how to proceed (Langley & Simon, 1995). The resulting tight coupling between humans and machines, with the two influencing one another, makes it increasingly difficult, or even impossible, to decouple their influence on the resulting model (Amershi et al., 2014).

Over time, this close collaboration with machines sometimes allows managers to identify rules and/or models that either optimize the utility function or come sufficiently close to an optimal solution to be practically useful.⁵ If these models are sufficiently robust, they can subsequently be used to automate a task. Managers are taken "out of the loop," which allows them to focus on more demanding and valuable tasks. Augmented learning thus aims to provide increasing levels of automation, replacing time-consuming human activity with automated processes that improve accuracy, efficiency, and/or effectiveness (Langley & Simon, 1995). Consequently, augmentation may enable a transition to automation over time.

To provide illustrations of such transitions from automation to augmentation, we briefly discuss two examples from managerial practice. Organizations are increasingly employing AI-based solutions in human resource (HR) management to acquire talent (Stephan, Brown, & Erickson, 2017). For example, JP Morgan Chase chose an augmentation approach to assess candidates. A team of experienced HR managers worked closely with an AI-based solution to identify reliable, firm-specific predictors of candidates' future job performance. It took a full year of intensive interaction between the human experts and the AI-based solution to remove statistically biased, or socially vexed predictors, and make the system robust. After the initial augmentation stage, JP Morgan Chase decided to automate the candidate assessment task on the basis of the identified criteria. By removing humans from this activity, the bank intends to

increase the candidate assessment's fairness and consistency, while also making the process faster and more efficient (Riley, 2018).

Product innovation is another key domain of AI application in management (Daugherty & Wilson, 2018: 67f). For example, Symrise, a major global fragrance company, adopted an augmentation approach to generate ideas. An AI-based solution helped the company's master perfumers identify correlations between specific customer demographics and different combinations of ingredients based on the company's database of 1.7 million fragrances. Subsequently, Symrise's master perfumers used their expertise to confirm or reject the possible connections, create additional ones, and refine them further. After two years of close interaction between the master perfumers and the machine, the resulting model was considered sufficiently robust to automate the idea generation task. Based on a customer's requirements, the AI-based system now searches for possible new fragrance formulas far more rapidly and comprehensively than humans can, which has helped increase these formulas' novelty, while simultaneously reducing the search cost and time significantly (Bergstein, 2019).

As the examples illustrate, organizations may initially choose augmentation to address a complex task, but this advanced interaction between managers and AI-based solutions helps them expand their understanding of the task over time, which sometimes allows subsequent automation. While such a transition relaxes the tension temporarily, it resurfaces when conditions change over time (Smith & Lewis, 2011). For example, digitalization is likely to significantly alter the skills that JP Chase Morgan's future talents need to be successful. Bankers will need advanced data-science skills, which did not play a role in the extant employee data. Such substantial changes therefore make the automated solutions function less effectively (Davenport & Kirby, 2016: 72). Organizations should therefore, at least temporarily, return to augmentation, which allows human and machines to jointly work through the changing situation and adjust their models accordingly.

We conclude that the two AI applications in management are not only contradictory, but also interdependent. Organizations may opt for one or the other application at a given point in time, which softens the underlying tension temporarily, but fails to resolve it. Eventually, organizations will face the same choice again, demonstrating the two applications' interdependent nature and cyclical relationship.

Increasing the spatial scale. Paradox theory explores tensions not only over time, but also across space. Paradoxical tensions are nested and interwoven across multiple levels of analysis (Andriopoulos & Lewis, 2009). Addressing a tension at one level of analysis may therefore simply reproduce the tension on a different level (Smith & Lewis, 2011). Elevating the spatial scale from one task to multiple tasks allows us to explore automation and augmentation's nested interdependence across levels of analysis.

Focusing our attention on the use of either one (i.e., automation) or the other (i.e., augmentation) solution for a specific task sets artificial boundaries, fosters distinctions, and fuels opposites (Smith & Tracey, 2016). However, in practice, managerial tasks rarely occur in isolation, but are generally embedded in a managerial process. There are interdependencies between the various tasks constituting this process. These interdependencies cause managerial interventions in one task to have ripple effects throughout the process (Lüscher & Lewis, 2008). If organizations automate a task hitherto reserved for humans, this change could affect other, closely related human tasks, and lead managers to start interacting with machines. Such interactions are often iterative, resulting in the augmentation of adjacent tasks.

For example, at Symrise, the automation of the "idea generation" task also affected the preceding "objective setting" and the succeeding "idea selection" tasks in the product innovation process. Consequently, in the initial, objective-setting stage, the company's master perfumers must now enter customers' objectives and constraints into the AI-based system to allow the automated generation of fragrance formulas matching these requirements in the

subsequent idea-generation stage (Goodwin et al., 2017). This is often an iterative process, with the master perfumers circling back to adjust the objectives and the constraints according to the system outputs. In the later idea-selection stage, the master perfumers continue using their human senses, expertise, and intuition to select one of the formulas that the machine proposed. They subsequently use the AI-based solution to further refine their chosen formula (Bergstein, 2019). For example, the master perfumers employ the machine to experiment with different dosages of the selected formula's ingredients. This refinement process can include hundreds of iterations between the machine and the master perfumers. This iterative process involving close human-machine interaction has led to the augmentation of the idea selection task.⁶

As this example illustrates, a task's automation can lead to human-machine interaction in the preceding and/or the succeeding tasks in the managerial process. Automation in one task "spills over," enabling adjacent tasks' augmentation. These spillovers are particularly fast in AI systems, which often rely on distributed computing and cloud-based solutions that make the knowledge gained from a given insight immediately accessible across the system (Benlian, Kettinger, Sunyaev, & Winkler, 2018). The two AI applications in management are therefore not only interdependent across time, but also across space. While automation and augmentation are distinct activities operating in different temporal or spatial spheres, they are nevertheless intertwined at a higher level of analysis. Viewed as a paradox, automation and augmentation are no longer separate, but are mutually enabling and constituent of one another.

Persistence of the tension. Paradox refers to a tension between interdependent elements, but this tension is only considered paradoxical if it *persists* over time (Schad et al., 2016). We argue that the emerging co-existence of interdependent automated and augmented tasks will persist in the management domain. Sometimes, highly visible advancements driven by machine-learning applications are misinterpreted and extrapolated to imply that we are on the threshold of advancing toward artificial general intelligence.⁷ However, there is widespread

agreement among computer scientists that we are actually far from machines wholly surpassing human intelligence (Brynjolfsson & Mitchell, 2017; Walsh, 2017). Technical and social limitations make the full automation of complex managerial processes impossible in the foreseeable future. Managers will therefore remain involved in these processes and interact with machines on a wide range of tasks.

A few of machines' limitations are worth pointing out here: First, machines have no sense of self or purpose, which means managers need to define their objectives (Braga & Logan, 2017). Setting objectives is closely related to taking responsibility for the associated tasks and outcomes; consequently, while organizations can extend accountability to machines, responsibility requires intentionality, which is an exclusively human ability (Floridi, 2008). In turn, humans can only take responsibility if they retain some level of involvement with and control over the relevant tasks. In our example of product innovation at Symrise, the perfumers set the objectives, remain involved throughout the innovation process, and take responsibility for its outcomes. The same is true of HR managers in the talent acquisition process.

Second, in respect of complex managerial tasks, machines can only provide a range of options that all relax certain real-life constraints.⁸ Managers need to use their intuition and common-sense judgment – reconciling the machine output with reality – to make a final decision about the most desirable option (Brynjolfsson & McAfee, 2014: 92). In our example of talent acquisition at JP Morgan Chase, the AI-based solution enabled the candidate assessment's automation, but HR managers are still needed for the subsequent candidate selection (Riley, 2018), because no model can cover this task's full complexity. Machines cannot fully capture ambiguous predictors, such as cultural fit or interpersonal relations, for which there simply are no codified data available. The same applies to the product development process at Symrise, where the master perfumers ultimately choose one of the machine's suggested fragrance formulas (Bergstein, 2019).

Third, machines are limited to the specific task for which they have been trained. They cannot take on other tasks, since they do not possess the general intelligence to learn from their experience in one domain to conduct tasks in other domains (Davenport & Kirby, 2016: 35).⁹ Managers therefore need to ensure contextualization beyond an automated task. For example, HR managers still need to spend hours coordinating meetings to ensure that their hiring decisions are aligned with the business strategy, and product developers need to continue interacting with marketing departments to align their products with the business models.

Fourth, machines do not possess human senses, perceptions, emotions, and social skills (Braga & Logan, 2017).¹⁰ For example, HR managers can use their emotional and social intelligence to provide a “human touch,” or the advanced communication required to build true relationships, entice talent to work for their firm, and convince others to support the decisions made (Davenport & Kirby, 2016: 74). In the Symrise case, machines can neither smell nor fully predict how humans will perceive new fragrances, or the emotions and memories they trigger. Master perfumers have these skills and can also use them to subsequently tell a compelling story about a fragrance and its meaning, which is important for its commercialization.

To conclude, the augmentation of a managerial task may enable its subsequent automation. Such automation can, in turn, trigger further augmentation in closely related managerial tasks. While these dynamics are likely to promote increasing augmentation and automation, technological and social limitations prevent progress toward the full automation of managerial tasks in organizations. This is particularly true of managerial task contexts characterized by high degrees of ambiguity, complexity, and rare events, which limit deterministic approaches’ applicability (Davis & Marcus, 2015). In such contexts, automation and augmentation provide different, partly conflicting, but also complementary logics and functionalities that organizations require. While the optimal balance between automation and augmentation depends on contingencies, such as organizations’ AI expertise and the nature of

the environmental contexts they face, organizations will experience a persistent tension between these interrelated applications of AI in management.

Management Strategies

Recent technological progress has made the AI tension salient for organizations. Organizations facing such a salient tension tend to apply management strategies to address it. According to paradox theory, these organizational responses fuel reinforcing cycles that can be either negative or positive (Smith & Lewis, 2011). If organizations are unaware of a tension's paradoxical nature, they risk applying partial strategies, which cause vicious cycles that escalate the tension. Conversely, organizations that accept a tension as paradoxical and pay attention to its competing demands could enable virtuous cycles (Schad et al., 2016).

Vicious cycles. Organizations are likely to prioritize automation due to its promise of short-term cost efficiencies (Davenport & Kirby, 2016: 204). This strategy forces organizations' competitors to also pursue automation in order to remain cost competitive. Consequently, the whole industry may be "entering (...) in a race toward the zero-margin reality of commoditized work" (Davenport & Kirby, 2016: 204). Over time, these organizations lose the human skills required to alter their processes (Endsley & Kiris, 1995). Human experts are either made redundant through automation, or they lose their specific skills regarding the tasks they no longer pursue. Prior research has shown that automation can deskill humans, make them complacent, and diffuse their sense of responsibility (Parasuraman & Manzey, 2010; Skitka, Mosier, & Burdick, 2000). Ultimately, organizations become entrenched in their automated processes, because automation is limited to specific tasks in well-understood domains and imposes formal rules that narrow organizations' choices and penalize deviation (Lindebaum et al., in press). To conclude, while automation can free up resources for potential search activities, it is also associated with short-term thinking, the loss of human

expertise, and lock-in effects that, together, fuel a reinforcing cycle, which makes it increasingly difficult for organizations to implement such search activities.¹¹

In contrast, organizations could follow the three books' combined advice and focus on augmentation. This AI application requires extensive resources to work through iterative cycles of human-machine learning. Contrary to automation, augmentation demands continued human involvement and experimentation (Amershi et al., 2014). Since emotions and other subjective factors affect humans, augmentation is difficult or even impossible to replicate, which means every augmentation initiative is a new learning effort (Holzinger, 2016). Owing to their inherent complexity and uncertainty, augmentation efforts often fail (Amershi et al., 2014). Furthermore, the continued human involvement implies that human biases persist, which means augmentation outcomes are never fully consistent, reliable, or persistent (Huang, Hsu, & Ku, 2012). To legitimize their large augmentation investments, organizations experiencing failure may be tempted to reinforce their augmentation efforts further, which could escalate their commitment (Sabherwal & Jeyaraj, 2015; Staw, 1981), with failure leading to continued augmentation, in turn leading to continued failure.

To conclude, one-sided orientations toward either automation or augmentation cause vicious cycles, because they neglect the dynamic interdependencies between AI's dual applications in management. Managers limiting their perspective to either automation or augmentation risk developing partial and incomplete managerial solutions. While these solutions may be appropriate within the strict boundaries that time and space impose, the use of AI in management causes an organizational tension that persists across time and space.

Virtuous cycles. Paradox theory offers a more constructive response to tensions by envisioning a virtuous cycle, with organizations overcoming their defensiveness to embrace these tensions and viewing them as an opportunity to find synergies that accommodate and transcend the opposing poles (Schad et al., 2016).

A first step toward enabling such a virtuous cycle is the *acceptance* of the tensions as paradoxical (Smith & Lewis, 2011). While managers initially perceive automation and augmentation as a trade-off, they may eventually recognize that they cannot simply choose between these dual AI applications, because either choice intensifies the need for its opposite. However, transitioning to a more encompassing paradox perspective requires cognitive and behavioral complexity (Miron-Spektor, Ingram, Keller, Smith, & Lewis, 2018). Stimulating an exchange between organizational actors with different perspectives, such as data scientists and business managers, could develop more complex understandings of the phenomenon. Once actors accept that automation and augmentation can and should co-exist, they can explore the dynamic relationship between them mindfully, which could be part of their organization's vision or guiding principles regarding the use of AI in management.

While acceptance lays the groundwork for virtuous cycles, it has to be complemented with a subsequent *resolution* strategy (Smith & Lewis, 2011). Resolution involves seeking responses to paradoxical tensions through a combination of differentiation and integration practices (Poole & van de Ven, 1989).

Differentiation allows organizations to recognize and appreciate automation and augmentation's distinctive benefits and leverage them separately. Organizations can purposefully iterate between distinct automation and augmentation tasks, allowing long-term engagement with both forces. For example, Symrise's master perfumers iterate between automation (i.e., when generating alternative fragrance formulas) and augmentation (i.e., when selecting and refining the most promising formula). The use of automation allows exploration beyond humans' abilities by searching through the whole landscape of possible options. Their cognitive limitations mean that humans' search field is restricted, while machines do not face such information-processing limitations (Davenport & Kirby, 2016: 17). Excluding humans at this stage may help break path dependencies and promote greater novelty. Switching to

augmentation allows machine limitations to be overcome by subsequently using humans' more holistic and intuitive information processing to choose between options and contextualize beyond the specific task at hand (Brynjolfsson & McAfee, 2014: 92).

While such differentiation allows for engaging in both automation and augmentation, integration enables the finding of linkages that transcend the two poles (Smith & Lewis, 2011). By switching, the machine's independent output can be used to challenge human intuition and judgment, with human feedback enabling further rounds of machine analysis (Hoc, 2001). At these transition points, automation and augmentation become mutually enabling. The two AI approaches' juxtaposition stimulates learning and fosters adaptability, allowing the combination of (machine) rationality and (human) intuition, which enables more comprehensive information processing and better decisions (Calabretta, Gemser, & Wijnberg, 2017). Through integration, automation and augmentation jointly generate outcomes that neither application can enable individually.

It is no easy feat to ensure such integration. As described above, the risks of organizations over-emphasizing either automation or augmentation are real. Integration therefore requires humans to retain overall responsibility for a managerial process. Prior studies have shown that maintaining overall human responsibility not only reduces human bias (Larrick, 2004), but also prevents human-machine collaboration biases (Skitka et al., 2000). As these studies show, assigning the overall responsibility for processes to humans leads to increased vigilance and verification behavior, the consideration of a wider range of inputs prior to making decisions, and the use of greater cognitive complexity when processing such information. Consequently, retaining human responsibility for managerial processes promotes integration, which transcends automation and augmentation.

Outcomes

Paradox theory suggests that managing tensions through the dynamic strategies of acceptance and resolution fosters sustainability (Smith & Lewis, 2011). By managing paradox, organizations enable learning and creativity, promote flexibility and reliance, and unleash human potential. However, paradox scholars also acknowledge that narrow organizational attention to just one of the tensions' poles can trigger unintended organizational and societal consequences (Schad & Bansal, 2018). We therefore conclude our analysis of the automation-augmentation paradox by assessing its organizational and societal outcomes.

Organizational outcomes. The three books we reviewed argue that organizations benefit greatly from using AI. In particular, the authors emphasize augmentation's potential to increase productivity, improve service quality, and foster innovation. Moreover, they assume that the combination of complementary human and machine skills will increase the quality, speed, and extent of learning in organizations (Brynjolfsson & McAfee, 2014: 182; Daugherty & Wilson, 2018: 106; Davenport & Kirby, 2016: 206). In contrast, we have argued that focusing on either automation or augmentation can lead to reinforcing cycles that harm long-term performance. We suggest that organizations benefit if they differentiate between and integrate across automation and augmentation.

Differentiation allows organizations to benefit from both AI applications' unique benefits. Automation enables organizations to drive cost efficiencies, establish faster processes, and ensure greater information-processing rationality and consistency. As described before, augmentation provides complementary benefits arising from the mutual enhancing of human and machine skills. The integration of automation and augmentation leads to additional benefits that accrue from the synergies between these interdependent activities. Automation could free up scarce resources for augmentation, which, in turn, could help identify the rules and/or models that enable automation. Balancing automation and augmentation helps prevent the

escalating cycles that focusing on just one of these AI applications could cause. Furthermore, the combination of automation and augmentation could enable new business models. AI is, for example, the major driver behind the current trend toward personalized medicine, with treatments being tailored to each patient's specific biological profile (Fleming, 2018; Lichfield, 2018). While augmentation allows for identifying patterns in large volumes of patient data, automation makes the design and manufacturing of tailored drugs economically viable.

These varied benefits suggest that automation and augmentation's combination creates complementary returns that lead to superior firm performance. Together, AI's dual applications in management provide organizations with a range of benefits that neither automation nor augmentation can provide alone. However, realizing these benefits is contingent upon organizations' active management of the automation-augmentation paradox.

Societal outcomes. Tensions' systemic nature is a central tenet of paradox theorizing (Jarzabkowski, Bednarek, Chalkias, & Cacciatori, 2019; Smith & Lewis, 2011). Paradoxes are embedded in open systems and their implications extend beyond a single organization's boundaries. Consequently, it is important to adopt a more systemic perspective of paradox, which not only takes the organizational outcomes into consideration, but also tensions' and their management's larger, system-wide or societal implications (Schad & Bansal, 2018).

While firms may gain profits from their use of AI in management, the three books – to a varying extent – also point out that the larger societal implications are less certain (e.g., Brynjolfsson & McAfee, 2014: 171). There is a risk that organizations could take a narrow perspective of either automation or augmentation, triggering unintended consequences that affect society negatively. However, if organizations adopt a more comprehensive perspective, the outcomes could be positive for both business and society. We explore these issues further by focusing our attention on two societal outcomes discussed extensively in the three books:

AI's labor market impact (e.g., Brynjolfsson & McAfee, 2014: 147ff), and its effects regarding social equality and justice (e.g., Daugherty & Wilson, 2018: 129ff).

First, a one-sided focus on automation could cause extensive job losses and result in the deskilling of managers who relinquish tasks to machines, which could lead to the further risks of rising unemployment and social inequality (Brynjolfsson & McAfee, 2014: 171; see also Autor, 2015). Conversely, one-sided augmentation is likely to cause another “digital divide” (Norris, 2001), with social tensions arising between the few who currently have the capabilities and resources for augmentation and those who do not (Brynjolfsson & McAfee, 2014: 134f; see also Brynjolfsson & Mitchell, 2017).

Balancing automation and augmentation could, however, enable a virtuous cycle of selective deskilling (i.e., humans offload tasks where their abilities are inferior to those of machines) and strategic requalification (i.e., humans stay ahead of machines in their core abilities), thereby gradually enhancing both human and machine capabilities. This virtuous cycle could help organizations reduce the digital transition's negative effects on their employees and the labor market at large. Employees and managers whose basic skills are made redundant by automation could be given the opportunity to gradually build higher-level augmentation skills that remain in demand. This skill-enhancement cycle could also help “rehumanize work” by gradually shifting the focus from repetitive and monotonous tasks to more creative and fulfilling tasks (Daugherty & Wilson, 2018: 214).

A recent initiative at UBS's investment-banking division illustrates this virtuous cycle. UBS used an AI-based solution to automate the task of reading and executing client demands for fund transfers, which previously took an investment banker 45 minutes per demand. Simultaneously, the bank implemented another AI-based solution to augment the development of trading strategies. The investment bankers use the time freed up by automation to collaborate closely with the AI-based tool to explore new strategies for adaptive trading. Data scientists

provide the investment bankers with organizational support to develop their augmentation skills. Consequently, the combined use of automation and augmentation allowed UBS to exploit cost efficiencies while exploring new client solutions (Arnold & Noonan, 2017).

Second, the use of AI in management could also have implications for social equality and fairness. Automation takes humans “out of the loop,” reducing human biases and, in turn, promising greater equality and fairness. For example, using automation for credit approval could reduce bankers’ bias that might previously have kept people from qualifying for credit due to their ethnicity, gender, or postal code (Daugherty & Wilson, 2018: 167). Similarly, automated candidate assessment based on pre-determined criteria and consistent machine processing could help eliminate humans’ implicit biases in their hiring decisions.

However, real-world applications show that machine biases caused by noisy data, statistical errors, and/or socially vexed predictors often lead to new, even more systematic discrimination. Daugherty and Wilson (2018: 179) cite the example of an automated AI system used to predict defendants’ future criminal behavior, which has been shown to be biased against black defendants. Another example involves Amazon, which discontinued the use of an automated AI-hiring tool found to discriminate against female applicants for technical jobs (Dastin, 2018). In contrast, augmentation is likely to reduce such machine biases through human backtesting and feedback. However, the intense interaction between managers and machines increases the risk of human biases being carried over to machines. This problem is particularly pernicious, since machines then confirm humans’ biased intuition, which makes humans even less likely to question their preconceived positions (Huang et al., 2012).

The solution could be once more to combine differentiation and integration practices to address the paradox. Differentiation allows for independent analyses with (i.e., augmentation) and without (i.e., automation) human involvement. Integration ensures that machine outputs are used to challenge humans, and human inputs to challenge machines (Hoc, 2001), allowing

for mutual learning (Panait & Luke, 2005) and debiasing (Larrick, 2004). Furthermore, humans “in the loop” could explain the system’s outputs, rendering algorithmic decisions more accountable and transparent (Binns et al., 2018).

For example, much like JP Morgan Chase, Unilever differentiates clearly between automation (used for the initial candidate assessment) and augmentation (used for the final selection) in their talent acquisition process. Integration is ensured by retaining the overall human responsibility for the entire process. Unilever reports that the combination of automation and augmentation has led to a 16 percent increase in new hires’ ethnic and gender diversity, making it the “most diverse class to date.” At the same time, the company managed to save 70 000 person-days of interviewing and assessing candidates, which resulted in annual cost savings of GBP 1 million and a reduction of the time-to-hire by 90 percent (from an average of four months to two weeks) (Feloni, 2017).

DISCUSSION

Inspired by three recent business books, we have explored the emergent use of AI-based applications in organizations to automate and augment managerial tasks. The central argument of this *review essay* is that automation and augmentation are not only separable and conflicting, but in fact fundamentally interdependent. We suggest that the prevailing tradeoff perspective is overdrawn; viewed as a paradox, augmentation is both the driver and outcome of automation, and the two applications of AI develop and fold into one another across space and time. The automation-augmentation model’s popularity stems largely from its clear boundaries and simplicity. However, adopting this intuitively appealing model uncritically obscures how automation and augmentation intertwine in managerial practice. This *review essay* therefore sheds light on automation and augmentation’s complementarities and identifies opportunities to transcend the paradoxical relationship between them.

In the following, we discuss our reflections' implications for organization and management research. We will argue that the ways in which scientific research on AI is currently conducted need to change to accurately capture and analyze its organizational and societal implications for managerial practice. Our discussion will be loosely structured along the famous "5w & 1h" questions (who, how, what, why, where, and when).¹²

Our analysis of the automation-augmentation paradox suggests that the jury is out on whether the use of AI in management will turn out to be a blessing or a curse. Scholars can still make a difference by exploring the topic and informing practice regarding the ways forward. However, this will require a change in *who* conducts research on AI. Currently, computer scientists, roboticists, and engineers are the scholars most commonly studying AI. Their primary objective is to automate as far as possible, because they often regard humans as "a mere disturbance in the system that can and should be designed out" (Cummings, 2014: 62). Computer scientists may be expert technologists, but they are not generally trained social or behavioral scientists (Rahwan et al., 2019). Instead, they tend to use laboratory settings for their research, which reduce the inherent variability that characterizes human behavior. These settings permit the use of methodologies aimed at maximizing algorithmic performance, but disregard the role of humans and the wider organizational and societal implications.

The three business books' augmentation perspective reveals the importance of inducing social scientists to participate in the debate. With respect to managerial tasks, humans will remain "in the loop" and interact closely with machines. Management scholars are particularly well-equipped to study these human-machine interactions in real-world settings, as well as to explore their organizational and societal implications. It is therefore no wonder that the three business books devote the bulk of their attention to augmentation. However, our analysis suggests that management research limited to augmentation may be as biased as computer science's traditional focus on automation. Research on AI in management would therefore

benefit greatly from more interdisciplinary efforts. The technological and the social worlds are merging (Orlikowski, 2007), which means that computer scientists and management scholars no longer study separate phenomena. By juxtaposing and integrating their different perspectives, theories, and methodologies, computer scientists and management scholars can jointly create a foundation for meaningful research on the use of AI in management.

Such efforts also require a change in *how* research on AI is conducted in the management domain. The limited work to date provides separate accounts with clear-cut contrasts between augmentation and automation. On the one side, the research and practice doomsayers warn us that automation will enslave humans, supervise and control them, and drive out every iota of humanity (e.g., Bostrom, 2014; Ford, 2015). For example, in their recent *AMR review essay*, Lindebaum et al. (in press) maintain that automation may lead to a technology-enabled totalitarian system with formal and oppressive rules representing the end of human choice. On the other side, there are technology utopians (e.g., Kurzweil, 2014; More & Vita-More, 2013), including the authors of two of the reviewed books (Daugherty & Wilson, 2018; Davenport & Kirby, 2016), who argue that humans will remain in control and use augmentation to create huge benefits for organizations and society.¹³

The automation-augmentation paradox suggests that both perspectives are equally biased. Automation and augmentation are not good or evil *per se*. Derrida (1967) argued that humans tend to construct binary oppositions in their narratives, with a hierarchy that privileges one side of the dichotomy while repressing the other. He warns that this approach is overly simplistic; it makes us forget that one side of a dichotomy cannot exist without the other. Similarly, researchers need to accept that automation and augmentation are interdependent AI applications in management that cannot be neatly separated and designated as either good or evil. These applications provide complementary functionalities that are both potentially useful

for organizations. The complex interaction between these varied AI applications over time could have *both* positive *and* negative organizational and societal implications.

Research on AI in management therefore needs to complexify its theorizing by moving from simple either/or perspectives to more encompassing both/and ones. Complexifying theories is essential for understanding complex phenomena (Tsoukas, 2017), because “it takes richness to grasp richness” (Weick, 2007: 16). More encompassing perspectives, such as paradox theory (Schad et al., 2016) or systems theory (Sterman, 2000), offer a vantage point from which researchers can observe the dynamic interplay between automation and augmentation. Accepting and embracing this complexity allows for studying AI and its managerial applications “in the wild.” This leads to a more comprehensive and, ultimately, more rigorous and relevant discussion of AI’s organizational and societal implications.

By working through this complexity, it becomes apparent *what* research needs to be conducted. At the most basic level, management scholars need to acknowledge that humans are no longer the sole agents in management, although most theories to date focus exclusively on human agency. Scholars need to overcome this human bias and integrate intelligent machines into their theories. The use of AI for managerial tasks implies that machines are no longer simple artifacts, but a new class of agents in organizations (Floridi & Sanders, 2004). While machines have fundamental limitations, their actions nevertheless enjoy far-reaching autonomy, because humans delegate knowledge tasks to these machine agents and allow them act on their behalf (Rai, Constantinides, & Sarker, 2019).

Such automation leads to machine behavior that deviates significantly from the human behavior that management theories traditionally describe. The bulk of current management research relies on the behavioral assumptions of boundedly rational human actors, who – due to their information-processing limits and cognitive biases – engage in satisficing rather than maximizing behavior (Argote & Greve, 2007; Cyert & March, 1963). However, intelligent

machines used for automation do not have these limitations; they have practically unlimited information-processing capacity and exhibit perfectly consistent behavior. Nonetheless, they can introduce statistical biases and have other limitations that humans do not have (Elsbach & Stigliani, 2019). These differences lead to entirely new machine behaviors (Rahwan et al., 2019). We could, for example, speculate that because machines do not have humans' limitations when searching, organizations using automation and augmentation could suffer less from learning myopia and path dependencies (Levinthal & March, 1993). Management scholars thus need to broaden their perspective to include human and machine agents and explore their distinct behaviors in organizational contexts.

If we increase the complexity further, we have to acknowledge that these human and machine agents do not simply co-exist in separate worlds (working on separate tasks), but are interdependent (interacting on the same or closely related tasks). Augmentation therefore implies close collaboration between humans and machines. Since automation and augmentation are interdependent, this interaction spreads across organizations. When addressing this human-machine interaction, management scholars need to first explore how machines shape managerial behavior. For example, Lindebaum et al. (in press) describe how autonomous algorithms can direct and constrain human behavior by imposing formal rationality. This perspective resonates with Foucault's (1977) concept of panoptic surveillance, characterizing IT as a form of omnipresent architecture of control that creates, maintains, and cements central norms of expected behavior (see also Lyon, 2003; Zuboff, 1988, 2019).

However, our broader paradox perspective of AI in management reveals that humans also shape machine behavior. Managers define the objectives, set constraints, generate and choose the training data, and provide machines with feedback. In machine-learning systems, humans shape and reshape algorithms through their daily actions and interactions (Deng, Bao, Kong, Ren, & Dai, 2017). In a management context, machines may influence human behavior,

but without setting a static norm of conduct or an unsurpassable rule (Cheney-Lippold, 2011). Managers participate in writing and rewriting the rules that shape behavior. Management research should thus explore both machines' influence on human behavior and humans' influence on machine behavior in the context of AI use in management.

If we increase the complexity even further, we finally see that humans and machines influence one another in an *iterative* process. While machine algorithms shape human actions, humans shape these algorithms through their actions, which create a recursive relationship (Beer, 2017). Through augmentation, humans and machines become so closely intertwined that they collectively exhibit entirely new, emergent behaviors, which neither show individually (Amershi et al., 2014). The use of AI in management leads to hybrid organizational systems manifesting collective behaviors. It may therefore be difficult, or even impossible, to distinguish between humans and machines or their respective learning and actions.

While it is convenient, and sometimes helpful, to separate research studies analyzing how machines influence managers and vice versa, studies examining hybrid organizational systems comprising both managers and machines should provide the greatest benefit. Only such studies can examine the feedback loops between human influence on machine behavior and machine influence on human behavior, which cause the emergent behaviors that are otherwise impossible to predict (Rahwan et al., 2019). Management scholars must provide further insights into how such hybrid organizational systems function by exploring their interactive behaviors. This systemic perspective also enables them to predict and/or explore the automation-augmentation paradox's systemic dynamics, which, as we have described before, can include unintended consequences and escalating cycles. Management research therefore has a crucial mandate to explore managers' continued interactions with machines, as well as the emergent behaviors and systemic outcomes they cause.

This discussion leads us to the heart of the fourth question, namely *why* research on AI in management is essential. As argued above, the emergent use of AI in management leads to iterative interactions between humans and machines. The resulting hybrid organizational systems exhibit behaviors and produce organizational, as well as societal, effects that are impossible to predict precisely and are often entirely unanticipated (O’Neil, 2016). No single actor in these systems has full control over these outcomes. Consequently, it is difficult to apportion accountability for outcomes to specific actors, which creates an “accountability gap” (Mittelstadt, Allo, Taddeo, Wachter, & Floridi, 2016). There is widespread fear of and the first empirical evidence that this lack of accountability can have harmful societal consequences regarding equality (Autor, 2015), privacy (Acquisti, Brandimarte, & Loewenstein, 2015), security (Brundage et al., 2018), and transparency (Castelvecchi, 2016).

Traditional managerial and organizational solutions may be inadequate for addressing such systemic problems sufficiently (Schad & Bansal, 2018). Management research should therefore contribute to the development of new organizational solutions that allow AI’s benefits to be realized, while mitigating the associated negative side effects. In order to fully understand the automation-augmentation paradox’s societal implications, scholars could adopt a relational ontology which accepts that, in the digital age, human and machine agents are so closely intertwined in hybrid collectives that their relations determine their actions (Floridi & Taddeo, 2016). Rather than focusing on individual actors, the interactions between these actors should be the unit of analysis. Such a perspective enables a discussion of “distributed morality” (Floridi, 2013), which relies on shared ethical norms developed through collaborative practices (Mittelstadt et al., 2016), and critically assesses whether, and to what extent, it replaces or complements our current focus on individual responsibility in the digital age.

With regard to the fifth question, *where* refers to the locus of management scholars’ research attention. AI is a particularly broad research field with a great variety of organizational

and societal implications. Accordingly, researchers from disciplines such as astronomy (Haiman, 2019), biology (Webb, 2018), law (Corrales, Fenwick, & Forgó, 2018), medicine (Topol, 2019), politics (Helbing et al., 2019), psychology (Jaeger, 2016), and sociology (McFarland, Lewis, & Goldberg, 2016) have all addressed and presented conceptual ideas. In this regard, our focus was exclusively on the emergent use of AI for managerial tasks in practice. While management scholars may (and should) become involved in the broader discussion of AI and its societal implications, the use of automation and augmentation in management relates to the core of management scholars' research. We therefore suggest that the management domain should be the specific focus of our attention.

In this *review essay*, our focus was predominantly on questions of organizational functioning, such as those pertaining to the effect of collaboration between multiple humans and machines on managerial tasks. While such meso-level topics will likely play a central role in future research, organizations' use of AI in management should also be explored on the micro and macro levels of analysis. Micro-level research could investigate how the emergence of AI-based solutions changes the role of managers in organizations. In the past, management theories emphasized managers' domain expertise, which granted them expert power and status in their organizations (Finkelstein, 1992). Although domain expertise remains relevant for managers regarding educating and challenging machines, automation and augmentation will lead to institutionalized knowledge – for example, in the form of algorithms – which is often superior to individual managers' expert knowledge. At the same time, general human skills that complement machines, such as creativity, common sense, and advanced communication (Davenport & Kirby, 2016: 30), as well as integration skills such as AI literacy, will gain further importance in an era of automation and augmentation (Daugherty & Wilson, 2018: 191). These developments could lead to important shifts in managers' roles, competencies, and status.

Macro-level research could explore how the emergence of automation and augmentation in management leads to institutional action and change. For example, AI is often applied in open systems, blurring organizational boundaries (Panait & Luke, 2005); data are collected widely, with diverse stakeholders updating them continuously and collectively through their actions; and inputs from agents within and outside the organization impact the automation and augmentation process, which, in turn, can have wide-reaching societal implications. A core focus of management scholars' future research attention should therefore be on studying how broader networks of actors, comprising activists, companies, governments, public institutions, and international organizations, collaborate to set standards, build institutions, and organize collective action to address issues pertaining to the use of AI in management.

This leaves the last remaining question of *when* scholars should address the phenomenon of emerging AI use in managerial practice. The answer is: "Immediately!" Managers in key organizational domains, including customer management, human resources, marketing, product innovation, sales, and strategy have already started working closely with intelligent machines on automated and augmented tasks. This introduction of AI in practice will profoundly change the nature of management. These developments offer many fruitful areas for scholarly research. Management scholars still have the opportunity to make a lasting impact on how organizations perceive and cope with the complex challenges they face. Our *review essay* shows that there is an urgent need for a better understanding, more reliable theories, and sustainable managerial solutions. We therefore close with a call to action and encourage our readers to embrace the topic of AI in management.

ENDNOTE

¹ This step can also be done by using unsupervised machine learning (Russell & Norvig, 2009), which allows the machine to induce rules directly from the data. If the task is deterministic, and the rules are simple and clear, these rules can be readily used for automation.

² A utility function represents the organization's preference ordering over a choice set, allowing it to assign a real value to each alternative. In the field of AI, utility functions are used to convey various outcomes' relative value to machines, which in turn allows them to propose alternatives that optimize the utility function (Russell & Norvig, 2009).

³ Consistent with the three books, we adopt a broad definition of AI, comprising both rule-based automation and machine learning. In rule-based automation, which is sometimes also called "robotic process automation" (RPA), the machine is static in the sense that it adheres to the explicit rules it has been given (Daugherty & Wilson, 2018: 50; Davenport & Kirby, 2016: 48). In contrast, machine learning gives the machine the ability to learn from experience without being explicitly programmed to do so (Mitchell, 1997).

⁴ Machine-learning solutions already employ measures against overfitting, such as cross-validation and regularization. However, these measures can only complement, not replace human responsibility and intervention in managerial tasks (Greenwald & Oertel, 2017).

⁵ The computer-science literature distinguishes between tractable or polynomial (P) problems and intractable or non-polynomial (NP) problems (Dean, 2016; Hartmanis & Stearns, 1965). Less complex (P) problems are amenable to optimization and rule-based automation. In contrast, machines working on more complex (NP) problems encounter optimization problems. While the optimal solution may be out of reach, machine-learning solutions can find models that approximate such a solution with certain accuracy. These solutions are therefore

suboptimal (given that they inevitably relax certain real-life constraints), but may be close enough to the optimal solution to be suitable for practical application (Fortnow, 2013).

⁶ The talent acquisition process in our JP Morgan Chase example functions similarly: HR managers now engage in augmentation to initially set the machine's objectives and constraints (objective setting) and to subsequently select from the candidates that the machine suggested (candidate selection) (Riley, 2018).

⁷ A recent example is the media hype around AlphaGo Zero, an AI-based system representing state-of-the-art Chess and Go play (Silver et al., 2017). AlphaGo Zero learned these games through trial-and-error (or reinforcement) without human guidance, only playing games against itself. However, people often overlook that programmers still needed to feed AlphaGo Zero an important piece of human knowledge: the rules of the game. Chess and Go have explicit, finite, and stable goals, rules, and reward signals, which allow machine learning to be optimized. Most real-world managerial problems are far more complex than games like Chess or Go. For example, the rules of managerial problems might not be known, might be ambiguous, and/or might change over time. While AlphaGo Zero is impressive, it represents little if any progress toward artificial general intelligence.

⁸ Intractable (NP) problems imply discrete optimization problems. While the optimal solution is out of reach, relaxation of the constraints allows these problems to be addressed (Fortnow, 2013). Humans use their experience to reduce the search space of exponential possibilities by means of heuristic selection. Machines can subsequently use approximation algorithms to provide a range of possible solutions that all relax certain real-life constraints.

⁹ The phenomenon of "catastrophic forgetting" explains this machine limitation; that is, having learned one task and subsequently transferred to another, a machine-learning system simply "forgets" how to perform the previously learned task (Taylor & Stone, 2009). Humans, on the

other hand, possess the capacity to transfer learning, allowing them to generalize from one task context to another (Parisi, Kemker, Part, Kanan, & Wermter, 2019).

¹⁰ Several current projects are aimed at deploying AI-based agents capable of perceiving and responding to emotional cues. However, these agents are very limited in their capabilities, because fundamental technical and ethical challenges limit their potential for human-level emotional sentience in the foreseeable future (McDuff & Czerwinski, 2018).

¹¹ Prior studies have shown that slack resources are a necessary, but insufficient prerequisite for organizational search. Slack resources can also induce complacency and inertia, especially if organizational factors work against leveraging slack resources for search (Desai, in press).

¹² The “5w & 1h” questions are a method used in areas as varied as journalism, research, and police investigations to describe and evaluate a subject comprehensively. The method’s origins have been traced to Aristotle’s *Nicomachean Ethics* (Sloan, 2010).

¹³ We acknowledge that Brynjolfsson and McAfee (2014) provide a far more balanced discussion of AI’s organizational and societal implications than the two more recent business books (Daugherty & Wilson, 2018; Davenport & Kirby, 2016).

REFERENCES

- Acquisti, A., Brandimarte, L., & Loewenstein, G. 2015. Privacy and human behavior in the age of information. *Science*, 347: 509–514.
- Amershi, S., Cakmak, M., Knox, W. B., & Kulesza, T. 2014. Power to the people: The role of humans in interactive machine learning. *AI Magazine*, 35(4): 105–120.
- Andriopoulos, C., & Lewis, M. W. 2009. Exploitation-exploration tensions and organizational ambidexterity: Managing paradoxes of innovation. *Organization Science*, 20: 696–717.
- Argote, L., & Greve, H. R. 2007. A behavioral theory of the firm - 40 years and counting: Introduction and impact. *Organization Science*, 18: 337–349.
- Arnold, M., & Noonan, L. 2017. Robots enter investment banks' trading floors. *Financial Times*. <https://www.ft.com/content/da7e3ec2-6246-11e7-8814-0ac7eb84e5f1>. July 7.
- Autor, D. H. 2015. Paradox of abundance: Automation anxiety returns. In S. Rangan (Ed.), *Performance and progress: Essays on capitalism, business, and society*: 237–260. Oxford: Oxford University Press.
- Bartunek, J. M., & Ragins, B. R. 2015. Extending a provocative tradition: Book reviews and beyond at AMR. *Academy of Management Review*, 40: 474–479.
- Beer, D. 2017. The social power of algorithms. *Information, Communication & Society*, 20: 1–13.
- Benlian, A., Kettinger, W. J., Sunyaev, A., & Winkler, T. J. 2018. The transformative value of cloud computing: A decoupling, platformization, and recombination theoretical framework. *Journal of Management Information Systems*, 35: 719–739.
- Bergstein, B. 2019. Can AI pass the smell test? *MIT Technology Review*, 122(2): 82–86.
- Binns, R., Van Kleek, M., Veale, M., Lyngs, U., Zhao, J., & Shadbolt, N. 2018. “It’s reducing a human being to a percentage”: Perceptions of justice in algorithmic

- decisions. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 377–391. New York: ACM Press.
- Bostrom, N. 2014. *Superintelligence: Paths, dangers, strategies*. Oxford: Oxford University Press.
- Braga, A., & Logan, R. 2017. The emperor of strong AI has no clothes: Limits to artificial intelligence. *Information*, 8: 156–177.
- Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., Dafoe, A., Scharre, P., Zeitzoff, T., Filar, B., Anderson, H., Roff, H., Allen, G. C., Steinhardt, J., Flynn, C., Ó hÉigearthaigh, S., Beard, S., Belfield, H., Farquhar, S., Lyle, C., Crootof, R., Evans, O., Page, M., Bryson, J., Yampolskiy, R., & Amodei, D. 2018. *The malicious use of artificial intelligence: Forecasting, prevention, and mitigation*.
<https://maliciousaireport.com>. February 20.
- Brynjolfsson, E., & McAfee, A. 2014. *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. New York: W.W. Norton.
- Brynjolfsson, E., & McAfee, A. 2017. The business of artificial intelligence. *Harvard Business Review*, July issue.
- Brynjolfsson, E., & Mitchell, T. 2017. What can machine learning do? Workforce implications. *Science*, 358: 1530–1534.
- Calabretta, G., Gemser, G., & Wijnberg, N. M. 2017. The interplay between intuition and rationality in strategic decision making: A paradox perspective. *Organization Studies*, 38: 365–401.
- Cariani, P. 2010. On the importance of being emergent. *Constructivist Foundations*, 5: 86–91.
- Castelvecchi, D. 2016. Can we open the black box of AI? *Nature*, 538: 20–23.
- Cheney-Lippold, J. 2011. A new algorithmic identity. *Theory, Culture & Society*, 28(6):

164–181.

Corrales, M., Fenwick, M., & Forgó, N. (Eds.). 2018. *Robotics, AI and the future of law*.

Singapore: Springer.

Cummings, M. M. 2014. Man versus machine or man + machine? *IEEE Intelligent Systems*,

29(5): 62–69.

Cyert, R. M., & March, J. G. 1963. *A behavioral theory of the firm*. Englewood Cliffs, NJ:

Prentice-Hall.

Dastin, J. 2018. Amazon scraps secret AI recruiting tool that showed bias against women.

Reuters. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>. October 10.

Daugherty, P., & Wilson, H. J. 2018. *Human + machine: Reimagining work in the age of*

AI. Boston, MA: Harvard Business Review Press.

Davenport, T. H., & Kirby, J. 2016. *Only humans need apply: Winners and losers in the*

age of smart machines. New York: HarperCollins.

Davis, E., & Marcus, G. 2015. Commonsense reasoning and commonsense knowledge in

artificial intelligence. *Communications of the ACM*, 58(9): 92–103.

Dean, W. 2016. Computational complexity theory. In E. N. Zalta (Ed.), *The Stanford*

encyclopedia of philosophy. Stanford, CA: Metaphysics Research Lab, Stanford University.

Deng, Y., Bao, F., Kong, Y., Ren, Z., & Dai, Q. 2017. Deep direct reinforcement learning for

financial signal representation and trading. *IEEE Transactions on Neural Networks and Learning Systems*, 28: 653–664.

Derrida, J. 1967. *De la grammatologie*. Paris: Les Éditions de Minuit.

Desai, V. M. In press. Can busy organizations learn to get better? Distinguishing between the

competing effects of constrained capacity on the organizational learning process.

Organization Science, doi: 10.1287/orsc.2019.1292.

Deutsche Telekom. 2018. Deutsche Telekom's guidelines for artificial intelligence. *Deutsche Telekom*.

<https://www.telekom.com/resource/blob/532446/f32ea4f5726ff3ed3902e97dd945fa14/d1-180710-ki-leitlinien-en-data.pdf>. April 24.

Elsbach, K. D., & Stigliani, I. 2019. New information technology and implicit bias. *Academy of Management Perspectives*, 33: 185–206.

Endsley, M. R., & Kiris, E. O. 1995. The out-of-the-loop performance problem and level of control in automation. *Human Factors*, 37: 381–394.

Fails, J. A., & Olsen, D. R. 2003. Interactive machine learning. *Proceedings of the 8th International Conference on Intelligent User Interfaces*, 39–45. New York: ACM Press.

Fan, J., Han, F., & Liu, H. 2014. Challenges of big data analysis. *National Science Review*, 1: 293–314.

Feloni, R. 2017. Consumer-goods giant Unilever has been hiring employees using brain games and artificial intelligence - and it's a huge success. *Business Insider*.
<https://www.businessinsider.com/unilever-artificial-intelligence-hiring-process-2017-6>.
June 28.

Finkelstein, S. 1992. Power in top management teams: Dimensions, measurement, and validation. *Academy of Management Journal*, 35: 505–538.

Fleming, N. 2018. How artificial intelligence is changing drug discovery. *Nature*, 557: S55–S57.

Floridi, L. 2008. Information ethics: A reappraisal. *Ethics and Information Technology*, 10: 189–204.

- Floridi, L. 2013. *The philosophy of information*. Oxford, UK: Oxford University Press.
- Floridi, L., & Sanders, J. W. 2004. On the morality of artificial agents. *Minds and Machines*, 14: 349–379.
- Floridi, L., & Taddeo, M. 2016. What is data ethics? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374: 1–4.
- Ford, M. 2015. *The rise of the robots: Technology and the threat of mass unemployment*. New York: Basic Books.
- Fortnow, L. 2013. *The golden ticket: P, NP, and the search for the impossible*. Princeton, NJ: Princeton University Press.
- Foucault, M. 1977. *Discipline and punish: The birth of the prison*. New York: Pantheon Books.
- Gillespie, T. 2014. The relevance of algorithms. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society*: 167–194. Cambridge, MA: The MIT Press.
- Goodwin, R., Maria, J., Das, P., Horesh, R., Segal, R., Fu, J., & Harris, C. 2017. AI for fragrance design. *Proceedings of the machine learning for creativity and design workshop at NIPS*. San Diego, CA: Neural Information Processing Systems Foundation.
- Greenwald, H. S., & Oertel, C. K. 2017. Future directions in machine learning. *Frontiers in Robotics and AI*, 3: 79.
- Haiman, Z. 2019. Learning from the machine. *Nature Astronomy*, 3: 18–19.
- Hartmanis, J., & Stearns, R. E. 1965. On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, 117: 285–306.
- Helbing, D., Frey, B. S., Gigerenzer, G., Hafen, E., Hagner, M., Hofstetter, Y., van den Hoven, J., Zicari, R. V., & Zwitter, A. 2019. Will democracy survive big data and

- artificial intelligence? In D. Helbing (Ed.), *Towards digital enlightenment*: 73–98. Cham, Switzerland: Springer.
- Hoc, J.-M. 2001. Towards a cognitive approach to human-machine cooperation in dynamic situations. *International Journal of Human-Computer Studies*, 54: 509–540.
- Holzinger, A. 2016. Interactive machine learning for health informatics: When do we need the human-in-the-loop? *Brain Informatics*, 3: 119–131.
- Huang, H.-H., Hsu, J. S.-C., & Ku, C.-Y. 2012. Understanding the role of computer-mediated counter-argument in countering confirmation bias. *Decision Support Systems*, 53: 438–447.
- IBM Think Blog. 2017. Transparency and trust in the cognitive era. *IBM Think Blog*. <https://www.ibm.com/blogs/think/2017/01/ibm-cognitive-principles>. January 17.
- Jaeger, H. 2016. Deep neural reasoning. *Nature*, 538: 467–468.
- Jarzabkowski, P., Bednarek, R., Chalkias, K., & Cacciatori, E. 2019. Exploring inter-organizational paradoxes: Methodological lessons from a study of a grand challenge. *Strategic Organization*, 17: 120–132.
- Jordan, M. I., & Mitchell, T. M. 2015. Machine learning: Trends, perspectives, and prospects. *Science*, 349: 255–260.
- Kellogg, K., Valentine, M., & Christin, A. In press. Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, doi: 10.5465/annals.2018.0174.
- Kurzweil, R. 2014. The singularity is near. In R. L. Sandler (Ed.), *Ethics and emerging technologies*: 393–406. London: Palgrave Macmillan.
- La Roche, J. 2017. IBM’s Rometty: The skills gap for tech jobs is “the essence of divide.” *Yahoo! Finance*. <https://finance.yahoo.com/news/ibms-rometty-skills-gap-tech-jobs-essence-divide-175847484.html>. November 16.
- Langley, P., & Simon, H. A. 1995. Applications of machine learning and rule induction.

Communications of the ACM, 38(11): 54–64.

Larrick, R. P. 2004. Debiasing. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making*: 316–338. Malden, MA: Blackwell.

Levinthal, D. A., & March, J. G. 1993. The myopia of learning. *Strategic Management Journal*, 14: 95–112.

Lichfield, G. 2018. The precision medicine issue. *MIT Technology Review*.

<https://www.technologyreview.com/s/612285/editors-letter-the-precision-medicine-issue>. October 23.

Lindebaum, D., Vesa, M., & den Hond, F. In press. Insights from *The Machine Stops* to better understand rational assumptions in algorithmic decision-making and its implications for organizations. *Academy of Management Review*, doi: 10.5465/amr.2018.0181.

Lüscher, L. S., & Lewis, M. W. 2008. Organizational change and managerial sensemaking: Working through paradox. *Academy of Management Journal*, 51: 221–240.

Lyon, D. 2003. *Surveillance as social sorting: Privacy, risk, and digital discrimination*. London: Routledge.

Marr, B. 2018. The amazing ways how Unilever uses artificial intelligence to recruit & train thousands of employees. *Forbes*.

<https://www.forbes.com/sites/bernardmarr/2018/12/14/the-amazing-ways-how-unilever-uses-artificial-intelligence-to-recruit-train-thousands-of-employees>. December 14.

McDuff, D., & Czerwinski, M. 2018. Designing emotionally sentient agents.

Communications of the ACM, 61(12): 74–83.

McFarland, D. A., Lewis, K., & Goldberg, A. 2016. Sociology in the era of big data: The ascent of forensic social science. *The American Sociologist*, 47: 12–35.

Miron-Spektor, E., Ingram, A., Keller, J., Smith, W. K., & Lewis, M. W. 2018.

- Microfoundations of organizational paradox: The problem is how we think about the problem. *Academy of Management Journal*, 61: 26–45.
- Mitchell, T. M. 1997. *Machine learning*. Boston, MA: McGraw-Hill.
- Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. 2016. The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2): 1–21.
- More, M., & Vita-More, N. 2013. *The transhumanist reader: Classical and contemporary essays on the science, technology, and philosophy of the human future*. Malden, MA: Wiley-Blackwell.
- Nadella, S. 2016. The partnership of the future. *Slate*.
<https://slate.com/technology/2016/06/microsoft-ceo-satya-nadella-humans-and-a-i-can-work-together-to-solve-societys-challenges.html>. June 28.
- Newell, A., Shaw, J. C., & Simon, H. A. 1959. Report on a general problem solving program. *International Conference on Information Processing*, 256–264. Santa Monica, CA: Rand Corporation.
- Newell, A., & Simon, H. 1956. The logic theory machine - A complex information processing system. *IRE Transactions on Information Theory*, 2: 61–79.
- Nilsson, N. J. 1971. *Problem-solving methods in artificial intelligence*. New York: McGraw-Hill.
- Norris, P. 2001. *Digital divide: Civic engagement, information poverty, and the internet worldwide*. Cambridge, UK: Cambridge University Press.
- O’Neil, C. 2016. *Weapons of math destruction: How big data increases inequality and threatens democracy*. New York: Crown.
- Orlikowski, W. J. 2007. Sociomaterial practices: Exploring technology at work. *Organization Studies*, 28: 1435–1448.
- Panait, L., & Luke, S. 2005. Cooperative multi-agent learning: The state of the art.

Autonomous Agents and Multi-Agent Systems, 11: 387–434.

Parasuraman, R., & Manzey, D. H. 2010. Complacency and bias in human use of automation:

An attentional integration. *Human Factors*, 52: 381–410.

Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. 2019. Continual lifelong

learning with neural networks: A review. *Neural Networks*, 113: 54–71.

Poole, M. S., & van de Ven, A. H. 1989. Using paradox to build management and

organization theories. *Academy of Management Review*, 14: 562–578.

Press, G. 2016. *Only Humans Need Apply* is a must-read on AI for Facebook executives.

Forbes. <https://www.forbes.com/sites/gilpress/2016/09/21/only-humans-need-apply-is-a-must-read-on-ai-for-facebook-executives>. September 21.

Putnam, L. L., Fairhurst, G. T., & Banghart, S. 2016. Contradictions, dialectics, and

paradoxes in organizations: A constitutive approach. *Academy of Management Annals*, 10: 65–171.

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. “Sandy,” Roberts, M. E., Shariff, A., Tenenbaum, J. B., &

Wellman, M. 2019. Machine behaviour. *Nature*, 568: 477–486.

Rai, A., Constantinides, P., & Sarker, S. 2019. Next-generation digital platforms: Toward

human-AI hybrids. *MIS Quarterly*, 43: iii–ix.

Raisch, S., Hargrave, T. J., & van de Ven, A. H. 2018. The learning spiral: A process

perspective on paradox. *Journal of Management Studies*, 55: 1507–1526.

Riley, T. 2018. Get ready, this year your next job interview may be with an A.I. robot.

CNBC. <https://www.cnbc.com/2018/03/13/ai-job-recruiting-tools-offered-by-hirevue-mya-other-start-ups.html>. March 13.

- Russell, S., & Norvig, P. 2009. *Artificial intelligence: A modern approach* (3rd ed.). Englewood Cliffs, NJ: Prentice-Hall.
- Sabherwal, R., & Jeyaraj, A. 2015. Information technology impacts on firm performance: An extension of Kohli and Devaraj (2003). *MIS Quarterly*, 39: 809–836.
- Schad, J., & Bansal, P. 2018. Seeing the forest and the trees: How a systems perspective informs paradox research. *Journal of Management Studies*, 55: 1490–1506.
- Schad, J., Lewis, M. W., Raisch, S., & Smith, W. K. 2016. Paradox research in management science: Looking back to move forward. *Academy of Management Annals*, 10: 5–64.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., & Hassabis, D. 2017. Mastering the game of Go without human knowledge. *Nature*, 550: 354–359.
- Simon, H. A. 1987. Two heads are better than one: The collaboration between AI and OR. *Interfaces*, 17: 8–15.
- Skitka, L. J., Mosier, K., & Burdick, M. D. 2000. Accountability and automation bias. *International Journal of Human-Computer Studies*, 52: 701–717.
- Sloan, M. C. 2010. Aristotle’s *Nicomachean Ethics* as the original *Locus* for the *Septem Circumstantiae*. *Classical Philology*, 105: 236–251.
- Smith, W. K., & Lewis, M. W. 2011. Toward a theory of paradox: A dynamic equilibrium model of organizing. *Academy of Management Review*, 36: 381–403.
- Smith, W. K., & Tracey, P. 2016. Institutional complexity and paradox theory: Complementarities of competing demands. *Strategic Organization*, 14: 455–466.
- Staw, B. M. 1981. The escalation of commitment to a course of action. *Academy of Management Review*, 6: 577–587.
- Stephan, M., Brown, D., & Erickson, R. 2017. Talent acquisition: Enter the cognitive

- recruiter. *Deloitte Insights*. <https://www2.deloitte.com/insights/us/en/focus/human-capital-trends/2017/predictive-hiring-talent-acquisition.html>. February 28.
- Sterman, J. 2000. *Business dynamics: Systems thinking and modeling for a complex world*. Boston, MA: Irwin/McGraw-Hill.
- Taylor, M. E., & Stone, P. 2009. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(1): 1633–1685.
- The Economist. 2017. Harvard Business School risks going from great to good. *The Economist*. <https://www.economist.com/business/2017/05/04/harvard-business-school-risks-going-from-great-to-good>. May 4.
- Topol, E. J. 2019. High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25: 44–56.
- Tsoukas, H. 2017. Don't simplify, complexify: From disjunctive to conjunctive theorizing in organization and management studies. *Journal of Management Studies*, 54: 132–153.
- Walsh, T. 2017. The singularity may never be near. *AI Magazine*, 38(3): 58–62.
- Webb, S. 2018. Deep learning for biology. *Nature*, 554: 555–557.
- Weick, K. E. 2007. The generative properties of richness. *Academy of Management Journal*, 50: 14–19.
- Westcott Grant, K. 2018. Netflix's data-driven strategy strengthens claim for “best original content” in 2018. *Forbes*. <https://www.forbes.com/sites/kristinwestcottgrant/2018/05/28/netflixs-data-driven-strategy-strengthens-lead-for-best-original-content-in-2018>. May 28.
- Wladawsky-Berger, I. 2018. *Human + Machine*: The impact of AI on business transformation. *Wall Street Journal*. <https://blogs.wsj.com/cio/2018/04/13/human-machine-the-impact-of-ai-on-business-transformation>. April 13.
- Zuboff, S. 1988. *In the age of the smart machine: The future of work and power*. New

York: Basic Books.

Zuboff, S. 2019. *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. New York: PublicAffairs.

Sebastian Raisch (sebastian.raisch@unige.ch) is Professor of Strategy at GSEM, University of Geneva, Switzerland. He received his Ph.D. in Management from the University of Geneva and his habilitation from the University of St. Gallen. His research interests include artificial intelligence, organizational ambidexterity, and organizational paradox.

Sebastian Krakowski (sebastian.krakowski@hhs.se) is a Ph.D. candidate at GSEM, University of Geneva, and a research associate at the House of Innovation, Stockholm School of Economics, Sweden. He was a visiting researcher at Warwick Business School. His research explores artificial intelligence's behavioral implications in organizations and society.