

December 2020

## Special Issue Editorial: Artificial Intelligence in Organizations: Current State and Future Opportunities

Hind Benbya

Thomas H. Davenport

Stella Pachidi

Follow this and additional works at: <https://aisel.aisnet.org/misqe>

---

### Recommended Citation

Benbya, Hind; Davenport, Thomas H.; and Pachidi, Stella (2020) "Special Issue Editorial: Artificial Intelligence in Organizations: Current State and Future Opportunities," *MIS Quarterly Executive*: Vol. 19 : Iss. 4 , Article 4.

Available at: <https://aisel.aisnet.org/misqe/vol19/iss4/4>

This material is brought to you by the AIS Journals at AIS Electronic Library (AISeL). It has been accepted for inclusion in MIS Quarterly Executive by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact [elibrary@aisnet.org](mailto:elibrary@aisnet.org).

# EDITORS' COMMENTS

## Special Issue Editorial: Artificial Intelligence in Organizations: Current State and Future Opportunities

### Introduction

Artificial intelligence (AI) is typically defined as the ability of machines to perform human-like cognitive tasks, including the automation of physical processes such as manipulating and moving objects, sensing, perceiving, problem solving, decision making and innovation.<sup>1</sup> AI is currently viewed as the most important disruptive new technology for large organizations.<sup>2</sup> However, the technology is still in a relatively early state in large enterprises and largely absent from smaller ones other than technology startups. Surveys<sup>3</sup> suggest that fewer than half of large organizations have meaningful AI initiatives underway, although the percentage is increasing over time.

For most organizations, AI projects remain somewhat experimental—undertaken as a pilot or proof-of-concept initiative. Relatively few organizations have deployed AI on a production basis, a problem that we describe in greater detail below. Of course, the experimental nature of its use means that many organizations have achieved little or no economic return on their AI investments. However, some analysts<sup>4</sup> suggest that AI adoption will eventually have a

considerable positive impact on company growth and profitability.

A 2020 Deloitte survey of executives revealed that AI is currently being applied in organizations to support the following diverse objectives:<sup>5</sup> making processes more efficient (28%), enhancing existing products and services (25%), creating new products and services (23%), improving decision making (21%) and lowering costs (20%). Although reducing headcount is a common objective cited in AI-oriented press, it was mentioned least in this survey (11%).

Executives initially focused on using AI technologies to automate specific workflow processes and repetitive work. Such processes were linear, stepwise, sequential and repeatable. Currently, however, firms are moving toward nonsystematic cognitive tasks that include decision making, problem solving and creativity, which, until recently, seemed beyond the scope of automation. AI technologies are also progressively enabling people and machines to work collaboratively in novel ways. In manufacturing, for example, in order to fulfill customized orders and handle fluctuations in demand, employees are partnering with robots to perform new tasks without having to manually overhaul any processes. AI technologies are also performing certain tasks autonomously, though complex tasks like driving a car in all conditions remain tantalizingly out of reach.

However, autonomous systems are beginning to appear that can perform tasks without any human involvement at all because systems are capable of training themselves and adjusting to new training data. Consider automated financial trading: because it depends entirely on algorithms, companies can complete transactions much more quickly with AI systems than with systems relying on humans. In a similar fashion, robots are performing narrow tasks autonomously in manufacturing settings.<sup>6</sup>

1 Innovation is defined here as the design, creation, development and/or implementation of new or altered products, services, systems, organizational structures, management practices and processes, or business models; see Benbya, H. and Leidner, D. (2018) "How Allianz UK Used an Idea Management Platform to Harness Employee Innovation," *MIS Quarterly Executive* (17:2), 2018, pp. 141-157; Yan, J. Leidner, D. and Benbya, H. "Differential Innovativeness Outcomes of User and Employee Participation in an Online User Innovation Community," *Journal of Management Information Systems* (35:3), 2018, pp. 900-933.

2 "Big Data and AI Executive Survey 2019, Executive Summary of Findings," *NewVantage Partners*, 2019, <https://newvantage.com/wp-content/uploads/2018/12/Big-Data-Executive-Survey-2019-Findings-Updated-010219-1.pdf>

3 "AI 360: Hold, Fold, or Double Down," *Genpact*, 2020, <https://www.genpact.com/uploads/files/ai-360-research-2020.pdf>

4 "Notes from the AI Frontier: Modeling the Impact of AI on the World Economy," *McKinsey & Co.*, 2018, <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy#>

5 "Thriving in the Era of Pervasive AI: Deloitte's State of AI in the Enterprise, 3rd Edition," *Deloitte Insights*, Deloitte, 2020, <https://www2.deloitte.com/us/en/insights/focus/cognitive-technologies/state-of-ai-and-intelligent-automation-in-business-survey.html>

6 Davenport, T., & Ronanki, R., Artificial Intelligence for the Real World. *Harvard Business Review* (96:1), 2018, 108-116.

Companies such as Amazon and Google have attempted to create highly ambitious applications of AI, including autonomous vehicles, unattended retail checkout and drone delivery. Some of these “moon shots” have been successful, but some highly ambitious projects, including cancer treatment, have been largely unsuccessful thus far despite considerable expenditures. Less ambitious “low-hanging fruit” projects have been more successful in many firms and are perhaps more consistent with the current narrow intelligence of AI systems.

Likewise, most autonomous AI applications remain limited to low-risk areas with limited costs associated with failure. Although many AI systems can do certain things better than humans, workers’ trust in AI technology is still limited because of the issues associated with this technology, such as algorithmic bias, unexplainable outcomes, privacy invasions and/or lack of accountability. Consumers are also skeptical about AI and surveys suggest that most or many would not want autonomous vehicles, dislike dealing with chatbots and so forth.

This December 2020 *MIS Quarterly Executive* special issue is titled “AI in Organizations: Current State and Future Opportunities.” It details current challenges and implications that may arise from AI applications and ways to overcome such challenges to realize the potential of this emerging technology. The collection of papers in this issue (December), combined with a forthcoming (March) article, offers insights to managers currently implementing digital transformation initiatives driven by AI technology, to practitioners considering implementing AI in their businesses and to research-oriented faculty and students. In this editorial, we first provide a brief history of AI and an overview of AI typologies. We discuss the current challenges, implications and future opportunities of AI so that readers are better equipped to understand the five papers in the special issue. Finally, we summarize the special issue articles and highlight the contributions each makes.

## Brief History of AI

AI as an academic field dates back to the 1950s. The term AI was first introduced during a multidisciplinary program presented

at Dartmouth College in 1956. The program aimed to study the possibility that machine intelligence could imitate humans and involved researchers from various fields including scientists, mathematicians and philosophers. Despite early promises of the practical usefulness of AI, it largely failed to deliver and faced several obstacles during the 1960s and 1970s, the biggest of which was the lack of computational power to do anything substantial. Research funding gradually stalled and the field lost momentum. During the 1980s and 1990s, governments and firms made significant investments in research on expert systems, which revived interest in AI. Machine learning and neural networks began to flourish as practitioners integrated statistics and probability into their applications. At the same time, the personal computing revolution began. Over the next decade, digital systems, sensors, and the internet proliferated, providing all kinds of data for machine-learning experts to use when training adaptive systems. Although the growth of AI and machine learning has been intermittent over the decades, unprecedented computing capacity and growing volumes of data have provided momentum for the recent development of artificial intelligence applications.

## AI Types and technologies

There are many types of AI systems. One typology differentiates AI systems based on the kind of intelligence they display. A second typology distinguishes AI applications based on the type of technology embedded into the AI system, whereas a third is based on the functions performed by the AI.

**Based on intelligence:** Philosophical debates on AI are centered on the notion of intelligent machines—that is, machines that can learn, adapt and think like people.<sup>7</sup> AI types based on such a notion fall, in general, into three categories: artificial narrow intelligence, artificial general intelligence and artificial superintelligence.

While narrow (or weak) AI is usually able to solve only one specific problem and is unable to transfer skills from domain to domain, general AI aims for a human-level skill set. Once general

<sup>7</sup> Lake, B., Ullman, T., Tenenbaum, J. and Gershman, J. “Building Machines that Learn and Think Like People,” *Behavioral and Brain Sciences* (40), 2017, e253

**Table 1: AI Technologies and Domains of Application**

Technology	Brief Description	Example Application
<b>Machine learning</b> <ul style="list-style-type: none"> <li>Reinforcement learning</li> <li>Supervised learning</li> <li>Unsupervised learning</li> </ul>	Learns from experience. Learns from a set of training data. Detects patterns in data that are not labeled and for which the result is not known.	Highly granular marketing analyses on big data.
<b>Deep learning</b>	A class of machine learning that learns without human supervision, drawing from data that is both labeled and unlabeled.	Image and voice recognition, self-driving cars.
<b>Neural networks</b>	Algorithms that endeavor to recognize underlying relationships in a set of data through a process that mimics the way the human brain operates.	Credit and loan application evaluation, weather prediction.
<b>Natural language processing</b>	A computer program able to understand human language as it is written or spoken.	Speech recognition, text analysis, translation, generation.
<b>Rule-based expert systems</b>	A set of logical rules derives from human experts.	Insurance underwriting, credit approval
<b>Robotic process automation</b>	Systems that automate structured digital tasks and interfaces.	Credit card replacement, validating online credentials.
<b>Robots</b>	Automatically operated machines that automate physical activity, manipulate and pick up objects.	Factory and warehouse tasks.

AI is achieved, it is believed that it might lead to superintelligence that exceeds the cognitive performance of humans in virtually all domains of interest.<sup>8</sup> This type of superintelligence could emerge following evolutionary and complex adaptive systems principles.<sup>9</sup> If humans could create AI intelligence at a roughly human level, so the argument goes, then this creation could, in turn, create yet higher intelligence and eventually evolve further.<sup>10</sup> AI enthusiasts have

provided estimates and outline scenarios for when technological growth will reach the point of singularity at which machine intelligence will surpass human intelligence. This raises philosophical arguments about the mind and the ethics of creating artificial beings endowed with human-like intelligence. Although the futuristic literature assumes that AI systems will be able to perform all tasks as well as or even better than humans, this type of artificial general intelligence does not yet exist. There are, however, some AI programs, such as the GPT-3 language prediction application, that are beginning to exhibit some aspects of general intelligence.<sup>11</sup>

**Based on technology:** A second typology differentiates between the technologies that

8 Bostrom, N. (2014). *Superintelligence Paths, Dangers, Strategies*. Oxford University Press.

9 See Benbya, H. Nan, N., Tanriverdi, H. and Yoo, Y. "Complexity and Information Systems Research in the Emerging Digital World," *MIS Quarterly* (44:1), 2020, pp. 1-18, for a recent article on evolutionary principles, and Benbya, H. and McKelvey, B. "Using Coevolutionary and Complexity Theories to Improve IS Alignment: A Multi-Level Approach," *Journal of Information Technology* (21:4), 2006, pp. 284-298 for an elaboration of such principles in IT management.

10 Hawking, S., Russell, S., Tegmark, M., & Wilczek, F. "Stephen Hawking: Transcendence Looks at the Implications of Artificial Intelligence: But Are We Taking AI Seriously Enough?," *The Independent*, May 5, 2014, <https://www.independent.co.uk/news/science/stephen-hawking-transcendence-looks-implications-artificial-intelligence-are-we-taking-ai-seriously-enough-9313474.html>.

11 GPT-3 stands for generative pre-trained transformer version three. It is a powerful machine-learning system that can rapidly generate text with minimal human input. After an initial prompt, it can recognize and replicate patterns of words to work out what comes next, see Thierry, G. "New AI Can Write Like a Human but Don't Mistake that for Thinking," *The Conversation*, September 17, 2020, <https://theconversation.com/gpt-3-new-ai-can-write-like-a-human-but-dont-mistake-that-for-thinking-neuroscientist-146082>

are embedded into AI systems, which include machine learning (and its subclasses deep learning and reinforcement learning), natural language processing, robots, various automation technologies (including robotic process automation) and rule-based expert systems (still in broad use although not considered a state-of-the-art technology). One recent survey<sup>12</sup> suggests that all contemporary AI technologies (machine learning, deep learning, natural language processing) are either currently being used or will be used within a year by 95% or more of large adopters of AI. Table 1 below provides brief definitions and the domain of AI technology applications.

**Based on function:** This distinction differentiates between four types of AI: conversational, biometric, algorithmic and robotic. These categories overlap somewhat; for example, conversational and biometric AI already make extensive use of algorithmic AI models, and robotic AI is increasingly doing so as well.

*Conversational AI* refers to the general capability of computers to understand and respond using natural human language. Such systems include both voice- and text-based technologies and vary largely based on their capability, domain and level of embodiment. Simple conversational AI is mainly used to handle repetitive client queries whereas smart conversational AI, enabled by machine learning and natural language processing, has the potential to undertake more complex tasks that involve greater interaction, reasoning, prediction and accuracy. Conversational AI has been used in many different fields, including finance, commerce, marketing, retail and healthcare. Although the technology behind smart conversational agents is continuously under development, they currently do not have full human-level language abilities, sometimes resulting in misunderstanding and user dissatisfaction.<sup>13</sup>

*Biometric AI:* Biometrics relies on techniques to measure a person's physiological (e.g., fingerprints, hand geometry, retinas, iris, facial image) or behavioral traits (e.g., signature, voice, keystroke rhythms). AI-powered biometrics

uses applications such as facial recognition, speech recognition and computer vision for identification, authentication, and security objectives in computer devices, the workplace, home security, etc. While fingerprints have the longest history as a marker of identity and continue to be used in a number of applications across the world,<sup>14</sup> other bodily markers, such as the face, voice and iris or retina, are proliferating, with significant research exploring their potential large-scale application. Meanwhile, the ubiquity of facial images and voice recordings tagged with people's names on the internet, alongside algorithms to transform such data into biometric recognition systems, has accelerated the use of such data at a larger scale—for example, to identify suspects, monitor large events and surveil protests. Such large-scale use has triggered calls for regulation to introduce new laws, reform existing laws, or ban the use of such data in some contexts.

*Algorithmic AI* revolves around the use of machine learning (ML) algorithms—a set of unambiguous instructions that a mechanical computer can execute. Some ML algorithms can be trained on structured data and are specific to narrow task domains, such as speech recognition and image classification. Other algorithms, especially deep learning neural networks, can learn from large volumes of labeled data, enhance themselves by learning, and accomplish a variety of tasks such as classification, prediction and recognition. For example, neural networks can analyze parameters of bank clients such as age, solvency and credit history, and decide whether to approve loan requests. Such networks can also use face recognition to let only authorized people into a building or predict outcomes such as the rise or fall of a stock based on past patterns and current data. Despite the potential of ML algorithms, there are concerns that, in some cases, it may not be possible to explain how a system has reached its output. Such algorithms may also be susceptible to introducing or perpetuating discriminatory bias.

*Robotic AI:* Physical robots have been used for many years to perform dedicated tasks in factory automation. Recently, AI (including ML

12 “Thriving in the Era of Pervasive AI,” op. cit., 2020.

13 “What is a Chatbot? All You Need to Know About Chatbots!” *Botpress: Open-Source Conversational AI Platform*, 2018, <https://botpress.io/learn/what-and-why/>

14 Amba K., ed. “Regulating Biometrics: Global Approaches and Urgent Questions,” *AI Now Institute*, September 1, 2020, <https://ainowinstitute.org/regulatingbiometrics.html>.

and NLP) has become increasingly present in robotic solutions, enabling robots to move past automation and tackle more complex and high-level tasks. AI-enabled robots are equipped with the ability to sense their environment, comprehend, act and learn. This helps robots perform many tasks by successfully navigating their surroundings, identifying objects around them and assisting humans with various tasks, such as robot-assisted surgeries.

## Current Challenges

### AI's Deployment Problem

One of the major concerns with AI in organizations at present is that many systems are only experimental and never deployed in production. Pilot AI projects are relatively easy to develop and are only intended to demonstrate that the technology is feasible in concept. Deployment, on the other hand, requires a variety of tasks and capabilities that may be in short supply—for example, integration with existing technology architectures and legacy infrastructure, change in business processes and organizational culture, reskilling or upskilling of employees, substantial data engineering and approaches to organizational change management. Full production deployment tends to take much longer than pilot projects and cost substantially more.

Surveys of organizations and market research reports in the U.S. and globally suggest that deployment challenges involving big data and AI are widespread. A 2019 survey<sup>15</sup> of large financial services and life sciences firms found that firms are actively embracing AI technologies and solutions, with 91.5% of firms reporting ongoing investment in AI. However, only 14.6% of firms reported that they have deployed AI capabilities for widespread production. A 2019 global McKinsey survey reported under the headline “AI Adoption Proves its Worth, but Few Scale Impact” indicates that between 12% (consumer packaged goods firms) and 54% (high tech firms) of firms have at least one machine learning application implemented in a process or product, but only 30% of respondents overall

reported using AI in products or processes across multiple business units and functions.<sup>16</sup>

In order to address such deployment concerns, companies need to plan for the possibility of deployment from the beginning. Some companies, such as Farmers Insurance, have a well-defined process that, when appropriate, seeks to move projects from the pilot phase to full deployment.<sup>17</sup> In a survey of early U.S. adopter organizations, 54% of executives said that their organization has a process for moving prototypes into production and 52% reported having an implementation road map. Such organizational approaches may facilitate moving more AI systems into deployment but attempts may be only in their early stages.

### AI Talent Issues

Securing a sufficient volume and level of human AI talent is a challenge for many organizations—particularly those that are not in the technology sector. Data scientists and AI engineers are still scarce, although many university programs have arisen to train them. Firms that are unable to pay high levels of compensation and are not located in technology centers are likely to have difficulty hiring the desired number of skilled employees. Many companies should attempt to not only hire new employees with AI skills but to retrain existing employees to the degree possible. Even when companies do manage to hire data scientists and other types of analytical and artificial intelligence talent, there is little consensus within and across companies about the qualifications for such roles. The term “data scientist” might mean a job with a heavy emphasis on statistics, open-source coding, or working with executives to solve business problems with data and analysis. Some view the role as focused only on developing models, while others view it as extending to the deployment of models in production. The idea of

<sup>16</sup> “State of AI in the Enterprise, 2nd Edition,” *Deloitte Insights*, 2018, [https://www2.deloitte.com/content/dam/insights/us/articles/4780\\_State-of-AI-in-the-enterprise/DI\\_State-of-AI-in-the-enterprise-2nd-ed.pdf](https://www2.deloitte.com/content/dam/insights/us/articles/4780_State-of-AI-in-the-enterprise/DI_State-of-AI-in-the-enterprise-2nd-ed.pdf)

<sup>17</sup> Davenport, T. and Bean, R. “Farmers Accelerates its Time to Impact with AI.” *Forbes*, August 1, 2018, <https://www.forbes.com/sites/tomdavenport/2018/08/01/farmers-accelerates-its-time-to-impact-with-ai/#51430150b672>

<sup>15</sup> “Big Data and AI Executive Survey,” op cit., 2019.

data scientist “unicorns” who possess high levels of all these skills has never been very realistic.<sup>18</sup>

The skills taught in university programs aimed at training AI-oriented workers vary widely and some universities offer multiple programs with different emphases. For both newly hired and experienced employees, titles such as data scientist and AI engineer are not likely to be a good guide for understanding actual capabilities. Further, activities involved in the deployment of AI systems and related organizational change issues may not be taught at all by many technically focused programs. There is an increasing need for a new type of professional who can understand business problems and translate them into algorithmic problems and, vice versa, explain technical insights to business managers.<sup>19</sup>

Initiatives<sup>20</sup> exist in the early stages to standardize the different types of data, analytics, AI roles and requisite skills across organizations; however, developing new standards typically takes many years. In the meantime, companies should devote significant attention to classifying and certifying the different types of AI and data science jobs needed in their organizations. Companies would also benefit from expanding their talent pool by working with universities directly on educational programs and by building and nurturing communities within their organizations for employees working in data teams. These steps are essential for companies looking to use AI to both improve current operations and expand opportunities for digital innovation.

### AI and Social Dysfunctions

Aside from deployment and talent challenges, there are a few other potential AI dysfunctions that managers should be aware of and make plans to avoid.

*Algorithmic bias:* The employment of AI systems in classification or prediction tasks often comes with the risk of algorithmic bias, which means that the outcomes of the machine

learning algorithm can put certain groups at a disadvantage.<sup>21</sup> This has already been observed in various cases, including algorithms that are used to score job applicants and appear to be racist, or algorithms that recommend sentences to judges and appear to propagate the preconceptions implicit in past sentencing decisions that were used as training data. Algorithmic bias can also lead to consequences distributed across large subsections of society by affecting the type of information to which people are exposed. This happens, for example, when machine learning algorithms behind social media propagate fake news or enable the targeting of individuals for political campaigns. To reduce potential algorithmic bias, managers will need to be proactive by performing small-scale experiments and simulations before implementing such algorithms, regularly evaluating the dataset used for training, and using human reviewers who regularly provide feedback to system designers. In politically and socially sensitive domains like judicial sentencing, firms may find it necessary to publish their algorithms to preclude accusations of bias.

*Unexplainable decision outcomes:* Potential social dysfunctions resulting from AI implementation may be caused by decision outcomes of some machine learning algorithms—deep learning in particular—that cannot be easily explained because of the vast number of feature layers involved in their production. This may lead to problematic situations, such as unexplainable evaluations of high school teachers, or parole decisions that cannot be justified and may appear to be unfair.<sup>22</sup> Organizations need to respond to regulators’ calls for explainability by avoiding “black box” AI applications and by choosing algorithms whose outcomes can be explained. Being open about the data that is used and explaining how the model works in nontechnical terms is also necessary to ensure customers’ trust and to avoid potential dysfunctions triggered by a lack of transparency. Indeed, in some industries such as banking, regulators often force firms to use explainable algorithms.

18 Davenport, T. “Beyond Unicorns: Educating, Classifying, and Certifying Data Science Talent,” *Harvard Data Science Review*, May 19, 2020, <https://hdsr.mitpress.mit.edu/pub/t37qjoi7/release/2>

19 Henke, N., Levine, J. and McNerny, P. “You Don’t Have to Be a Data Scientist to Fill this Must-Have Role,” *Harvard Business Review*, Feb. 5, 2018, <https://hbr.org/2018/02/you-dont-have-to-be-a-data-scientist-to-fill-this-must-have-analytics-role>

20 See, for example, <https://www.iadss.org/>

21 Davenport, T. *The AI Advantage: How to Put the Artificial Intelligence Revolution to Work*, 2018, MIT Press.

22 O’Neil, C. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, 2016, Broadway Books.

*Blurring accountability boundaries:* As AI is used to enhance or even automate decision-making procedures, the issue of accountability arises. Who is responsible if a traffic accident involving a driverless car occurs? Who is responsible for approving parole to a criminal who eventually commits another crime? Who is responsible for big financial losses incurred by algorithmic trading? These are only a few examples of cases in which accountability boundaries are blurred. Managers will need to proactively focus on the reasons and processes that may lead to potential harm. They should also carefully consider how they engage the different actors that directly or indirectly interact with outcomes produced by AI systems (AI developers and designers, business users, institutions) and clarify responsibility and legal liability upfront.<sup>23</sup>

*Invaded privacy:* Ethical issues arise even before any action is recommended or performed by AI systems, with privacy being reported as one of the main ethical considerations underlying AI implementation.<sup>24</sup> Data is the primary resource fed into AI systems and is often perceived to be a source of competitive advantage. AI's need to process increasingly large amounts of data thus conflicts with the right to maintain control over data and its use in order to preserve privacy and autonomy. Organizations need to ensure that their data practices comply with relevant policies on the use of personal data (e.g., GDPR in EU countries) and avoid potential privacy violations. Developing auditable algorithms and performing algorithmic audits on them to identify what data is used and what variables feed into decision-making processes represent helpful solutions that can be used to increase transparency regarding how consumer data is processed and used.<sup>25</sup> Overall, openness about how data is handled is essential for ensuring customer trust.

23 Dourish, P. "Algorithms and their Others: Algorithmic Culture in Context," *Big Data & Society*, 3(2), 2016, 1-11.

24 Kinni, T. "Ethics Should Precede Action in Machine Intelligence." *MIT Sloan Management Review*, June 29, 2017, <https://sloanreview.mit.edu/article/ethics-should-proceed-action-in-machine-intelligence>

25 Mittelstadt, B. "Automation, Algorithms, And Politics | Auditing for Transparency in Content Personalization Systems," *International Journal of Communication* (10), 2016, Article 12.

## Implications

Despite the existing challenges, AI has the potential to dramatically change how the workforce is structured, how jobs are designed, how knowledge is managed and how decisions are made. Such changes will have broader implications for organizations and societies, many of which have yet to be understood or realized. However, the most common effects are likely to be on how work is conducted in the future.

**AI and the future of work:** Recent developments in AI are already affecting the workplace in different ways.

*Automating work tasks:* AI will have a significant impact on several occupations by automating mundane tasks and rendering various human skills obsolete. Given that AI can perform tasks that previously required human judgment, the effects of AI-enabled automation differ from those of past technologies, particularly regarding their impact on knowledge workers.<sup>26</sup> AI introduces new threats to the authority of professionals such as physicians, lawyers, consultants and architects, whose expertise, judgment and creativity have thus far been highly valued and considered irreplaceable. While the need for such professions will not disappear in the near future, the changing nature of such work is already a reality. There are many predictions about how AI will impact such work, but thus far job losses have been relatively minor.<sup>27</sup>

*Changing expertise:* AI technology that is capable of automating some tasks is already active in the workplace. In law firms, for example, a plethora of applications have been developed for automating due diligence and contract review tasks that were previously performed by junior lawyers. In sales, conversational AI can now automate various tasks that previously had

26 See Davenport, T. *Thinking for a Living: How to Get Better Performance and Results from Knowledge Workers*, 2015, Harvard Business School Press; Benbya, H., "Knowledge Management Systems Implementation: Lessons from the Silicon Valley," 2018, Neal-Schuman Publishers; and Faraj, S., Pachidi, S. and Sayegh, K. "Working and Organizing in the Age of the Learning Algorithm," *Information and Organization* (28:1), 2018, pp. 62-70.

27 For one prominent prediction, see Frey, C. B. and Osborne, M. *The Future of Employment: How Susceptible Are Jobs to Computerization*, Oxford Martin School working paper, 2013, <https://www.oxfordmartin.ox.ac.uk/downloads/academic/future-of-employment.pdf>

to be carried out by account managers. While such automation can increase the efficiency of operations and decrease labor costs, they threaten to create gaps for professionals by automating processes that had previously been used to acquire knowledge about customers or to developed expertise. This will eventually lead to changes in the knowledge base required for affected occupations and could potentially even trigger their restructuring. For example, in the legal profession, law graduates now often find it necessary to develop data science skills and technical skills instead of following the traditional career path of a lawyer.

*Augmenting professionals:* In many cases, AI systems are not yet able to replace human experts but can augment human work by supporting experts' decision-making processes. For example, in the medical field, while there was initially some concern that AI could fully replace radiologists, it is now clear that the role of AI will be to augment the work of radiologists.<sup>28</sup> Nevertheless, as AI systems are introduced to support the diagnostic capabilities of radiologists, several unintended consequences have arisen—for example, radiologists may confront communication barriers in interactions with data scientists, and conflicts between AI and their own diagnoses may cause radiologists to question their own capabilities versus those of the AI.<sup>29</sup> This becomes even more complicated considering the potential inscrutability of machine learning algorithm functions, meaning that specific outcomes often cannot be easily traced or explained. In any case, the nature of work is changing dramatically and, while many observers predict that the combination of human and machine intelligence will always reign supreme,<sup>30</sup> we have yet to see how “augmented professionals” will carry out their work and what further implications will arise for the workplace, organizations and institutions.

28 Davenport, T. and Dreyer K. “AI will Change Radiology, but it Won't Replace Radiologists”, *Harvard Business Review*, March 27, 2018, <https://hbr.org/2018/03/ai-will-change-radiology-but-it-wont-replace-radiologists>

29 Lebovitz, S. Lifshitz-Assaf, H. and Levina, N. *To Incorporate or Not to Incorporate AI for Critical Judgments: The Importance of Ambiguity in Professionals' Judgment Process*, NYU Stern School of Business, January 15, 2020, available at SSRN: <https://ssrn.com/abstract=3480593>

30 Brynjolfsson, E. and McAfee, A. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, 2014, Norton.

## Organizational Implications

The introduction of AI is associated with significant changes in how organizations are managed.

*Changing authority arrangements:* Unavoidably, as discussed above, expertise is redefined and the knowledge and skills of technology practitioners such as machine learning experts, data scientists and data analysts will become increasingly valued in the workplace, which may lead to a restructuring of authority arrangements across all levels of an organization. On a tactical level, technology practitioners will likely gain authority and control over work design and decision-making processes, given that they have the ability to prescribe how AI systems will affect operations. On a more strategic level, new roles will be added to boards and upper echelons, triggering questions related to the established regime: For example, where does the CIO's jurisdiction end and the CDO's start when it comes to planning a major digital transformation through implementing AI technology?

*Changing coordination:* The use of AI to algorithmically manage work will lead to fundamental changes in organizational design and coordination. For instance, work tasks must be redefined so that they can be broken down into smaller subtasks that can then be algorithmically assigned to workers on digital labor platforms such as UpWork or Amazon MTurk.<sup>31</sup> Machine learning algorithms can be used to coordinate more proactively by analyzing historical data to predict the need for skills and expertise in future projects. Furthermore, practitioners and managers will need to collaborate with new experts entering the workplace in data processing, algorithm development, data visualization and so forth. Collaboration among individuals with different types of expertise can make work coordination more challenging and may result in a substantially different execution of an organization's operations and services.

*Changing valuation schemes:* The way that performance is evaluated is also changing substantially, as employees are now being

31 Faraj, S., Pachidi, S. and Sayegh, K. “Working and Organizing in the Age of the Learning Algorithm,” *Information and Organization* (28:1), 2018, pp. 62-70.

assessed by machine learning algorithms, often with little understanding of the variables included in the inherent model or the extent to which a specific variable might determine a specific outcome. Even product quality checks are increasingly becoming an automated task performed by robots.<sup>32</sup> Such fundamental changes in the values that matter to an organization not only substantially impact how firms manage their employees but can also lead to countereffects on the employee side. For example, Amazon delivery drivers have reportedly begun hanging phones on trees outside of dispatch stations because the closer proximity to the dispatch network allows these drivers to receive delivery requests a split-second sooner than other drivers.<sup>33</sup> This is a fascinating demonstration of how employees try to game algorithmic systems in their effort to maintain some control over their work.

*Industrial transformations:* AI technology is currently enabling significant digital transformations that not only redefine what an organization does but also blur industry boundaries. Traditional manufacturing organizations are currently taking advantage of machine learning technology to transform their focus from the production of goods to the provision of services. GE's digital transformation effort is a well-known example of such an attempt, with AI being the driving force behind their predictive maintenance services. Such developments invoke a number of questions that must be addressed: Who are the new competitors of such digitally transformed organizations? How should such organizations be regulated? How do relationships with customers change? Who are the new partners of such organizations?

## Future Opportunities

As companies continue to use AI, they will explore a variety of different directions. We suggest several of them below.

*Management and governance mechanisms:* Leading companies using AI already have

32 Mahdawi, A. "The Domino's 'Pizza Checker' Is Just the Beginning: Workplace Surveillance Is Coming for You," *The Guardian*, October 15, 2019.

33 Soper, S. "Amazon Drivers Are Hanging Smartphones in Trees to Get More Work," *Bloomberg*, September 1, 2020, <https://www.bloomberg.com/news/articles/2020-09-01/amazon-drivers-are-hanging-smartphones-in-trees-to-get-more-work>.

management and governance mechanisms in place. We have already mentioned those related to deployment. In addition, organizations in various states of adoption have put in place a wide range of internal organizational structures and roles to manage and govern AI projects. The results of a 2018 survey<sup>34</sup> illuminate a number of governance mechanisms used to manage AI projects: respondents indicated that 45% of firms had appointed AI champions, 37% of firms had created an AI center of excellence,<sup>35</sup> and 37% of firms had developed a comprehensive strategy for AI.<sup>36</sup>

*Democratization of data science and AI:* Tools like automated machine learning<sup>37</sup> can structure and automate the workflow of creating and implementing a machine learning model. Such tools can be employed to improve the productivity of professional data scientists or to enable less highly educated "citizen data scientists" to complete data science and AI projects. Several startups and large cloud vendors have made such capabilities available and it seems likely that the democratization of data science and AI development—i.e., the notion that anyone, even those with little to no expertise, can perform data science if provided ample data and user-friendly analytics tools—will continue to advance.

*Ongoing model improvement:* Companies that are heavily committed to AI often find that they have many models and algorithms in place, some of which are in production processes and systems. Since businesses are, to some extent, dependent on the accuracy of these models, it is important to monitor them for "drift" (i.e., inaccuracy of predictions) and improve them over time. Vendors are developing tools to support this process under the banner of MLOps (machine learning operations), which are most widely used in data- and analytics-dependent industries like financial services.

34 Davenport, T. *The AI Advantage: How to Put the Artificial Intelligence Revolution to Work*, 2018, MIT Press.

35 Davenport, T. and Dasgupta, S. "How to Set Up an AI Centre of Excellence," *Harvard Business Review*, January 16, 2019, <https://hbr.org/2019/01/how-to-set-up-an-ai-center-of-excellence>

36 Davenport, T. and Mahidhar, V. "What's Your Cognitive Strategy?" *MIT Sloan Management Review*, 2018, <https://sloanreview.mit.edu/article/whats-your-cognitive-strategy/>

37 Sharma, M. "Navigating the New Landscape of AI Platforms," *Harvard Business Review*, March 10, 2020, <https://hbr.org/2020/03/navigating-the-new-landscape-of-ai-platforms>

*AI explainability and transparency:* As outlined above, it is now widely known that AI models can be biased against certain groups and individuals. Some firms have established AI ethics organizations or “algorithm review boards” to assess transparency issues. Complex models, such as those implicit in deep learning neural networks, may be impossible to interpret or explain. Some vendors provide “prediction explanations” that point out influential variables or features and their direction of influence, but this is not yet possible for the most complex models. Many organizations and researchers are now working on new approaches to improve explainability, but we are only in the early stages of addressing this issue successfully.

*Reduced requirements for data:* Many AI models, particularly deep learning neural networks, require large amounts of data in order to be trained effectively. A new deep learning-based natural language generation model called GPT-3, for example, used billions of words to train the model and has 175 billion variables and parameters. Some researchers have argued that the trend toward such volumes of data is unsustainable and that new approaches to AI should aim to use less data. However, this trend is also in its nascent stages.

## Special Issue Papers

This special issue started out as a conversation between the guest senior editors and the editors-in-chief of two journals—*MISQ Executive* (MISQE) and the *Journal of the Association of Information Systems* (JAIS)—on the need to create concerted efforts to contribute to both IS theory and practice. This special issue is the outcome of this dialogue, which began at the 2019 Pre-ICIS Special Issue Workshop in Munich, Germany. We received over 50 extended abstracts and selected 30 submissions for discussion; we received early feedback from the special issue editorial board and the participating senior editors from both journals.

The special issue received a total of 50 submissions. About half of the submissions were sent out for review after the initial screening and, after three rounds, five articles were accepted for publication in the MISQE special issue. The first four papers appear in the December 2020 issue;

The last paper will be published in the March 2021 issue.

Table 2 maps the contributions that each paper makes to the special issue along with the type of AI technology it covers. We then briefly discuss each of the papers, outline the challenges faced by firms in adopting AI technologies and offer guidelines to manage such challenges.

The first paper in the special issue, “Addressing the Key Challenges of Developing Machine Learning AI Systems for Knowledge-Intensive Work,” is by Zhewei Zhang, Joe Nandhakumar, Jochem Thomas Hummel and Lauren Waardenburg. The paper discusses how a machine learning AI for a legal practice firm (LegalTechCo) was developed to help legal professionals make faster and better-informed decisions. The authors studied the development of the AI system at LegalTechCo over a couple of years. They identified three challenges involved in developing machine learning systems. The challenges are related to how to define ML problems, how to manage the training of ML models and how to evaluate ML AI performance. The authors propose three guidelines (and twelve recommendations) that executives can use to address the various challenges. The guidelines include: 1) co-formulate the appropriate machine learning AI problems; 2) develop machine learning AI through iterative refinement; and 3) go beyond the numeric measurements and ask for clues.

The second paper, “Unintended Consequences of Introducing AI Systems for Decision Making,” is by Anne-Sophie Mayer, Franz Strich and Marina Fiedler. This paper focuses on the unintended consequences of introducing an autonomous AI system in the banking industry. It draws on a case study from one of the largest banks in Germany (Main Finance). Main Finance confronted several issues in the small loan segment, including: 1) increased competition from new market participants due to digitization; 2) mismatched personnel resources; 3) high default rates; and 4) a decline in profitability. To address the issues faced, the firm introduced an AI system based on ML to make decisions about who is qualified for loans. The authors document the implementation of the AI system and its consequences from the perspective of both frontline workers and senior management. While the introduction of the AI system

**Table 2: The Focus and Contributions of the Special Issue Papers**

Paper	Authors	AI technology	Industry	Contribution
1	Zhang, Nandhakumar, Hummel and Waardenburg	Machine Learning Algorithmic AI	Legal	Covers challenges related to developing machine learning systems
2	Mayer, Strich and Fiedler	Machine learning Algorithmic AI	Banking	Discusses intended and unintended consequences of introducing an autonomous AI system
3	Asatiani, Malo, Nagbøl, Penttinen, Rinta-Kahila and Salovaara	Machine learning Algorithmic AI	Government	Offers ways to address explainability issues
4	Reis, Maier; Mattke; Creutzenberg and Weitzel	Machine Learning, Natural Language Programming	Healthcare	Explains physician resistance to an AI virtual agent
5	Schuetzler, Grimes, Rosser and Giboney	Conversational AI	Multiple examples	Offers guidelines to design conversational AI systems

enhanced profitability and helped address the main challenges facing loan management, it also resulted in employees' perceived loss of competence and reputation and unpredictability of decisions. From senior management's perspective, the AI system resulted in employees' loss of critical thinking and expertise and in the misuse of the system. The authors offer several guidelines to prevent related consequences: 1) maintain employees' abilities to reflect and understand underlying processes; 2) understand and guide the shift of employees' roles; 3) make the AI system as transparent and explainable as possible; and 4) reconsider customer groups excluded from the AI.

The third paper, "Challenges of Explaining the Behavior of Black-Box AI Systems," is by Aleksandre Asatiani, Pekka Malo, Per Rådberg Nagbøl, Esko Penttinen, Tapani Rinta-Kahila and Antti Salovaara. The authors document ways used by the Danish Business Authority (DBA)—an agency under Denmark's Ministry of Industry, Business, and Financial Affairs—to deal with challenges associated with explainability. The availability of large volumes of data-enabled DBA to pursue machine learning for core tasks such as supporting companies' legal compliance, checking annual reports for signs of fraud, and identifying companies early enough on their route to distress to ensure that timely support can be given. The organization has been able to implement AI responsibly and legally even

though the inner workings are not always entirely explainable. The authors build on a six-dimensional framework of an intelligent agent to discuss explainability challenges at DBA: 1) the model; 2) the goal; 3) training data; 4) input data; 5) output data and 6) environment. They further offer guidelines for managers to address explainability issues: 1) use modular design to increase AI explainability; 2) avoid online learning if explainability is a priority; and 3) facilitate continuous open discussion between stakeholders.

The fourth paper, "Addressing User Resistance Would Have Prevented a Healthcare AI Project Failure," is by Lea Reis, Christian Maier, Jens Mattke, Marcus Creutzenberg and Tim Weitzel. The authors discuss a case of AI implementation failure in a German hospital. The hospital decided to integrate AI to improve their anamnesis-diagnosis-treatment-documentation process with the intent of giving physicians more time to care for patients and reducing process costs. A virtual agent based on machine learning and natural language processing was developed to support different activities: 1) the cognitive agent engages with patients to perform anamnesis, collects data and provides structured documentation; 2) the cognitive agent applies decision support algorithms to suggest a diagnosis based on the structured recorded data; and 3) the cognitive agent engages with the physician to provide treatment options.

However, after nine months of developing the use case and the test version and six months of technological testing, the project team realized that the hospital's physicians did not want to use the system. While the physicians acknowledged that complementary knowledge supporting the diagnosis decision was valuable to themselves and the patients, they refused to approve the project. The team decided to postpone the project indefinitely until they could better understand the reasons for the physicians' rejection and what steps should be taken to ensure future project success. The authors document the reasons behind the physicians' rejection of the cognitive agent and offer recommendations to address them.

The fifth paper, "Your Agent Is Ready: Guidance for Designing Conversational Agents," is by Ryan Schuetzler, Mark Grimes, Holly Rosser and Justin Giboney. The paper focuses on chatbot design. Chatbots are used by organizations to improve business processes, automate routine interactions and provide an automated social touchpoint for customers. The authors build on their experience with chatbot design and use examples of chatbots across industries to offer a decision guide about when and how chatbots should be deployed. The framework presented in the paper asks questions and offers considerations that should be discussed early on in the bot development process, and offers a number of implicit signals that bots can use to create natural, human-like conversations.

## Conclusion

The five papers selected for this special issue along with this editorial provide a variety of examples of AI applications across industries and discuss challenges and implications for organizations. As AI technology is still maturing, awareness regarding the new management challenges it poses and the implications it raises for the workplace and the organization are still emerging, but the most common effect will likely be on how work is conducted in the future. Therefore, companies need to begin work now on developing AI applications that create economic value and that lead to new ways of orchestrating work by humans and machines. Leaders will need to understand and prepare for how AI will impact their workforce by, for

example, upskilling workers to do existing jobs with AI and retraining and hiring other workers to fulfill the new roles that AI will demand.

## About the Special Issue Editors

### Hind Benbya

Hind Benbya is a professor and head of IS and Business Analytics at Deakin University and is a visiting policy fellow at the Oxford Internet Institute at the University of Oxford. Hind's research expertise includes digital innovation, IT-enabled transformation and artificial intelligence. Her work has appeared in *MIS Quarterly*, the *Journal of Management Information Systems*, *MIT Sloan Management Review*, *MISQ Executive* and *Decision Support Systems*, among other outlets. She is currently a senior editor for *MISQ Executive*, guest senior editor for the *Journal of the Association of Information Systems* and a member of the editorial board of the *Journal of Strategic Information Systems*. Hind has been a visiting professor at Cambridge Judge Business School, UCLA Anderson School and the London School of Economics. She has received several best paper awards, regularly works with leading firms in Europe, the U.K. and the U.S. and has presented her research at premier academic and practitioner venues.

### Thomas H. Davenport

Tom Davenport is the President's Distinguished Professor of Information Technology and Management at Babson College, visiting professor at the Oxford Saïd Business School, fellow at the MIT Initiative on the Digital Economy and senior advisor to Deloitte Analytics. He teaches analytics/big data in executive programs at Babson, Harvard Business School and School of Public Health and MIT Sloan School. Davenport pioneered the concept of competing on analytics with his best-selling 2006 Harvard Business Review article and 2007 book. His most recent book is *The AI Advantage: How to Put the Artificial Intelligence Revolution to Work*. He has written or edited nineteen other books and over 200 articles for *Harvard Business Review*, *MIT Sloan Management Review*, *The Financial Times* and many other publications. He

is a regular contributor to the *Wall Street Journal* and *Forbes*. He has been named one of the top 25 consultants by *Consulting News*, one of the 100 most influential people in the IT industry by Ziff-Davis and one of the world's top fifty business school professors by *Fortune*.

### **Stella Pachidi**

Dr. Stella Pachidi is a lecturer of information systems at Judge Business School, University of Cambridge. Her research interests lie in the intersection of technology, work and organizing. Currently, her research projects include the introduction of artificial intelligence technologies in organizations, managing challenges in the workplace during digital transformation, and practices of knowledge collaboration across boundaries. She holds a Ph.D. in business administration from VU University Amsterdam, an M.Sc. in business informatics from Utrecht University and an M.Sc. in electrical and computer engineering from National Technical University of Athens. Dr. Pachidi has articles in information systems and organization journals and books including *Organization Science*, *Information and Organization* and *Computers in Human Behavior*. She has presented her work at various major conferences in the fields of technology and organizations including Academy of Management Meeting, International Conference on Information Systems, European Group for Organizational Studies Colloquium, Process Symposium and others.