

Artificial Intelligence in Organizations: Current State and Future Opportunities

Hind Benbya, Deakin University and Oxford Internet Institute (Australia and U.K.)
Thomas H. Davenport, Babson College and Oxford Saïd Business School (U.S.)
Stella Pachidi, Cambridge University (U.K.)

Recommended citation: Benbya, Hind; Davenport, Thomas H.; and Pachidi, Stella (2020) "Artificial Intelligence in Organizations: Current State and Future Opportunities," *MIS Quarterly Executive*: 19 (4).
Available at: <https://aisel.aisnet.org/misqe/vol19/iss4/4>

Introduction

Artificial intelligence (AI) is typically defined as the ability of machines to perform human-like cognitive tasks. These can include automation of physical processes such as manipulating and moving objects, sensing, perceiving, problem-solving, decision-making and innovation¹. AI is currently viewed as the most important and disruptive new technology for large organizations.² However, the technology is still in a relatively early state in large enterprises, and largely absent from smaller ones other than technology startups. Surveys³ 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. Relatively few organizations have deployed AI on a production basis, a problem that we describe in greater detail below. The experimental use, of course, means that many organizations have achieved little or no economic return on their AI investments. However, some analysts⁴ suggest that AI adoption will eventually have considerable positive impact on company growth and profitability.

AI is being applied in organizations for diverse objectives⁵: to make processes more efficient (28% as one of top two), to enhance existing products and services (25%), to create new products and services (23%), to improve decision-making (21%), and to lower costs (20%). Although a common theme in the AI-oriented press is related to reducing headcount, this objective got the lowest number of mentions at 11%.

Executives initially focused on using AI technologies to automate specific workflow processes and repetitive work. Such processes were linear, stepwise, sequential and repeatable. But now, firms are

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; 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), pp. 900-933.

² NewVantage (2019) "Big data and AI executive survey 2019, executive summary of findings," NewVantage Partners, <https://newvantage.com/wp-content/uploads/2018/12/Big-Data-Executive-Survey-2019-Findings-Updated-010219-1.pdf>

³ Genpact, 2020. "AI 360: Hold, fold, or double down," <https://www.genpact.com/uploads/files/ai-360-research-2020.pdf>

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

⁵ Deloitte (2020) "Thriving in the era of pervasive AI: Deloitte's state of AI in the enterprise, 3rd edition," Deloitte Insights, <https://www2.deloitte.com/us/en/insights/focus/cognitive-technologies/state-of-ai-and-intelligent-automation-in-business-survey.html>

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 ones like driving a car in all conditions remain tantalizingly out of reach.

We are beginning, however, to see autonomous systems that can perform tasks without any human involvement at all, as the system can train itself and adjust to new training data. Consider automated financial trading. Because it depends entirely on algorithms, companies can complete transactions much faster with it than with systems relying on humans. In a similar fashion, robots are performing narrow tasks autonomously in manufacturing settings.⁶

Some companies, such as Amazon.com 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 expenditures⁷. Less ambitious “low hanging fruit” projects have been more successful in most firms and are perhaps more consistent with the narrow intelligence possessed by AI systems at the moment.

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

This December 2020 Special Issue of the MIS Quarterly Executive is titled “AI in organizations: current state and future opportunities.” It details current challenges and implications that might arise from AI applications, and the 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 will be insightful to managers who are currently running digital transformation initiatives driven by AI technology, to practitioners who are 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 current challenges, implications and future opportunities regarding AI to enable readers to better 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 date back to the 1950s. The term AI was first introduced during a multidisciplinary program presented at Dartmouth 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.

⁶ Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108–116.

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 rejuvenated interest in AI. Machine learning and neural networks began to flourish as its 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 would become more common, 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 given momentum to 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 function 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 people⁸. AI types based on such a notion fall in general into three categories: artificial narrow intelligence, artificial general intelligence and artificial super intelligence.

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 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.⁹ This type of superintelligence can emerge following evolutionary and complex adaptive systems principles¹⁰. The argument states that if humans could create AI intelligence at a roughly human level, then this creation could, in turn, create yet higher intelligence and eventually evolve further¹¹. AI enthusiasts are providing estimates and outline scenarios for when technological growth will reach the point of singularity, where 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 just as well as, or even better than, humans, this type of artificial general intelligence does not exist yet. There are, however, some AI programs, such as the GPT-3 language prediction application, that are beginning to exhibit some aspects of general intelligence.¹²

⁸ Lake, B., Ullman, T., Tenenbaum, J. and Gershman, J. "Building machines that learn and think like people," *Behavioral and Brain Sciences*, 2017

⁹ Bostrom, N. (2014). *Superintelligence Paths, Dangers, Strategies*. Oxford, Oxford University Press.

¹⁰ See Benbya, H. Nan, N., Tanriverdi, H. and Yoo, Y. (2020), "Complexity and information systems research in the emerging digital world," *MIS Quarterly* (44: 1), pp. 1-18, 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.

¹¹ Hawking, S., Russell, S., Tegmark, M., & Wilczek, F. (2014). Transcendence looks at the implications of artificial intelligence - but are we taking AI seriously enough? *The Independent*, 01.05.2014.

¹² 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 recognise and replicate patterns of words to work out what comes next, see Thierry, 2020. New AI can write like a human but don't mistake that for thinking, *The Conversation*, Sept. 17., 2020, <https://theconversation.com/gpt-3-new-ai-can-write-like-a-human-but-dont-mistake-that-for-thinking-neuroscientist-146082>

Based on technology: A second typology differentiates between the technologies embedded into the AI systems which include machine learning, (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 survey¹³ suggests that all the 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 domain of applications of AI technologies.

Technology	Brief Description	Example Application
Machine learning <ul style="list-style-type: none"> • Reinforcement learning • Supervised learning • Unsupervised learning 	Learns from experience Learns from a set of training data Detects patterns in data that aren't labeled and for which the result isn't known	Highly granular marketing analyses on big data
Deep Learning	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
Neural Networks	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
Natural Language Processing	A computer program able to understand human language as it is written or spoken	Speech recognition, text analysis, translation, generation
Rule-based expert systems	A set of logical rules derives from human experts	Insurance underwriting, credit approval
Robotic process automation	Systems that automate structured digital tasks and interfaces	Credit card replacement, validating online credentials
Robots	Automatically operated machines that automate physical activity, manipulate and pick up objects	Factory and warehouse tasks

Table 1: AI technologies and domains of application

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 with 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 are mainly used to handle repetitive client queries whereas smart conversational AI, enabled by machine learning and natural language processing, have the potential to undertake more complex tasks that involve greater interaction, reasoning, prediction, and accuracy. Conversational AI have been used in many different fields, including finance, commerce, marketing, retail, and healthcare. Although the technology behind

¹³ Deloitte (2020) "Thriving in the era of pervasive AI: Deloitte's state of AI in the enterprise, 3rd edition," Deloitte Insights, <https://www2.deloitte.com/us/en/insights/focus/cognitive-technologies/state-of-ai-and-intelligent-automation-in-business-survey.html>

smart conversational agents is continuously under development, they currently do not have full human-level language abilities, sometimes resulting in misunderstanding and users' dissatisfaction.¹⁴

Biometric AI: Biometrics relies on techniques to measure a person's physiological (fingerprints, hand geometry, retinas, iris, facial image) or behavioral traits (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, workplace, home security among others. While fingerprints have the longest history as a marker of identity and continue to be used in a number of applications across the world¹⁵, other bodily markers like face, voice, and iris or retina are proliferating, with significant research exploring their potential large-scale application. Meanwhile, the ubiquity of face images and voice recordings tagged with people's names on the Internet alongside algorithms to transform such data into biometric recognition systems has accelerated their use at a larger scale. Examples include identifying suspects, monitoring large events and surveilling protests. Such large-scale use has triggered calls for regulation to introduce new laws, reform existing laws, or ban their use 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, are able to 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, a neural network can analyze the parameters of bank clients such as age, solvency, and credit history, and decide whether to approve a loan request. It can use face recognition to only let authorized people into a building. And it can 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. They 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 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 do a lot of tasks from successfully navigating their surroundings, to identifying objects around the robot or assisting humans with various tasks such robotic-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. A pilot AI project is relatively easy to develop and is only meant 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. These can include, for example, integration with existing technology architectures and legacy infrastructure, change in business processes and organizational culture, reskilling or upskilling of employees, substantial data

¹⁴ What is a Chatbot? All You Need to Know About Chatbots!. Botpress: Open-Source Conversational AI Platform. 2018. URL: <https://botpress.io/learn/what-and-why/>

¹⁵ Amba Kak, ed., "Regulating Biometrics: Global Approaches and Urgent Questions" AI Now Institute, September 1 2020, <https://ainowinstitute.org/regulatingbiometrics.html>.

engineering, and approaches to organizational change management. Full production deployment tends to take much longer than pilots and cost substantially more.

Surveys of organizations and market research reports in the US and globally suggest that deployment challenges are widespread with big data and AI. A survey¹⁶ (of large financial services and life sciences firms) found that firms were actively embracing AI technologies and solutions, with 91.5% of firms reporting ongoing investment in AI. But only 14.6% of firms reported that they have deployed AI capabilities into widespread production. In a 2019 global McKinsey [survey](#) with the headline “AI adoption proves its worth, but few scale impact,” between 12% (in consumer packaged goods) and 54% (in high tech firms) had at least one machine learning application implemented in a process or product. Only 30% of respondents overall were using AI in products or processes across multiple business units and functions.¹⁷

In order to address these 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 seeks to move projects when appropriate from the pilot phase to full deployment¹⁸. In a survey of US early adopter organizations, 54% of executives said their organization has a process for moving prototypes into production; and 52% had an implementation road map. These organizational approaches would seem to be an aid to getting more AI systems into deployment, but they may only be in the 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 can’t pay high levels of compensation and aren’t based in technology centers are likely to have difficulty hiring the desired number of skilled employees. Many companies should attempt not only to 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 only as developing models, others as extending to deployment of the models in production. The idea of data scientist “unicorns” who possess all these skills at high levels was never very realistic¹⁹.

Within university programs to train AI-oriented workers, the skills taught in such programs 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 to their actual capabilities. Further, activities involved in deployment of AI systems and related

¹⁶ NewVantage (2019) “Big data and AI executive survey 2019, executive summary of findings,” NewVantage Partners, <https://newvantage.com/wp-content/uploads/2018/12/Big-Data-Executive-Survey-2019-Findings-Updated-010219-1.pdf>

¹⁷ Deloitte (2018) “State of AI in the enterprise, 2nd edition,” Deloitte Insights, 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

¹⁸ Davenport, T. and Bean, R., 2018. “Farmers accelerates its time to Impact with AI.” *Forbes*, August 1. <https://www.forbes.com/sites/tomdavenport/2018/08/01/farmers-accelerates-its-time-to-impact-with-ai/#51430150b672>

¹⁹ Davenport, T., (2020). “Beyond unicorns: educating, classifying, and certifying data science talent,” *Harvard Data Science Review*, May 19, <https://hdsr.mitpress.mit.edu/pub/t37qjoi7/release/2>

organizational change issues may not be taught at all by many technically focused programs. There is an increasing need of a new type of professionals who can understand the business problems and translate them into algorithmic problems, and vice versa to explain the technical insights to business managers.²⁰

There are initiatives²¹ in the early stages to standardize the different types of data, analytics, and AI roles and requisite skills across organizations. This is an excellent idea, but developing new standards typically takes many years.

In the meantime, companies need to devote considerable attention to classifying and certifying the different types of AI and data science jobs needed in their organizations. Companies also would 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 on their data teams. These steps are essential for companies looking to use AI to improve both current operations and opportunities for digital innovation.

AI and social dysfunctions

Aside from the deployment and talent challenges, there are a few other potential dysfunctions from AI that managers need to be aware of and plan 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 disadvantage²². 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 infused in past sentencing decisions that were used as training data. Algorithmic bias can also have consequences distributed across large subsections of society by affecting the type of information that people are exposed to. This happens, for example when the machine learning algorithms behind social media propagate fake news or enable the targeting of individuals for political campaigns. To reduce any potential algorithmic bias, managers will need to be proactive by performing small-scale experiments and simulations before implementing such algorithms; regularly evaluate the dataset used for training; and involve human reviewers who regularly provide feedback to the 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: The possible social dysfunctions from AI implementation can increase if one considers the fact that the decision outcomes of some machine learning algorithms—deep learning in particular—cannot be easily explained due to the vast amount of feature layers involved in their production. This could lead to problematic situations, such as unexplainable evaluations of high school teachers, or parole decisions that cannot be justified and may cause rage when they also appear unfair²³. 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 non-technical terms is

²⁰ Henke, N., Levine, J., and McInerney, P., (2018) "You don't have to be a data scientist to fill this must-have role," *Harvard Business Review*, Feb. 5, <https://hbr.org/2018/02/you-dont-have-to-be-a-data-scientist-to-fill-this-must-have-analytics-role>

²¹ See for example, <https://www.iadss.org/>

²² Davenport, T. H. (2018). *The AI advantage: How to put the artificial intelligence revolution to work*. MIT Press.

²³ O'Neil (2016) *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books.

also necessary to ensure customers' trust and to avoid potential dysfunctions triggered by lack of transparency. In some industries such as banking, regulators sometimes force firms to use explainable algorithms.

Blurring accountability boundaries: As AI is used to enhance or even automate decision making procedures, the issue of accountability arises. Who is responsible in the case of a traffic accident with a driverless car? Who is responsible for approving parole to a criminal who eventually commits another crime? Who is responsible for a big financial loss in algorithmic trading? These are only few of the cases where the accountability boundaries are blurred. Managers will need to proactively focus on the reasons and processes that may lead to potential harm. They also need to carefully consider how they engage the different actors that directly or indirectly interact with the outcomes produced by the AI system (AI developers and designers, business users, institutions), and clarify responsibility and legal liability upfront²⁴.

Invaded privacy: Ethical issues arise even before any action is recommended or performed by the AI system, with privacy being reported as one of the main ethical considerations behind AI implementation²⁵. Data is the primary resource that is fed into the AI systems, and quite often is seen as a source of competitive advantage. AI's need to process increasingly large amounts of data thus conflicts with people's right to maintain control over their data and its use in order to preserve their privacy and autonomy. Organizations need to ensure that their data practices comply with the relevant policies on the use of personal data (e.g. GDPR in EU countries) and avoid any possible privacy violation. Developing auditable algorithms and performing algorithmic audits on them to identify what data is used and what variables feed into a decision-making procedure are helpful solutions to increase transparency of how consumers' data are processed and used²⁶. Overall, being open about how the data is handled is necessary to ensure customers' trust.

Implications

Despite 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. These changes will have broader implications on organizations and societies, many of which have yet to be understood or realized. But 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 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, as for the first time they get to affect knowledge workers²⁷. Professions such as doctors, lawyers, consultants or architects, whose expertise, judgment and creativity have thus far been highly

²⁴ Dourish, P. (2016). Algorithms and their others: Algorithmic culture in context. *Big Data & Society*, 3(2), 1-11.

²⁵ Kinni (2017) Ethics Should Precede Action in Machine Intelligence. *MIT Sloan Management Review*

²⁶ Mittelstadt, B. (2016). Automation, algorithms, and politics| auditing for transparency in content personalization systems. *International Journal of Communication*, 10, 12.

²⁷ See Davenport, T. (2005), *Thinking for a Living: How to Get Better Performance and Results from Knowledge Workers*, Harvard Business School Press ; Benbya, H. (2008), *Knowledge Management Systems Implementation : Lessons From the Silicon Valley*, Neal-Schuman Publishers, and Faraj, S., Pachidi, S., and Sayegh, K. 2018. "Working and Organizing in the Age of the Learning Algorithm," *Information and Organization* (28:1), pp. 62-70.

valued and considered irreplaceable, for the first time in history appear threatened. And while the end of those professions is not for the near future, the changing nature of their work is already a reality. There are many predictions about how much job loss from AI will take place, but thus far it has been relatively small.²⁸

Changing expertise: AI technology that is able to automate some of the workers' tasks is already in the workplace. In law firms for example, a plethora of applications have been developed for automating the 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 to be carried out by account managers. While such automations can increase efficiency of operations and decrease labor costs, they leave professionals with voids in the processes they used to acquire knowledge about their subjects or customers, or the ways through which they would develop their expertise. This will eventually lead to changes in the knowledge of the affected occupations and could potentially even trigger their restructuring. For example, in the legal profession there is already the tendency for various law graduates to develop data science skills and engage with legal tech instead of following the traditional career path of a lawyer.

Augmenting professionals: In several of the cases, AI systems are not yet able to replace human experts, but they can augment their work by supporting experts' judgment and decision-making processes. For example, the debate has now moved away from the "end of radiologists" focus, and acknowledges that radiologists will not be replaced by AI tools any time soon, but they will be augmented by them.²⁹ Yet, as AI systems are introduced in the radiology profession to support the radiologists' diagnosis process, we begin to see several unintended consequences on their everyday work: from having to overcome communication barriers in their unavoidable interactions with data scientists, to even doubting the prediction of the AI system or questioning their own diagnosis³⁰. This becomes even more complicated if we consider that most often, the way in which a machine learning algorithm functions and comes to render a specific outcome cannot be easily traced or explained.

Thus, the nature of work is changing dramatically, and while many observers predict that the combination of human and machine intelligence will always be the winning one³¹, we have yet to see how the "augmented professionals" will carry out their work, and with what further implications for the workplace, the organization and institutions.

Organizational implications

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

Changing authority arrangements: Unavoidably, as we have discussed above, expertise is redefined and the knowledge and skills of technology practitioners such as machine learning experts, data scientists or data analysts become increasingly valued in the workplace. This can lead to restructured authority arrangements across all levels of hierarchy. On a tactical level, technology practitioners will

²⁸ For one prominent prediction, see Carl Benedikt Frey and Michael Osborne (2013), *The future of employment: how susceptible are jobs to computerization*, Oxford Martin School working paper, <https://www.oxfordmartin.ox.ac.uk/downloads/academic/future-of-employment.pdf>

²⁹ Davenport, T. and Dreyer K. (2018), "AI will change radiology, but it won't replace radiologists", *Harvard Business Review*, March 27. <https://hbr.org/2018/03/ai-will-change-radiology-but-it-wont-replace-radiologists>

³⁰ Lebovitz, S. Lifshitz-Assaf, H. and Levina, N. 2020, *To Incorporate or Not to Incorporate AI for Critical Judgments: The Importance of Ambiguity in Professionals' Judgment Process* (January 15, 2020). NYU Stern School of Business, Available at SSRN: <https://ssrn.com/abstract=3480593>

³¹ Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company.

gain authority and control over the work design and decision-making procedures, given that they have the ability to prescribe how the introduced AI systems will affect the operations and work. But even on a more strategic level, new roles get to join the board, triggering questions towards the established regime: e.g. where does the jurisdiction of the CIO end and where does the jurisdiction of the CDO start when it comes to planning a major digital transformation with the implementation of AI technology?

Changing coordination: The use of AI to manage work algorithmically leads to fundamental changes in organizational design and coordination. Work tasks are redefined so that they can be broken down into smaller sub-tasks and are then algorithmically assigned to workers of digital labor platforms such as UpWork or Amazon MTurk³². Machine learning algorithms can be used to coordinate more on a proactive basis, by analyzing historical data to predict the needs in skills and expertise for future projects. Furthermore, practitioners and managers need to collaborate with new experts who enter the workplace with expertise in data processing, algorithm development, data visualisation, etc. Collaboration amongst people with different types of expertise can make work coordination more challenging. Changing coordination results in a substantially different execution of an organization's operations and services

Changing valuation schemes: The way in which performance is evaluated is also changing substantially, as employees are assessed by machine learning algorithms most often without knowing what variables are included in the inherent model, or the extent to which a specific variable contributes to the production of a specific outcome. Even the quality check of products becomes an automated task itself with the use of robots³³. Those fundamental changes in the values that matter in the organization substantially impact how firms manage their employees but can also lead to counter performances from the employees' side. For example, delivery drivers hanging their phones on trees outside Whole Foods in order to get more rides was a fascinating yet ironic demonstration of how people try to game the system in their effort to maintain some control over their work.³⁴

Industrial transformations: AI technology is currently enabling significant digital transformations that are not only redefining what an organization does, but even result in blurring industry's boundaries. A lot of traditional manufacturing organizations for example are 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 popular example of such an attempt, with AI being the driving force behind their predictive maintenance services. Who are the new competitors of such a digitally transformed organization? How should it be regulated? How do the relationships with their customers change? Who are their new partners? A lot of new questions arise that need to be addressed.

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 management and governance mechanisms in place. We've already mentioned those related to deployment. In addition, organizations in various states of adoption, have put in place a wide range of internal

³² Faraj, S., Pachidi, S., and Sayegh, K. 2018. "Working and Organizing in the Age of the Learning Algorithm," *Information and Organization* (28:1), pp. 62-70.

³³ Mahdawi A. (2019). The Domino's 'pizza checker' is just the beginning - workplace surveillance is coming for you. *The Guardian*, October 15 2019.

³⁴ Soper S. (2020). Amazon Drivers Are Hanging Smartphones in Trees to Get More Work. *Bloomberg*, September 1, 2020.

organizational structures and roles to manage and govern AI projects. A survey³⁵ suggested that the governance mechanisms used to manage AI projects include appointing AI champions, which 45% of respondents said their firms had already done; creating an AI center of excellence³⁶, which 37% had done; and developing a comprehensive strategy for AI³⁷, which 37% had also adopted.

Democratization of data science and AI—Tools like automated machine learning³⁸ can structure and automate the workflow of creating and implementing a machine learning model. These 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, the notion that anyone, with little to no expertise, can do 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 in production processes and systems. Since their business is dependent in part on the accuracy of these models, it’s important to monitor them for “drift”—inaccuracy of predictions—and improve them over time. Vendors are developing tools to support this process under the banner of “MLOps”—machine learning operations—and they are most widely used in data- and analytics-dependent industries like financial services.

AI explainability and transparency—As outlined above, it is now widely known that AI models can be biased against certain groups and individuals³⁹. Some firms have established AI ethics organizations⁴⁰ or “algorithm review boards” to assess transparency issues. Complex models, such as those in deep learning neural networks, may be impossible to interpret or explain⁴¹. Some vendors provide “prediction explanations” that point out influential variables or features and their direction of influence, but this isn’t yet possible for the most complex models. Many organizations and researchers are now working on new approaches to explainability, but we are only in the early stages of addressing the issue successfully.

Reduced requirements for data—Many AI models, particularly deep learning neural networks, require large amounts of data 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 researchers⁴² have argued that the trend to such volumes of data is unsustainable, and that new approaches to AI can use less data. This is another trend in the early stages, however.

³⁵ Davenport, T. H. (2018). *The AI advantage: How to put the artificial intelligence revolution to work*. MIT Press.

³⁶ Davenport, T. and Dargupta, S. (2019) “How to set up an AI centre of excellence,” *Harvard Business Review*, January 16, <https://hbr.org/2019/01/how-to-set-up-an-ai-center-of-excellence>

³⁷ Davenport, T. and Mahidhar, V. (2018) “What’s your cognitive strategy?” *MIT Sloan Management Review*, Summer. <https://sloanreview.mit.edu/article/whats-your-cognitive-strategy/>

³⁸ Sharma, M. (2020) “Navigating the new landscape of AI platforms,” *Harvard Business Review*, March 10. <https://hbr.org/2020/03/navigating-the-new-landscape-of-ai-platforms>

³⁹ Li, M. (2019) “Are your algorithms upholding your standards of fairness?” *Harvard Business Review*, Nov. 5. <https://hbr.org/2019/11/are-your-algorithms-upholding-your-standards-of-fairness>

⁴⁰ Davenport, T., 2019. “What does an AI ethicist do?” *MIT Sloan Management Review*, June 24. <https://sloanreview.mit.edu/article/what-does-an-ai-ethicist-do/>

⁴¹ Royal Society (2019). Explainable AI: the basics; policy briefing. <https://royalsociety.org/-/media/policy/projects/explainable-ai/ai-and-interpretability-policy-briefing.pdf>

⁴² Wilson, J., Daugherty, P., and Davenport, C. (2019) “The future of AI will be about less data, not more.” *Harvard Business Review*, Jan. 24. <https://hbr.org/2019/01/the-future-of-ai-will-be-about-less-data-not-more>

Special issue papers

This special issue started as a conversation between the guest senior editors and the editors in chief of two journals: the *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 such dialogue, a dialogue started at the pre-ICIS special issue Workshop held in Munich. We received over 50 extended abstracts, 30 submissions were selected for discussion and 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 to 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 each paper makes to the special issue along with the type of AI technology it covers. We then briefly discuss each of the papers and outline the challenges firms faced while adopting AI technologies and guidelines to manage such challenges.

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's resistance to an AI virtual agent
5	Schuetzler, Grimes, Rosser, and Giboney	Conversational AI	Multiple examples	Offers guidelines to design conversational AI systems

Table 2: Special issue papers focus and contribution

The first paper in the special issue, “*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 proposed three guidelines (and twelve recommendations) for executives 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; 3) Go beyond the numeric measurements and ask for clues.

The second paper, *“Well-meant is not well done: unintended consequences of introducing AI”* 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)⁴³. 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 a loan. The authors document the implementation of the AI system and its consequences from the perspective of both front-line workers and senior management. While the introduction of the AI system enhanced profitability and helped address the main challenges faced with 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; 4) Reconsider customer groups excluded from the AI.

The third paper, *“Implementing Black-Box Artificial Intelligence: Lessons from Tackling Explainability Issues at the Danish Business Authority”* 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. Availability of large volumes of data enabled DBA to pursue machine learning for such core tasks such as supporting companies’ legal compliance, checking annual reports for signs of fraud, and identifying companies early enough in their route to distress 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) goal, (3) training data, (4) input data, (5) output data, (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, 3) Facilitate continuous open discussion between stakeholders.

The fourth paper, *Resistance to AI: A case study exploring the implementation failure of cognitive agents in healthcare* is by **Lea Reis; Christian Maier; Jens Mattke; Marcus Creutzenberg, Tim Weitzel**. The authors discuss a case of AI implementation failure in a German hospital. The hospital has 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

⁴³ A pseudonym

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 do not want to use the system. While the physicians acknowledge that complementary knowledge supporting the diagnosis decision is valuable to themselves and the patients, they refuse to approve the project. The team decided to postpone the project indefinitely until they can better understand the reasons for the physicians' rejection and what steps to take to ensure future project success. The authors document the reasons behind the physician's 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, or provide an automated social touchpoint for customers. The authors build on their experience with chatbot design and use examples of chatbot across industries to offer a decision guide about when and how chatbots should be deployed. The framework presented in the paper provides questions and considerations that should be discussed early in the bot development process. And offers a number of implicit signals bots can use to create natural, humanlike conversations.

Conclusion

The five papers selected for this special issue along with this editorial provide a variety of examples of AI applications across industries, challenges and implications for organizations. Table 1 summarizes the AI technology and industry covered by each paper, as well as the main contribution each paper makes. We then briefly discuss each of the papers and outline the challenges firms faced while adopting AI technologies and the recommendations offered to manage such challenges.

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 is likely to be on how work is conducted in the future. Thus, 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. And leaders will have to understand how AI will impact their workforce, then get them prepared : upskill some workers to do existing jobs with AI, and retrain and hire others for the new roles that AI will demand.

About the Special Issue Guest Editors

Hind Benbya

Hind Benbya is Professor and Head of IS and Business Analytics at Deakin University and 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 others. She is currently senior editor of the *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 received several best paper awards, regularly works with leading firms in Europe, UK and the US and presents 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 wrote 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 in 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 PhD in Business Administration from VU University Amsterdam, a MSc in Business Informatics from Utrecht University, and a MSc 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 in various major conferences in the fields of technology and organizations including the Academy of Management Meeting, the International Conference on Information Systems, the European Group for Organizational Studies Colloquium, the Process Symposium and other.