

Artificial Intelligence and the World of Work, a Co-Constitutive Relationship

Carsten Østerlund, Syracuse University; costerlu@syr.edu

Mohammad Hossein Jarrahi¹, University of North Carolina; jarrahi@unc.edu

Matthew Willis, University of Michigan; mandwill@umich.edu

Karen Boyd, University of Maryland; klboyd@umd.edu

Christine T. Wolf, IBM Research – Almaden; ctwolf@us.ibm.com

Abstract: The use of intelligent machines-- digital technologies that feature data-driven forms of customization, learning, and autonomous action-- is rapidly growing and will continue to impact many industries and domains. This is consequential for communities of researchers, educators, and practitioners concerned with studying, supporting, and educating information professionals. In the face of new developments in artificial intelligence (AI), the research community faces three questions: 1) How is AI becoming part of the world of work? 2) How is the world of work becoming part of AI? and 3) How can the information community help address this topic of work in the age of intelligent machines? This opinion piece considers these three questions by drawing on discussion from an engaging 2019 iConference workshop organized by the NSF supported Work in the Age of Intelligent Machines (WAIM) research coordination network (RCN)².

Introduction

Artificial intelligence (AI) and its relation to work has become central in our cultural discourse, clear to even a casual reader of contemporary news and media outlets. Technological breakthroughs in the field of AI promise to change the way we organize work (Davenport & Kirby, 2016). The artful integration of AI and work, however, remains an open challenge. There is currently limited empirical understanding and research to guide the information community in this area—e.g., labor, motivation, cognition, machine learning, data science, human-computer interaction, and information science, among others—in coherent ways. Such interdisciplinarity is necessary if we are to push beyond assumptions and hype to open up possibilities for diverse and inclusive AI futures.

This article draws on a 2019 iConference workshop about work in the age of intelligent machines. An interactive workshop technique helped harness the collective insights of a multidisciplinary group of researchers. A key theme that emerged from the workshop made it clear that to push beyond obvious connections and commonsense notions associated with new technologies, we must stay open to

¹ Corresponding author

² <https://waim.network>

numerous possible AI futures. That is, embedding AI with people's practices, infrastructures and organizational life could follow more than one trajectory and history tells us that we cannot expect any easy causal connections working one way or the other. In short, our exploration should ask not only how AI is part of the world of work, but also how the world of work is part of AI (Haraway, 1991). Based on these inputs, this paper aims to provide an impetus for disciplinary convergence around the interplay between AI and the world of work.

AI in the World of Work

The information field has long been interested in workplace changes associated with rapidly increasing technological capabilities (e.g., Zuboff, 1988). Given this extant topic of concern, workshop participants wondered: what is new about AI and work that extends beyond IT and work? To what degree will AI in work reproduce existing structures or introduce new dynamics relative to the concepts of labor, agency, sensemaking, system fluidity, and opacity?

Labor

Deskilling and job replacement are two central concerns around adoption of intelligent systems (Gray & Suri, 2019). AI algorithms are making inroads into the realm of perception and cognition, domains less susceptible to technology-enabled automation in the past (Brynjolfsson & Mitchell, 2017). This could lead to deskilling as tasks transfer from human to AI or augment existing human skills. Autonomous vehicles are a canonical example affecting truck drivers, one of the largest job categories in the labor market. Will truck drivers still be needed for some sub-tasks? Will a new division of labor emerge involving both humans and AI? At higher-end pay scales, new job categories appear to emerge such as "data scientist" and "machine learning engineer," but it is unclear if these truly constitute new jobs or are merely renamed positions. Many workshop participants agreed we need to better understand AI in the world of work by asking: what types of labor incorporate and make use of AI? What types of actors (human and non-human) are involved? The answers to these questions are likely to depend on the specific forms of work. For instance, the role of AI in crowdsourcing or gig work arrangements may differ radically from work settings governed by professional organizations, such as medicine or law.

Agency

Understanding AI in work involves questions about the types of agency these systems will take. AI advances into perception and cognition have reignited the debate about agency and information systems (Neff & Nagy, 2018). Will we, in fact, be seeing a qualitative shift in agency and control to AI-supported systems, workshop participants wondered. On one hand, information systems have for a long time given orders, e.g. through factory scheduling systems. On the other hand, some AI applications now act more independently. For instance, neural net training involves back-propagation where the system finetunes weights based on error rates generated during previous iterations to learn over time.

Sensemaking

The issue of agency gets further muddled by changes to how we encounter new systems. Personal intelligent assistants, such as Siri and Alexa, have captured the public imagination and raised questions about sensemaking (Maedche et al., 2019). Smart machines are increasingly acting like humans and interfacing with them through visual and voice-based channels. This is not entirely new (for example, ELIZA was introduced with similar affordances in 1966); however, new AI algorithms create more seamless digital ecosystems, which enable users to practice natural interactions with the machine. A challenge of designing personal assistant systems could be that the system renders itself invisible by closely simulating human behaviors, leading users to overestimate the system's capabilities.

Fluidity

Emerging AI systems tend to be dynamic, personal, and customizable. The rise of unsupervised learning algorithms provides unprecedented opportunities for discovering new patterns in data and often more effective predictions (Jarrahi, 2019). However, the same characteristics make it difficult for managers, users, and even designers to make sense of the system and its inferences, especially if the system constantly changes with new data. The fluid nature of these systems sets them apart from previous generations of intelligent systems.

AI fluidity goes hand in hand with the personalization built into machine learning. A system evolving to suit individual user behavior promotes seamless interactions, but also raises the danger of echo chambers where people only get exposed to a narrow sliver of information shaped by their prior interactions. Cross comparisons become challenging when the system's behavior and output becomes individualized.

Opacity

The intrinsic opacity of many systems powered by machine learning algorithms makes it difficult even for AI experts to explain the rationale behind an AI model's decision making processes (Brynjolfsson and McAfee 2017). It is not uncommon to hear sentiments like "the output looks optimal, but I don't know why?" A lack of transparency and explainability can hamper trust if humans cannot readily decipher how their technological partners make decisions. Furthermore, this can affect the reliability of decision making. For example, when scientists rely upon ML techniques to create new scientific practices, there may not be adequate chains of reasoning that ensure that actions taken are valid, reliable, and replicable. Together, these changes in the nature of technology raise the importance of data and algorithmic literacy. Soon, many workers may need to use intelligent assistants in daily work practices. In addition to common competencies that have undergirded knowledge work for many years, workers have to develop an understanding of this emerging class of technologies, their affordances, their limits, and how to optimize human-AI partnerships in practice. In addition to the daily interactions with workers, AI systems' functioning and decision making will need to be made visible to stakeholders inside and outside of organizations, whose daily lives will be affected by decisions enabled by these systems. The transparency of AI in the world of work goes beyond academic or engineering concerns— it also matters to users, shareholders, regulators, and citizens.

The World of Work in AI

Exploring AI futures and affordances entails an understanding of not only AI in the world of work but equally important, the world of work in AI. The world of work sneaks into AI in the form of big data feeding the algorithms, data management activities, the bias of past practices, ethics, and governance.

Big data

AI performance goes hand in hand with computationally accessible (big) data. The two reinforce one another. Digitalization and datafication of individual and organizational practices power AI inferences. At the same time, algorithms cohesively bring together data assemblages scattered across sources, geographies, and systems. This infusion of user-generated and often personal data into algorithms can have unintended legal and privacy implications that large companies like Facebook and Alphabet still struggle to address.

Data Management

Additionally, we need to ask ourselves who maintains or should maintain big data. Where does it take place? What are the stages in and histories of its production? For instance, instead of programming a computer to recognize a human face, it can be taught how to identify a face from a collection of images or video available on social media, private databases, surveillance systems, or publicly available images and video. We need to better understand the integration of sources and data signals that enable machines to become capable of sophisticated function mapping so that we can assess the credibility of its outputs and intervene if necessary.

Bias

All data management practices come with inherent biases that can become masked by the sheer complexity of mapping labeled and unlabeled data of various types and sources. Humans who label data introduce their own interpretations, context, and biases into the system's feedback loop. These considerations of data context and bias become magnified with systems that run at scale. Researchers need to develop research programs and actionable strategies to help detect and control sources of biases in AI. However, it is difficult to decide whether and when to intervene in a dataset or algorithm. For example, some people believe that intervening in training data is the most effective way to reduce bias in algorithms, where others see manipulating training data as compromising the power of patterns "naturally" emerging from the data.

Governance

To face the biases and ethical concerns associated with data management practices in AI we need to ask what kinds of local, national, and international bodies do or should claim jurisdiction over AI. Policy makers enter uncharted territory given the lack of understanding about how AI systems actually function, difficulty in auditing them, and lack of visibility to users, the public, regulators, and often even algorithm builders. As such, researchers should help decision-makers at different levels and across

boundaries (e.g. enterprises, technology firms, government, and international organizations) develop policies and regulation that extinguish potential negative consequences of AI uses and invite fair, ethical implementations that benefit individuals, institutions, and society at large.

With the implementation of AI systems, legions of workers will see their jobs changed. As a result, new labor policies can help these workers learn new skill sets and compete in a changing labor market. Reskilling not only applies to these workers, but also those who do mundane, “ghost work” behind the scenes without many opportunities for professional growth (Gray & Suri, 2019).

Discussions at the workshop also highlighted a local-global dynamic, pointing to the impacts of local decisions on global issues, which requires a global policy perspective. For example, even if the techniques used in China (a country of particular interest among participants, given its heavy investment in AI development) are the same, the training data ML learns from; the market needs it fills; the work environment that produces it; the humans who build, use, interpret, and are affected by its outcomes; the definitions and priorities of relevant values; and the government that regulates it may be different. These factors are difficult to study and intervene in. It also makes local interventions in, for example, the US or European Union potentially less effective in a global market and less appealing to regulators, who may believe that regulation will allow China to “win,” strategically speaking.

Implications and Future Directions

By now it should be clear that AI in the world of work and the world of work in AI do not represent a dichotomy with opposing forces. Each side collapses into the other under close scrutiny. Nevertheless, the distinction forces us to keep an open mind, avoiding one-directional or one-dimensional explorations of AI trajectories and possible futures; instead it encourages an exploration of the specificity and situatedness of AI and work. Below, we discuss implications of these questions for different industries as well as how they may impact our future educational, system design, and methodological practices.

Implications for Different Industries and Sectors

One finds a range of potential AI uses in different sectors. The variations across industries appear closely related to not only sector-specific drivers of value, but equally important the availability of data, how it suits existing technologies, and the applicability of various algorithmic solutions (Chui et al 2018). Below, we briefly broach implications of AI technologies in three major industries.

Service Industry

AI is expected to usher in extensive transformation in several service industries, specifically financial, consulting, and IT services (Nhat et al., 2020). Over the past decades a growing part of the service lifecycle has already become digitized. Still, AI holds unique potential to transform many touchpoints in service production and delivery including marketing strategies, sales processes, and customer service (T. Davenport et al., 2020). For example, AI’s predictive power can help organizations better gauge

demands and adjust business processes and service delivery accordingly. AI may also reshuffle customer service or technicians' schedules based on a more timely and concise analysis of the ebbs and flows of service requests as well as historical precedents. Furthermore, intelligent assistants are already helping both customers and workers to more easily interact with information at different points of service (e.g., via more natural conversations) (Maedche et al., 2019).

It is assumed that AI will disrupt key elements of the service value chain such as consumers' journeys; however, for decision-making practices in the enterprise, a human-in-the-loop approach is still needed. This approach requires not only a new division of labor between human actors and artificial agents, but a new set of competencies on the part of workers. Roles may shift from handling monotonous and repetitive activities towards more human-centred tasks (such as exception handling) which requires social and emotional intelligence as well as a contextual perspective. For instance, a vast majority of underwriting processes are already automated and performed by intelligent systems; however, a more subjective, manual review and subsequent rehashing decision would still need to be made by credit analysts (Smith, 2020).

Manufacturing and Agriculture

AI's greatest potential in manufacturing appears to be associated with supply chain, logistics and maintenance (Atkinson, 2019; Chui et al., 2018; Muro et al., 2019). For instance, predictive maintenance combines machine learning and Internet of Things (IoT) sensors on machinery to detect anomalies in e.g temperature, sound or vibration. Applying deep learning techniques on large amounts of high dimensional sensory data will allow manufactures to predict failures and plan repairs in ways that reduce downtime and operational cost.

In agriculture, AI may help increase productivity and yield by detecting poor plant nutrition, weeds, pests, water needs, or diseases in plants and pets. Such analysis relies on hundreds of thousands of data points detailing weather conditions, temperatures, soil, water usage and in some cases real time images gathered by drones flying over the fields. With these tools, farmers hope to improve harvest quality, better predict their yield and even the protein content of specific crops such as wheat.

AI has also fueled an increased deployment of robots (Czarniawska & Joerges, 2020). In manufacturing this process has been taking place over the past three decades and today the assembly of a BMW involves up to 1000 robots. These automation processes appear to be reshaping supply chains by allowing some organizations to reshore by moving manufacturing capabilities from developing countries with low labor costs back to developed nations. In farming, one finds an increased use of milking robots, robotic fencers and mobile barn cleaning robots (Frey & Osborne, 2017).

Healthcare

The opportunities of AI to transform healthcare are plentiful. AI may enable greater analysis and synthesis of the growing amount of health information than is currently possible. Clinical decision making and automation of mundane tasks could also make room for additional time in the patient-

provider consolation. However, the application of AI in this work domain is fraught with challenges and unintended consequences (Quiroz et al., 2019). AI may increase the complexity of data management practices and potentially introduce error (Challen et al., 2019). In an already overtaxed system practicing clinicians have little time for training on new systems or to conduct additional data entry. The use of any AI system has to be proven as there is little room for error and no slack in the system.

For several years now, radiology and medical imaging have seen the rollout of AI and its influence on the work of professionals (Nichols et al., 2019). Rapid image analysis and interpretation have created a fluid working environment—no longer having to wait for images to process and receiving diagnostic and interpretive guidance and automatic artifact recognition and highlighting. Recent comparisons of AI and radiologist performances showed that deep learning algorithms suffered from limited specificity for categorizing findings from a chest radiograph (Singh et al., 2018). These and similar findings in medical imagery yield two broader points for the future of AI and work in healthcare. First, AI should be thought of as a supportive tool, working in symbiosis with clinical professionals, administrative staff, and patients. In the case of radiology—AI tools should augment analysis and reduce the burden of cognitive load on radiologists, but not removing human judgment and critical appraisal from the clinical workflow. Second, if AI tools interrupt clinician analysis and interpretation they can lead to misjudgments and errors in patient care. To augment, not hamper clinical knowledge and intuition the field needs to address AI the lack of transparency associated with many AI systems and their potential introduction of biases (Topol, 2019).

Cutting across all sectors, ongoing data acquisition and management will become a challenge (Chui et al., 2018). To realize AI promises, service providers, manufacturers, farmers and healthcare providers will need to gather thousands of data records and in some cases millions before AI models can perform at the level of humans. For many organizations, it can be difficult to obtain and label data sets of this magnitude. This will require companies to develop strategies that allow them to collect, integrate, and process vast amounts of data on an ongoing basis. Some AI potentials can only be realized with a diverse range of data types spanning from numerical and text-based, to images, video, and audio files. Many models necessitate retraining on a regular basis to adjust for potential changes in conditions. Training data will thus need to be refreshed regularly and sometimes daily, for models to retain currency.

Educational Practices

If we hope to train the general public and future experts to design, analyze, and use AI, we must teach data and algorithmic literacy. Data often appears to be neutral; exposing students to the way data is value-laden will help them navigate a labor market shaped by AI. They should be able to make sense of AI systems, navigate their opacity, and be able to deal with situations where the system breaks down, requiring improvisational problem-solving. Developing algorithmic literacy calls for knowledge about interfaces, data management, design, ethics, and governance as described above.

Workers and management alike must understand the capabilities of AI and its limits. At the workshop, participants pointed to concepts of “learning” and “intelligence” as perpetuating the misleading belief

that a machine “understands” its domain. This can lead to overestimating AI’s capabilities; unawareness of the potential for errors and discriminatory bias; and the inability to critically assess outputs, participate in making the algorithm more accurate, or intervene in job tasks that are influenced by ML-driven tools. Beyond understanding the system itself, workers must understand the assemblages and larger ecosystems of which the systems are part. When biases seep up from social systems into AI systems (like hiring algorithms biased against women and minorities) what should HR workers do? The prospective employee? These remain to be seen.

Design Practices

Our design practices face new challenges and questions in the emerging AI era. New interactional modes, such as voice or gestural commands, call for new approaches to enable dynamic relationships between machine and human learning (Maedche et al., 2019). This may require efforts to develop interfaces that engender trust through algorithmic access and transparency. Centering questions of opacity early in the design process could be one way to help ensure that AI’s effects on work are visible, intelligible, and corrigible.

Designing AI systems requires a sociotechnical perspective - AI does not exist in a vacuum. An automated vehicle, for example, may not get far without instrumented roads, readily-available cellular network coverage, and unperturbed Bluetooth signal. How do infrastructural assumptions around access and connectivity shape the future of work? People with dial up, intermittent, or no internet access may be unable to benefit from some AI tools. This may also preclude people in low-income or rural areas from opportunities to develop important skills for jobs that include AI systems.

Integrating AI systems into everyday work practices also requires alignment between professional and occupation logics and values and machine actions and capabilities. For example, engineers and developers seeking to create a solution as close to optimal as possible, and thus may be satisfied with an AI system that assists directly with the creation of the output if they believe the solution will be better. Artists or other content creators, on the other hand, may be less interested in optimizing final outcomes than enhancing and enriching their creative practice. An AI-driven tool that will help inspire (e.g., suggesting novel source content for the artist to consider or proposing unique linkages among content) or guide (e.g., highlighting certain pixels in an image editing tool) could support artistic logics. Stock traders may desire some balance between a system that supports their intent (developing insight about the market) and optimization (analyzing individual stocks and executing trades as fast as possible). Autopilot systems must prioritize safety over optimal routing and support human intent and intuition primarily in service of safety. Understanding the importance of human intent and the risk of human error in a work task can inform design decisions about how work can be distributed between human and AI in systems tailored to professions.

The fluidity of many AI systems complicates user experience (UX) design. Your Spotify or Netflix recommendations, for example, are not the same after your teenage child uses the account for a couple of weeks. Neither are the recommendations solely for “you,” when the family watches content together

as a social unit. The UX constantly shifts and tactics are needed to help people build personal connections, tell stories, and see themselves in the spaces and narratives shaped by AI. This is equally important in workplace settings - where careers and livelihoods are entangled with the data-driven, sociotechnical systems that we use to get work done. The future of work calls for a participatory approach where workers play a central role in the design and ongoing alignment of AI systems in their everyday work practices (Wolf & Blomberg, 2019).

Methodological Practices

As researchers, we face many of the same challenges as policy makers, educators and designers around AI and work. How do we study what is opaque and fluid? How do we deal with access where data and algorithms fall under organizational control? The field calls for new methods that consider and can describe personalized systems that change over time (Østerlund et al., 2020). This may lead to new epistemological and ontological considerations about the nature of algorithms and new frameworks for describing and evaluating training data. AI may also become a tool in the study of its own phenomena. Adversarial learning is beginning to be used to explore biases, training one ML model to offset the biases of another. It may be possible to explore Google or Amazon algorithms through innovative uses of ML that experiment with different types of participation.

Current Literature Landscape

The current state of research on AI and work is multidisciplinary and dispersed across disciplinary lines. Table 1 provides a few examples of research communities interested in the topic and some exemplary publications from each community. It is noteworthy that cross-cutting conferences and networks have recently provided researchers with opportunities to convene in workshops at conferences such as HICSS³ or WAIM Convergence Conferences⁴. Beyond the communities listed in Table 1 one finds an abundance of research exploring AI adoption and implications. These efforts often focus on non-work contexts such as the use of social media. For example, communication scholars have provided useful conceptualization of humans' interactions with 'AI-driven media' (Sundar, 2020) or 'communicative AI' (Guzman & Lewis, 2020), and have formed an interest group around communication between human and intelligent machine⁵.

³ <https://hicss.hawaii.edu>

⁴ <https://waim.network/events>

⁵ <https://www.icahdg.org/group/hmc>

Research community	Example of primary focus	Example of relevant publications and survey of literature
HCI/CSCW	Algorithmic decision making Explainability Algorithmic bias Human-Algorithm collaboration	(Lee, 2018) (Wolf, 2019) (Robert et al., 2020) (Gray & Suri, 2019)
Information and library science	AI anxiety Knowledge management Accountability AI and transforming libraries	(Johnson & Verdicchio, 2017) (Pee et al., 2019) (Griffey, 2019) (Hoffmann et al., 2018)
Organization/management science	Transformation of organizing practices AI and strategy HR practices Organizational decision making	(Metcalf et al., 2019) (Kietzmann & Pitt, 2020) (Kellogg et al., 2020) (Glikson & Woolley, 2020) (Haenlein & Kaplan, 2019)
Information systems	Datafication of work activities Human-AI hybrid tasks AI-based automation	(Gal et al., 2020) (Rai et al., 2019) (Gal et al., 2020) (Seeber et al., 2020)
Medical informatics	Accountability HR in medical contexts diagnosis and treatment recommendations	(T. Davenport & Kalakota, 2019) (Becker, 2019) (Maddox et al., 2019) (Meskó et al., 2018)

Table 1. Research on the issues of work and AI and examples of relevant publications

JASIST Special Issue

In closing, AI and the future of work stands as a central theme for educators, designers, and researchers in the information field (Duan et al., 2019). As suggested by the iConference workshop participants, we must consider both AI in the world of work and the world of work in AI to describe and intervene in the many possible futures of work. A forthcoming JASIST special issue will provide a venue for researchers hoping to further such an approach and help address the real need for new theoretical perspectives, methodological capacities and analytical tools that can integrate and extend insights on AI and work.

References

- Atkinson, R. D. (2019). *Robotics and the Future of Production and Work*. Information Technology and Innovation Foundation. Accessed April 26, 2020. <https://itif.org/publications/2019/10/15/robotics-and-future-production-and-work>
- Atkinson, R. D. (2019). *Robotics and the Future of Production and Work*. Information Technology and Innovation Foundation. <https://itif.org/publications/2019/10/15/robotics-and-future-production-and-work>
- Becker, A. (2019). Artificial intelligence in medicine: What is it doing for us today? *Health Policy and Technology*. <https://www.sciencedirect.com/science/article/pii/S2211883718301758>
- Brynjolfsson, E., & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370), 1530–1534.
- Challen, R., Denny, J., Pitt, M., Gompels, L., Edwards, T., & Tsaneva-Atanasova, K. (2019). Artificial intelligence, bias and clinical safety. *BMJ Quality & Safety*, 28(3), 231–237.
- Chui, M., Manyika, J., Miremadi, M., Henke, N., Chung, R., Nel, P., & Malhotra, S. (2018). Notes from the AI frontier: Insights from hundreds of use cases. *McKinsey Global Institute*. <http://straty.com/wp-content/uploads/2018/04/McKinsey-AI-Frontier-paper-April-2018.pdf>
- Czarniawska, B., & Joerges, B. (2020). *Robotization of Work?: Answers from Popular Culture, Media and Social Sciences*. Edward Elgar Publishing.
- Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48(1), 24–42.
- Davenport, T. H., & Kirby, J. (2016). *Only humans need apply: winners and losers in the age of smart machines*. Harper Business New York, NY.
- Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. In *Future Healthcare Journal* (Vol. 6, Issue 2, pp. 94–98). <https://doi.org/10.7861/futurehosp.6-2-94>
- Duan, Y., Edwards, J. S., & Dwivedi, Y. K. (2019). Artificial intelligence for decision making in the era of Big Data – evolution, challenges and research agenda. In *International Journal of Information Management* (Vol. 48, pp. 63–71). <https://doi.org/10.1016/j.ijinfomgt.2019.01.021>
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280.
- Gal, U., Jensen, T. B., & Stein, M.-K. (2020). Breaking the vicious cycle of algorithmic management: A virtue ethics approach to people analytics. *Information and Organization*, 30(2), 100301.
- Glikson, E., & Woolley, A. W. (2020). Human Trust in Artificial Intelligence: Review of Empirical Research. *Annals*. <https://doi.org/10.5465/annals.2018.0057>
- Gray, M. L., & Suri, S. (2019). *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.
- Griffey, J. (2019). *Artificial Intelligence and Machine Learning in Libraries*. American Library Association.

- Guzman, A. L., & Lewis, S. C. (2020). Artificial intelligence and communication: A Human–Machine Communication research agenda. *New Media & Society*, 22(1), 70–86.
- Haenlein, M., & Kaplan, A. (2019). A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence. *California Management Review*, 61(4), 5–14.
- Haraway, D. J. (1991). *Simians, cyborgs and women: the reinvention of nature*. Free Association Books.
- Hoffmann, A. L., Roberts, S. T., Wolf, C. T., & Wood, S. (2018). Beyond fairness, accountability, and transparency in the ethics of algorithms: Contributions and perspectives from LIS. *Proceedings of the Association for Information Science and Technology*, 55(1), 694–696.
- Jarrahi, M. H. (2019). In the age of the smart artificial intelligence: AI's dual capacities for automating and informing work. In *Business Information Review* (Vol. 36, Issue 4, pp. 178–187).
<https://doi.org/10.1177/0266382119883999>
- Johnson, D. G., & Verdicchio, M. (2017). AI Anxiety. *Journal of the Association for Information Science and Technology*, 68(9), 2267–2270.
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at Work: The New Contested Terrain of Control. *Annals*, 14(1), 366–410.
- Kietzmann, J., & Pitt, L. F. (2020). Artificial intelligence and machine learning: What managers need to know. *Business Horizons*, 63(2), 131–133.
- Lee, M. K. (2018). Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. In *Big Data & Society* (Vol. 5, Issue 1, p. 205395171875668).
<https://doi.org/10.1177/2053951718756684>
- Maddox, T. M., Rumsfeld, J. S., & Payne, P. R. O. (2019). Questions for Artificial Intelligence in Health Care. *JAMA: The Journal of the American Medical Association*, 321(1), 31–32.
- Maedche, A., Legner, C., Benlian, A., Berger, B., Gimpel, H., Hess, T., Hinz, O., Morana, S., & Sollner, M. (2019). AI-Based Digital Assistants: Opportunities, Threats, and Research Perspectives. *Bus Inf Syst Eng*, 61(4), 535–544.
- Meskó, B., Hetényi, G., & Györfy, Z. (2018). Will artificial intelligence solve the human resource crisis in healthcare? *BMC Health Services Research*, 18(1), 545.
- Metcalf, L., Askay, D. A., & Rosenberg, L. B. (2019). Keeping Humans in the Loop: Pooling Knowledge through Artificial Swarm Intelligence to Improve Business Decision Making. *California Management Review*, 61(4), 84–109.
- Muro, M., Whiton, J., & Maxim, R. (2019). *What jobs are affected by AI?: better-paid, better-educated workers face the most exposure*. <https://www.voced.edu.au/content/ngv:85736>
- Neff, G., & Nagy, P. (2018). Agency in the digital age: using symbiotic agency to explain human--technology interaction. In *A Networked Self and Human Augmentics, Artificial Intelligence, Sentience* (pp. 113–123). Routledge.
- Nhat, L. V., Wirtz, J., Kunz, W. H., Paluch, S., Gruber, T., Martins, A., & Patterson, P. G. (2020). Service robots, customers and service employees: what can we learn from the academic literature and where are the gaps? *Journal of Service Theory and Practice*, ahead-of-print(ahead-of-print).
<https://doi.org/10.1108/JSTP-04-2019-0088>
- Nichols, J. A., Herbert Chan, H. W., & Baker, M. A. B. (2019). Machine learning: applications of artificial intelligence to imaging and diagnosis. *Biophysical Reviews*, 11(1), 111–118.
- Østerlund, C., Crowston, K., & Jackson, C. (2020). Building an Apparatus: Refractive, Reflective, and Diffractive Readings of Trace Data. *Journal of the Association for Information Systems*, 21(1), 10.
- Pee, L. G., Pan, S. L., & Cui, L. (2019). Artificial intelligence in healthcare robots: A social informatics study of knowledge embodiment. *Journal of the Association for Information Science and Technology*, 70(4), 351–369.
- Quiroz, J. C., Laranjo, L., Kocaballi, A. B., Berkovsky, S., Rezazadegan, D., & Coiera, E. (2019). Challenges of developing a digital scribe to reduce clinical documentation burden. *NPJ Digital Medicine*, 2, 114.

- Rai, A., Constantinides, P., & Sarker, S. (2019). *Editor's comments: next-generation digital platforms: toward human–AI hybrids*. <https://dl.acm.org/doi/abs/10.5555/3370135.3370136>
- Robert, L. P., Pierce, C., Marquis, L., Kim, S., & Alahmad, R. (2020). Designing fair AI for managing employees in organizations: a review, critique, and design agenda. *Human–Computer Interaction*, 1–31.
- Seeber, I., Bittner, E., Briggs, R. O., de Vreede, T., de Vreede, G.-J., Elkins, A., Maier, R., Merz, A. B., Oeste-Reiß, S., Randrup, N., Schwabe, G., & Söllner, M. (2020). Machines as teammates: A research agenda on AI in team collaboration. *Information & Management*, 57(2), 103174.
- Singh, R., Kalra, M. K., Nitiwarangkul, C., Patti, J. A., Homayounieh, F., Padole, A., Rao, P., Putha, P., Muse, V. V., Sharma, A., & Digumarthy, S. R. (2018). Deep learning in chest radiography: Detection of findings and presence of change. *PLoS One*, 13(10), e0204155.
- Smith, P. (2020). *Human AI Partnership In Underwriting: A Task-Centered Analysis of the Division of Work* (M. H. Jarrahi (ed.)) [Undergraduate Honor]. University of North Carolina at Chapel Hill.
- Sundar, S. S. (2020). Rise of Machine Agency: A Framework for Studying the Psychology of Human–AI Interaction (HAI). *Journal of Computer-Mediated Communication: JCMC*, 25(1), 74–88.
- Topol, E. J. (2019). High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56.
- Wolf, C. T. (2019). Explainability Scenarios: Towards Scenario-based XAI Design. *Proceedings of the 24th International Conference on Intelligent User Interfaces*, 252–257.
- Wolf, C. T., & Blomberg, J. L. (2019). Evaluating the Promise of Human-Algorithm Collaborations in Everyday Work Practices. *PACMHCI*, 3 (CSCW), Article No 143.
- Zuboff, S. (1988). *In the age of the smart machine: The future of work and power* (Vol. 186). Basic books New York.