

RESEARCH PRIMER

Open Access

Primer on artificial intelligence and robotics



Manav Raj and Robert Seamans* 

* Correspondence: rseamans@stern.nyu.edu
NYU Stern School of Business, 44
West 4th Street, New York, NY
10012, USA

Abstract

This article provides an introduction to artificial intelligence, robotics, and research streams that examine the economic and organizational consequences of these and related technologies. We describe the nascent research on artificial intelligence and robotics in the economics and management literature and summarize the dominant approaches taken by scholars in this area. We discuss the implications of artificial intelligence, robotics, and automation for organizational design and firm strategy, argue for greater engagement with these topics by organizational and strategy researchers, and outline directions for future research.

Keywords: Automation, Artificial intelligence, Robotics, Future of work, Organizational design

Introduction

Artificial intelligence (AI) and robotics have become increasingly hot topics in the press and in academia. In October 2017, *Bloomberg* published an article claiming that artificial intelligence is likely to be the “most disruptive force in technology in the coming decade” and warning that firms that are slow to embrace the technology may risk extinction.¹ Similarly, the following month, the *Financial Times* declared that the “robot army” is transforming the global workplace.² This interest is likely due to the rapid gains that artificial intelligence has been making in some applications, such as image recognition and abstract strategy games, and that advanced robotics has been making in labs, even though widespread commercial applications may be lagging (Felten et al. 2018).

Scholars have been increasingly interested in the economic, social, and distributive implications of artificial intelligence, robotics, and other types of automation. For example, over the past 2 years, economists at the University of Toronto have convened conferences around the economics of artificial intelligence, which have been attended by a dazzling array of economics scholars from diverse point of views including Nobel Prize winners Edmund Phelps, Paul Romer, Joseph Stiglitz. Some research has taken a morez, and others.³ There are a number of well-attended conferences for legal, manufacturing, technical, and general-interest communities such as the World Conference on Robotics and Artificial Intelligence, WeRobot, and AI Now.

Organizational scholars are a bit late to the game and have only just started to focus on the organizational implications of artificial intelligence, robotics, and other types of advanced technologies. However, as we describe in this primer, we believe that these technologies present a unique opportunity for organizational scholars. Periods of great

technological change can bring about great progress but also great turmoil. For example, while the steam engine led to great economic growth (see, e.g., Crafts 2004) it also led to job displacement. It is important for organizations to understand and anticipate the effects that artificial intelligence, robotics, and other types of automation may have, and design themselves accordingly. While many lessons can be drawn from prior episodes of automation, it is possible that artificial intelligence and robotics may have unique consequences. Differences from prior episodes of automation include that (1) the nature of business activity has shifted dramatically over the past decade such that many businesses now rely on platform (i.e., 2-sided market) business models, (2) artificial intelligence is likely to affect white-collar workers more so than blue-collar workers (while perhaps robotics may affect blue-collar workers more than white-collar workers), and (3) artificial intelligence may affect the links between establishments and firms (e.g., monitoring and firm scope).

This article is a primer on artificial intelligence, robotics, and automation. To begin, we provide definitions of the constructs and describe the key questions that have been addressed so far. We discuss implications of these technologies on organizational design, then describe areas in which organizational scholars can make substantial contributions to our understanding about how artificial intelligence and robotics are affecting work, labor, and organizations. We also describe ways in which organizational scholars have been using artificial intelligence tools as part of their research methodology. Finally, we conclude with a call for more research in this fertile area.

Artificial intelligence, robotics, and automation: definitions and key questions

Definitions

Studies of artificial intelligence and robotics base their theory and analysis on constructs of automation, robotics, artificial intelligence and machine learning, and automation. In this body of literature, use of robotics, artificial intelligence, and machine learning technologies can be used both as independent and as dependent variables—as dependent variables to examine factors that encourage or discourage the adoption and use of these technologies and independent variables to see how the use of these technologies impacts a variety of outcomes, such as effects on labor, productivity, growth, and firm organization. It is important that organizational scholars carefully define any such constructs in their studies and to avoid confusing these related but distinct constructs. The definitions below are meant to be a helpful first step in such an endeavor.

Robotics

The International Federation of Robots (IFR), an international industrial group focused on commercial robotics, defines an industrial robot as an “automatically controlled, reprogrammable, multipurpose manipulator, programmable in three or more axes, which can be either fixed in place or mobile for use in industrial automation applications.”⁴ While this definition is a starting point, other roboticists may differ on dimensions such as whether a robot must be automatically controlled or could be autonomous or whether a robot must be reprogrammable. At a broader level, any machine that can be

used to carry out complex actions or tasks in an automatic manner may be considered a robot.

Artificial intelligence and machine learning

Similar to robotics, artificial intelligence is a construct with varying definitions and potentially broad interpretations. For starters, it is useful to distinguish between general and narrow artificial intelligence (Broussard 2018). “General artificial intelligence” refers to computer software that can think and act on its own; nothing like this currently exists. “Narrow artificial intelligence” refers to computer software that relies on highly sophisticated, algorithmic techniques to find patterns in data and make predictions about the future. In this sense, the software “learns” from existing data and hence is sometimes referred to as “machine learning” but this should not be confused with actual learning. Broussard (2018) writes that “machine ‘learning’ is more akin to a metaphor...: it means that the machine can improve at its programmed, routine, automated tasks. It doesn’t mean that the machine acquires knowledge or wisdom or agency, despite what the term *learning* might imply [p. 89].”

Many applications of machine learning focus on prediction and estimation of unknowns based on a given set of information (Athey 2018; Mullainathan and Spiess 2017). There are a variety of algorithms that can be used for this machine learning. Some of these techniques are relatively straightforward uses of logit models which would be familiar to most organizational scholars, whereas others involve highly sophisticated algorithms that attempt to mimic how a human brain looks for patterns in data (the latter are called “neural networks”). Artificial intelligence technology can be used towards a variety of purposes, including playing abstract strategy games such as Chess or Go; to playing real-time video games such as Atari, Asterix, or Crazy Climber; to image or street number recognition; to natural language translation; and many other uses.

Automation

Automation refers to the use of largely automatic, likely computer-controlled, systems and equipment in manufacturing and production processes that replace some or all of the tasks that previously were done by human labor. Automation is not a new concept, as innovations such as the steam engine or the cotton gin can be viewed as automating previously manual tasks. One of the concerns for scholars in this area revolves around how and in what contexts increased use of robotics and artificial intelligence technology may lead to increased automation, and the impact that this form of increased automation may have on the workforce and the design of organizations.

Disentangling artificial intelligence, robotics, and automation

While artificial intelligence, robotics, and automation are all related concepts, it is important to be aware of the distinctions between each of these constructs. Robotics is largely focused on technologies that could be classified as “manipulators” as per the IFR definition, and accordingly, more directly relates to the automation of physical tasks. On the other hand, artificial intelligence does not require physical manipulation, but rather computer-based learning. The distinction between the two technologies can

become fuzzier as applications of artificial intelligence may involve robotics or vice versa. For example, “smart robots” are robots that integrate machine learning and artificial intelligence to continuously improve the robots’ performance.

Both artificial intelligence and robotics technologies are capable of automation. However, an open question is how and whether the effects of automation may differ across the two technologies. Some scholars contend that computerization and the increased use of artificial intelligence have the potential to automate certain non-routine tasks compared to the more rote tasks previously subjected to automation (Frey and Osborne 2017; Autor et al. 2006). Accordingly, it is possible that technologies incorporating artificial intelligence may be able to automate far more tasks than pure robotics-based technologies.

Importantly, even though a technology such as artificial intelligence or robotics may automate some of the tasks previously done by human labor, it does not necessarily imply that the human has been automated out of a job. In many cases, a computer or robot may be able to complete relatively low-value tasks, freeing up the human to focus efforts instead on high-value tasks. In this sense, artificial intelligence and robotics may *augment* the work done by human labor.

Distinction from information and communication technology

In addition to the distinction across the concepts of robotics, artificial intelligence, and automation, we additionally draw readers’ attention to the contrast between artificial intelligence and robotics, and computerization and information technologies more generally. Similarly to robotics and artificial intelligence, information and communication technology (ICT) has been of interest to researchers and policymakers with regards to both its potential to increase productivity and its ability to affect labor (e.g., Autor et al. 2003; Bloom et al. 2014; Akerman et al. 2015). However, while artificial intelligence and robotics may reduce the cost of storing, communicating, and transmitting information much like ICT, they are distinct. ICT can refer to any form of computer-based information system (Powell and Dent-Micallef 1999), while artificial intelligence and robotics may be computer-based but are not necessarily information systems. This distinction can be especially difficult to navigate given the broadness and variation in the definitions used for robotics and artificial intelligence in the literature. Again, we urge organizational scholars to carefully define any of these constructs in their studies.

Key questions and areas of interest

Extant work on artificial intelligence and robotics addresses a number of major questions regarding the effect of these technologies on firms and individuals.

Artificial intelligence, robotics, and productivity

Research on robotics and artificial intelligence builds off of the substantial body of literature surrounding innovation and technological development. Innovation is a key factor in contributing to economic growth (Solow 1957; Romer 1990) and has been an area of interest for both theorists and policymakers for decades. Literature on robotics and automation has pointed to the impressive potential of these new technologies.

Brynjolfsson and McAfee (2017) claim that artificial intelligence has the potential to be “the most important general-purpose technology of our era.” Graetz and Michaels (2018) suggest that robotics added an estimated 0.37 percentage points to annual GDP growth for a panel of 17 countries from 1993 and 2007, an effect similar to that of the adoption of steam engines on economic growth during the industrial revolution.

Artificial intelligence, robotics, and labor

Historically, excitement around radical new technologies is tempered by anxieties regarding the potential for labor substitution (Mokyr et al. 2015). A body of work has shown that automation spurred by innovation can both complement and substitute for labor. Acemoglu and Restrepo (2018) examine how increased industrial robotics usage has impacted regional US labor markets between 1990 and 2007. Their findings suggest that the adoption of industrial robotics is negatively correlated with employment and wages—specifically that each additional robot reduced employment by six workers and that one new robot across a thousand workers reduced wages by 0.5%. Graetz and Michaels (2018) find that while wages increase with robot use, on average, hours worked drops for low- and middle-skilled workers. A similar study in Germany suggests that each additional industrial robot leads to a loss of two manufacturing jobs, but these jobs are offset by newly created roles in the service industry (Dauth et al. 2017).

Increasingly, work on automation considers or focuses on artificial intelligence rather than just robotics. Frey and Osborne (2017) predict how increased computerization, in particular, machine learning technologies, will affect non-routine tasks. Based on the tasks most involved in an occupation, the authors propose which occupations may be more or less at risk of automation in the future. Their results suggest that 47% of employment in the USA is at high risk of computerization. Frey and Osborne’s work has been applied by researchers in other countries. Using the same methodology, Brzeski and Burk (2015) suggest that 59% of the German workforce may be highly susceptible to automation, while Pajarinen and Rouvinen (2014) suggest that 35% of Finnish jobs are at high risk. Similar to the task-based approach utilized by Frey and Osborne, Brynjolfsson et al. (2018b) take a task-based approach to assess occupations’ suitability for machine learning. They show that occupations across the wage and wage bill spectrum are equally susceptible, suggesting that machine learning will likely affect different parts of the workforce than earlier waves of automation.

Work on automation and labor has focused on different units of analysis. Much of the existing work in economics has focused on the economy as a whole. For example, Frey and Osborne (2017) measure the risk of automation on an occupation by occupation level but consider the occupations at a global level. Similar work by McKinsey Global Institute (MGI 2017) does the same, and recent work by Accenture considers these at the country level (Accenture 2018). US-specific work has been done by Brynjolfsson et al. (2018b) and Felten et al. (2018). Some research has taken a more focused approach and highlights the effect of artificial intelligence and automation on specific sectors of the economy. For example, Acemoglu and Restrepo (2018) highlight that the largest effects of technology adoption will occur in manufacturing, especially among manual and blue-collar occupations and for workers without a college degree.

Distributional effects of artificial intelligence and robotics

Existing work on artificial intelligence and robotics has also attempted to identify “winners” and “losers” and to understand the distributional effects of these new technologies. A body of this work looks at cross-industry effects. Autor and Salomons (2018) show that industry-specific productivity increases are associated with a decrease of employment within the affected industry; however, positive spillovers in other sectors more than offset the negative own-industry effect. Similarly, Mandel (2017) examines brick-and-mortar retail stores during the rise of e-commerce and finds that new jobs created at fulfillment and call centers more than make up for job losses at department stores.

Other work looks at how skill composition can affect the potential complementary or substitution effects of these new technologies. A recent working paper by Choudhury et al. (2018) looks at performance effects of the use of artificial intelligence by workers with different types of training. They find productivity with artificial intelligence technology is highly affected by an individual’s background with computer science and engineering. Individuals who have requisite computer science or engineering skills are better able to unlock superior performance using artificial intelligence technologies than individuals without those skills. Felten et al. (2018) use an abilities-based approach to assess the link between recent advances in artificial intelligence and employment and wage growth. They find that occupations that require a relatively high proportion of software skills see growth in employment when affected by artificial intelligence, while other occupations do not see a meaningful relationship between the impact of artificial intelligence and employment growth.

Algorithmic decision-making and bias

There is a growing literature in economics, strategy, and information systems that studies the use of machine learning algorithms in decision-making. Some of the authors in this literature use disaggregated, micro-level data to draw insights as to how artificial intelligence affects firms or individuals differently depending on their background. Some of this work examines whether and how the use of artificial intelligence and machine learning tools affects individual biases. For example, machine-based algorithms appear to outperform judges in making decisions regarding potential detainment pre-trial and also reduce inequities (Kleinberg et al. 2018). Hoffman et al. (2017) find that managers who choose to hire against recommendations constructed by machine-based algorithms choose worse hires. Together, these results appear to suggest that machine learning algorithms may have potential in improving decision quality and equity.

However, other research cautions that machine learning algorithms often contain their own form of bias. For example, a machine learning algorithm designed to deliver advertisements for Science, Technology, Engineering, and Math occupations targeted men more than women, despite the fact that the advertisement was explicitly intended to be gender-neutral (Lambrech and Tucker 2018); Google’s Ad Settings machine learning algorithm displays fewer advertisements for high-paying jobs to females than to males (Datta et al. 2015); and artificial intelligence-based tools used in judicial decision-making appear to display racial biases (Angwin et al. 2016). While these biases are troubling, some argue that compared to the counterfactual of human decision-making, algorithmic processes offer improvements in quality and fairness, and

in particular, machine learning tools are best able to mitigate biases when human decision-makers exhibit bias and high levels of inconsistency (Cowgill 2019).

Recommender systems are a common tool on e-commerce platforms and frequently incorporate machine learning or artificial intelligence algorithms in the creation of their recommendations (Adomavicius and Tuzhilin 2005). Barach et al. (2018b) show that the use of recommendation systems for sellers can substitute for explicit monetary incentives in online marketplaces, highlighting one method by which firms can use artificial intelligence technologies to cut costs. Barach et al. (2018a, 2018b) study recommendation systems in online labor marketplaces and find that firms use AI-driven recommendations to identify an initial set of generally acceptable partners before relying on internal capabilities to select the best match. In particular, the use of the recommendation system is used less for specialized jobs and for experienced employees.

Other areas of interest

In addition to the above areas, research on artificial intelligence and robotics has started to examine a broader range of questions, such as how artificial intelligence may help stimulate innovation (Cockburn et al. 2018), the role of policy in an economy featuring artificial intelligence (Goolsbee 2018), and the role of artificial intelligence in international trade (Brynjolfsson et al. 2018a; Goldfarb and Treffer 2018). There are other important firm strategy and policy questions left to answer in this space such as the impact of artificial intelligence on firm structure, the factors that lead to increased adoption of these technologies, and distributional effects of artificial intelligence across industries, geographies, and occupations. However, aside from literature studying machine learning algorithms, research in this area has been slowed by a lack of available data, especially at the firm level. We discuss future directions of research below.

While there are some data sets containing information on the diffusion of robotics, it is largely at an aggregate level which does not allow for detailed microanalysis and differences across industries and regions can be obscured. There are currently no public data sets on the utilization or adoption of artificial intelligence at either the micro or the macro level, as the most complete sources of information are proprietary and inaccessible to the general public and the academic community (Raj and Seamans 2018; McElheran 2019). Despite these limitations, scholars studying management and organizations have constructed data sets and conducted research using trade magazines and other industry-specific resources. For example, using the industrial robotics industry as a setting, scholars have established that prior technological experience and technological knowledge are associated with greater innovative behavior following the introduction of a disruptive technology (Roy and Sarkar 2016; Roy and Islam 2017). Researchers have also used the industrial robotics industry as a setting to study organizational search and identify two distinct dimensions of search—search scope and search depth (Katila and Ahuja 2002). Nevertheless, the next stage in the evolution of research in this area should involve a proliferation of data to conduct a more focused and rigorous analysis of important questions regarding these technologies, firm adoption, and its consequences in an empirical manner.

Implications of artificial intelligence and robotics for organizational design

Historically, advances in technology have reshaped the workforce and our work habits and required organizations to adjust their design paradigms in dramatic ways. For example, in the last two decades, the rise of the Internet has led firms to increasingly embrace remote work and virtual teams which can cross geographic boundaries and use virtual means to coordinate actions (Kirkman and Mathieu 2005). A significant challenge for firms lies in recognizing when this reorganization is beneficial and what are the boundaries to adjusting to the new technology. Kirkman and Mathieu (2005) note the importance of weighing the “presses” that operate on real-world teams that influence the effectiveness of face-to-face interaction compared to virtual interactions.

Similarly, artificial intelligence and robotics technology have the capacity to reshape firms and change the structure of organizations dramatically. As discussed above, the adoption of artificial intelligence and robotics technologies will likely alter the bundle of skills and tasks that many occupations are comprised of. By that aspect alone, these technologies will reshape organizations and force firms to restructure themselves to account for these changes. Boundaries between occupations within firms are likely to shift as some tasks are automated, and individuals within firms that choose to adopt these technologies are likely to have greater exposure to computer technologies. In addition, the composition of the labor force may change to adopt to the new set of skills that are most valued. These changes are also likely to be reflected in the design of organizations as they seek configurations to get the most value out of their human capital.

Interfirm boundaries are also likely to shift as robotics and artificial intelligence technologies are adopted more widely. In a seminal article, Coase (1937) argues that firms will expand until the cost of organizing an additional transaction within the firm equals the cost of carrying out the same transaction on the market. Increased usage of artificial intelligence and robotics technology has the potential to greatly reduce costs within firms, potentially leading to fewer transactions on the market. Tasks that previously had to be contracted to other firms may now be able to be transferred in-house, or alternatively, firms may find that tasks that were done within the firm can be more efficiently done by other organizations with greater access and facility with these technologies. In addition, a firm may avoid adopting newer technologies such as robotics if the technology is highly specific to the firm and the firm faces risk of hold-up from an opportunistic downstream customer (Williamson 1985).

Regardless of what form the effect takes, the strategy literature consistently presents evidence that incumbent firms struggle during technological discontinuities (e.g., Tushman and Anderson 1986; Henderson and Clark 1990). Despite the challenges presented by radical innovation, incumbents can be successful when they are “pre-adapted,” and their historical capabilities and assets can be leveraged to take advantage of the new technology (Klepper 2002; Cattani 2006). In the specific context of robotics technology, Roy and Sarkar (2016) present evidence that the presence of in-house users of robots and access to scientific knowledge will best prepare firms to be flexible and adapt to new, “smarter” robotics technology. To the extent that this finding is generalizable, firms may consider employing individuals with experience with these technologies and increase their facility with scientific knowledge in the area to best be able to take advantage of potential benefits from adoption.

Future research directions for organizational scholars

There are a number of topics related to robotics, artificial intelligence, and automation that would benefit from research by organizational scholars. For example, the popular press tends to associate artificial intelligence and robotics with substitution, in part because of an assumption that productivity gains are at the expense of labor. The evidence does not support this conclusion, however. For example, Furman and Seamans (2019) show that there is no correlation between a country's labor force and its productivity, and Autor and Salomons (2018) show that while productivity growth may have a negative employment effect on the sector that experiences the growth, this is more than made up for by employment gains in related sectors.

More generally, there are reasons to expect that artificial intelligence will have complementary effects on labor. This has been the case for prior episodes of automation—for example, Bessen (2015) describes how the adoption of ATM machines by banks was associated with an increase in bank employment—and early evidence suggests it will be the same for artificial intelligence. Bessen et al. (2018) provide survey evidence that software sold by artificial intelligence startups is designed in most cases to augment the work that humans do. According to their findings, artificial intelligence startups are most likely to provide technology that helps their customers “make better predictions or decisions”, “manage and understand data better”, and “gain new capabilities to improve services or provide new products.” It is notable that these are all related to management and strategy. Given the dramatic impact that these technologies could have on labor and society, it is vitally important to have a clear understanding of the relationship between artificial intelligence, robotics, and labor. This is one area that we believe would greatly benefit from research by organizational scholars, who are adept at describing mechanisms affecting the organization of work.

There are a variety of other questions surrounding artificial intelligence and robotics that we encourage organizational scholars to turn to. One topic that has yet to be explored in much detail surrounds the establishment and firm-level consequences for adoption of artificial intelligence and robotics technology. Research could examine performance consequences as well as outcomes related to firm organization and strategy. Scholars can study in what circumstances and in what kinds of firms such adoption has the greatest impact. Additionally, adoption of these technologies within a firm may have consequences for the adopting firm as well as other firms in the industry, including firms upstream and downstream from the focal firm. The adoption of the technology itself can be viewed as an outcome, and scholars can examine what circumstances and factors encourage or discourage the use of these technologies. Certain industries, management styles, or organizational forms may be particularly quick to adopt, and market level forces may also impact the adoption decision. Industry and organizational factors may play a role as well as the backgrounds of individuals and managers within organizations. Greater work can be done to identify what factors contribute to adoption and differential effects once technology is adopted.

Further, more specific to management scholars, we need a detailed understanding about how artificial intelligence and robotics affect the nature of work. This includes not only how artificial intelligence and robotics change a given type of work or occupation (for example, by changing the relative importance of skills and tasks required for an occupation), but also how artificial intelligence and robotics affect the way in which

individuals interact with each other in the workplace. That is, we suspect that these technologies will change the type of work that we do, and also how that work is designed and organized as part of a larger production system.

To put this into perspective, in the early 2000s, online communication allowed for the creation of “virtual teams” (Jarvenpaa and Leidner 1999). Organizational scholars have highlighted many of the ways in which virtual teams need to be managed differently than non-virtual teams (e.g., Gilson et al. 2015; Kirkman and Mathieu 2005). Relatedly, we believe that a deeper understanding of how artificial intelligence is affecting workplace organization will help inform some of the economic studies of the effect of artificial intelligence on labor. More broadly, artificial intelligence and robotics are likely to substitute for labor in some cases, but complement labor in other cases. A better understanding of how work is done in the future will help inform conditions under which we can expect these technologies to be complementary to labor and when we should expect labor substitution.

The adoption and use of artificial intelligence and robotics technology also raises important questions with policy implications. Researchers can begin to examine the distributional effects of technology adoption across different demographics and regions. Feldman and Kogler (2010) show that industries, and even occupations within industries, tend to be geographically clustered. Because of that, the consequences of artificial intelligence and robotics may be far more pronounced in some geographies compared to others. In addition, to industry- and occupation-based differences, other factors may influence a company’s ability to take advantage of these technologies. For example, these new technologies may have significant implications for entrepreneurs. Entrepreneurs may lack knowledge of how best to integrate robotics with a workforce and often face financing constraints that make it harder for them to adopt capital-intensive technologies. In the case of artificial intelligence, entrepreneurs may lack data sets on customer behavior, which are needed to train artificial intelligence systems.

In the case that artificial intelligence and robotics do substitute for labor in certain industries or occupations, the labor market may look dramatically different from how it does now, and significant work will need to be done to help prepare the next generation of workers to adapt to the new environment. There will be a need to evaluate what skills and tasks are still valuable in the labor market compared to skills and tasks that can now be fully automated. This calls for a greater understanding of the worker experience in firms and occupations affected by artificial intelligence and robotics to craft appropriate worker education, job training, and re-training programs.

Machine learning and artificial intelligence tools for organizational research

In addition to being a subject of future research, machine learning and artificial intelligence technologies also offer potential as tools to be used by researchers in examining a broad range of questions. The computational abilities of artificial intelligence technologies open the door to analyses that were not previously feasible due to computational complexity.

Machine learning tools make fewer a priori assumptions regarding data when fitting models, and tools such as decision trees, random forests, K-nearest neighbors, and neural networks thus allow for the recognition of complex patterns and offer potential

for inductive theory building (Choudhury et al. 2019). Further, machine learning models can under certain conditions also improve causal inference with high-dimensional data and can help sort meaningful variables from confounding information (Belloni et al. 2014). As an example of a potential application, Guzman (2019) uses the double Least Absolute Shrinkage and Selection Operator (LASSO) methodology discussed by Belloni et al. (2014) to construct an appropriate counterfactual in his analysis of the effect of migration on start-up performance.

Additionally, machine learning tools allow for the analysis of unstructured textual data through natural language processing techniques, such as vector space models and topic modeling, and provide an opportunity to explore novel, difficult-to-measure constructs (Menon et al. 2018). For example, Furman and Teodoridis (2018) use topic modeling techniques to measure scientists' research trajectories, Kaplan and Vakili (2014) use topic modeling to construct a measure of novelty in patent applications, and Feng (2019) uses topic modeling to construct measures of local knowledge spillovers.

Conclusion

Artificial intelligence and robotics have experienced dramatic increases in performance, and this has led to greater funding for artificial intelligence and robotics start-ups, more popular press articles on how these technologies will change the world, and a recent increase in academic research around the consequences of these technologies for firms, workers, and economies. In this Primer, we define the key concepts, review the existing literature, identify implications for organizational design, and describe opportunities for organizational and strategy scholars.

Much of the literature that has been undertaken in this area focuses on how the adoption of robotics and artificial intelligence technologies affects economic growth and labor markets. This is still a prime area for further research given the important implications for social welfare. In addition, a lack of comprehensive data on the adoption and use of artificial intelligence and robots means that much of the existing work relies on expert or crowd-sourced opinions rather than empirical evidence (e.g., Frey and Osborne 2017; Brynjolfsson et al. 2018b; Felten et al. 2018). In the future, better collection and organization of data will allow for more direct empirical studies and will allow scholars to examine adjacent questions, such as differences in terms of performance and labor market consequences for different types of robotics or artificial intelligence technologies. We need evidence-based research on how artificial intelligence affects firm-level productivity, employment, and wages, as well as research on how artificial intelligence may affect economic outcomes with distributional consequences, such as innovation, business dynamism, and inequality.

There are multiple opportunities for organizational and strategy scholars to contribute to our understanding about how these technologies are affecting our society. In particular, we highlight the following questions as those that organizational and strategy scholars may be particularly well-suited to address:

- Which types of firms are more likely to adopt artificial intelligence and robotics technologies? Are there certain management styles or organizational forms that

may be particularly quick to adopt? Are there market level forces that impact the adoption decision?

- Do artificial intelligence and robotics increase or decrease inequality within an occupation, firm or region? Are there certain management or regulatory policies that can mitigate or exacerbate any detrimental effects of artificial intelligence and robotics?
- How does one firm's adoption of artificial intelligence and robotics affect its competitors in the same industry or market, as well as upstream suppliers and downstream customers? Under what conditions does use of artificial intelligence or robotics help new entrants compete with established incumbents?
- How do artificial intelligence and robotics affect the nature of work? In what ways do artificial intelligence and robotics change the relative importance of skills and tasks required for an occupation? How do artificial intelligence and robotics affect the way in which individual workers interact with each other in the workplace? Under what organizational conditions do artificial intelligence and robotics substitute or complement for labor?

Given the broad range of potential research questions, the far-reaching consequences of these technologies, and important practical and policy implications that may spill out of future work in this area, we believe that this is an exciting and fertile field for future research in organizations and management. It is our hope that this Primer can serve as a resource for organizational scholars who hope to build on this literature in the future.

Endnotes

¹<https://www.bloomberg.com/professional/blog/new-era-artificial-intelligence-now-biggest-tech-disrupter/>

²<https://www.ft.com/content/f04128de-c4a5-11e7-b2bb-322b2cb39656>

³<https://www.economicsofai.com/nber-conference-toronto-2017/>.

⁴<https://ifr.org/standardisation>

Authors' contributions

MR and RS equally co-authored the manuscript. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 24 October 2018 Accepted: 6 May 2019

Published online: 29 May 2019

References

- Accenture (2018) It's learning. Just not as we know it. Available at: <https://www.accenture.com/us-en/insights/future-workforce/transforming-learning>
- Acemoglu D, Restrepo P (2018) Robots and jobs: Evidence from US labor markets. NBER Working Paper
- Adomavicius G, Tuzhilin A (2005) Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. *IEEE Trans Knowl Data Eng* 17(6):734–749
- Akerman A, Gaarder I, Mogstad M (2015) The skill complementarity of broadband internet. NBER Working Paper
- Angwin J, Larson J, Mattu S, Kirchner L (2016) Machine bias: there's software used across the country to predict future criminals. And it's biased against blacks. ProPublica
- Athey S (2018) The impact of machine learning on economics. In: Agarwal AK, Gans J, Goldfarb A (eds) *The economics of artificial intelligence: an agenda*. University of Chicago Press

- Autor DH, Katz LF, Kearney MS (2006) The polarization of the US labor market. *AEA Papers and Proceedings* 96(2):189–194
- Autor DH, Levy F, Murnane RJ (2003) The skill content of recent technological change: an empirical exploration. *Q J Econ* 118(4):1279–1333
- Autor DH, Salomons A (2018) Is automation labor-displacing? Productivity growth, employment, and the labor share. National Bureau of Economic Research, Cambridge, p w24871
- Barach M, Golden JM, Horton JJ (2018b) Steering in online markets: the role of platform incentives and credibility. Working Paper
- Barach M, Kaul A, Leung M, Lu S (2018a) Small numbers bargaining in the age of big data: evidence from a two-sided labor matching platform. SSRN Working Paper
- Belloni A, Chernozhukov V, Hansen C (2014) High-dimensional methods and inference on structural and treatment effects. *J Econ Perspect* 28(2):29–50
- Bessen JE (2015) How Computer Automation Affects Occupations: Technology, Jobs, and Skills. Boston University School of Law, Law and Economics Working Paper No. 15-49.
- Bessen JE, Impink SM, Seamans RC, Reichensperger L (2018) The Business of AI Startups. Boston University School of Law, Law and Economics Research Paper No. 18-28.
- Bloom N, Garicano L, Sadun R, Van Reenen J (2014) The distinct effects of information technology and communication technology on firm organization. *Manag Sci* 60(12):2859–2885
- Broussard M (2018) Artificial unintelligence: how computers misunderstand the world. MIT Press, Cambridge
- Brynjolfsson E, Hui X, Lu M (2018a) Does machine translation affect international trade? Evidence from a large digital platform. National Bureau of Economic Research Working Paper (No. w24917)
- Brynjolfsson E, McAfee A (2017) The Business of Artificial Intelligence. *Harvard Business Review*, July 2017.
- Brynjolfsson E, Mitchell T, Rock D (2018b) What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings* 108:43–47
- Brzeski, C, Burk, I (2015) Die Roboter Kommen. ING DiBa Report
- Cattani G (2006) Technological pre-adaptation, speciation, and emergence of new technologies: how coming invented and developed Fiber optics. *Ind Corp Chang* 15(2):285–318
- Choudhury P, Allen R, Endres M (2019) Developing theory using machine learning models. SSRN Working Paper
- Choudhury P, Starr E, Agarwal R (2018) Machine learning and human capital: experimental evidence on productivity complementarities. SSRN Working Paper
- Coase RH (1937) The nature of the firm. *Economica* 4(16):386
- Cockburn I, Henderson R, Stern S (2018) The impact of artificial intelligence on innovation. NBER Working Paper
- Cowgill B (2019) Bias and productivity in humans and machines. Columbia Business School Working Paper
- Crafts N (2004) Globalisation and economic growth: a historical perspective. *World Econ* 27(1):45–58
- Datta A, Tschantz MC, Datta A (2015) Automated experiments on ad privacy settings. *Proceedings on Privacy Enhancing Technologies* 2015(1):92–112
- Dauth W, Findeisen S, Suedekum J, Woessner N (2017) German robots – the impact of industrial robots on workers. IAB Discussion Paper
- Feldman MP, Kogler DF (2010) Stylized facts in the geography of innovation. *Handbook of the Economics of Innovation* 1(C):381–410
- Felten EW, Raj M, Seamans RC (2018) A Method to Link Advances in Artificial Intelligence to Occupational Abilities. *AEA Papers and Proceedings* 108:54–57
- Feng FS (2019) The proximity of ideas: an analysis of patent text using machine learning. NYU Stern Working Paper
- Frey CB, Osborne MA (2017) The future of employment: how susceptible are jobs to computerisation? *Technol Forecast Soc Chang* 114(January):254–280
- Furman J, Seamans RC (2019) AI and the Economy. *Innovation Policy and the Economy* 19:161–191
- Furman J, Teodoridis F (2018) The cost of research tools and the direction of innovation: evidence from computer science and electrical engineering. SSRN Working Paper
- Gilson LL, Maynard MT, Jones Young NC, Vartiainen M, Hakonen M (2015) Virtual teams research: 10 years, 10 themes, and 10 opportunities. *J Manag* 41(5):1313–1337
- Goldfarb A, Trefler D (2018) AI and international trade. NBER Working Paper
- Goosbee A (2018) Public policy in an AI economy. National Bureau of Economic Research, Cambridge, p w24653
- Graetz G, Michaels G (2018) Robots at work. *Rev Econ Stat.* 100(5):753–768
- Guzman J (2019) Go west young firm: agglomeration and embeddedness in startup migrations to Silicon Valley. SSRN Working Paper
- Henderson RM, Clark KB (1990) Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. *Adm Sci Q* 35(1):9–30
- Hoffman M, Kahn LB, Li D (2017) Discretion in hiring. *Q J Econ* 133(2):765–800
- Jarvenpaa SL, Leidner DE (1999) Communication and trust in global virtual teams. *Organ Sci* 10(6):791–815
- Kaplan S, Vakili K (2014) The double-edged sword of recombination in breakthrough innovation. *Strateg Manag J* 36(10):1435–1457
- Katila R, Ahuja G (2002) Something old, something new: a longitudinal study of search behavior and new product introduction. *Acad Manag J* 45(6):1183–1194
- Kirkman BL, Mathieu JE (2005) The dimensions and antecedents of team virtuality. *J Manag* 31(5):700–718
- Kleinberg J, Lakkaraju H, Leskovec J, Ludwig J, Mullainathan S (2018) Human decisions and machine predictions. *Q J Econ* 133(1):237–293
- Klepper S (2002) Firm survival and the evolution of oligopoly. *RAND J Econ* 33(1):37–61
- Lambrecht A, Tucker CE (2018) Algorithmic Bias? An empirical study into apparent gender-based discrimination in the display of STEM career ads. SSRN Working Paper
- Mandel M (2017) How ecommerce creates jobs and reduces income inequality. Progressive Policy Institute
- McElheran K (2019) Economic measurement of AI. NBER Working Paper
- McKinsey Global Institute (MGI). 2017. "A future that works: automation, employment, and productivity." Available at: <https://www.mckinsey.com/~/media/mckinsey/featured%20insights/Digital%20Disruption/Harnessing%20automation%20for%20a%20future%20that%20works/MGI-A-future-that-works-Executive-summary.ashx>

- Menon A, Choi J, Tabakovic H (2018) What you say your strategy is and why it matters: natural language processing of unstructured text. *Acad Manag Proc* 2018(1)
- Mokyr J, Vickers C, Ziebarth NL (2015) The history of technological anxiety and the future of economic growth: is this time different? *J Econ Perspect* 29(3):31–50
- Mullainathan S, Spiess J (2017) Machine learning: an applied econometric approach. *J Econ Perspect* 31(2):87–106
- Pajarinen M, Rouvinen P (2014) Computerization threatens one third of Finnish employment. ETLA Brief from The Research Institute of the Finnish Economy
- Powell TC, Dent-Micallef A (1999) Information technology as competitive advantage: the role of human, business, and technology resources. *Strateg Manag J* 18(5):375–405
- Raj M, Seamans RC (2018) AI, Labor, Productivity, and the Need for Firm-Level Data. In: Agrawal A, Gans JS, Goldfarb A (eds) *NBER Economics of Artificial Intelligence*. University of Chicago Press, Chicago
- Romer PM (1990) Endogenous technological change. *J Polit Econ* 98(5):S71–S102
- Roy R, Islam M (2017) Nuanced role of relevant prior experience: sales takeoff of disruptive products and product innovation with disrupted technology in industrial robotics. In: Furman J, Gawer A, Silverman BS, Stern S (eds) *Advances in strategic management*, vol 37. Emerald Publishing Limited, pp 81–111
- Roy R, Sarkar MB (2016) Knowledge, firm boundaries, and innovation: mitigating the Incumbent's curse during radical technological change: mitigating Incumbent's curse during radical discontinuity. *Strateg Manag J* 37(5):835–854
- Solow RM (1957) Technical change and the aggregate production function. *Rev Econ Stat* 39(3):312–320
- Tushman ML, Anderson P (1986) Technological discontinuities and organizational environments. *Adm Sci Q* 31(3):439–465
- Williamson OE (1985) *The economic institutions of capitalism*. Macmillan, New York

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)
