

Working and Organizing in the Age of the Learning Algorithm

Samer Faraj

McGill University

Stella Pachidi

University of Cambridge

Karla Sayegh

McGill University

January 2018

For Information and Organization

Abstract

Learning algorithms, technologies that generate responses, classifications, or dynamic predictions that resemble those of a knowledge worker, raise important research questions for organizational scholars related to work and organizing. We suggest that such algorithms are distinguished by four consequential aspects: black-boxed performance, comprehensive digitization, anticipatory quantification, and hidden politics. These aspects are likely to alter work and organizing in qualitatively different ways beyond simply signaling an acceleration of long-term technology trends. Our analysis indicates that learning algorithms will transform expertise in organizations, reshape work and occupational boundaries, and offer novel forms of coordination and control. Thus, learning algorithms can be considered performative due to the extent to which their use can shape and alter work and organizational realities. Their rapid deployment requires scholarly attention to societal issues such as the extent to which the algorithm is authorized to make decisions, the need to incorporate morality in the technology, and their digital iron-cage potential.

Introduction

Recent technological developments have given rise to algorithms that are capable of learning from data to undertake tasks that previously required human judgement. Such a development could be the most consequential to work and organizing since the beginning of the digital age sixty years ago. Algorithms have historically been understood as a program containing a fixed sequence of instructions executed until a solution is reached (Hopcroft and Ullman, 1983). Algorithms function by taking in data that has been restructured, formatted, and prepared for processing. With continuous advances in programming science, coupled with Moore's law (the continuous increase in hardware computing power), algorithms are becoming more predominant and performant, reaching a crucial stage of evolution. With the digitization of most economic and social interactions, algorithmic technologies are now being used on a hitherto unforeseen and unknown scale. Driving the change is the emergent set of learning technologies that were developed in the field of artificial intelligence. For the first time since the rise of the field of artificial intelligence half a century ago, learning algorithms are able to perform reliably, and potentially exceed, an increasing array of tasks that historically had been the domain of humans. We use the term learning algorithms to refer to an emergent family of technologies that build on machine learning, computation, and statistical techniques, as well as rely on large data sets to generate responses, classifications, or dynamic predictions that resemble those of a knowledge worker.

Beyond the recognition that a growing part of humanity now collaborates and communicates digitally, it is important to recognize that such technologies are increasingly impacting broader aspects of society. For example, algorithms change newsroom culture by providing a measure of the readership of online newspaper articles, to emphasize stories likely to be popular (Christin, 2014). They alter legal procedure and judicial discretion by introducing statistical models of crime prediction in legal proceedings and identifying perpetrators (Christin, 2016). Algorithms that manage engine functioning have been used by car manufacturers to defeat engine pollution tests (Ramsey, 2015) or by Uber to avoid city inspections (Krisher, 2017). Algorithms are also constitutive for the modern electoral campaigns, as they help sway votes by targeting individual citizens with customized cross-platform communications that emphasize select memes, images, keywords, hashtags, analyses and stories, based on an analysis of the content of a person's corpus of online communication (Cadwalladr, 2017; Shaffer, Carey, and Starling, 2017).

For many scholars, algorithmic technologies do not seem to qualitatively differ from previous developments of the digital age (e.g., the personal computer, the Internet, Web 2.0) that have each heralded novel ways of organizing or new business models. They enable digital platforms and the shift to the short-term employment. Algorithm developments positively increase economic output by broadly increasing consumption (Brynjolfsson and McAfee, 2014). For firms, they increased economic gain via improved targeting capacities, reduced costs, and enhanced productivity (Newell and Marabelli, 2015). Beyond economic value, however, predictions of future consequences are mixed. Some analyses emphasize that learning algorithms are likely to take over major aspects of the knowledge economy, threatening to extinguish up to half of full-time jobs in the United States within two decades (Frey and Osborne, 2017). Despite the increasing prevalence of learning algorithms in every facet of work and life, and the multitude of associated predictions surrounding the future of work, scholars and practitioners alike are only now beginning to understand their transformative effects on work and organizing. The debate centers on the extent to which algorithms can take over many aspects of human work. What will matter is the capacity of contemporary workers to adapt their ways of knowing and working and embrace novel technologies, with

augmentative effects (Brynjolfsson and Mitchell, 2017; Nelson and Irwin, 2014). In the balance of the paper, we examine emerging characteristics of algorithmic technologies and discuss how they might transform work and organizing, with implications for future research

Consequential aspects of learning algorithms

In this section we examine certain emergent aspects of learning algorithms that make them highly consequential for work and organizing. More specifically, we distinguish four consequential aspects: black-boxed performance, comprehensive digitization, anticipatory quantification, and hidden politics.

Black-boxed performance

In today's digital era, users of web technologies are constantly presented with search results, ads, or personalized recommendations that cannot be easily explained, redressed, or adjusted because the algorithms behind them are opaque, inaccessible, and unmodifiable. For one, sophisticated algorithms such as Google's search, Facebook's newsfeed, or Amazon's recommendation algorithm constitute the "secret sauce" that is crucial to business success. They are worth billions of dollars and are tightly protected as intellectual property (O'Neil 2017). Design choices are internal to the firm and unavailable for public scrutiny for both commercial and competitive reasons. As a result, unknown design choices that reflect social assumptions and economic valuation become inscribed into the software. Different people putting in the same search term are likely to generate different recommendations because the algorithm takes into account additional individual and contextual factors. Even if auditing were to be applied to such algorithms, understanding them would be limited only to a select professional class of people with highly specialized skills and technical training for comprehending code of immense size and logical complexity (Dourish, 2016).

While the above sources of inscrutability could be said to apply to other more traditional types of algorithms as well, learning algorithms are even more opaque because they do not rely on pre-specified instructions, but on evolving weights and networks of connections that get refined with each additional data point (Burrell, 2016; Michalski, Carbonell, and Mitchell, 2013). As a matter of fact, even if a learning algorithm is coded in a very simple and comprehensible way, understanding how it arrived at its result can still be unclear, even to its developers, due to the quantity and complex interaction of the data fed into the algorithm. For example, when activists in the Occupy Wall Street movement accused Twitter of censoring their activities from becoming a trending topic, the engineers at Twitter insisted that no censorship had taken place, and while they could explain the code that they had written, they were unable to explain why the Occupy Wall Street activities did not become a trending topic (Dourish, 2016).

To some extent, the inscrutability of learning algorithms could prove to be a positive aspect because the less people understand how the learning algorithm comes to a certain outcome, the less they will try to "game" the system (O'Neil, 2016). This could result in decreasing corruption in various instances such as credit scoring in loan application processes, competitive tendering, or evaluation of work performance. From that perspective, learning algorithms can function as panopticons, imposing discipline that is often desired in the workplace (Burton-Jones, 2014; Zuboff, 1988). However, black-boxed performance may also result in various problematic situations. In the case of legal learning algorithms that apply high legal standards to the judgement of legal cases, if defendants ask why they were convicted, it would not be fair

to get a response from the system's developers that 'the system trained on tons of data and thus came to this decision' (Tegmark, 2017). Similarly, in an accident with a self-driving car, it could be difficult to identify the specific reasons that led the self-driving car to act in the way it did, so as to allocate legal responsibility accordingly (Knight, 2017). Overall, learning algorithms increasingly produce unknown and unexpected outcomes that cannot be explained by humans in a straightforward manner (Dourish, 2016). As we discuss later in greater detail, this poses significant consequences for various facets of work and organizing.

Comprehensive Digitization

The performance of learning algorithms improves with the provision of very large and accurate data sets. An important aspect of today's algorithms is the combining of comprehensive data extracted in one context with data from other contexts and possibly of different nature. For example, the digitization of various markers of the physical and built environment provides digital maps and GPS based location data. The digitization of devices and machines (known as the Internet of Things) adds another source of important data. Adding to these the digitization of economic transactions and human interactions, along with the digital tracing of human activities through mobile devices, we accomplish "the digitization of just about everything" (Brynjolfsson and McAfee 2014, p. 65).

This amalgamation of previously separate facets of life allows for algorithms to use a much richer set of inputs with the goal of arriving at more precise predictions. For instance, an Amazon bookstore algorithm with access to various facets of a professor's social life, published research, and utterances on Twitter or on various social networks can potentially provide more accurate book suggestions. Companies such as Facebook collect approximately 100 different dimensions of user activity on their platform and combine them with third-party databases that provide complementary information regarding income, net worth, house value, banking and financial information, charity donation, digital radio listening, podcast downloads, and pharmacy purchases. Such a mashup allows the targeting of consumers with more precisely appealing advertising and offers (Bergstein, 2017).

However, this algorithmic dependence on ubiquitous data is limited by its sensitivity to what data is available irrespective of the veracity of the content. For example, search engine optimization allows the skewing of algorithmic results in order to ensure certain results are listed first. Autocompletion suggestions given by Google and other search engines can be extremely useful for users. However, these same results can manifest racism, gender bias, and other non-neutral views due to the sensitivity of the learning algorithm to historical search activities (Cadwalladr, 2016). As seen in the recent example of the Microsoft TAY social media chatbot that learned from the social media streams and from interactions, algorithms can generate results that may be racist and misogynist (Shah and Chakkattu, 2016). For example, as seen in the recent 2016 US election, the logic of sharing information that aligns with the users' expressed opinions and preferences (induced from their tracked behavior on social media) can lead to a narrowing of intellectual horizon and a propagation of the phenomenon of "fake news" (Olson, 2016). Furthermore, manipulating algorithms to show alternative viewpoints can lead to deeper polarization or a more complete rejection of other perspectives (Zuckerberg, 2017).

Anticipatory quantification

The output of learning algorithms is principally a prediction or likelihood of an outcome. Given

the lowering costs and easy availability and processing of vast amounts of trace data, it comes as no surprise that our society is becoming dominated by the logic of quantification (Espeland and Stevens, 2008). With most objects being tagged, locations identified, people's attributes marked, behavior traced, and interactions mapped, all aspects of working and living can be digitally represented and quantified (Wagner-Pacifici et al., 2015). This quantitative turn is based on the assumption that digital data can stand for social life; that is to say, the detailed measurement of different aspects of a person's activities and utterances get to be constitutive of that person.

Algorithms can plow through an immense quantity and breadth of data to identify patterns and correlations. By studying a large enough number of people and by covering enough of their actions/utterances, algorithms seem to know people in counterintuitive, but uncannily accurate ways. For example, an analysis of less than a hundred of a person's Facebook "likes" can lead to the identification with high accuracy of a person's skin color, sexual orientation, or political affiliation (Grassegger and Krogerus, 2017). More impressively, an algorithmic analysis of less than 300 of such "likes" can often predict an individual's personality better than their friends, family, or even spouse (Youyou, Kosinski and Stillwell, 2015). In short, this ability to build statistical models that target individuals via customized messages that leverage aspects of personality, political leanings, and affective proclivities is the culmination of the digital quantification logic.

Learning algorithms can be reductionist because they inscribe predictive models of action based on a limited set of predictive variables emerging from the model's underlying data. Predictive modeling builds on correlational analysis of a large number of dimensions to generate a predictive model that focuses on future actions. A possible impact is to reduce individuals to a set of measured dimensions and such a representation avoids dealing with a person's (or object's) evolution and discourages alternative explorations that may explain how one ends up in a specific category (Ananny, 2016). For example, is the Amazon "customers who bought this item also bought" algorithm adequately accurate in predicting similar reading preferences between two adults because they both purchased books for young children? This focus on assigning people to categories is somewhat similar to the thought process of an individual who sees others primarily through the lens of ethnic categories and predictively assigns stereotypical actions/opinions to them. Learning algorithms thus cannot, as currently constituted, take into account the complex nature of individuals, their intentionality, or their desires, leading to a necessarily skewed, imperfect, and incomplete digital representation.

As the regime of quantification enters multiple fields, from journalism and academic practice to the management of human resources, such algorithmic outcomes appear to be more objective than humans (Boyd and Crawford, 2012), reducing the output of decision making processes to the output of a predictive model. As a result, people may be evaluated based on their propensity to act rather than on their actions (Mayer-Schönberger and Cukier, 2012). For example, several states in the US are already allowing a proprietary algorithm to decide on whether inmates should be granted parole, based on the likelihood of not committing a crime (Israni, 2017). Furthermore, as algorithms are automatically executed, they can now operate with no human judgment to mitigate their operation (Introna, 2016). Yet, by processing people and situations into specific buckets, the algorithm is seldom neutral. Some predictions, such as when a company tries to determine which customers are most likely to be receptive to its salespeople, are less consequential. Other predictions, however, such as no-fly lists or predictive policing watch lists, are less benign for the "false-positives" on the list. This unbound empiricism seems to lack feedback mechanisms that allow users

to review their publicly available data records for errors/omissions or to explain certain associations that put them in one bucket or another or even choose the buckets they wish to associate with (Tufekci, 2014). Individuals are effectively reduced to how the algorithm categorizes them, with serious implications such as erroneously being assigned a high probability of committing an undesirable act.

Hidden politics

As with other technologies, algorithms are political in their design (Akrich, 1992; Winner, 1980). For one, algorithms are imbued with the value choices of their designers, whether these have been made implicitly or explicitly (Introna, 2016). The criteria included in the algorithmic processes reflect valuation schemes, beliefs, and ethical standards. Those design decisions are not necessarily the result of a collective decision, or are written in official codes of conduct. Instead, they are often made in an informal, intuitive, and idiosyncratic way by the individual who designs the algorithm (Annany, 2016; Kitchin, 2014). As a result, whether intendedly or unintendedly, algorithms become artifacts with political qualities, giving certain people, objects, ideas, and events higher power, status, visibility, or importance than others (Annany, 2016). For example, search engine algorithms select and rank web pages while using a set of decision criteria (such as keywords, stop words, or advertising bids) that have been pre-defined by their designers (Introna and Nissenbaum, 2000). It would come as no surprise then that news about social movements are often excluded from the search engines of certain countries, or that Google would eventually be accused of prioritising its own related websites in relevant search queries (European Commission, 2017).

Beyond designers' intentions, algorithms may have political consequences due to the datasets that they are entwined with. Politics play out in the classification, selection, and pre-processing of the data that is fed into an algorithmic technology. Such categorizations, inclusions, and exclusions constitute a significant political and semantic parameter that influences the algorithmic outcome (Bowker and Star, 2000). For example, in credit-scoring loan applications, the classification process that determines what constitutes a credit-worthy applicant will inevitably include the values and beliefs of the practitioners who do the pre-processing of the data and the pre-classification of the training dataset, thus possibly introducing bias (O'Neil, 2016). If the algorithm prioritizes a negative correlation between living in a low income neighbourhood and loan repayment likelihood, the algorithm could end up discriminating against minority applicants in ways contrary to existing laws. Human supervision and monitoring may become necessary when the algorithm has legal or political consequences. Finally, as the proliferation of fake news on social media demonstrates, there is no easy way for an algorithm that prioritizes the sharing of links by users to identify in an "objective" manner the veracity of news and what to do about it (Dewey, 2016). Considering data to be raw is an "oxymoron" (Bowker, 2005; Gitelman, 2013), meaning that data is never pure, raw, or objective. When data is personal, sensitive, or can lead to economic/reputational consequences, it can no longer be considered as apolitical.

Consequences for work and organizing

Learning algorithms are likely to transform knowledge work and its organization in qualitatively different ways than historical technology advances have done. On the surface, we are likely to continue observing increases in automation, gains in productivity, job gains/losses, and the emergence of novel occupational categories, as we have done in the past. However, the combined effects of black-boxed performance, comprehensive digitization, anticipatory quantification, and hidden politics that characterize

learning algorithms are likely to raise unexpected challenges that organizations, occupations, and professions must confront. In the following sections, we discuss several consequential areas of inquiry that can serve as points of departure for future studies of algorithms and organizing.

Transforming expertise

Routine cognitive tasks that are well-defined and can be reduced to sets of rules and patterns are likely to be taken over by learning algorithms. Advancements in deep learning approaches coupled with comprehensive digitization mean that learning algorithms can now emulate the ways in which tacit knowledge is acquired by workers. For example, machine learning techniques involving image recognition, a process that is difficult to explicitly articulate and explain but that we, as humans, develop through experience, can now discriminate with an accuracy that surpasses our own (Esteva et al., 2017). For most jobs with a tacit knowledge component, up to 30 percent of tasks could be directly taken over by algorithms (Manyika et al., 2017). The challenge for workers and organizations is to understand which parts of jobs are most likely to be automated, how it affects the rest of the job, and how to manage the tension inherent in the emergent entwining between people and the algorithm.

As with other technologies that threaten the legitimacy of professional expertise, learning algorithms are likely to be resisted by occupational members. As learning algorithms enter jurisdictional spaces traditionally controlled by knowledge workers, their coexistence with humans is likely to lead to clashes over what type of expertise is valued and whose expertise has primacy. For example, a traditional salesperson who is used to relational ways of knowing that involve personal and frequent contact with the customer and detailed knowledge of the client's context, will likely need to shift to new ways of knowing built around an analytics algorithm that constructs a model of sales and then identifies sales opportunities based on prediction modelling (Pachidi et al., 2014). These algorithmic ways of knowing may be difficult to adopt if knowledge workers are unable to learn, manage, and control the technology's black-boxed performance. In cases in which the accuracy of the algorithmic output proves to be higher and leads to devaluation of human expertise, it is likely that the affected occupations will devote more of their identities and tasks to monitoring, improving upon, leveraging, and auditing algorithms.

For skilled professions whose expertise and training are dependent upon tasks suitable for learning algorithms (image recognition tasks for example), the reliance on such technologies by incumbents for routine tasks may threaten the development of the profession's future experts. If trainees must make do with decreasing participation in the work that they must ultimately perform, legitimate peripheral participants will face challenges in progressing to the center of their practice communities (Lave and Wenger, 1991). Further, over time, as incumbent experts retire, the replenishment of the occupational expertise that understands the tasks taken over by the algorithm is in question. The medical profession is a case in point. Intelligent heart and lung image recognition algorithms are already guiding radiologists toward certain diagnoses, by identifying suspicious tissue and attaching specific probabilities of malignancy to them, and compare favorably with board-certified doctors (Mukherjee, 2017). A recent study about the introduction of robotic surgery in teaching hospitals showed how trainees, no longer able to assist surgeons directly at the operating table, worked around norms and conventions of the surgical profession to learn in vicarious, self-directed, and less legitimate ways with mixed outcomes (Beane, 2018). For other occupational groups, algorithms may be leveraged to train newcomers more efficiently. For example, Udacity has already developed automated sales support systems built on chat room data designed to

conversationally coach in real time, rather than taking over the work of customer service representatives (Brynjolfsson and McAfee, 2017). Thus, while short-term efficiency, productivity, and value gains from learning algorithms are evident, they must be balanced with long-term consequences for the quality and quantity of professional expertise and the future of skilled work.

Reshaping occupational work and boundaries

Beyond taking over routine components of knowledge work, the deployment of learning algorithms may foster the need for a changed form of occupational accountability. If algorithmic predictions were automatically turned into action without additional scrutiny or human monitoring, who becomes accountable for “mistakes,” misdiagnoses, or for when erroneous convictions become contested. Simply, learning algorithms require humans to ensure accountability (Lupton and Jutel, 2015). There could be extremely atypical cases that the algorithm may miss and that a human could still recognize. In our radiology example, failure of the algorithm to identify a rare form of cancer or to misdiagnose a scan because of an interfering condition not previously encountered by the algorithm, can have life-altering consequences for the patient. In addition, the algorithm assumes a systematic process by which a human tacitly performs every “routine task” (Bergstein, 2018). For example, in predicting how a judge may rule in a criminal case (through patterns of complex relationships between the facts of a case and prior decisions for similar cases), the algorithms would assume that the judge would come up with the same results for all cases with the same characteristics (Remus and Levy, 2017). But the hallmark of human decision making, especially in highly consequential situations, is to rely not just on previously encountered patterns, but also on a highly diverse set of contextual factors. Thus, the judge in practice can vary his/her decision-making “steps” for cases with similar characteristics based on factors not captured in the training data but germane to the case, limiting the predictive power of the algorithm.

Much attention has been focused on the ability of learning algorithms to take over many of the repetitive and routine tasks that comprise most knowledge jobs (Frey and Osborne, 2017; Manyika et al., 2017). It is unclear whether professions will be unable to adapt, as they have always done in periods of technological change (Anteby et al., 2017). A more likely outcome is for algorithms to augment the importance of human judgement in knowledge work. The case of radiology could be illustrative of how knowledge work may evolve. On one hand, some computer scientists contend that hospitals “should stop training radiologists now” (Mukherjee, 2017, p. 12), because algorithms are starting to recognize malignancies with the same level of accuracy as the typical radiologist. However, a more likely scenario is for the radiologists to happily relinquish the preliminary scanning of medical images to the algorithm in order to focus on the more complicated cases, or those that the algorithm finds ambiguous. Another possible outcome is the migration of image reading to other professionals such as emergency department doctors who can now rely on the algorithm to identify urgent medical issues without waiting for the radiologist report. While it is possible that there will be a need for a lower number of radiologists overall, those who remain in the profession are likely to become even more specialized, engaging in work requiring complex interpretation and judgement.

Learning algorithms are likely to usher a transformation of many professions away from their current focus and priorities toward new identities and ways of working. Some occupations may leverage the capacities of the technology to provide informational insights for learning to repurpose their expertise and re-conceptualize their identities. For example, in the past librarians redefined their domains of expertise

following the introduction and development of internet search. Over time, and as search technologies became more sophisticated, librarians became “connectors of people and ideas” by helping information-seekers sort through vast volumes of information, as opposed to simply performing searches for them (Nelson and Irwin, 2014). More recently, a study of technology-enabled crowdsourcing at NASA traced how resident scientists redefined themselves as solution-seekers rather than problem solvers. In so doing, they did not attempt to acquire and understand the necessary scientific knowledge required to solve the complex scientific problems they put to the crowd, but rather used the technologies to become experts at finding the rare others who could (Lifshitz-Assaf, 2017). Similar changes are likely to emerge for all professions affected by learning algorithms and will be accompanied by shifting identities and renegotiated occupational boundaries.

Providing new forms of control

Learning algorithms are gradually producing a new form of organizational control. Across a variety of jobs with structured, decomposable sub-tasks or with concrete deliverables that are subject to ratings or evaluation by the crowds of users, the anticipatory quantification aspect of algorithms is bounding the action possibilities available to individual workers and reducing their discretion (Orlikowski and Scott, 2014; Annany, 2016). For example, Amazon workers (cynically referred to as “Amabots”) are often being evaluated primarily by the time it takes them to locate and secure packages off the shelf for shipping. Similarly, journalists are increasingly being ranked and compared based on algorithm-generated metrics of online article popularity. Uber drivers are dynamically receiving different income for the same work based on time-based analytics with little forewarning, explanation, or redress (O’Connor, 2016). A recent enterprise-based communication application, Slack, is using learning algorithms to predict which messages are of higher importance and give a sense of the topics each employee participates in (Woyke, 2018). Thus, knowledge workers face an interesting paradox in today’s algorithm-imbued work settings: the recorded actions that may comprise the basis for their agency and control might be used to regulate or even erode this control. The availability of traces of human behavior can constrain agency in subtle ways that are not readily discernible to users shut out of the design, deployment, and evolution of the learning algorithms (Tufekci, 2014). Thus, the learning algorithm has the potential to revive the piecework system, one that is minutely monitored, only this time with an algorithmic managerial span of control whose logic is opaque.

To counteract these structuring forces, knowledge workers will likely engage in counter-performances designed to game algorithmic directives and their anticipatory quantification: they may employ gaming strategies, limit or censor their participation, or use political resources to alter the basis for their evaluation. First, gaming strategies are already widespread. For example, in order to produce substantive articles of their choosing, journalists have been reported to buy ratings performance by producing several quick-and-dirty clickbait articles to “reset the score” (Christin, 2014). Recent media coverage of the FitBit wearable exposed how workers gamed the device and manipulated their fitness data to accumulate rewards, with consequences for organizations that promote their use to keep down health insurance premiums (Bachman, 2016; Weiczner, 2016). Second, users may limit their participation or participate selectively. In the Slack example, because utterances/actions are heavily analyzed, knowledge workers may choose not to engage in debates or questioning and to contribute curated comments that they believe will be assessed favorably by the algorithm only. Lastly, to counteract the monitoring of the algorithms, occupational communities may engage in political framing strategies from inside organizations to convince their hierarchies to change the value-imbued assumptions behind the metrics (Pine and

Liboiron, 2015). For journalists, this may mean lobbying for readership engagement metrics rather than the number of page views. As sophisticated actors encounter algorithmic structures in their everyday work, they are likely to use the cultural tools at their disposal as resources for action (Swidler, 1986). In order to retain control, knowledge workers are likely to navigate a constant tension between “working to the algorithm” and “working the algorithm.” Their resultant actions and improvisations are likely to generate performative struggles that propel work trajectories and possibilities in ways yet unexplored.

Augmenting Coordination

Learning algorithms will enable organizations to divide and allocate tasks as well as to integrate efforts in novel ways. Learning algorithms are already automating administrative coordination by managing task decomposition and integration (Zammuto et al., 2007). In online communities, such as Amazon Mechanical Turk or Upwork, platforms that attract very large numbers of contract workers or volunteer contributors from all over the world, algorithms are used to solve the organizational problem of recruitment, negotiating contracts, and motivating employees. The tasks are then decomposed by the algorithm into smaller ones that allow a divide-and-assign strategy so that they can be allocated to individual workers (Faraj et al., 2011; Boudreau and Lakhani, 2013). Once subtasks are completed, they are algorithmically re-composed into a working whole. Likely outcomes of the growing capability of learning algorithms include their ability to assemble high-functioning teams by matching expertise, the flagging of individuals who cause bottlenecks or are not working to expectation, the performance of quality control and task rerouting, and finally the dynamic reconfiguration of roles and workflows (Faraj et al., 2018; Valentine et al., 2017). For example, corporate applications, such as Slack, are dynamically tracing evolving knowledge conversations and are matching the “right people with the right information to do the right tasks” (Woyke, 2018). These emergent AI-powered technologies point to the increasing reliance on learning algorithms as mechanisms for new forms of coordination that reduce the transaction, search, and coordination costs of matching a distributed population of workers to specific tasks. Algorithms also make work more visible to co-workers, sending pings for completed tasks or reminders to slow performers (Schildt, 2017).

Algorithmic coordination that is focused on the minutia of tasks is likely to limit knowledge recombination and the emergence of serendipitous or novel solutions. For example, one feature of the Slack application previously discussed is that it is able to scan a user’s unread messages and to flag a handful of messages that are deemed as most important by its algorithms. In this example, the learning algorithm “curates” what knowledge workers pay attention to, making its own inclusion/exclusion decisions based on an analysis of the content of a person’s corpus of online communication. Workers that are primarily exposed to knowledge related to a narrow expertise domain are likely to become confined to a limited way of thinking and reasoning. Instead, they may prefer to understand and appreciate other epistemic viewpoints and may value the possibility of novel knowledge combination (Faraj et al., 2016). While learning algorithms may improve the coordination of work, they may inadvertently limit the knowledge that employees are exposed to, which holds detrimental effects for knowledge diversity and organizational innovation (Dougherty and Dunne, 2012; Dahlander and Frederiksen, 2012).

Discussion & Conclusion

Learning algorithms can be considered performative due to the extent to which their use can shape and alter work and organizational realities. They raise important research questions for organizational

scholars related to expertise, occupational boundaries, control, and the coordination of knowledge work. Our evaluation of these questions points to the various ways in which learning algorithms are likely to alter work and organizing in qualitatively different ways beyond simply signaling an acceleration of long-term technology trends. Beyond the direct impact of organizing, these technologies foreground digital data in all its forms and are implicitly dismissive of the value of the situated experience held by embedded, embodied social actors (Constantiou and Kallinikos, 2014). With little societal debate or oversight, learning algorithms produce “the possibility of modifying the behaviors of persons and things for profit and control” (Zuboff, 2015 p. 85). Yet, “who” exactly does the modifying? That remains the core political and societal question that algorithm researchers need to address.

An emergent issue for both academics and policy makers is the extent to which algorithms are authorized to make individual and organizational choices, for what types of decisions, and with what implications. Already, a number of algorithmic decisions that rely primarily on digital data at the exclusion of more qualitative and complex types of parameters have shown the difficulty of assigning the authority to the algorithm. For example, an investigation performed for a lawsuit by two teachers who were terminated as a result of an algorithmic decision, based on student quantitative test results, exposed the inefficiency of the algorithm in holistically assessing teaching performance (Webb, 2017). Thus, learning algorithms come with highly performative effects that arise through their usage and bring about a novel form of evaluation and performance that emphasizes quantified comparison and ranking –based on easily available digital data (Orlikowski and Scott, 2014). Users may become increasingly lured by the inscribed technological preferences and alter their behaviours accordingly. Over time, this influences social expectations and norms and gives rise to new institutions that shape how citizens are rated, the kind of stories they have access to, or even which political candidate they may find themselves swayed to vote for (Grassegger and Krogerus, 2017). More than with previous technological change moments, society needs to implement the right policies that regulate the designer goals and the corporate control behind the development of such highly performative technologies.

To break with the existing *laissez-faire* approaches to “algorithmification,” radical proposals, such as those advocating the taxation of technologies that replace the work previously done by a human worker (Schiller, 2017), represent markers in a societal debate that is yet to take place. Other proposals are being put forth regarding governance mechanisms to manage the full transparency and accessibility of individual data. An example is the provision of individual approval mechanisms for publicly-available personal data in order to control the data that third parties are given access to (Stohl, Stohl, and Leonardi, 2016). Currently, individual users on social networks are not afforded control over their own data and their digital traces remain exploitable. Their youthful and not so youthful indiscretions surface as part of their digital persona and become increasingly consequential in a world governed by algorithms.

By challenging our existing frameworks that ontologically separate between people and things, learning algorithms bring to the fore issues of technology morality. Given that algorithms blur the line between user and technology, human and technological agency become entangled. Learning algorithms that imitate cognition and behavior cannot grasp the ethical choices that humans are faced with. For example, a self-driving car algorithm can observe the occupant driving too fast or in a disorderly manner. The issue of what the algorithm should do, or should be allowed to do, raises complex ethical issues: can it “overpower” the driver whenever they drive in a seemingly dangerous manner (they could be in a hurry) or only under select circumstances (being drunk). These technologies, due to their active mediation of human

intent and action, may need to incorporate a moral dimension (Verbeek, 2009). A societal evaluation of the extent to which the technology can influence or even control the user are crucial to ensure a balance between society and the developers of the technology. Leaving the decision of what the algorithm can or cannot do in the hands of the commercial developer may be too shortsighted from a societal point of view.

In conclusion, learning algorithms offer a surface-level similarity to the Weberian bureaucracy in that they create a new digital iron-cage, but whose bars are not readily graspable for bending – where rules are not readily understood or available for interpretation and scrutiny. Algorithms are already changing work so that all aspects of one's performance are quantified, compared to others, and managed against an algorithmic model that is likely to usher in unexpected transformations that need to be better understood (Orlikowski and Scott, 2016). By raising dilemmas regarding expertise, occupational boundaries, coordination, and control, learning algorithms are profoundly consequential and compel researchers and practitioners to expand our understanding of their transformative effects and their broader social and ethical implications. Answering the question of the consequences of learning algorithms for work and organizing is likely to alter our most basic understanding of organizational theory.

References

- Akrich, M. (1992). The De-scription of Technical Objects. In W. E. Bijker & J. Law (Eds.), *Shaping Technology / Building Society: Studies in Sociotechnical Change* (pp. 205-224). Cambridge, MA: MIT Press.
- Ananny, M. (2016). Toward an ethics of algorithms: Convening, observation, probability, and timeliness. *Science, Technology, & Human Values*, 41(1), 93-117.
- Anteby, M., Chan, C. K., & DiBenigno, J. . (2016). Three lenses on occupations and professions in organizations: Becoming, doing, and relating. *The Academy of Management Annals*, 10(1), 183-244.
- Bachman, R. (June 9, 2016). Want to cheat your Fitbit? Try a puppy or a power drill, *Wall Street Journal*.
- Beane, M. (2018). Shadow Learning: Building Robotic Surgical Skill When Approved Means Fail. *Administrative Science Quarterly*, 0001839217751692.
- Bergstein, B. (2017). We need alternatives to Facebook. *MIT Technology Review* 120(3), 86-90.
- Bergstein, B. (2018). The Great AI Paradox. *MIT Technology Review* 121(1), 76-80.
- Bowker, G. C. (2005). *Memory practices in the sciences* (Vol. 205). Cambridge, MA: MIT Press.
- Boyd, D., & Crawford, K. (2012). Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon. *Information, Communication & Society*, 15(5), 662-679.
- Boudreau, K. J., & Lakhani, K. R. (2013). Using the crowd as an innovation partner. *Harvard Business Review*, 91(4), 60-69.
- Brynjolfsson, E., & McAfee, An. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*: WW Norton & Company.
- Brynjolfsson, E., & McAfee, A. (2017). The business of artificial intelligence. *Harvard Business Review*.
- Brynjolfsson, E. & Mitchell, T. (2017). What can machine learning do? Workforce implications. *Science*, 358(6370): 1530-1534.
- Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1-12.
- Burton-Jones, A. (2014). What have we learned from the Smart Machine?. *Information and Organization*, 24(2), 71-105.
- Cadwalladr, C. (December 4, 2016). Google, democracy and the truth about internet search, *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2016/dec/04/google-democracy-truth-internet-search-facebook>
- Christin, A. (August 28, 2014). When it comes to chasing clicks, journalists say one thing but feel pressure to do another, *Neiman Lab*. Retrieved from <http://www.niemanlab.org/2014/08/when-it-comes-to-chasing-clicks-journalists-say-one-thing-but-feel-pressure-to-do-another/>
- Christin, A. (June 20, 2016). The hidden story of how metrics are being used in courtrooms and newsrooms to make more decisions, *Ethnographymatters.net*. Retrieved from <http://ethnographymatters.net/blog/category/editions/co-designing-with-machines/>.
- Constantiou, I. D., & Kallinikos, J. (2015). New games, new rules: Big data and the changing context of strategy. *Journal of Information Technology*, 30(1), 44-57.
- Dahlander, L., & Frederiksen, L. (2012). The Core and Cosmopolitans: A Relational View of Innovation in User Communities. *Organization Science*, 23(4), 988-1007. doi: 10.1287/orsc.1110.0673
- Dewey, C. (October 12, 2016). Facebook has repeatedly trended fake news since firing its human editors, *Washington Post*. Retrieved from https://www.washingtonpost.com/news/the-intersect/wp/2016/10/12/facebook-has-repeatedly-trended-fake-news-since-firing-its-human-editors/?utm_term=.09aed4963adb
- Dougherty, D., & Dunne, D. D. (2012). Digital science and knowledge boundaries in complex innovation. *Organization Science*, 23(5), 1467-1484.

- Dourish, P. (2016). Algorithms and their others: Algorithmic culture in context. *Big Data & Society*, 3(2), 2053951716665128.
- Espeland, W. N., & Stevens, M. L. (2008). A sociology of quantification. *European Journal of Sociology*, 49(03), 401-436.
- Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*, 542(7639): 115-118.
- European Commission. (June 27, 2017). Antitrust: Commission fines Google €2.42 billion for abusing dominance as search engine by giving illegal advantage to own comparison shopping service. from http://europa.eu/rapid/press-release_IP-17-1784_en.htm
- Faraj, S., Jarvenpaa, S. L., & Majchrzak, A. (2011). Knowledge Collaboration in Online Communities. *Organization Science*, 22(5), 1224-1239.
- Faraj, S., Sayegh, K., & Rouleau, L. (2018). Knowledge collaboration in organizations: From information processing to social knowing. In R. D. Galliers & M. Stein (Eds.), *The Routledge Companion to Management Information Systems* (pp. 370-386). London: Taylor and Francis.
- Faraj, S., Von Krogh, G., Monteiro, E., & Lakhani, K. R. (2016). Online community as space for knowledge flows. *Information Systems Research*.
- 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.
- Grassegger, H., & Krogerus, M. (2017) The Data That Turned the World Upside Down, Motherboard Vice. Retrieved from https://motherboard.vice.com/en_us/article/mg9vvn/how-our-likes-helped-trump-win.
- Gitelman, L. (2013). *Raw data is an oxymoron*. MIT Press.
- Hopcroft, J.E., & Ullman, J. D. (1983). *Data structures and algorithms*. Boston, MA: Addison-Wesley
- Introna, L. D. (2016). Algorithms, governance, and governmentality: On governing academic writing. *Science, Technology, & Human Values*, 41(1), 17-49.
- Introna, L. D., & Nissenbaum, H. (2000). Shaping the Web: Why the politics of search engines matters. *The information society*, 16(3), 169-185.
- Israni, E.. (October 26, 2017). When an algorithm helps send you to prison, New York Times. Retrieved from <https://www.nytimes.com/2017/10/26/opinion/algorithm-compas-sentencing-bias.html>
- Kitchin, R. (2014). Big Data, new epistemologies and paradigm shifts. *Big Data & Society*, 1(1).
- Knight, Will (2017). The Dark Secret at the Heart of AI. *MIT Technology Review*, 120(3), 55-63.
- Krisher, Tom. (May 5, 2017). Feds probe Uber's use of fake app to stymie city inspectors, *The Washington Post*.
- Lave, J., & Wenger, E. (1991). *Situated Learning: Legitimate Peripheral Participation*: Cambridge University Press.
- Lifshitz-Assaf, H. (2017). Dismantling Knowledge Boundaries at NASA: From Problem Solvers to Solution Seekers. *Administrative Science Quarterly*, Forthcoming.
- Lupton, D., & Jutel, A. (2015). 'It's like having a physician in your pocket!' A critical analysis of self-diagnosis smartphone apps. *Social Science & Medicine*, 133, 128-135.
- Manyika, J., Lund, S., Chui, M., Bughin, J., Woetzel, J., Batra, P., & Ko, R. (2017). *Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation* (pp. 1-148): McKinsey Global Institute.
- Mayer-Schönberger, V., & Cukier, K. (2012). *Big Data: A revolution that transforms how we work, live, and think*. Houghton Mifflin Harcourt.
- Michalski, R. S., Carbonell, J. G., & Mitchell, T. M. (Eds.). (2013). *Machine learning: An artificial intelligence approach*. Springer Science & Business Media.
- Mukherjee, S. (April 3, 2017). AI v. MD. *New Yorker*.
- Nelson, A., & Irwin, J. (2014). "Defining What We Do--All Over Again": Occupational Identity, Technological Change, and the Librarian Internet-Search Relationship. *Academy of Management Journal*, 57 (3), 892-928.

- O'Connor, S. (September 8, 2016). When your boss is an algorithm, *Financial Times*. Retrieved from <https://www.ft.com/content/88fdc58e-754f-11e6-b60a-de4532d5ea35>
- O'Neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*: Broadway Books.
- O'Neil, C. (2016). How Algorithms Rule our Working Lives, *The Guardian*. Retrieved from <https://www.theguardian.com/science/2016/sep/01/how-algorithms-rule-our-working-lives>
- Olson, Parmy (November 6, 2016). How Facebook helped Donald Trump become President, *Forbes*. Retrieved from <https://www.forbes.com/sites/parmyolson/2016/11/09/how-facebook-helped-donald-trump-become-president/#1e1b8da059c5>
- Orlikowski, W. J., & Scott, S. V. (2014). What happens when evaluation goes online? Exploring apparatuses of valuation in the travel sector. *Organization Science*, 25(3), 868-891.
- Orlikowski, W. J., & Scott, S. (2016). Digital Work: A Research Agenda. In B. e. Czarniawska (Ed.), *A Research Agenda for Management and Organization Studies*. Cheltenham, UK: Edward Elgar.
- Pachidi, S., Berends, H., Faraj, S., Huysman, M., & van de Weerd, I. (2014). What Happens When Analytics Lands in the Organization? Studying Epistemologies in Clash. In *Academy of Management Proceedings* (Vol. 2014, No. 1, p. 15590). Academy of Management.
- Pine, K. H., & Liboiron, M. (2015). The politics of measurement and action. Paper presented at the Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems.
- Ramsey, M. (September 23, 2015). Volkswagen emissions problem exposed by routine university research, *The Wall Street Journal*.
- Remus, D., & Levy, F. (2017). Can Robots Be Lawyers: Computers, Lawyers, and the Practice of Law. *Geo. J. Legal Ethics*, 30, 501-558.
- Schildt, Henri. (2017). Big data and organizational design—the brave new world of algorithmic management and computer augmented transparency. *Innovation*, 19(1), 23-30.
- Schiller, R. (March 22, 2017). Why robots should be taxed if they take people's jobs. *The Guardian*. Retrieved from <https://www.theguardian.com/business/2017/mar/22/robots-tax-bill-gates-income-inequality>
- Shaffer, K., Carey, C.E., & Starling, B. (May 5, 2017). Democracy Hacked: A Massive, Pro-Le Pen Disinformation Campaign Hits Twitter, 4chan, and the Mainstream Media, *Democracy Now*. Retrieved from <https://medium.com/data-for-democracy/democracy-hacked-a46c04d9e6d1#.vqrp0z4at>
- Shah, S., & Chakkattu, J. (2016). Microsoft kills AI chatbot Tay (twice) after it goes full Nazi. *Digital Trends*. <https://www.digitaltrends.com/social-media/microsoft-tay-chatbot/>
- Stohl, Cynthia, Stohl, Michael, & Leonardi, Paul M. (2016). Managing opacity: Information visibility and the paradox of transparency in the digital age. *International Journal of Communication*, 10, 123-137.
- Swidler, Ann. (1986). Culture in action: Symbols and strategies. *American Sociological Review*, 273-286.
- Tegmark, M. (2017). *Life 3.0: Being Human in the Age of Artificial Intelligence*. Knopf.
- Tufekci, Z. (2014). Engineering the public: Big data, surveillance and computational politics. *First Monday*, 19(7).
- Tufekci, Z. (Writer). (2017). We're building a dystopia just to make people click on ads. In N. Y. TED, New York (Producer). YouTube: TED.
- Valentine, M. A, Retelny, D., To, A., Rahmati, N., Doshi, T., & Bernstein, M. S. (2017). *Flash organizations: Crowdsourcing complex work by structuring crowds as organizations*. Paper presented at the Proceedings of the 2017 Conference on Human Factors in Computing Systems.
- Verbeek, P. (2009). Ambient Intelligence and Persuasive Technology: The Blurring Boundaries Between Human and Technology. *Nanoethics*, 3, 231-242.
- Wagner-Pacifci, R., Mohr, J. W., & Breiger, R. L. (2015). Ontologies, methodologies, and new uses of Big Data in the social and cultural sciences.

- Webb, Shelby (2017). Houston teachers to pursue lawsuit over secret evaluation system System shaped firings, bonuses between 2011-15, Houston Chronicle. Retrieved from <http://www.houstonchronicle.com/news/houston-texas/houston/article/Houston-teachers-to-pursue-lawsuit-over-secret-11139692.php>.
- Weiczner, J. (June 10, 2016). Fitbit Users are finding creative ways to cheat, *Forbes*.
- Woyke, E. (2018). Slack Hopes Its AI Will Keep You from Hating Slack. MIT Technology Review. <https://www.technologyreview.com/s/609319/slack-hopes-its-ai-will-keep-you-from-hating-slack/>
- Winner, L. (1980). Do artifacts have politics?. *Daedalus*, 121-136.
- Youyou, W., Kosinski, M., & Stillwell, D. (2015). Computer-based personality judgments are more accurate than those made by humans. *Proceedings of the National Academy of Sciences*, 112(4), 1036-1040.
- Zammuto, R. F., Griffith, T. L., Majchrzak, A., Dougherty, D., & Faraj, S. (2007). Information Technology and the Changing Fabric of Organization. *Organization Science*, 18(5), 749-762. doi: 10.1287/orsc.1070.0307
- Zuboff, S. (1988). *In the Age of the Smart Machine: The Future of Work and Power*. New York: Basic Books.
- Zuboff, S. (2015). Big other: Surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30(1), 75-89.
- Zuckerberg, Mark. (February 16, 2017). Building global community. Retrieved from <https://www.facebook.com/notes/mark-zuckerberg/building-global-community/10154544292806634/>