A more-than-human approach to researching AI at work: Alternative narratives for AI and networked learning

Terrie Lynn Thompson

Faculty of Social Sciences, University of Stirling, terrielynn.thompson@stir.ac.uk

Bruce Graham

Computing Science and Mathematics, Faculty of Natural Sciences, University of Stirling, bruce.graham@stir.ac.uk

Abstract

Artificial intelligence (AI) is increasingly manifest in everyday work, learning, and living. Reports attempting to gauge public perception suggest that amidst exaggerated expectations and fears about AI, citizens are sceptical and lack understanding of what AI is and does (Archer et al., 2018). Professional workers practice at the intersection of such public perceptions, the AI industry, and regulatory frameworks. Yet, there is limited understanding of the day-to-day interactions and predicaments between workers, AI systems, and the publics they serve. This includes how these interactions and predicaments generate opportunities for learning and highlight new digital fluencies needed. We bring social and computing science perspectives to begin to examine the prevailing AI narratives in professional work and learning practices. Some AIs (such as deep machine learning systems) are so sophisticated that a human-understandable explanation of how it works may not be possible. This raises questions about what professional practitioners are able to know about the AI systems they use: their new digital co-workers. We argue that a co-constitutive human-AI perspective could provide useful insights on questions such as: (1) How is professional expertise and judgment re-distributed as workers negotiate and learn with AI systems? (2) What trust and confidence in new AI-infused work practices is needed or possible and how is this mediated? (3) What are the implications for professional learning: both learning within work and the workplace and more formal curriculum? Given the early stages of this field of inquiry, our aim is to evoke discussion of alternative human-AI narratives suited for the messy—and often unseen—realities of everyday practices.

Keywords

artificial intelligence, professional learning, professional work, ethics of technology, public understanding of technologies

Introduction

As artificial intelligence (AI) weaves its way into everyday work, learning and living, labour is being re-distributed between workers and their new digital counterparts. National policies globally present ambitious aspirations for rapid uptake of AI, positioned as a key driver of innovation, labour productivity, and economic growth that needs to be advanced swiftly in order to attain global competitiveness and leadership. AI is also seen as key to finding solutions for critical societal challenges including the UN Sustainable Development Goals.

However, it is not clear what impact AI has, and should have, on workers, particularly professional workers. Or what work-related policies and organizational practices are needed to address these changes. Largely thought to be immune from automation, professional work is now challenged as AI increasingly adds advanced data analytics to augment complex professional decisions and automates tasks (Susskind & Susskind, 2015). Following other approaches (European Commission (EC), 2019; Nilsson, 2010), our working definition of AI is any computational system that carries out a task that is normally associated with a degree of intelligence when performed by humans. The rising prominence of complex AI systems in the workplace challenge roles and skills as new decision-making processes distribute professional judgment and responsibility across AI-human systems. Coming to the fore is the trustworthiness of AI outputs, as emphasized in recent policy recommendations by The High-Level Expert Group on Artificial Intelligence of the European Commission (2019).

Increased use of AI to deliver professional services depends on an informed, critical, and willing public. However, the escalating debate about the incursion of AI into the workplace remains stubbornly polarized. Recent reports attempting to gauge public perception suggest that amidst exaggerated expectations and fears about AI, citizens are sceptical, believe “it won’t happen to me”, and lack understanding of what AI is and does (Archer et al., 2018). Others point to the divergence between the AI hype and the views of experts (e.g., Bristows, 2018). AI narratives have long been influenced by fiction, which fan the fear of robots replacing humans and depict versions of AI that are well beyond current or even near future reality. These narratives are important (The Royal Society (RS), 2018). Informed and positive, they drive investment and innovation at all stages of development from research to commercialization and inspire a future generation of students to become creators of AI. However, negative perceptions fuelled by spurious narratives could lead to public backlash that curtails AI development.

Professional workers practice at the intersection of such public perceptions and prevailing narratives about AI, professional regulatory frameworks, the fast-paced AI industry, and their own competencies and degree of trust regarding AI systems. We take a broad view of the professional worker: a member of an occupational group “that defines itself as collectively sharing particular knowledges and practices, and that is publicly accountable for its service” (Fenwick & Nerland, 2014, p. 2). Although the impact of AI on work is far-reaching, much of the current focus is on macro-employment trends: jobs gained/lost, what work can be automated, and re-skilling the workforce for the “jobs of tomorrow”. In this paper, we argue the importance of attending to professional workers’ day-to-day experiences and interactions with AI systems to provide the urgently needed fine-grained analysis to study what trust and confidence in new AI-infused work practices is required and how this can be generated in new arrangements of work and work-related learning. This evidence will allay concerns that designers and policy makers often make interventions for change in everyday contexts with little understanding of how people produce and experience such algorithmic systems (Pink et al., 2017).

Because work and work-based learning are often inextricably linked it is generative to look at both in order to understand the implications of these technology-mediated practices for workers and their networked ways of learning and working. We begin with an initial exploration of issues around how professional expertise and judgment is re-distributed as workers negotiate and learn with AI systems. We raise questions about what trust and confidence in new AI-infused work practices is needed (or possible) and how is this mediated. Finally, given the early stages of this field of inquiry, our aim is to evoke discussion of alternative human-AI narratives suited for the messy—and often unseen—realities of everyday practices and consider implications for researching these practices.

Negotiating with AI: Re-distribution of professional work-learning

The rapid pace of recent AI advances is driven by machine-learning algorithms including deep learning; exponential increase in computing capacity which can train larger and more complex models much faster; and vast amounts of data (Manyika et al., 2017). Such shifts are shaping assertions that “we are on the cusp of a new automation age in which technologies not only do things we thought only humans could do, but can increasingly do them at a superhuman level” (Manyika et al., 2017, p. 24). However, current discourse on AI and its impact on provision of professional services suggest that AI debate and research is in the early stages and does not yet untangle important distinctions and complexities. For example, much of the rhetoric focuses on the broad trope of jobs lost or gained through automation. Necessary to inform next steps in AI-related development and policy is an understanding of the significant changes in work itself and the learning opportunities inherent in arrangements of work.

Given the range of tasks AI can do (intelligent decision support, classification, prediction, visual object recognition and image processing, speech recognition, natural language processing, and natural language generation) changes to work are complex. There is limited evidence of how AI is being used now and how workers’ tasks have changed where this has happened (Frontier Economics, 2018). Professional bodies responsible for profession-specific regulations and codes of conduct, within workplaces and at the national level, are grappling with drastically changing professional work landscapes, ethical dilemmas, and a desire to seize opportunities afforded by AI while also minimizing risk.

Edwards and Fenwick (2016) ask how we think about professional responsibility and accountability when decisions are delegated to complex digital systems or what it means to consider a professional as a responsible agent when capability is distributed across human and digital actors. Evidence is needed of how AI-mediated work practices are changing decision-making processes, the valuing of professional judgment, and newly distributed responsibilities for algorithmic-influenced decisions. Allert and Richter (2018) highlight a profound shift: as automation and algorithmitization of knowledge work turn data into a resource for, and product of, computation, certain regimes of knowledge that replace subjective experience with objectified data come to the fore. In addition to delegating routine tasks to AI, complex decisions are increasingly based on computational analysis of big data raising questions about the capacity and need for human judgment. Although decision makers may be reluctant to depart from algorithmic recommendations (thus further undermining individual judgment and discretion), others argue that not all decisions can be coded (Agrawal et al., 2019).

As professionals undertake new and different responsibilities for knowledge, understandings of where "expert" knowledge resides becomes blurred. The outsourcing of work activities to, and with, algorithms is leading to new forms of "algorithmic management": prolific data collection and surveillance, transfer of performance evaluations to rating systems or other metrics, and the use of "nudges" and penalties to indirectly incentivize worker behaviors (Kolbjørnsrud et al., 2016; Mateescu & Nguyen, 2019). The following examples highlight some of the complexity of these shifts in responsibility and control.

As reported by Tromans (2019), the recent ban obtained by France's judges on the use of public court data for the statistical analysis and prediction of their decisions in court (i.e., legal predictive analytics) has led the French National Bar Council to demand that lawyers should also be excluded from statistical analysis of their actions in court. France may be the "first country in the world where litigation analysis and predictive modelling face such a comprehensive ban" (para 6). In light of France's "Open Data" movement, intended to make all public data available online, Tromans (2019) points to contradictions in the emergence of a "two-tier" public data system: "citizens can know some things, but not others, even when the underlying information is public" (para 13); and the work of legal professionals and court practices are further obscured with some lawyers claiming this move as "irreconcilable with their mission to represent and defend their clients" (para 15).

The tensions evident in the French court system relate to the openness of AI systems and the data upon which they build. Further concerns arise when AI systems move from merely informing to prescribing professional decisions and actions. In the case reported by the AI Now Institute, the use of student test data to make teacher employment decisions including promotions and terminations revealed, in a subsequent law suit, that no one in the school district could explain or even replicate the determinations made by the system even though the district had access to all the underlying data (Whittaker et al., 2018). The teachers who contested the AI outputs were told that system was simply to be believed and could not be questioned. After the vendor fought against providing access to detailed information on how its system worked, and a ruling that such an AI system could contravene constitutional due process protections, the school district eventually abandoned the third-party AI system in question.

The private-public partnerships that often sustain extensive use of AI systems in the provision of professional service are potentially problematic as decision-making, responsibility, accountability, and the underlying data are not only increasingly distributed across a range of actors but sometimes "black boxed". Predictive algorithms are often used in advanced in criminal justice systems to inform decision-making in policing patterns as well as bail and sentencing decisions. Described in a recent Council of Europe (2018) motion as effective systems valued by the authorities that use them, they nevertheless urge attention to: (1) how such systems are usually provided by private companies and not subject to public scrutiny; and (2) how police departments may lose control over their own data and become dependent on the private companies that have acquired this data.

Our recent work has highlighted contradictions in the current rhetoric about AI and its actual level of uptake in provision of professional work and services. This is consistent with an ethnographic study on the use of AI-mediated risk-assessment tools in the USA criminal court system. Christin's (2017) analysis suggests that such AI systems are often actively resisted in criminal courts and far less powerful and persuasive than suggested in the current narratives extolling widespread AI deployment. She notes that because the judges and prosecutors in her study did not trust the algorithms (they did not know the companies they come from, they did not understand their methods, and often found them useless), the AI outputs often went unused (para 12). Christin (2017) describes how the software was used, score sheets printed out and added to the defendants' files, after which the "scores then seemed to disappear and were rarely mentioned during hearings" (para 12). Christin's (2017) study foregrounds the importance of attending to actual everyday practices: she found that the issue creating resistance was not the transition to complex AI risk-assessment tools per se but rather the more basic transition to paperless case-management systems.

Perhaps the best way to describe the current situation is an uneasy alliance: there are many aspects of work that can be done better and in ways that do not minimize or devalue the human but there are also many potential uncertainties and dilemmas. It is possible to build on the opportunities created by the current wave of AI systems. Polonski (2018) provides examples of how police forces use AI to map when and where crime is likely to occur and how doctors can use it to predict when a patient is most likely to have a heart attack or stroke. There is evidence of significant economic benefits when AI is used to optimize production processes, especially when coupled with suitable workforce retraining in the AI technologies to avoid staff layoffs (PAI, 2018). Image processing by deep neural networks (Le Cun, Bengio & Hinton, 2015) is a strong success story for AI and promises to at least take the drudgery out of examining large volumes of medical imaging data for signs of disease. And it shows promise to be able to find disease indicators in such data that are not evident to human experts. AI developed in-house by Zymergen, a start-up company in the USA to automate laboratory services, found that close collaboration with laboratory scientists during the AI development was crucial to establishing trust in the end systems (PAI, 2018).

Such collaboration between AI developers and workers is extremely important. Deepening involvement with AI systems not only distributes, but also amplifies, workers' implicatedness (Thiele, 2014) and thus expands their ethical responsibilities. Questions of how professional work is valued within the new algorithmic culture (Bayne, 2015) are extremely timely. Workers therefore need to be part of the design and development of responsible human-AI interaction in ways that do not minimize human intelligence. Research evidence is needed to inform how workers can be more purposefully and thoughtfully implicated, valued, and involved in the development of AI systems.

Trusting AI co-workers

Setting out a framework for achieving trustworthy AI, the EC (2019) identifies trustworthy AI as a foundational ambition: not only the technology's inherent properties, but adopting a socio-technical approach that attends to both human and technology actors throughout the AI ecosystem and life cycle. There are good reasons for caution. Bias and lack of transparency in how algorithms work are shortcomings in current AI systems and an active area of research. The AI Now Institute point to widespread testing of AI systems "in the wild" in which AI systems with significant decision making are tested on live populations, often with little oversight (Whittaker et al., 2018). Addressing these issues is crucial for developing AI systems that workers and the public trust.

If people do not know how AI arrives at decisions, they will not trust it; an issue attributed to the failure of IBM Watson for Oncology, an AI system designed to assist doctors with cancer diagnoses. Polonski (2018) highlights the tensions that emerged in the deployment of IBM's AI system:

> If Watson provided guidance about a treatment that coincided with their own opinions, physicians did not see much value in Watson's recommendations. The supercomputer was simply telling them what they already know, and these recommendations did not change the actual treatment. … [If] Watson generated a recommendation that contradicted the experts' opinion, doctors would typically conclude that Watson wasn't competent. And the machine wouldn't be able to explain why its treatment was plausible because its machine learning algorithms were simply too complex to be fully understood by humans. Consequently, this has caused even more mistrust and disbelief, leading many doctors to ignore the seemingly outlandish AI recommendations and stick to their own expertise. (paras 5-6)

Adding to the challenge of understanding how this trust develops is that AI is often invisible, making it difficult for people to understand how and when they interact with it (Bristows, 2018). The problem is exacerbated by the increasing availability of (relatively) easy-to-use software tools for creating data-trained AI systems (e.g., deep neural networks), enabling AI systems to be built by people who have little or no understanding of the inner workings of such systems and their limitations.

Nevertheless, Bunz (2017) states that if we do not want to live with blackboxed technologies, it is essential to learn how to interact with them more attentively (p. 253). Without this attentiveness, there will be repercussions. For example, consider Uber's deliberate obscuring of the algorithms that determine demand and supply pricing of fares, which led to drivers to "game the system" in order to control and create price surges (Rosenblat & Stark, 2015). The efforts of these workers to address information asymmetries highlight the consequences of imposed algorithms that are not transparent or trusted by workers.

Explainable AI (XAI) is seen as essential if workers are to "understand, appropriately trust, and effectively manage an emerging generation of artificially intelligent machine partners" (Gunning, 2017) and is meant to afford humans a degree of functional understanding of AI outputs. However, a fundamental question arises about worker and public expectations of an AI system: Is the expectation to replicate human expertise and/or to improve upon it? If it is the former, then we would likely expect to be able to interrogate the AI to understand how it has arrived at an output, in the same way we could ask a human expert. That said, if we can accept that the AI system may work differently from human reasoning and potentially with higher performance, could workers and the broader publics accept that a human-understandable explanation of how the AI works may not be possible?

The operation of many AI technologies, such as rule-based systems, case-based reasoning, and decision trees, is transparent to humans. An approach to XAI is to try to use these technologies to model the performance of non-transparent AI systems, such as deep neural networks (Ribeiro, Singh & Guestrin, 2016). The downside is that any "explanation" that arises is still only an approximation to what the AI is really doing (though the same may be true for a human expert asked to give an explanation of how they reached a conclusion). At best, the explanation such an approach provides may, to a degree, match the way the neural network has arrived at a decision.

In this situation, the important factor in deploying AI in the workplace is whether adding such a level of explanation provides increased and necessary trust in the AI, whether or not the explanation is strictly accurate. Ultimately, truly powerful AI systems may not be understandable and therefore the entire AI ecosystem (which includes designers, industry, policy makers, workers, researchers, and the public) needs to find other ways of establishing trust in such systems. This could include continual monitoring of the utility of the outcomes produced by the AI so that trust is established via increasing confidence in the robustness and performance of the AI. Deployed AI systems should come under critical performance appraisal in the same way as a human employee. For example, a recent large-scale study of existing published research concludes that current AI systems perform only as well as humans (Liu et al., 2019). However, Liu et al. (2019) add a caveat that the quality of most of these studies is still poor, with only 14 of 82 providing a robust comparison between the AI and human doctors.

One challenge to the development of trustworthy AI is built-in bias. Because humans exhibit bias in decision-making either consciously or unconsciously, a potential selling point for AI decision support systems is their lack of bias. Unfortunately, this is difficult to achieve in practice, as it requires large and truly representative data sets to underpin the training of the AI. For example, Hao (2019) explains how risk assessment tools used in the justice system are designed to generate a recidivism score (a single number estimating the likelihood that a person will reoffend) that is then used by a judge to help determine what type of rehabilitation services particular defendants should receive. However, Hao (2019) points out that such tools are often driven by algorithms trained on historical crime data, which means that populations that are historically disproportionately targeted by law enforcement (e.g., low-income and minority communities) are at risk of high recidivism scores. These algorithms may in fact "amplify and perpetuate embedded biases and generate even more bias-tainted data to feed a vicious cycle" (Hao, 2019, para 10).

The issue of bias in datasets and algorithms is now widely recognised by AI developers and is rightly part of the public AI narrative on the limitations of AI systems. The onus is therefore on a the range of actors involved in the AI ecosystem to understand and to identify—in practice—the limitations and biases of the system and to work towards generating genuinely unbiased—trustworthy—datasets for use in training AI. This is a hidden unappreciated cost in AI deployment.

More-than-human sensibilities

Bucher (2016) asks: "When confronted with the seemingly obscure and hidden, what are our methodological options?" (p. 82). Sociomaterial and more-than-human sensibilities provide a way to conceptualize and study the complex interactions that unfold between AI systems, workers, ways of working, workplaces, policies, and public discourse in the delivery of professional services. Work practices are seen as distributed across a network and changes to work and professionalism as a series of complex social and material (digital) relations. Such distinctions have long been the hallmark of networked learning theorizing and research. Taking a relational view of learning, networked learning focuses on connections among learners, other people, learning resources and technologies (Goodyear et al., 2004). AI systems introduce a myriad of new actors and connections into these networks.

This more co-constitutive perspective helps to avoid over-simplistic deterministic stances and instead brings complex objects (such as AI) out of the background and into critical inquiry, thus offering more inclusive accounts of what it means to be human in an increasingly technologized world (Barad, 2003). Much of the current discourse around AI systems reinforces the binary of human-machine, worker-AI, and human vs. artificial intelligence. Workers and AI systems are often described and portrayed as somehow connected, yet separate, entities. And yet, many current and promising uses of AI systems in professional work and provision of professional services seem to be about how AI systems and humans work together (i.e., AI co-workers and job sharing) and less so on outright replacement by robots or algorithms.

Understanding the larger social changes and the ethical implications around work and workers demands sensibilities, theory, and methodologies to look beyond the human vs. technology by seeing the human-technology together as the phenomena of interest. AI-mediated work practices are not one thing performed by two actors but rather redolent with what Mol (2002) describes as manyfoldedness: “different versions, different performances, different realities, that co-exist in the present” (p. 79). More-than-human perspectives acknowledge the performativity of AI systems and take into account the myriad of foldings and unfoldings between human and nonhuman actors: meshworks (e.g., Ingold, 2005) which are both performed into being and performative.

Ingold (2005) writes that people increasingly find themselves in environments “built as assemblies of connected elements” (p. 46). Yet in practice they continue to thread their own ways through these environments, tracing paths as they go. Ingold (2012) writes about improvising passages: as beings thread their way through and among the ways of others (human and material), they must “improvise a passage” (p. 49). Each new passage lays a new line in the meshwork: “the trails along which a life is lived” (Ingold, 2005, p. 47). Considering the intra-actions (Barad, 2003) between AI system, professional workers, and work could be considered a way of improvising passages through the AI-mediated work landscape: the laying of lines. In so doing, human beings do not merely interact with their materials (aka data and outputs of AI systems) in pre-determined ways but rather co-respond with them in creative and improvisational modes.

Indeed, this is a very important contribution to be made by social sciences within interdisciplinary research endeavours. The emphasis on a socio-technical approach that attends to both human and technology actors advocated by the EC (2019) and its work on trustworthy AI is a promising beginning, although it too seems to still place human and AI actors in separate camps. Bucher (2016) advises that the uptake of algorithms as a field of research within the social science and humanities is very recent, which creates openings for innovative ways to conceptualize and undertake this research.

Changing AI narratives

One challenge is the small number of similar and potentially misleading narratives that dominate public debate, in part generated by a global confluence of powerful AI knowledge brokers and mediators that include government, industry, research institutions, the media, and the 3rd sector. The narratives about AI prevalent in public discourse inevitably shape the deployment of AI in the workplace.

Public perception of AI is shaped by hundreds of years of stories that people have told about humans and machines, often of a dystopian nature. In these stories, AI is embodied (a robot) and super-intelligent, a trope that leads to inflated expectations and fears about the technology and influences the way the technology is portrayed in popular culture and the media. It is important to recognize that AI deployment in various work sectors is currently performed in the context of workers and publics who bring expectations and beliefs about AI: accurate or not. A recent report by the Royal Society (RS, 2018) summarises the common narratives and their drivers. As an easy target for sensationalism and hype, stories about AI often reinforce fears and/or hope for its future potential of AI and muddy the waters as to its immediate possibilities (e.g., if and when the “AI singularity” will happen). Understanding, acknowledging, and then pedagogically addressing these perceptions in order to clarify and educate workers and the publics they serve about the realistic nature of AI in the provision of professional services is vital to successful deployment.

There is an urgent need to utilize more innovative participatory research methods to enable new AI narratives to emerge through two-way public dialogues. Hauert (2015), robotics researcher and co-founder of Robohub (https://robohub.org/), an online community of robotics experts dedicated to connecting the robotics community to the public in order to demystify robotic technologies, spur innovation, and raise ethical and legal questions that require discussion, writes:

> Irked by hyped headlines that foster fear or overinflate expectations of robotics and artificial intelligence (AI), some researchers have stopped communicating with the media or the public altogether. But we must not disengage. The public includes taxpayers, policy-makers, investors and those who could benefit from the technology. They hear a mostly one-sided discussion that leaves them worried that robots will take their jobs, fearful that AI poses an existential threat, and wondering whether laws should be passed to keep hypothetical technology 'under control'. My colleagues and I spend dinner parties explaining that we are not evil but instead have been working for years to develop systems that could help the elderly, improve health care, make jobs safer and more efficient, and allow us to explore space or beneath the oceans. (paras 15-16).

It is possible to re-craft compelling narratives about AI that accurately reflect, as emphasized by the RS (2018), "the underlying science and its possibilities while acknowledging scientific and social uncertainties" (p. 20). In this paper, we hope to have sparked discussion and thinking about alternative human-AI narratives and ways of conceptualizing research suited for the messy—and often unseen—realities of everyday AI-mediated professional work practices.

References

Agrawal, A., Gans, J. S., & Goldfarb, A. (2019). Exploring the impact of artificial Intelligence: Prediction versus judgment. Information Economics and Policy, (47), 1–6.

Archer, H., Writer-Davies, R., & McGeoghegan, M. (2018). AI, automation, and corporate reputation. https://www.ipsos.com/en/ai-automation-and-corporate-reputation.

Allert, H., & Richter, C. (2018). Perspectives on data and practices. Paper presented at the WebSci'18: 10th ACM Conference on Web Science, Amsterdam.

Barad, K. (2003). Posthumanist performativity: Toward an understanding of how matter comes to matter. Signs: Journal of Women in Culture and Society, 28(3), 801-831.

Bayne, S. (2015). Teacherbot: Interventions in automated teaching. Teaching in Higher Education, 20(4), 455-467. doi: 10.1080/13562517.2015.102078

Bristows. (2018). Artificial Intelligence: Public perception, attitude and trust. https://www.bristows.com/viewpoint/articles/artificial-intelligence-public-perception-attitude-and-trust/

Bucher, T. (2016). Neither black nor box: Ways of knowing algorithms. In S.

Kubitschko & A. Kaun (Eds.), Innovative method in media & comm research (pp. 81-98). Cham: Palgrave.

Bunz, M. (2017). The deed for a dialogue with technology. In M. T. Schafer & K. van Es (Eds.), The datafied society: Studying culture through data (pp. 249-254). Amsterdam: Amsterdam University Press.

Christin, A. (2017). The mistrials of algorithmic sentencing. Logic, Issue 3 Justice. Retrieved from https://logicmag.io/03-the-mistrials-of-algorithmic-sentencing/

Council of Europe (2108). Justice by algorithm – the role of artificial intelligence in policing and criminal justice systems. Motion for a recommendation (Doc. 14628). Retrieved from http://assembly.coe.int/nw/xml/XRef/Xref-XML2HTML-en.asp?fileid=25062&lang=en

Edwards, R., & Fenwick, T. (2016). Digital analytics in professional work and learning. Studies in Continuing Education, 38(2), 213-227. doi:10.1080/0158037X.2015.1074894

European Commission. (2019). Ethics guidelines for trustworthy AI. Report by the High-Level Expert Group on Artificial Intelligence. https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai

Fenwick, T., & Nerland, M. (2014). Reconceptualising professional learning: Sociomaterial knowledges, practices and responsibilities (pp. 1-7). London: Routledge.

Frontier Economics. (2018). The Impact of Artificial Intelligence on Work. https://royalsociety.org/topics-policy/projects/ai-and-work/

Goodyear, P., Banks, S., Hodgson, V., & McConnell, D. (2004). Research on networked learning: An overview. In P. Goodyear, S. Banks, V. Hodgson, & D. McConnell (Eds.), Advances in research on networked learning (pp. 1–11). Dordrecht: Kluwer Academic Publishers.

Gunning, D. (2017). Explainable Artificial Intelligence (XAI) Programme. Published by Defence Advanced Research Projects. https://www.darpa.mil/attachments/XAIProgramUpdate.pdf

Hao, K. (2019, January 21). AI is sending people to jail—and getting it wrong. MIT Technology Review. Retrieved from https://www.technologyreview.com/s/612775/algorithms-criminal-justice-ai/
Hauert, S. (2015, May 27). Robotics: Ethics of artificial intelligence. Nature, 521(7553). Retrieved from https://www.nature.com/news/robotics-ethics-of-artificial-intelligence-1.17611
Ingold, T. (2005). Up, Across and along. Retrieved from http://spacesyntax.tudelft.nl/media/Long%20papers%20I/tim%20ingold.pdf
Ingold, T. (2012) Looking for lines in nature. Earthlines 3, 48-51.
Kolbjørnsrud, V., Amico, R., & Thomas, R. J. (2016). How artificial intelligence will redefine management. Harvard Business Review, Nov, 1-6.
Le Cun, Y, Bengio, Y., & Hinton, G. (2015). Deep Learning. Nature, 521, p. 436-444.
Liu, X., Faes, L., Kale, A., et al. (2019). A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. Lancet Digital Health, 1, e271–97
Manyika, J., Lund, S., Chui, M., Bughin, J., Woetzel, J., Batra, P…. Sanghvi, S. (2017). Jobs lost, jobs gained: Workforce transitions in a time of automation. Retrieved from McKinsey & Company / McKinsey Global Institute website https://www.mckinsey.com/featured-insights/future-of-work/jobs-lost-jobs-gained-what-the-future-of-work-will-mean-for-jobs-skills-and-wages /
Mateescu, A., & Nguyen, A. (2019). Explainer: Algorithmic management in the Workplace. Retrieved from Data & Society website http://datasociety.net
Mol, A. (2002). The body multiple: Ontology in medical practice. London: Duke University Press.
Nilsson, N. J. (2010). The quest for artificial intelligence: A history of ideas and achievements. Cambridge: Cambridge University Press.
Partnership on AI. (PAI, 2018). AI, Labor, and Economy Case Studies. Retrieved from PAI website https://www.partnershiponai.org/compendium-synthesis/
Pink, S., Sumartojo, S., Lupton, D. & Heyes La Bond, C. (2017). Mundane data: The routines, contingencies and accomplishments of digital living. Big Data & Society, Jan-June, 1-12.
Polonski, V. (2019, January 9). People don't trust AI – here's how we can change that. The Conversation. Retrieved from https://theconversation.com/people-dont-trust-ai-heres-how-we-can-change-that-87129
Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. Retrieved from: arXiv:1602.04938v1
Rosenblat, A., & Stark, L. (2015). Uber's drivers: Information asymmetries and control in dynamic work. SSRN Electronic Journal. Retrieved from https://doi.org/10.2139/ssrn.2686227
Royal Society. (2018). AI narratives: Portrayals and perceptions of AI and why they matter. https://royalsociety.org/topics-policy/projects/ai-narratives/
Susskind, R., & Susskind, D. (2015). The future of the professions: How technology will transform the work of human experts. Oxford: Oxford University Press.
Thiele, K. (2014). Ethos of diffraction: New paradigms for a (post)humanist ethics. Parallax, 20(3), 202–216.
Tromans, R. (2019, July 1). Now French lawyers demand statistical data ban, following judges' move. Retrieved from Artificial Lawyer website https://www.artificiallawyer.com/2019/07/01/now-french-lawyers-demand-statistical-data-ban-following-judges-move/
Whittaker, M. et al. (2018). AI Now Report 2018. Retrieved from AI Now website http://ainowinstitute.org/AI_Now_2018_Report.pdf