

Proposal for Presenter Symposium

Organizational Implications of Artificial Intelligence

Co-Organizers:

Ruth Aguilera

D'Amore-McKim School of Business
Northeastern University
360 Huntington Ave, Boston, MA, 02115
r.aguilera@northeastern.edu

Deepika Chhillar

Gies College of Business
University of Illinois Urbana-Champaign
1206 S. Sixth, Champaign, IL, 61820
deepika4@illinois.edu

Discussants:

Robert Wayne Gregory

McIntire School of Commerce
University of Virginia
125 Ruppel Drive, Charlottesville, VA,
22903
rg7cv@comm.virginia.edu

Ivanka Visnjic

ESADE Business School
Ramon Llull University
Av. Torre Blanca, 59
08172 Sant Cugat, Barcelona, Spain
ivanka.visnjic@esade.edu

Presenters (*) and co-authors:

1. Addressing the Key Challenges of Developing Machine Learning AI Systems for Knowledge-Intensive Work

Zhewei Zhang
Warwick Business School
University of Warwick
zhewei.zhang@wbs.ac.uk

Joe Nandhakumar
Warwick Business School
University of Warwick
joe.nandhakumar@wbs.ac.uk

Jochem Thomas Hummel*
Warwick Business School
University of Warwick
jochem.hummel@wbs.ac.uk

Lauren Waardenburg
School of Business and Economics
Vrije Universiteit Amsterdam
l.waardenburg@vu.nl

2. The Politics of Visibility: Digital Prisms and Coded Visions

Mikkel Flyverbom*
Copenhagen Business School
mf.msc@cbs.dk

Frederik Schade
Copenhagen Business School
fsc.msc@cbs.dk

3. Who is Responsible for AI Decisions?

Kirsten Martin*
Mendoza College of Business
University of Notre Dame
kmarti33@nd.edu

4. Morality in the age of Artificially Intelligent Algorithms

Christine Moser
Vrije Universiteit Amsterdam
c.moser@vu.nl

Frank den Hond
Hanken School of Economics
Vrije Universiteit Amsterdam
frank.denhond@hanken.fi

Dirk Lindebaum*
Grenoble Ecole de Management
dirk.lindebaum@grenoble-em.com

Sponsor Divisions:

Strategic Management (STR); Organization and Management Theory (OMT); Technology & Innovation Management (TIM)

OVERVIEW OF SYMPOSIUM

Artificial Intelligence (AI) refers to “a system’s ability to correctly interpret external data, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation” (Haenlein and Kaplan, 2019, p.5). Since the beginning of the 21st century, we are transitioning from employing reactive *automatic* systems to proactive *autonomous* systems at workplaces (Tschang and Almira, 2020). Over the last decade, there has been an exponential growth in interest in adoption of big data and AI technologies by organizations. Strategic industries, such as healthcare, transportation, and energy are becoming increasingly intelligent, efficient, and accurate owing to applications of artificial intelligence such as robotics and autonomous vehicles, computer vision, virtual agents, and ML (Kellogg, Valentine, and Christin, 2020).

Despite the budding and promising extant literature on AI, there is still a large gap in our understanding of the multifarious implications of employing big data, algorithmic decision making and artificial intelligence in workplaces (Fleming, 2019; Kellogg et al., 2020). We believe that our field’s collective understanding may benefit from attention to organizational implications of AI-adoption for several reasons. One of the key reasons is the changing employment opportunities and job quality across industries with AI-based businesses. An interesting theoretical perspective in the realm of AI and the future of work is the human-in-the-loop perspective of algorithmic-adoption (Metcalf, Askay, and Rosenberg, 2020) where the main argument is that substantive human judgement and intelligence is imperative for algorithmic decision-making systems to function effectively. This perspective further leads us to contemplate what power structures emerge as a result of an algorithmically managed world. From a governance perspective, there are real concerns with organizational adoption of algorithmic management and delegating agency in organizations

(How much, and which tasks should be delegated?), bearing caution that we are still in an age of employing weak AI and that most ML tools learn from existing workplace routines.

Organizational solutions to some of these governance challenges may come with new challenges in terms of economic value capture. For instance, organizations with augmented teams, where machines and humans are integrated to complement each other will need to adapt their business models to accommodate additional costs of keeping a human-in-the-loop. With organizations increasingly using data and algorithms to conduct intellectual tasks previously performed by managers, they are also likely to undergo transformation with respect to their organizational decision-making, attention, control, and learning. Firms may also need to rethink traditional workplace hierarchies resulting in extensions to our current theoretical assumptions of organization design.

Algorithmic management has the potential to dramatically change the overall workforce structure and organizational control. Via this symposium, we aim to learn and discuss the importance of various organizational transformations that AI brings. The symposium has papers and discussants drawing on diverse disciplinary perspectives, that will help us gain a better understanding of these complex phenomena. Our goal is to push the boundaries of our collective understanding of the managerial implications of AI-adoption in organizations, and to that effect, our symposium is well aligned with this year's theme of bringing the manager back in management. In addition, we aim to discuss novel methods and theoretical perspectives in the field that may help in developing a better understanding of the adoption and spread of AI across organizations.

Presentations

In the first presentation, Jochem Hummel and his co-authors, identify three key challenges encountered by developers and users. They provide actions and recommendations for addressing

the challenges of developing machine learning AI systems for knowledge-intensive work. Authors propose that AI has great potential to change the way businesses operate. However, developers of AI systems face many challenges because of the significant differences compared to traditional systems. Zhang et al. (2020) describe the development of the AI system at LegalTechCo and the key challenges encountered. Their empirical findings illustrate how a machine learning AI system for legal practice was developed to help legal professionals make faster and better-informed decisions.

Next, in her presentation, Kirsten Martin argues for the responsibility of firms for the value-laden decisions of (employed/designed) AI - i.e., decisions including which algorithms to deploy, which target or outcome to use, and which outliers are appropriate. As more decisions are augmented with AI, firms are grappling with whether and how to take responsibility for the moral implications of those AI-enabled decisions. As such, the author argues that transparency is a means to identify biases and issues of fairness, to make the decision contestable, and for organizations to be held accountable and not as a goal itself: mere transparency is not enough. She suggests that focusing on firms as responsible for not only data, but also the algorithms, biases, and outliers, we can look at five areas for firms developing and adopting AI: data, algorithms, power, governance, and purpose.

In the next presentation Mikkel Flyverbom, in his co-authored work, continues to explore the phenomenon of AI governance by offering a conceptual and illustrative take on the forces and dynamics at play in new politics of visibility. The authors focus on two particular instances of data-driven, algorithmic approaches to governance, namely predictive policing and anti-radicalization efforts. The authors use these illustrations from policing and anti-radicalization to articulate the technological underpinnings and orientations of two forms of political interventions that they term – ‘politics of prediction’ and a ‘politics of conversion’. They empirically illustrate the need for a

renewed focus on rights and democratic values in the age of algorithms and AI and highlight questions about visibility and the reconfiguration of politics.

Finally, Dirk Lindebaum and his co-authors shine a light on the unexamined processes through which humanity risks to lose control over artificially intelligent technology with dire consequences for our morality. The authors argue that human morality is already co-constituted with algorithmic reckoning. Against current enthusiastic trends professing otherwise, they suggest that we need to recognize this co-constitution to avoid the risk of ending up in a situation in which our judgment is impoverished by being informed – largely or even completely so – by algorithmic reckoning. The authors emphasize that AI systems do have agency, which – when unrecognized and unchecked – enables them to inform, guide, and steer human judgment in decision making. Not only do algorithms already make decisions that involve moral judgment, they are inextricably intertwined with the way that our morality develops. In their call for (in)action, they highlight the need to reconsider how we develop, understand and apply algorithms in our daily lives and business.

RELEVANCE TO DIVISIONS

Organization and Management Theory (OMT)

The contemporary topics around the firm-adoption of AI in the paper presentations of our symposium, will add to the rich theoretical heritage of organizational theory. The theme discussed in our symposium is likely of high interest for scholars of OMT. Specially those interested in organizational decision-making and performance will find discussion of such research to enhance our understanding of current decision-making processes at firms. Presentations in our symposium bring forth major limitations and boundaries to organizational use of algorithmic decision-making. For example, tasks such as building and maintaining a positive organizational culture, inspiring co-

workers or coming up with out-of-the-box solutions, fundamentally require human intervention. Organizational actors attempting to cross-over those boundaries are liable to jeopardize organizational performance, institutional trust and employee motivation. There exists a promising stream of research, within OMT, some of which will be presented in our symposium, on challenges of sustainable value capture from AI adoption in traditional organizational activities. Interactive discussion along with these presentations in our symposium hold the potential to improve our current understanding of benefits and perils of AI adoption by organizations.

Strategic Management (STR)

If data are the new oil in an information age, then algorithms are the refineries that make those data valuable. However, to deliver to and share value among organizational stakeholders in a sustainable and accountable manner comes with strategic barriers and institutional changes. Through our symposium we aim to engage scholars of STR division, who are in a unique position to examine, evaluate, recommend strategic frameworks and guide policy making for a fair and just value distribution. With an explosion of data in the recent decade and multiple pipelines in the form of digital platforms, the time is certainly ripe for possibilities creating new business models (Gregory, Henfridsson, Kaganer, and Kyriakou, 2020). The questions then are - how are firms creating and capturing value from AI differently from the past, and what constitutes a new basis for a sustained competitive advantage? Our symposium aims to present an interesting avenue to hold such discussions and with these efforts, STR researchers and practitioners could benefit from the research findings and the interactive discussion in this event.

Technology and Innovation Management (TIM)

Innovations in AI lie at the heart of Technological and Innovation Management (TIM) scholars' interests. Additionally, our symposium discusses issues at the intersection of multiple scholarly

interests. Some illustrative research questions include – How is AI different from previous technological innovations that attempt to deal with bounded rationality in organizational decision-making process? Two presentations in this symposium illustrate that regulation as well as ethics of AI is different from that of earlier technological innovations, in that we have not yet built functional accountability mechanisms for AI-based technologies in most parts of the world today. Governance of artificial intelligence guarantees responsible innovation, leading to increased human trust and user adoption. With these efforts, we believe that TIM researchers and practitioners could benefit from the research findings and the interactive discussion in this event.

PROPOSED FORMAT OF SYMPOSIUM

Length: 90 minutes

Minutes 0-5: Welcome and introduction to the symposium

- Presenter: *Deepika Chhillar*

Minutes 5-65: Paper presentations (15 minutes each)

- Addressing the Key Challenges of Developing Machine Learning AI Systems for Knowledge-Intensive Work. *Presented by Jochem Thomas Hummel*
- Who is Responsible for AI Decisions? *Presented by Kirsten Martin*
- The Politics of Visibility: Digital Prisms and Coded Visions. *Presented by Mikkel Flyverbom*
- Morality in the age of Artificially Intelligent Algorithms. *Presented by Dirk Lindebaum*

10 mins each for the main idea of the paper followed by **5 mins** of discussion (research comments and theoretical implications) by *Ivanka Visnjic* and *Robert Gregory*

Minutes 65-85: Q&A from audience, moderated by *Ruth Aguilera*

Minutes 85-90: Concluding remarks

- Presenter: *Ruth Aguilera*

PRESENTATION SUMMARIES

Addressing the Key Challenges of Developing Machine Learning AI Systems for Knowledge-Intensive Work

Zhewei Zhang, Joe Nandhakumar, **Jochem Thomas Hummel**, and Lauren Waardenburg

A search of the websites of the top ten Fortune 500 firms shows that they have plans to deploy artificial intelligence (AI) systems and, in a recent survey of CEOs, 85% agreed that AI has the potential to significantly alter the way their businesses operate in the next five years. Today, large organizations that are not considering whether and how they can use “predictive algorithms,” “machine learning technologies” or “neural networks” are in the minority. AI has a long history dating back to the 1950s but, until now, it promised much but delivered less. With recent significant advancements in computational algorithms and computing capability, the time may have finally come for AI to fulfill its promise.

However, ready-to-use AI solutions are still scarce, especially for knowledge-intensive work. Moreover, an off-the-shelf AI system is unlikely to offer the kind of competitive advantage that organizations hope to gain from AI. Davenport suggests that for complex work settings, where there are no previous AI solutions, “*you really don’t have much choice but to do it yourself.*” To deploy AI solutions, organizations must therefore develop their own AI systems to suit their unique situations. Building an AI system in-house also provides greater control over the data and intellectual property that become core parts of the system.

The discipline of AI has evolved significantly since the 1950s. Many of the early forms of AI, which are now called rule-based systems, automated cognitive tasks through a computer executing

hard-coded rules derived from human expertise, such as performing routine processes or accessing predefined knowledge bases. The most recent advance in AI is *machine learning systems*, which involve computers autonomously developing algorithmic models to solve problems without explicit instructions from humans. We believe that machine learning systems are the key to automating complex tasks in knowledge-intensive work. In this article, our focus is on machine learning AI systems rather than rule-based systems. There are key differences between the two types:

Rule-based system: An information system that is created based on human knowledge hard-coded as rules. Many early forms of AI were often created as rule-based systems (e.g. IBM's Deep Blue).

Machine learning AI system: A form of AI that relies on large datasets to train machine learning models to handle complex tasks that used to require human knowledge (e.g. DeepMind's AlphaGo). Although the development of machine learning AI systems shares some similarities with rule-based systems, several central characteristics of machine learning systems make their development highly challenging. Rule-based system development relies on existing and codified human knowledge that flows from domain experts to developers, who translate that knowledge into rule-based algorithms for execution by a computer. In contrast, machine learning systems learn from large datasets to create their own logic, relying on opaque models derived via mathematical computing. The underlying knowledge and algorithms of rule-based systems can be analyzed and understood by humans. In contrast, however, machine learning systems are black boxes whose internal workings are incomprehensible to humans. This characteristic of machine learning AI systems creates significant development challenges.

To cast light on these challenges and derive actions and recommendations for executives, we studied the development of an AI system for legal practice at "LegalTechCo." As described in

Appendix, we collected data over 22 months (between June 2018 and March 2020) on the development of an AI system at Legal Tech Co. Because the project was based on a university-industry partnership involving two of the authors, we had unprecedented access to the system development process used by computer scientists (i.e., developers) and legal analysts (i.e., domain experts) working at Legal TechCo. This enabled us to collect data by participating in and recording all meetings, interacting freely with developers and domain experts, and accessing all project documentation. To avoid potential bias, data collection and analysis were carried out by authors not involved in the development of the AI system. Our qualitative data analysis (see Appendix) led to a richer understanding of AI development in organizations with knowledge-intensive work. In this article, we describe the development of the AI system at LegalTechCo and the key challenges encountered.

Who is Responsible for AI Decisions?

Kirsten Martin

As more decisions are augmented with artificial intelligence, firms are grappling with whether and how to take responsibility for the moral implications of those AI-enabled decisions. Whether using AI for who to hire, who to promote, how to price, or who to vaccinate, the use of AI can momentarily provide a smoke screen for accountability.¹ People focus on the role of AI in the decision and mistakenly believe assigning responsibility for the decision's outcomes and moral implications is in question.

The question is not if firms are responsible for their AI-augmented decisions (they are), but rather what are they responsible for? A growing number of scholars point to the biased and unfair outcomes, the privacy violations of data collected, and the degree the AI decision is transparent as objects of responsibility for firms. I have previously argued that data scientists are responsible for the allocation of responsibility for tasks between the human users and the AI program. Purposefully designing inscrutable AI-decision systems, where the users are unable to identify, judge, and fix mistakes, would render the users unable to take responsibility for the moral implications of their decisions and would place the responsibility of the designer of the AI program (Martin, 2019a). This makes accountability a design choice of the data scientist.² In addition, scholars are even moving away from transparency towards the obligation of contestability (Mulligan, Kluttz, & Kohli, 2020) and providing technical solutions for interpreting 'black box' AI (Bastani, Kim, & Bastani, 2017). Transparency is a means to identify biases and issues of fairness, to make the decision contestable,

¹ "The Stanford vaccine algorithm failed to prioritize house staff". (Bernstein, Beachum, & Knowles, n.d.)

² By data scientist, I mean the firm that employs the data scientist and developed the AI program. This could be the same company that uses the AI program for decisions such as when Google develops their own programs.

and for organizations to be held accountable and not as a goal itself. Mere transparency is not enough.

I wish to focus on the responsibility of firms for the value-laden decisions of the designed AI: decisions including which algorithms to deploy, which target or outcome to use, and which outliers are appropriate. A common misnomer is that machine learning (and all its subsets) ‘learns on its own’ using data from past behavior; as if the data scientist has no role in the development of an AI program. This throw-away line³ diminishes the important role of the data scientist and hides the many choices data scientists make *outside the choice of the type of data* to include the type of algorithm and the best predicted outcome to use for a given decision. For example, within machine learning, the data scientist can choose between supervised learning, reinforcement learning, and unsupervised learning; data scientists make assumptions about the type of relationship between the data and the type of outliers that are acceptable in making choices as to the type of algorithm to use. Data scientists also make tradeoffs around efficiency and accuracy and explain-ability and even fairness. All are value-laden decisions with moral implications for which the firm developing the AI is responsible.

The type of algorithm chosen also dictates who is considered an outlier. The point of artificial intelligence or machine learning is to find clusters of factors that are related to each other (unsupervised learning) or related to an outcome (supervised learning) *for most of the people represented in the data*. Outliers – those individuals that do not neatly fit within the hypothesized relationship between data – are inherent to the statistical underpinnings AI. In other words,

³ Another throwaway line claims that AI/ML/Neural Networks have advantages “over us such as speed, accuracy and lack of bias”. <https://www.forbes.com/sites/bernardmarr/2016/12/06/what-is-the-difference-between-artificial-intelligence-and-machine-learning/?sh=188516bd2742>

outliers always exist, are determined by the choice of data and algorithm, and are also individuals in many cases of organizations using AI (Aggarwal, 2017).

In focusing on firms as responsible for not only data, but also the algorithms, biases, and outliers, we can look at five types of questions for firms developing and adopting AI:

1. **Data:** Who is represented accurately in the dataset being used and who is not? For whom do the chosen factors work and for whom are the factors not predictive? (Gebru et al., 2018; Paullada, Raji, Bender, Denton, & Hanna, 2020)
2. **Algorithms:** What assumptions are the data scientist making when choosing an algorithm? Who are they prioritizing in choosing the algorithm and approach to machine learning? What are the alternatives? How are outliers treated?
3. **Power:** For whom is value created or destroyed with the use of the AI program? Are current power structures being reinforced? Are currently marginalized communities being disenfranchised or being recognized? (Gebru, 2019; Leavy, O'Sullivan, & Siaper, 2020)
4. **Governance:** How are mistakes being identified, judged, and corrected going forward? What mechanisms are in place to ensure mistakes and AI are being governed? (See also Martin, 2019b)
5. **Purpose:** What is the purpose of the use of AI in this instance? Does it align with the mission and values of the organization? Does the use of AI engender trust with users, employees, customers, etc? (see also Freeman, Parmar, & Martin, 2020)

6.

The Politics of Visibility: Digital Prisms and Coded Visions

Mikkel Flyverbom and Frederik Schade

Data analytics and automated forms of pattern recognition are a new frontier of action and influence, particularly in the shape of digital advertising, business intelligence and predictions. Data enthusiasts and big tech companies are busy rolling out data-driven approaches to security and police work, consumer analytics, human resource management and other forms of governance. The goal is more accurate, more proactive, and more objective methods for the anticipation of developments on the horizon – the opportunity to shape futures from the vantage point of the present. We argue that such datafied, algorithmic visibilities, or *coded visions*, constitute new forms of political interventions that are worth exploring if we want to make sense of the politics of visibility that this volume is about.

The chapter articulates how attempts to make social phenomena visible and governable through data and algorithmic sorting involve complex forms of knowledge work and socio-material entanglements that deserve more attention. To understand how futures are produced and steered through data and algorithms, we need to return to fundamental questions about how social worlds become *seeable*, *knowable* and *governable* through situated processes of knowledge production. Focusing on policing and anti-radicalization efforts as two compelling cases that serve to highlight our arguments about the consequences of digital transformations for politics and visibility, we illustrate how human actions are turned into digital information, and how such operations come to shape political interventions. By unpacking the technical, material and practical *production of visibilities* in the context of governance efforts, we highlight the political consequences of bringing

social phenomena into the register of the visible or turning things into objects of visibility. The chapter contributes to emergent research on how digital transformations and processes of datafication (Mejias and Couldry, 2019) condition and relate to contemporary attempts to frame and govern societal challenges and opportunities. The ambition is to conceptualize and illustrate the workings of digital visibility as a distinct mode of anticipatory governance and – by extension – a novel contemporary shape taken by politics and political interventions. The kinds of coded visions that we explore pave the way for novel kinds of politics, what we term a *politics of prediction* and a *politics of conversion*, both of which have far-reaching and potentially problematic ramifications.

Seeing, knowing, governing: An analytical framework

The production, circulation and management of visibilities involves a particular set of dynamics that have consequences for the kinds of governance and politics they enact and produce (Flyverbom 2019). In the following, we develop an analytical vocabulary that focuses on the relationship between ‘seeing’, ‘knowing’ and ‘governing’. We use each of these three terms to highlight key dynamics of the workings of visibilities in the context of governance.

Empirical illustrations

We focus on two particular instances of data-driven, algorithmic approaches to governance, namely predictive policing and anti-radicalization efforts. The main consideration guiding this selection has been one of locating two different and relatively widespread technologies, each representing distinct coded visions, which produce particular and distinct forms of politics. We demonstrate how predictive policing constitutes a particular ‘politics of prediction’, and how anti-radicalization (using techniques from digital advertising) constitutes an equally distinct ‘politics of conversion’. Both of these future-oriented forms of governance rely on digital technologies and data analytics to

construct visible objects and seek to anticipate and shape human behaviour through data-driven practices. At the same time, they represent distinct coded visions oriented towards anticipation (Anderson, 2010) and information reductionism (Tsoukas, 1997), which sets them apart from other modes of governance. In our subsequent discussion, we compare these seemingly distinct political formations while also attempting to unpack the dimensions they may have in common to locate a more general political formation related to the production of visibilities enabled by data-driven technologies per se. We start, however, by considering each technology in turn with regard to how they make their objects seeable, knowable and governable in novel ways.

Conclusion

In this chapter, we have offered a conceptual and illustrative take on the forces and dynamics at play in new politics of visibility. Our response has taken the shape of a conceptual and analytical vocabulary that places the production of visibilities at its core. The analytical distinction between *seeing*, *knowing* and *governing* paves the way for research into concrete work with digital traces, algorithms and visualizations carried out to prevent crime and radicalization. In combination with our interest in the performativity of digital visibility efforts, this conceptualization offers a starting point for investigations of the possible formation(s) of a new politics of visibility in the context of digital transformations. We use illustrations from policing and anti-radicalization to articulate the technological underpinnings and orientations of two forms of political interventions that we term a ‘politics of prediction’ and a ‘politics of conversion’. Locating the particular politics attached to distinct technologies – along with the coded visions signalling their directionality – we propose, can highlight both some fundamental dynamics at play in datafied, algorithmic forms of governance, and some emergent and troubling developments in governance and politics that may require a renewed focus on rights and democratic values.

Morality in the age of Artificially Intelligent Algorithms

Christine Moser, Frank den Hond, and **Dirk Lindebaum**

We are at risk, now, of artificially intelligent algorithms changing our morality in fundamental ways, and perhaps irreversibly so. We develop and use such algorithms to facilitate and enhance decision making, harboring the illusion that we, human beings, are in control and can develop algorithms to be aligned with, reflect, and espouse our morality. However, as we increasingly rely on artificially intelligent algorithms in decision making, we risk mistaking ‘reckoning’ – that is, the way how algorithms process data to inform decision making – for judgment. When we see reckoning and judgment as equivalent, or ontologically similar whereas they are not, reckoning may transform judgment by ontologically assimilating and subsequently substituting it, to the consequence of impoverishing our morality. In light of this risk, we are, in fact, far less in control than currently recognized. In this essay, against current enthusiastic trends professing otherwise, we set out to shine light on the unexamined processes through which we risk to lose control over artificially intelligent technology with dire consequences for our morality.

It is now widely recognized that artificially intelligent algorithms rapidly infiltrate and control many applications that affect the lives of millions of people. As a recent whitepaper published by the European Commission states:

“Artificial Intelligence is developing fast. It will change our lives by improving healthcare (e.g. making diagnosis more precise, enabling better prevention of diseases), increasing the efficiency of farming, contributing to climate change mitigation and adaptation, improving the efficiency of production systems through predictive maintenance, increasing the security of Europeans, and in many other ways that we can only begin to imagine” (EC, 2020: 1).

Nevertheless, the European Commission also recognizes that “(AI) entails a number of potential risks, such as opaque decision-making, gender-based or other kinds of discrimination,

intrusion in our private lives or being used for criminal purposes” (EC, 2020: 1). In order to reap the benefits from AI while minimizing its risks, the European Commission announces to develop a set of rules and regulations on the development and use of AI, such that it can give “citizens the confidence to take up AI applications and give companies and public organizations the legal certainty to innovate using AI” (EC, 2020: 3). Likewise, Virginia Dignum – a prominent writer on ‘responsible AI’ – argues that “perhaps the most important message of this book is that responsible AI is not about the characteristics of AI systems, but about our own role. We are responsible for how we build systems” (Dignum, 2019: 7, emphasis added). Overall, the European Commission, Dignum, and many other analysts and authorities, seem to have confidence in the possibility that individuals and organizations will benefit from AI while remaining in control of the technology. But even if we could remain in control, we are at risk of inflicting on ourselves a process by which we impoverish human morality as we know it and use it. This risk is of far greater magnitude than currently recognized in the rush toward adopting this technology.

Therefore, in this essay, we examine in three scenarios how we lead ourselves into the illusion of being in control, when, in fact, control may already be slipping through our fingers. Along the way, we closely examine human judgment and algorithmic reckoning in decision making in relation to morality. The current conviction is that morality will remain unaffected by algorithmic decision making, because the latter can be developed in the service of morality. However, this conviction of being and remaining in control represents dangerous naivety and myopia, because it overlooks the extent to which judgment is assimilated to algorithmic reckoning and the implications thereof for human morality. The naivety of this view is in the failure to see that decision making based on algorithmic reckoning is already co-constituting morality; its myopia is in the failure to see

how a progressive reliance on algorithmic reckoning in decision making may result in a future condition of algorithmic [morality].

In advancing our argument, we embrace an ontological vantage point because it enables us to better understand how phenomena interact (Lawson, 2019). Thus, we examine the ontological assumptions underlying algorithmic reckoning and human judgement, and interpret the mechanism through which judgment is assimilated to algorithmic reckoning as ontological assimilation. Recognizing the current lack of ontological understanding in relation to AI makes it not only more likely for structures of technological domination to go undetected and unchallenged, but also for social interventions to be less likely to succeed.

REFERENCES

- Aggarwal, C. C. 2017. An introduction to outlier analysis. *Outlier analysis*: 1–34. Springer.
- Anderson, B. 2010. Preemption, precaution, preparedness: Anticipatory action and future geographies. *Progress in Human Geography*, 34(6), 777-798.
- Bastani, O., Kim, C., & Bastani, H. 2017. Interpreting blackbox models via model extraction. *ArXiv Preprint ArXiv:1705.08504*.
- Bernstein, L., Beachum, L., & Knowles, H. n.d. Stanford apologizes for coronavirus vaccine plan that left out many front-line doctors. *Washington Post*.
<https://www.washingtonpost.com/health/2020/12/18/stanford-hospital-protest-covid-vaccine/>, December 20, 2020.
- Dignum, V., 2019. *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way*. Springer Nature.
- European Commission, 2020. White paper on artificial intelligence—a European approach to excellence and trust.
- Fleming, P., 2019. Robots and organization studies: Why robots might not want to steal your job. *Organization Studies*, 40(1), pp.23-38.
- Flyverbom, M. and Murray, J., 2018. Datastructuring—Organizing and curating digital traces into action. *Big Data & Society*, 5(2), p.2053951718799114.
- Freeman, R. E., Parmar, B. L., & Martin, K. 2020. *The Power of And: Responsible Business Without Trade-Offs*. Columbia University Press.
- Gebru, T. 2019. Oxford Handbook on AI Ethics Book Chapter on Race and Gender. *ArXiv Preprint ArXiv:1908.06165*.
- Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., et al. 2018. Datasheets for datasets. *ArXiv Preprint ArXiv:1803.09010*.
- Gregory, R.W., Henfridsson, O., Kaganer, E. and Kyriakou, H., 2020. The Role of Artificial Intelligence and Data Network Effects for Creating User Value. *Academy of Management Review*, (ja).
- Haenlein, M. and Kaplan, A., 2019. A brief history of artificial intelligence: On the past, present, and future of artificial intelligence. *California management review*, 61(4), pp.5-14.
- Kellogg, K.C., Valentine, M.A. and Christin, A., 2020. Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), pp.366-410.
- Lawson, T., 2019. *The Nature of Social Reality: issues in social ontology*. Routledge.
- Leavy, S., O’Sullivan, B., & Siapera, E. 2020. Data, Power and Bias in Artificial Intelligence. *ArXiv Preprint ArXiv:2008.07341*.
- Martin, K. 2019a. Ethical Implications and Accountability of Algorithms. *Journal of Business Ethics*, 160(4): 835–850.
- Martin, K. 2019b. Designing Ethical Algorithms. *MIS Quarterly Executive*, 18(2).
- Mejias, U.A. and Couldry, N., 2019. Datafication. *Internet Policy Review*, 8(4).
- Mulligan, D. K., Kluttz, D., & Kohli, N. 2020. Shaping Our Tools: Contestability as a Means to Promote Responsible Algorithmic Decision Making in the Professions. In K. Werbach (Ed.), *After the Digital Tornado*. Cambridge University Press.
- Paullada, A., Raji, I. D., Bender, E. M., Denton, E., & Hanna, A. 2020. Data and its (dis) contents: A survey of dataset development and use in machine learning research. *ArXiv Preprint ArXiv:2012.05345*.
- Perrault, R., Shoham, Y., Brynjolfsson, E., Clark, J., Etchemendy, J., Grosz, B., Lyons, T., Manyika, J., Mishra, S. and Niebles, J.C., 2019. The AI Index 2019 Annual Report. *AI*

Index Steering Committee, Human-Centered AI Institute, Stanford University, Stanford, CA.

Tschang, F.T. and Mezquita, E.A., 2020. Artificial Intelligence as Augmenting Automation: Implications for Employment. *Academy of Management Perspectives*, (ja).

Tsoukas, H., 1997. The tyranny of light: The temptations and the paradoxes of the information society. *Futures*, 29(9), pp.827-843.