On Machine Learning Systems and Production Harms: the case of Pig Farms

AGATHE BALAYN, Microsoft Research / Delft University of Technology, USA / Netherlands
SEDA GÜRSES, Delft University of Technology, Netherlands

Abstract

We wish to introduce a site of investigation that remains understudied in computer science (CS) research: how the way we *develop and deploy machine learning* (ML) systems into organizations requires and leads to transformations of these organizations' *operations* with broader social, political and economic effects. Current CS research focuses on the design of ML algorithms with high-accuracy outputs and of well-calibrated interactions between the users and these outputs. By doing so, this research abstracts the broader production and organizational contexts of ML systems. We posit that it is in studying how we bring ML systems *into the world* –in this case through how we produce such ML systems– that we can come to better understand the impact of these ML systems *on the world*. We propose a conceptual framework to do so using "production" and "operations" as key notions guiding our analysis. In particular, we study how deploying ML systems into operations in the domains of agriculture and manufacturing transforms these operations in the image of software production, with impact on workers, organizations, and the broader domains in which these organizations exist. To conclude, we explore potential explanations for why researchers are more likely to treat ML systems as *output*-generating machines, whose value and risks arise from the potential low quality of these outputs.

1 Introduction

Today, computer science research on machine learning (ML) is centered around the output-generating power of ML systems. Algorithm-oriented researchers develop ML systems that produce high-quality outputs from a bounded set of inputs (e.g., [7, 12]), and Human-Computer Interaction (HCI) researchers design well-calibrated modes of interactions between the end-users and these outputs (e.g., [24]). What researchers articulate as potential risks of ML systems stems from this conceptualization. For example, depending on the domain where an ML system is introduced, inaccurate outputs might lead to accidents and low business value [22], and biased outputs might lead to discrimination of certain populations and may cause reputational damage for ML providers [13]. Interdisciplinary scholars have enriched this understanding of the impact of using an ML system within various domains, by also pointing to the questionable goals of the systems [14], the resulting substitution of human workers and its impact on the labour market [1], or the resulting transformation of ML users' skills and labor conditions [16]. However, for an ML system to be used, it first has to be developed and deployed. Developing and deploying digital systems in target domains has always transformed organizations and their operations [2, 26]. In the context of ML systems, researchers have shown that the software production is organized into "supply chains" that reinforce power asymmetries across ML providers [23], and that it requires the exploitation of resources and low-wage labor that disproportionately exploit people from the Global South [5, 15]. Extending these findings, we argue that the way ML systems are produced in terms of, e.g., software architectures (e.g., services), development and deployment processes (e.g., agile, CI/CD), and economic models (e.g., subscription contracts) determines how ML systems redefine domain operations and reshapes associated practices, working conditions, and the social, political and economic makeup of these domains.

Authors' Contact Information: Agathe Balayn, balaynagathe@microsoft.com, Microsoft Research / Delft University of Technology, USA / Netherlands; Seda Gürses, f.s.gurses@tudelft.nl, Delft University of Technology, Netherlands.

2 Method

Analytical lenses. To aid us conceptually, we use an *operational* lens that we borrow from prior literature in engineering, organization, and science and technology studies (e.g., [4, 9, 17]). With operations, we refer to any activity conducted within an organization to produce (economic) value. We assume the digitalization of operations brings about intended and unintended social and economic transformations to organizations, matters we study empirically and critically.

To do so, we borrow from scholars, including Agre [2], Zuboff [26], Guerses and van Hoboken [10], and Poon [19] who have written about how artifacts composing digital systems (e.g., the data model of any software system) and the mundane *production processes* for these systems (e.g., agile software development) reconfigure operations in the deployment domain (e.g., the commodification of operations). We take methodological inspiration from these prior works to study ML systems and their effects.

Data. To understand the impact of ML production on a domain, we (1) investigate how production of ML systems is organized in practice, and (2) how the operations in a domain are organized before and after deployment of such systems. We conducted two sets of interviews to uncover the practices of ML providers (e.g., ML startups) and ML clients (e.g., the bank, the farmers) who develop, deploy, and use ML systems. First, we recruited 22 data scientists, ML engineers, data engineers, and product managers working within *ML provider* organizations that develop computer vision systems for various kinds of domains related to manufacturing (e.g., detecting defects, actuating robots on conveyor belts). Second, to analyze the effects of the ML systems on their deployment domains, we conducted 20 interviews in one single domain: pig farming. The deployment of computer vision systems in farms is expanding, with applications such as disease detection in barns, pig weight and carcass monitoring, etc. Some of our participants work for ML providers as CEOs, managers, technical developers, expert data annotators, zoological consultants, veterinarians, or researchers. The others are *clients* who may adopt the ML system or interact with them as *end-users* (e.g., farm workers, veterinarians, slaughterhouse managers, product testers). In the interviews, we inquired about ML system development, deployment, and use, as well as the primary challenges and other factors that shape such activities. We also reflected on the potential impacts these activities might have on the deployment domains, from which we derived potential harms.

3 ML production and organizational operations

When deploying ML systems in pig farms, we found that ML providers often request the ML clients to transform the existing physical structures in the farms. This can serve to orchestrate the material infrastructure of the ML systems in the barns, e.g., several ML providers had to restructure the barn cells to secure the cameras and Internet cables from pigs who like to chew on them. This can also serve to improve the accuracy of the ML system. Indeed, ML systems have difficulty producing accurate labels for a wide diversity of input vectors (i.e., they have difficulty generalizing over widely diverse operational domains [8, 18]). In response to this challenge, ML providers turned to increasing their control over the physical environment. For instance, by standardizing the barn structures, they narrowed down the space of environment states that the ML system should learn to draw inferences on, making it easier (and less costly) to develop an ML model with higher accuracy.

Such transformation can improve performance of the system while negatively impacting the deployment domains. (1) The transformations of the environment can be consequential for the pigs, potentially impacting their welfare [6]. Pigs would need to accommodate their habits, like spatial preferences, to the new conditions within the barns [11], they could also lose their enrichments and be burdened by the new lighting conditions in cases where the environment would need to be made more barren for the ML models to work better [21]. (2) By adopting these systems, ML clients create economic and technical dependency on their ML providers. Almost all providers offered ML systems as a service, asserting a continuous relationship in which farm operations are adapted to evolving service offerings. This service model differs economically from purchasing products, e.g., machinery. Adapting farm operations to keep up with a specific provider increases opportunity costs to switch to another provider. Unsubscribing from ML services may require switching back to human driven operation of farms, which may require reconfiguring barns optimized for the ML system, a complex reversal. (3) The high risks and costs associated with adopting ML systems to increase operational efficiency can bolster asymmetries in-between (potential) clients. The costly transformations described above can give advantages to larger players who can afford the costs and risks of adopting ML systems in contrast to smaller players. For instance, one medium-size farmer explained being in a tense economic situation (pig farming is a volatile market [3]), and struggling to convince their bank to adjust their business plan and integrate R&D costs to afford experimenting with the ML system. The experimentation risk that providers expect clients to take when adopting the ML systems is a challenge specific to these systems, that also accentuates inequalities between players in the domain with more or less risk capital. What's more, smaller players are less likely to already operate with a standardized physical structure (which could have facilitated the application of the transformations), making the transformations relatively more costly.

4 Discussion

Our findings reveal that operational transformations are not only driven by "values" in the design of these services. Rather, a lot of the issues arise because of the way software production is organized: when ML systems are offered 'as a service' that is heavily experimental, continuously evolves and binds organizations to associated subscription models, they introduce powerful dynamics within targeted organizations and domains. These findings underline that solely focusing on improving the quality of ML systems' outputs is not sufficient for addressing their harmful impact in deployment domains.

To obtain these findings, we studied the production of ML systems and its resulting operational transformations in tandem. This differentiates our approach from work in the social sciences that provide rich accounts of how the design of ML systems transforms organizations, but not how the design of these systems is shaped by how software production is organized. We invite scholars to study other domains where ML systems are introduced to transform operations and reconfigure the production of value. Similarly, there is great value in studying other production aspects (e.g., the scaling of the ML systems from one to multiple farms) that could reveal further power asymmetries and associated harms.

Operational harms deriving from how ML systems are produced are typically unattended in both computer science and social science literature on machine learning. Beyond presenting this knowledge gap, at the conference, we would like to interrogate the reasons for this gap. We can conjecture that some of this is due to the narrow epistemological focus of scholars on ML outputs, ML capabilities, and responsible AI issues that conveniently contribute to the hype around artificial intelligence and machine learning [25]; the circumscribe methodological interest in data and algorithm-centered activities that the community tends to reward more than grounded but non-generalizable findings [20]; the prioritization of societal harms over harms to non-human beings and political economic harms [22]; and the assumption that software production is not relevant to how computer science manifests itself in the world. Social scientists and STS scholars, on the other hand, provide rich empirical accounts of implementations of ML systems, but typically focus on the design (e.g., data structures, algorithms) or design activities (e.g., user participation) rather than the *conditions of software production* that inform how ML transforms operations in target domains. We posit that paying more attention to software production would help understand how this model of production is neither inevitable nor always desirable, and deserves greater engagement and contestation from computer and social scientists alike.

References

[1] Daron Acemoglu and Pascual Restrepo. 2018. Artificial intelligence, automation, and work. In *The economics of artificial intelligence: An agenda*. University of Chicago Press, 197–236.

[2] Philip E Agre. 1994. Surveillance and capture: Two models of privacy. *The information society* 10, 2 (1994), 101–127.

[3] Thomas M Banhazi, H Lehr, JL Black, H Crabtree, P Schofield, M Tscharke, and D Berckmans. 2012. Precision livestock farming: an international review of scientific and commercial aspects. *International Journal of Agricultural and Biological Engineering* 5, 3 (2012), 1–9.

[4] Rajiv D Banker, Robert J Kauffman, and Richard C Morey. 1990. Measuring gains in operational efficiency from information technology: a study of the Positran deployment at Hardee's Inc. *Journal of Management Information Systems* 7, 2 (1990), 29–54.

[5] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 610–623.

[6] Alex Blanchette. 2020. *Porkopolis: American animality, standardized life, and the factory farm*. Duke University Press.

[7] Terrance De Vries, Ishan Misra, Changhan Wang, and Laurens Van der Maaten. 2019. Does object recognition work for everyone?. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*. 52–59.

[8] Gintare Karolina Dziugaite, Alexandre Drouin, Brady Neal, Nitarshan Rajkumar, Ethan Caballero, Linbo Wang, Ioannis Mitliagkas, and Daniel M Roy. 2020. In search of robust measures of generalization. *Advances in Neural Information Processing Systems* 33 (2020), 11723–11733.

[9] Harun Farocki. 2022. War at a distance / erkennen und verfolgen (2003) dir. Harun Farocki. https://www.youtube.com/watch?v=9iomr1LPWMQ

[10] Seda Gurses and Joris Van Hoboken. 2017. Privacy after the agile turn. (2017).

[11] Dae Yeong Kang, Byeong Eun Moon, Myeong Yong Kang, Jung Hoo Kook, Nibas Chandra Deb, Niraj Tamrakar, Elanchezhian Arulmozhi, and Hyeon Tae Kim. 2025. Analysis of Pig Tendencies to Stay Specific Sections Within the Pig Barn According to Environmental Parameters and Facilities Features. *Agriculture* 15, 12 (2025), 1282.

[12] Kimmo Karkkainen and Jungseock Joo. 2021. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 1548–1558.

[13] Jared Katzman, Angelina Wang, Morgan Scheuerman, Su Lin Blodgett, Kristen Laird, Hanna Wallach, and Solon Barocas. 2023. Taxonomizing and measuring representational harms: A look at image tagging. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 14277–14285.

[14] Os Keyes, Jevan Hutson, and Meredith Durbin. 2019. A mulching proposal: Analysing and improving an algorithmic system for turning the elderly into high-nutrient slurry. In *Extended abstracts of the 2019 CHI conference on human factors in computing systems*. 1–11.

[15] Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. Quantifying the carbon emissions of machine learning. *arXiv preprint arXiv:1910.09700* (2019).

[16] Hao-Ping Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. 2025. The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI conference on human factors in computing systems*. 1–22.

[17] Brad Myers. 1994. Challenges of HCI design and implementation. *interactions* 1, 1 (1994), 73–83.

[18] Magdalini Paschali, Sailesh Conjeti, Fernando Navarro, and Nassir Navab. 2018. Generalizability vs. robustness: investigating medical imaging networks using adversarial examples. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 493–501.

[19] Martha Poon. 2007. Scorecards as devices for consumer credit: The case of Fair, Isaac & Company Incorporated. *The sociological review* 55, 2_suppl (2007), 284–306.

[20] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI. In *Proc. of the CHI Conference on Human Factors in Computing Systems*. 1–15.

[21] Juliette Schillings, Richard Bennett, and David Christian Rose. 2021. Animal welfare and other ethical implications of Precision Livestock Farming technology. *CABI Agriculture and Bioscience* 2, 1 (2021), 17.

[22] Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Rostamzadeh, Paul Nicholas, N'Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, et al. 2023. Sociotechnical harms of algorithmic systems: Scoping a taxonomy for harm reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 723–741.

[23] David Gray Widder and Dawn Nafus. 2023. Dislocated accountabilities in the "AI supply chain": Modularity and developers' notions of responsibility. *Big Data & Society* 10, 1 (2023), 20539517231177620.

[24] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the effect of accuracy on trust in machine learning models. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–12.

[25] Meg Young, Michael Katell, and PM Krafft. 2022. Confronting power and corporate capture at the FAccT Conference. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 1375–1386.

[26] Shoshana Zuboff. 1988. *In the age of the smart machine: The future of work and power*. Basic Books, Inc.