Regulations for Pigs Biting AI infrastructures!? A case for empirically studying AI production to inform and enrich AI governance

Agathe Balayn
agatheb.research@gmail.com
Microsoft Research / Delft University of Technology
USA / Netherlands

Seda Gürses
f.s.gurses@tudelft.nl
Delft University of Technology
Netherlands

Abstract

In this paper, we present the preliminary results of an in-depth empirical study about the ways in which AI-based services are developed and deployed in target domains (e.g., pig farms). We discuss the challenges our findings may pose to harm-based governance efforts. Prior work has proposed looking at harms related to the political economy of AI supply chains. Building on this work, we explore three use cases to reveal how current approaches to producing AI-based services transform the operations of organizations in target domains. Our study reveals novel concerns due to the ways in which these services are zipped into these operations. We close with a discussion on the potential and limitation of harm-based governance approaches to mitigate these more structural concerns we identify.

ACM Reference Format:
Agathe Balayn and Seda Gürses. 2025. Regulations for Pigs Biting AI infrastructures!? A case for empirically studying AI production to inform and enrich AI governance. In *Proceedings of Sociotechnical AI Governance (STAIG - CHI '25).* ACM, New York, NY, USA, 4 pages. https://doi.org/0

1 Proposition: Empirically studying how AI-based services transform the operations of a target domain can uncover new challenges to AI governance

Current AI governance primarily relies on accounting for and mitigating harms [5]. Harms-based governance approaches aim to identify, assess, and mitigate those harms caused by humans relying on AI capabilities. They promote proactivity for instance through the emphasis on design solutions that would support harm mitigation [18].

Researchers and policymakers have contributed to this approach, first through the development of *responsible AI principles* [6], *taxonomies of AI harms* [26], and *algorithmic metrics and methods* for harm mitigation [10], structuring practitioners' and policymakers' reflections and actions around harm. The harms they consider revolve around AI systems' outputs (e.g., biases, lack of explainability), around certain negative impact of their development (e.g., environmental impact), or around job replacement post-deployment. Following these efforts, researchers also increasingly explore ways to support practitioners in *operationalizing* techniques for harm mitigation. For that, they conduct *empirical studies* to investigate the concerns, needs, and challenges faced by individual stakeholders involved in the AI supply chain —primarily developers [4], product managers [24], and UX designers [15], or practitioners with designated "responsible AI" roles [9]. Insights from these works often result in practitioner-centered guidelines and practical tools to foster the development of responsible AI systems [8], or in organizational-level governance frameworks for AI providers [21]. They also inform harms-based AI-governance approaches, be it through regulation, standards, or responsible AI initiatives in organizations.

Another strand of work centers political economic aspects of software systems (AI-based or not) that may lead to harms. Researchers study how service architectures and iterative development processes create an (AI) supply chain, whose potential harms are hard to detect or pin onto a single stakeholder in the chain [3, 7, 28]. They investigate the types of AI models [17], AI development tools [27] and cloud infrastructures [16] that BigTech organizations make available to other providers of software services, and how they come to reinforce the power of the former. Ultimately, these researchers encourage policymakers to explore governance mechanisms that go beyond individual organizations, stretching across supply chains.

Building on such prior work, we present preliminary results of an empirical study to reflect on the potentials and limitations of such a harm-based AI governance approach. The study looks at the production of AI-based services to be introduced into target domains, like factories and farms. We focus on how the production of these services aims to transform the operations of these domains–the daily inner workings of an organization or business–, and how these transformations might ultimately impact humans, non-humans and their environments in complex ways. The results reveal the dynamics AI-based services are bringing to organizations, with insights into potential negative affects. This in turn calls for either extending the scope of existing harm-based policies or exploring governance approaches that address political economic concerns that fall beyond harm considerations.

STAIG - CHI '25, Yokohama, Japan

https://doi.org/0

2 Brief description of the methodology followed to empirically study AI-based services and AI supply chains

We conducted multiple studies interviewing individuals active in the production of AI-based services to be deployed in target domains, and individuals working within these target domains. Interlocutors could be doing in-house research, engineering data infrastructure, instrumenting the physical target environment, etc. We prepared these semi-structured interviews using grey literature relevant to the development of the AI-based services at hand or to their usage in practice.

These studies all followed an exploratory approach. At the start, our aim was broadly to understand how the production practices of the AI-based services might affect the operations of a target domain in ways that are still misunderstood, and how this might ultimately harm humans or their environment. For this reason, we asked an initial round of participants open questions about their production practices and the (envisioned) usages of their services. We identified potential consequential patterns that had not been discussed in prior literature. We then refined the orientation of the studies by adjusting the questions to target these potential patterns.

In terms of use-cases, we looked into the use of AI-based services in different target domains (enterprise, industrial, agricultural), and we focused on multiple AI applications that require different types of underlying AI algorithms. See Table 1 for more information.

Usecase	Target domain	AI-based service	Section of the supply chain
UC1	Operations at a pig farm, e.g., detecting early pig diseases, remote pig weighing.	Deep learning for computer vision tasks, e.g., video classification	Development, deployment, and usage of the AI system, e.g., by the farmers, veterinarians, slaughterhouses.
UC2	Operations in manufacturing, e.g., detecting defects, actuating robots on conveyor belts.	Deep learning for computer-vision tasks, e.g., object detection.	Development of the AI system and deployment in its physical environment, e.g., in the robotic arm's factory
UC3	Operations at an enterprise, e.g., answering employees HR questions.	Large language models; retrieval-augmented generation.	Development of the AI system by the provider and adoption decisions by the deployer.

Table 1: Synthesis of the AI use-cases we explored, and of the variations in our approaches for exploring these use-cases.

3 Illustration of a few findings from these empirical studies (focus on UC1)

The objective of the AI-based services we studied is typically to optimize the operations of a target domain, often to ultimately increase efficiency, scale, or speed, to reduce costs, or to increase revenue. Observations made in the target domain of pig farms (UC1) reveal how such an *aim* might be problematic in itself. Similarly to prior works that hypothetically discussed how and why AI-based services are introduced in a domain [14], several interview participants pointed out that the service clients aim to increase the production of pigs in already controversial intensive farming settings. This might not only bear implications for the pigs' populations (e.g., life conditions and mortality of the pigs [19]), but also for the environment (due to the resources the pig production requires and the waste it produces) [25], as well as the spread of diseases [1, 23].

Our observations also show that *producing* AI-based services is based on seemingly-inconspicuous transformations of the target domain, service architectures, and a continuous relationship between the AI service provider (e.g., the startup that builds the AI system) and the AI client (e.g., the farmer) in the target domain. We describe the consequences of deploying AI systems in such a way:

- **Impact on farmers' practices.** We found how the impact of AI systems on *jobs* is complex and reflects the production needs and purpose of AI-based services. In our interviews, producing AI services did not 'simply' require employing crowd workers with poor labor conditions [20], or replacing human-workers in their jobs [2]. Instead, actors in the farming sector get sucked into the production of AI-based services. This has to do with the way AI systems change how services are produced. Rather than augmenting or automating farmers' practices, AI-based services aim to mimic outcomes of workflows using predictions on input data. Once the service is deployed, farmers are increasingly asked to monitor it for potential mistakes in its predictions and to provide feedback to the AI provider. While their own work becomes partially automated, they become service "care-givers". Additionally, these actors as well as zoologists who typically conduct research about the farm animals might become AI "feeders" or even "designers". They are employed by the AI providers to support service design, be it for defining the labels to be predicted, determining the data to collect for AI training, or annotating videos of pigs. Importantly, in the long term, these farm actors might lose their domain knowledge that they acquire working closely with pigs [29], while transferring some of their knowledge to the AI provider who gains an upper hand.
- **Impact on the eco-system.** We found how the economic constraints on AI-based services bear *"social" implications*, that typically remain neglected. For instance, AI providers may optimize their services for particular pig breeds, leaving out breeds for which the services cannot deliver economically viable operational outcomes (e.g., pigs of darker colors are more difficult to detect in barns compared to pigs of lighter complexion). This might encourage the farm owners that want to adopt such AI-based services to concentrate their production around the breeds that AI providers can handle. Ultimately, this risks further decreasing biodiversity, which in the case of intensive pig farming is already an issue with broader health and environmental concerns [12].
- **Impact on the physical environment.** We found a tendency to transform the environments where the AI-based service is deployed to ensure better predictions by the AI system. AI-based services require a complex set-up within the farms (UC1) and an important reorganization of the physical infrastructure of the barns, e.g., to bring the Internet and Wifi connections, to run cables in between the pig cells, to install a well-protected server room, etc. AI providers face difficulties in building robust AI models that can adapt to any farm condition, and services that can resist the normal wear-and-tear provoked by the harsh farm conditions (e.g., pigs biting the camera cables, dust from the barns, etc.). To mitigate these problems across multiple farms, they experiment with standardizing cameras, changing light installments in barns, as well as constraining the movement of the animals. These changes can be costly for small farmers, and more

easily delivered by industrial farms with standardized physical installations.

- **Impact on economic condition and power relations.** We found that the AI deployment model gives the AI provider an upper hand in the transformation of farm operations. UC1 showed that AI providers (typically startups) rarely provide shrink-wrapped products, mostly producing services that establish a continuous economic relationship with the farmers and the provider companies. As part of their service, providers maintain the hardware placed in the farm, making switching providers cumbersome. Providers can use the data streams to update the internal AI algorithms or develop new algorithms, which they can sell as improved services, or to other farms. Farmers, in adopting services, depend on the economic success of the providers for their operational functioning. Besides, farmers' operations get coupled with the economic interests of the providers. If the providers need to cut costs, or decide to focus on a breed of pigs, they may transform the AI-based services in a way that no longer aligns with the farmers. This means, once a provider is successfully zipped into the operations of sufficient farms, they have an upper hand on a great number of farms, with the potential to push their operations to better fit their own economic interests. Once in this situation, farmers may not easily reverse back to not using the technology, e.g., in case they cannot afford it anymore, in case it doesn't work well enough to increase their productivity, or it requires greater physical changes than they can handle.

4 Examples of governance-related questions stemming from these findings

One way to analyze the results of our empirical study is to identify the harms hidden in the interviews, be it harms to the farmers, to the pigs, or to the broader farm eco-systems. These types of harm remain under-explored and would possibly constitute a new category within harm taxonomies [22, 26]. We explained in Sec. 1 that most AI researchers focus on understanding and tackling harms resulting from the output predictions of AI–based services, especially representational and allocative harms to individuals or groups [13], that have also been tackled in recent AI regulations. Such harms would not apply to our use cases, contrary to the harms we actually identified.

Methodologically, we were able to explore these impacts and harms by grounding our studies in specific domains of exploration, centering on specific AI systems with seemingly low risk (e.g., according to the AI Act), and adopting an analytical lens than includes software engineering processes with an attention to how companies providing and adopting AI systems are organizing their production. This is in contrast to prior work that focuses on datasets, as well as the inputs and outputs of an AI algorithm. To better understand the transformations that we touch on here, and whether they may lead to systemic harms that cannot be easily mitigated, we call on researchers to continue exploring various perspectives from which harm can be understood. This would include studying impact on different stakeholders, adopting different lenses (e.g., political economic, social, environment, etc.), and exploring in-depth specific domains where AI-based services are introduced.

How can we tackle these harms of AI-based services? We could find ways to address these harms individually. This could include ensuring that farmers are paid sufficiently for data annotation work, documenting and archiving richer accounts of farm work, pushing for biodiversity as an optimization goal, or developing AI systems that are robust to different environmental conditions. However, similar to prior studies on annotators, gig workers, and AI supply chains, our studies also reveal structural transformations that may occur due to how we produce AI-based services. The political economic conditions under which AI-based services are produced, as well as the existing conditions in the target domains suggest the existence of structural issues Specifically, we sketch how target organizations are transformed once they adopt AI-based services to enhance their operations, zipping them into long-term economic relationships with the service providers. We show how the economic and technical conditions in which these services come to be may establish systemic inequalities between service providers and organizations in the target domain, with potential harms to humans, non-humans and environments.

These issues may push the limits of harm-based approaches to govern them. The loss of farming skills, the diminution of the biodiversity, or the further reduction of small farms in the face of a large AI push and other power asymmetries, are transformations that are hard to mitigate and reverse without broader efforts. While we could go on to list the harms that stem from the transformations we described, our focus remains on these systemic transformations. They require us to question their root cause, i.e., the production model intrinsic to AI-based services. We identified that AI harms result from the requirements for producing AI technologies, the limitations of these technologies in terms of performance, and the subsequent way in which AI providers deliver the AI systems to the users (especially as services). Hence, we call on researchers to pay attention not only to individual stakeholders' work or organizational challenges, but also to production processes and their materiality (e.g., the re-organization of the barns), and to investigate how these aspects of AI might cause deep transformations of target domains. During the workshop, we could explore these systemic issues and what other governance approaches may be appropriate to govern them.

Services have not started from AI technologies: prior software systems were also offered as a service as this is more profitable to the software providers [11]. They can continuously obtain information about the service users, refine and extend their offering, continuously control pricing, etc. Having in mind this history that shapes the production of information technologies and especially AI-based services nowadays enables us to pose larger questions. Governance efforts might want to start by disentangling the factors that cannot be circumvented and that are inherent to AI technologies, from those factors that are only a feature of avoidable practical challenges. To what extent does service-based procurement stem from a desire for profitability for the AI provider or also from AI-specific challenges? Investigating what other modes of AI procurement (e.g., as a shrink-wrapped product) would be feasible could be generative to avoid many AI harms. More broadly, to what extent has the development of AI techniques been influenced by AI providers' desire to create services? Can we envision other techniques that would circumvent the problematic transformations?

References

[1] 2025. https://en.wikipedia.org/wiki/Intensive_pig_farming

[2] 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.

[3] Agathe Balayn and Seda Guerses. 2024. Misguided: AI regulation needs a shift in focus. *Internet Policy Review* 13, 3 (2024). https://policyreview.info/articles/news/misguided-ai-regulation-needs-shift/1796

[4] Agathe Balayn, Mireia Yurrita, Jie Yang, and Ujwal Gadiraju. 2023. "Fairness Toolkits, A Checkbox Culture?" On the Factors that Fragment Developer Practices in Handling Algorithmic Harms. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. 482–495.

[5] David Baldridge, Beth Coleman, and Jamie Amarat Sandhu. 2024. The terminology of AI Regulation: Preventing "harm" and mitigating "risk". https://srinstitute.utoronto.ca/news/terminology-regulation-risk-harm

[6] Roger Clarke. 2019. Principles and business processes for responsible AI. *Computer Law & Security Review* 35, 4 (2019), 410–422.

[7] Jennifer Cobbe, Michael Veale, and Jatinder Singh. 2023. Understanding accountability in algorithmic supply chains. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1186–1197.

[8] Wesley Hanwen Deng, Manish Nagireddy, Michelle Seng Ah Lee, Jatinder Singh, Zhiwei Steven Wu, Kenneth Holstein, and Haiyi Zhu. 2022. Exploring how machine learning practitioners (try to) use fairness toolkits. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 473–484.

[9] Advait Deshpande and Helen Sharp. 2022. Responsible ai systems: who are the stakeholders?. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*. 227–236.

[10] Hadi Elzayn, Shahin Jabbari, Christopher Jung, Michael Kearns, Seth Neel, Aaron Roth, and Zachary Schutzman. 2019. Fair algorithms for learning in allocation problems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 170–179.

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

[12] Juan Manuel Herrero-Medrano, Hendrik-Jan Megens, Martien AM Groenen, Mirte Bosse, Miguel Pérez-Enciso, and Richard PMA Crooijmans. 2014. Whole-genome sequence analysis reveals differences in population management and selection of European low-input pig breeds. *BMC genomics* 15 (2014), 1–12.

[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] Q Vera Liao, Hariharan Subramonyam, Jennifer Wang, and Jennifer Wortman Vaughan. 2023. Designerly understanding: Information needs for model transparency to support design ideation for AI-powered user experience. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–21.

[16] Dieuwertje Luitse. 2024. Platform power in AI: The evolution of cloud infrastructures in the political economy of artificial intelligence. *Internet Policy Review* 13, 2 (2024), 1–44.

[17] Dieuwertje Luitse and Wiebke Denkena. 2021. The great transformer: Examining the role of large language models in the political economy of AI. *Big Data & Society* 8, 2 (2021), 20539517211047734.

[18] Tambiama Madiega. 2021. Artificial intelligence act. *European Parliament: European Parliamentary Research Service* (2021).

[19] Steven McCulloch. 2024. The veterinary profession must condemn intensive pig farming. *Veterinary Record* 195, 8 (2024), 343–343.

[20] Milagros Miceli, Julian Posada, and Tianling Yang. 2022. Studying up machine learning data: Why talk about bias when we mean power? *Proceedings of the ACM on Human-Computer Interaction* 6, GROUP (2022), 1–14.

[21] Bogdana Rakova, Jingying Yang, Henriette Cramer, and Rumman Chowdhury. 2021. Where responsible AI meets reality: Practitioner perspectives on enablers for shifting organizational practices. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–23.

[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] Hugo Francisco de Souza. 2023. Intensive pig farming linked to dangerous zoonotic pathogens, study warns. https://www.news-medical.net/news/20231116/Intensive-pig-farming-linked-to-dangerous-zoonotic-pathogens-study-warns.aspx

[24] Hariharan Subramonyam, Jane Im, Colleen Seifert, and Eytan Adar. 2022. Solving separation-of-concerns problems in collaborative design of human-AI systems through leaky abstractions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–21.

[25] International Environmental Technology. [n. d.]. Intensive pig farming is contributing to air pollution in the UK, says New report. https://www.envirotech-online.com/news/air-monitoring/6/international-environmental-technology/intensive-pig-farming-air-pollution-uk-report/63227

[26] Risto Uuk, Carlos Ignacio Gutierrez, Daniel Guppy, Lode Lauwaert, Atoosa Kasirzadeh, Lucia Velasco, Peter Slattery, and Carina Prunkl. 2024. A Taxonomy of Systemic Risks from General-Purpose AI. *arXiv preprint arXiv:2412.07780* (2024).

[27] Fernando van der Vlist, Anne Helmond, and Fabian Ferrari. 2024. Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. *Big Data & Society* 11, 1 (2024), 20539517241232630.

[28] 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.

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