

How Could AI Supply Chain Research Shape HCI Inquiries And Vice-Versa?

Inha Cha
Georgia Institute of Technology
Atlanta, Georgia, USA
icha9@gatech.edu

David Gray Widder
School of Information
University of Texas at Austin
Austin, Texas, USA
dw658@cornell.edu

Blair Attard-Frost
University of Toronto
Toronto, Ontario, Canada
blair@blairaf.com

Jatinder Singh
Department of Computer Science &
Technology
University of Cambridge
Cambridge, United Kingdom
Research Center Trustworthy Data
Science and Security (RC-Trust) of the
UA Ruhr
University Duisburg-Essen
Duisburg, Germany
jatinder.singh@cl.cam.ac.uk

Agathe Balayn
Microsoft Research
New York City, New York, USA
balaynagathe@microsoft.com

Abstract

HCI research on AI has largely focused on end-user interactions or the practices of developers working to improve system performance. This meetup proposes to broaden that scope of interaction to encompass the material, political, and economic dynamics that shape AI across its lifecycle. Drawing on STS, we recognize AI as the outcome of extended supply chains involving practitioners, organizations, infrastructures, and governance, all entangled with power asymmetries and economic incentives. Much of this supply chain remains invisible within HCI, yet it profoundly influences how AI systems are designed, deployed, and experienced. The meetup will bring together HCI scholars to explore how supply chain perspectives can enrich HCI research. We plan to gather researchers of the HCI community and beyond, in order to discuss how they could integrate the AI supply chain in their research and what would be the implications, and how HCI can further contribute to supply chain studies.

CCS Concepts

• **Human-centered computing**; • **Social and professional topics** → **Computing / technology policy**; • **Computing methodologies** → **Artificial intelligence**; • **Applied computing** → **Law, social and behavioral sciences**;

Keywords

AI, Supply Chain, Value Chain, Research Agenda, AI Lifecycle, Infrastructure, Sociotechnical Systems

ACM Reference Format:

Inha Cha, David Gray Widder, Blair Attard-Frost, Jatinder Singh, and Agathe Balayn. 2026. How Could AI Supply Chain Research Shape HCI Inquiries And Vice-Versa?. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3772363.3778800>

1 Motivation and goal for the meetup

Research in HCI has often conceptualized artificial intelligence (AI) through the lens of human–AI interaction. For instance, HCI scholars have worked on the design of interfaces that allow end-users to interact with AI systems appropriately, and conducted research on the impact that these interactions have on end-users [1, 17, 23]. Additionally, HCI scholars have started to examine the AI lifecycle [13] with, e.g., discussions of stakeholder inclusion across design and evaluation [11], and infrastructural studies on the human and computational resources that underpin AI systems [15, 24]. Resonating with critical AI studies and STS, which conceptualize AI as an assemblage spanning data, models, compute, infrastructures, labor, and governance [2, 9, 20], this work underscores that such AI supply chains are not neutral backdrops but active forces shaping practices and interactions [14, 16].

Objectives of the meetup. This meetup seeks to start consolidating and expanding such a new lens on AI systems, and explore the many research opportunities that remain. We aim to foster a discussion that will expand the scope of what constitutes an "interaction" with an AI system to include the material, political, and economic interactions that exist throughout the system lifecycle

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI EA '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2281-3/26/04

<https://doi.org/10.1145/3772363.3778800>

and are essential for the system's development and operation. Beyond advancing this research, the goal of this meetup is to initiate a community and bring together HCI researchers interested in AI infrastructures, materiality, and political economy to foster connections across subfields. In doing so, the meet-up will appeal to researchers of AI supply chains, pipelines, and value chains; those working in responsible AI, governance, and participatory design; sub-communities focused on contestability, infrastructural perspectives, and domain-specific AI applications; and scholars building on STS and CSCW traditions. In conclusion, we see the session as both a networking opportunity and a facilitated discussion to collectively imagine how HCI can better engage with AI supply chains.

Examples of the ways in which the supply chain perspective interacts with HCI research. We now provide examples of research pathways that could be discussed during the meetup as they expose synergies between HCI research and research on AI supply chains.

First, the supply chain perspective invites us to revise our understanding of the interactions between the end-users and the AI systems. Rather than focusing only on the interactions between the end-users and the AI systems that is produced by the supply chain, this perspective invites HCI researchers to unpack how end-user experiences are also shaped by the infrastructural dependencies and procedural activities happening within the supply chain. For instance, although the *deployment* of the AI systems is often seen as a mere engineering challenge, application-specific literature shows that it can also be a source of harm for the end-users interacting with the AI systems. E.g., deploying computer-vision-based AI systems in farms creates economic dependencies between the farmers (end-users) and the AI providers (developers) who have to continuously pay a subscription for the material infrastructure (e.g., cameras in barns) to be maintained –maintenance being a procedural aspect of the supply chains– [7], and it might also diminish the welfare of the farmed animals whose barns have to be restructured to let such material infrastructure in [10]. From a supply-chain perspective, these consequences of interacting with AI are not “side effects”, but integral parts of the human–AI interaction: the farmer's experience is mediated as much by subscription contracts, infrastructure maintenance, and the welfare of animals, as by the interface with the AI system.

Second, the supply chain perspective invites us to reframe the HCI works that have analyzed the interactions between the AI developers and AI systems [5, 12]. On the one hand, this perspective invites us to consider more comprehensively the complex social and political interactions that unfold across layers of policy, compute allocation, and/or system design. This enables interrogation about how infrastructures, policy frameworks, and social factors interact to create AI systems. On the other hand, this perspective highlights the lack of studies on the ways AI developers further engineer the material infrastructures underlying AI systems and impact governance layers.

Third, the supply chain perspective opens up new research questions that HCI scholars are well-placed to tackle. The growing debate on AI sovereignty often framed at the national level, exemplifies this trend [18, 19, 21]. While sovereignty has been extensively analyzed in AI policy as part of supply chain concerns, it remains

still underexplored in HCI. Engaging with it offers HCI researchers an opportunity to investigate how sovereignty is imagined, enacted, and contested in specific practices and contexts. This line of inquiry resonates with prior HCI work on how collective visions of technology are translated and materialized in development practice and extends it to the AI development contexts. Questions of why sovereignty is pursued, by whom, which aspects of AI are subject to sovereign claims, and how such claims are enacted reveal sovereignty as an ongoing negotiation between autonomy, interdependence, and competing values. For HCI, this opens pathways to examine how sovereignty discourses shape sociotechnical practices across technical, organizational, and political layers. Crucially, these inquiries can inform how AI applications are developed and how interactions are designed for end users. For example, co-design, participatory, and speculative methods could be mobilized to surface community values that guide not only system-level decisions but also the design of concrete interaction mechanisms, interfaces, and practices of control.

2 Structure of the meetup sessions

We envision the meetup to be organized around three discussions, to kick start our workstream on AI supply chains. These discussions would happen both in small groups and with all the participants together, in order to facilitate multiple modes of participation. Discussion (1) would invite participants to reflect on how they relate to the concept of the supply chain, and reflect on potential future ways in which their research could benefit from the supply chain perspective. Discussion (2) would encourage participants to discuss what might be drivers and blockers of work around AI supply chains at CHI. This could be practical, organizational, or disciplinary challenges, that could potentially be solved by joining forces as a new interdisciplinary research community. Discussion (3) would center around the next steps for this research community. Led by the meetup organizers, this discussion would first consist in a synthesis of prior discussions and of the common themes that arose (to facilitate this process, the organizers will take notes all along the meetup on a common document), and then open to a larger interrogation about what the participants think could be the most appropriate next steps (e.g., developing a common research agenda, setting up case-studies).

3 Organizers

We have assembled a team of five organizers who all have conducted work around AI supply chains in the past (e.g., [2, 4, 8, 22]), have a foot in various disciplines (e.g., HCI, ML, systems, security, law, governance, STS, political economy), and have experience organizing workshops at various conferences (e.g., [3],[6]).

- **Inha Cha** is a PhD Student at Georgia Tech. Her research critically examines the sociotechnical dynamics, institutional logics, and geopolitical forces that shape how AI is developed and deployed.
- **David Gray Widder** is an assistant professor in the School of Information at the University of Texas at Austin, and an affiliate of Data and Society. He studies how people creating AI systems think about the downstream harms their systems make possible,

and the wider cultural, political, and economic logics which shape these thoughts.

- **Blair Attard-Frost** is an assistant professor in the Department of Political Science at University of Alberta, and a research fellow at the Alberta Machine Intelligence Institute (Amii). Her research addresses challenges of power, participation, and equity in AI governance and policymaking.
- **Jat Singh** is a professor at the Research Centre Trustworthy Data Science & Security (RC-Trust, Germany) and the Department of Computer Science and Technology, University of Cambridge. He leads the Compliant & Accountable Systems group, taking a socio-technical systems approach to governance, accountability, and compliance, with current research examining AI supply chains and service models.
- **Agathe Balayn** is a postdoctoral researcher at Microsoft Research. She has conducted extensive qualitative work in organizations producing and consuming AI systems, to understand the concerns of stakeholders along the AI supply chain, the factors impacting their practices, and the harms that might arise from their activities.

References

- [1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.
- [2] Blair Attard-Frost and David Gray Widder. 2025. The ethics of AI value chains. *Big Data & Society* 12, 2 (2025), 20539517251340603.
- [3] Agathe Balayn, Yulu Pi, David Gray Widder, Kars Alfrink, Mireia Yurrita, Sohini Upadhyay, Naveena Karusala, Henrietta Lyons, Gagatay Turkay, Christelle Tesono, et al. 2024. From stem to stern: Contestability along AI value chains. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing*. 720–723.
- [4] Agathe Balayn, Mireia Yurrita, Fanny Rancourt, Fabio Casati, and Ujwal Gadiraju. 2025. Unpacking Trust Dynamics in the LLM Supply Chain: An Empirical Exploration to Foster Trustworthy LLM Production & Use. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–20.
- [5] 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.
- [6] Eric P. S. Baumer, Inha Cha, Vera Khovanskaya, Rosemary Steup, Janet Vertesi, and Richmond Y. Wong. 2025. Exploring Resistance and Other Oppositional Responses to AI. In *Companion Publication of the 2025 Conference on Computer-Supported Cooperative Work and Social Computing (CSCW Companion '25)*. Association for Computing Machinery, New York, NY, USA, 156–160. doi:10.1145/3715070.3748295
- [7] Alberto Cavazza, Francesca Dal Mas, Maura Campa, and Valerio Brescia. 2023. Artificial intelligence and new business models in agriculture: the “ZERO” case study. *Management Decision* (2023).
- [8] 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.
- [9] Kate Crawford. 2021. *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- [10] Marian Stamp Dawkins. 2021. Does smart farming improve or damage animal welfare? Technology and what animals want. *Frontiers in Animal Science* 2 (2021), 736536.
- [11] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The Participatory Turn in AI Design: Theoretical Foundations and the Current State of Practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization* (Boston, MA, USA) (EAAMO '23). Association for Computing Machinery, New York, NY, USA, Article 37, 23 pages. doi:10.1145/3617694.3623261
- [12] 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.
- [13] Shipi Dhanorkar, Christine T. Wolf, Kun Qian, Anbang Xu, Lucian Popa, and Yunyao Li. 2021. Who needs to know what, when?: Broadening the Explainable AI (XAI) Design Space by Looking at Explanations Across the AI Lifecycle. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference* (Virtual Event, USA) (DIS '21). Association for Computing Machinery, New York, NY, USA, 1591–1602. doi:10.1145/3461778.3462131
- [14] W. Keith Edwards, Mark W. Newman, and Erika Shehan Poole. 2010. The infrastructure problem in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) (CHI '10). Association for Computing Machinery, New York, NY, USA, 423–432. doi:10.1145/1753326.1753390
- [15] Gabriel Grill. 2024. Constructing Capabilities: The Politics of Testing Infrastructures for Generative AI. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FAccT '24). Association for Computing Machinery, New York, NY, USA, 1838–1849. doi:10.1145/3630106.3659009
- [16] Steven J Jackson, Paul N Edwards, Geoffrey C Bowker, and Cory P Knobel. 2007. Understanding infrastructure: History, heuristics and cyberinfrastructure policy. *First monday* (2007). <https://firstmonday.org/ojs/index.php/fm/article/download/1904/1786>
- [17] 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.
- [18] Daniel Mügge. 2024. EU AI sovereignty: For whom, to what end, and to whose benefit? *Journal of European Public Policy* 31, 8 (2024), 2200–2225.
- [19] Huw Roberts, Emmie Hine, and Luciano Floridi. 2023. Digital Sovereignty, Digital Expansionism, and the Prospects for Global AI Governance. In *Quo Vadis, Sovereignty? New Conceptual and Regulatory Boundaries in the Age of Digital China*. Springer, 51–75.
- [20] Lucy Suchman. 2023. The uncontroversial ‘thingness’ of AI. *Big Data & Society* 10, 2 (2023), 20539517231206794.
- [21] UCL. 2024. Reclaiming digital sovereignty. <https://www.ucl.ac.uk/bartlett/public-purpose/publications/2024/dec/reclaiming-digital-sovereignty>
- [22] 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.
- [23] Qian Yang, Aaron Steinfeld, Carolyn Rosé, and John Zimmerman. 2020. Re-examining whether, why, and how human-AI interaction is uniquely difficult to design. In *Proceedings of the 2020 chi conference on human factors in computing systems*. 1–13.
- [24] Ben Zefeng Zhang, Tianling Yang, Milagros Miceli, Oliver L. Haimson, and Michaelanne Thomas. 2025. The Making of Performative Accuracy in AI Training: Precision Labor and Its Consequences. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 1203, 19 pages. doi:10.1145/3706598.3713112