

Community-Centered AI Governance

Bhavana Bheem

Massachusetts Institute of Technology
bbheem@mit.edu

The rapid expansion of artificial intelligence (AI) has intensified debates surrounding governance, accountability, and democratic participation in technological decision-making. Existing AI governance and policy research primarily focuses on national regulation and industry self-governance, where decision-making processes are often distant from the communities most affected by AI deployment (Bakiner 2023; Bareis and Katzenbach 2022; Hjaltalin and Sigurdarson 2024). In parallel, human-computer interaction (HCI) research has focused on design and audit methods for enabling users to exercise greater control, transparency, and agency over the technologies they use (Shen et al. 2021; DeVos et al. 2022; Lam et al. 2022). Yet, this work positions individuals primarily as end-users, consumers, and bystanders of technology, with less attention to how they participate in shaping and contesting technology beyond their role as users. Emerging HCI research has begun to challenge this framing by recognizing communities as stakeholders in technological decision-making and by examining technology as part of civic and political life (Eslami et al. 2025; Chan 2025). Thus, my dissertation research builds on this perspective by studying how communities engage with technology as civic actors, particularly through their participation in local governance and community organizing. I argue that examining these forms of civic engagement can reveal community priorities and values that may be difficult to capture through conventional user-centered methods, creating opportunities to bring those priorities into the design and governance of more accountable AI systems.

In this extended abstract, I focus on two projects within my dissertation research that examine local governance and community participation in AI policymaking and infrastructure development. These projects explore how communities articulate alternative visions of AI governance that challenge dominant national and industry sociotechnical imaginaries of technological development. The first project, which was accepted at AIES 2026, examines how local governments in the United States regulate artificial intelligence through municipal ordinances, procurement policies, executive directives, registries, and oversight mechanisms. With federal AI regulation largely non-existent, local governments have increasingly emerged as important sites of contestation, nego-

tiation, governance, and resistance surrounding AI systems and infrastructure (Marcucci, Kalkar, and Verhulst 2022). Drawing on theories of AI localism and counter-imaginaries (Marcucci, Kalkar, and Verhulst 2022; Kazansky and Milan 2021), this project investigates how local governments articulate values and priorities through AI governance. Specifically, the study addresses two research questions: (1) what domains and applications of artificial intelligence are targeted by local government policies; and (2) what values, priorities, and motivations are reflected in local government approaches to AI governance?

Using qualitative content analysis with inductive coding (Zhang, Wildemuth et al. 2009), I analyzed 77 AI policies enacted by US local governments before October 2025. Thematic analysis of the codes reveals three major counter-imaginaries within local AI governance efforts: democratic oversight, rights-based governance, and place-based alignment. First, many local policies foreground mechanisms for transparency, public participation, and contestation, emphasizing democratic oversight as a core feature of AI governance. Second, local governments frequently frame AI regulation through rights-based approaches grounded in civil rights, privacy, equity, and nondiscrimination. Third, many policies emphasize place-based alignment, positioning AI governance as an extension of local civic priorities and community values rather than adopting universalized approaches to technological governance. These findings suggest that local governments not only fill regulatory gaps left by federal inaction, but also articulate alternative governance models that foreground democratic accountability and situated community priorities.

While the first project examines how local governments formalize alternative AI governance frameworks through policy, the second examines how communities organize against AI infrastructure development before it becomes embedded within their local environments. Specifically, this project focuses on proposed hyperscale data center development in a small, semi-rural community in central Illinois, where residents and local organizations are navigating competing interests around land use, energy demand, water resources, environmental impacts, and the concentration of technological and economic power. As Illinois becomes a major site of data center development (Seets and Radde 2026), these conflicts provide an opportunity to examine

how communities participate in decisions about whether and how such development should proceed.

The project asks how community members understand and respond to the social, environmental, and political implications of data center development, and how their concerns enter—or fail to enter—local decision-making processes. While local governments have introduced mechanisms such as moratoriums and advisory task forces to facilitate public participation, community concerns are often incorporated only superficially into these processes (Gilman 2022). To examine these dynamics, we are conducting participatory workshops in collaboration with local organizations and in public community spaces, grounded in critical community practice and community-engaged scholarship (Butcher 2007; Blanchard and Furco 2021). The workshops seek to surface lived experiences and community concerns, support collective reflection on the political and economic forces shaping AI infrastructure development, and create space for communities to develop responses to local governance processes.

We will analyze workshop discussion notes and co-created materials using qualitative thematic analysis alongside critical analysis of infrastructural power. By centering rural and micro-urban perspectives, this project aims to show how communities develop situated understandings of AI infrastructure and make visible concerns that may be overlooked by formal planning and governance processes. It also seeks to extend research on the material and environmental politics of AI by examining infrastructure development not only as a technical or economic process, but as a site of community participation and contestation.

These two projects demonstrate how AI governance increasingly unfolds through local policymaking and community contestation. The first project shows how local governments articulate alternative AI imaginaries through enacting local policies that foreground democratic oversight, rights-based protections, and place-based priorities. The second examines how communities contest the material development of AI infrastructure, center localized expertise, and make visible social, environmental, and political concerns. These projects contribute to AI ethics and governance scholarship by expanding attention beyond national regulation and institutional governance toward the local contexts in which AI systems are negotiated, contested, and experienced. They also contribute to HCI by reframing communities not only as users or stakeholders of technological systems, but as civic actors whose collective practices can surface priorities and values relevant to the design and governance of those systems. This work connects AI ethics, technology policy research, and community-centered participatory research to develop approaches for designing and governing more accountable AI technologies that are responsive to the communities in which these systems are developed and deployed.

Acknowledgments

I would like to thank my advisor, Dr. Karrie Karahalios, as well as my collaborators Dr. Madelyn Rose Sanfilippo, Dr. Anita Chan, and Dr. Camille Cobb for their guidance, mentorship, and support throughout my dissertation research.

References

- Bakiner, O. 2023. Pluralistic sociotechnical imaginaries in Artificial Intelligence (AI) law: the case of the European Union's AI Act. *Law, Innovation and Technology*, 15(2): 558–582.
- Bareis, J.; and Katzenbach, C. 2022. Talking AI into being: The narratives and imaginaries of national AI strategies and their performative politics. *Science, Technology, & Human Values*, 47(5): 855–881.
- Blanchard, L.; and Furco, A. 2021. Faculty engaged scholarship: Setting standards and building conceptual clarity.
- Butcher, H. 2007. What is critical community practice? Case studies and analysis. In *Critical community practice*, 33–50. Policy Press.
- Chan, A. S. 2025. *predatory DATA*. University of California Press.
- DeVos, A.; Dhabalia, A.; Shen, H.; Holstein, K.; and Eslami, M. 2022. Toward User-Driven Algorithm Auditing: Investigating users' strategies for uncovering harmful algorithmic behavior. In *Proceedings of the 2022 CHI conference on human factors in computing systems*, 1–19.
- Eslami, M.; Fox, S.; Shen, H.; Fan, B.; Lin, Y.-R.; Farzan, R.; and Schwanke, B. 2025. From margins to the table: Charting the potential for public participatory governance of algorithmic decision making. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, 2657–2670.
- Gilman, M. E. 2022. Beyond window dressing: Public participation for marginalized communities in the datafied society. *Fordham L. Rev.*, 91: 503.
- Hjaltalin, I. T.; and Sigurdarson, H. T. 2024. The strategic use of AI in the public sector: A public values analysis of national AI strategies. *Government Information Quarterly*, 41(1): 101914.
- Kazansky, B.; and Milan, S. 2021. “Bodies not templates”: Contesting dominant algorithmic imaginaries. *New Media & Society*, 23(2): 363–381.
- Lam, M. S.; Gordon, M. L.; Metaxa, D.; Hancock, J. T.; Landay, J. A.; and Bernstein, M. S. 2022. End-user audits: A system empowering communities to lead large-scale investigations of harmful algorithmic behavior. *proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2): 1–34.
- Marcucci, S.; Kalkar, U.; and Verhulst, S. 2022. AI localism in practice: Examining how cities govern AI. *Available at SSRN 4284013*.
- Seets, S.; and Radde, K. 2026. Most new data centers in the U.S. are coming to rural areas. Accessed: 2026-05-21.
- Shen, H.; DeVos, A.; Eslami, M.; and Holstein, K. 2021. Everyday algorithm auditing: Understanding the power of everyday users in surfacing harmful algorithmic behaviors. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2): 1–29.
- Zhang, Y.; Wildemuth, B. M.; et al. 2009. Qualitative analysis of content. *Applications of social research methods to questions in information and library science*, 308(319): 1–12.