



COMMUNITY IN THE LOOP:

A Community-Led Framework for Municipal AI Governance in Austin, Texas

In partnership with the City of Austin, the Austin AI Alliance, and a network of community organizations that helped convene residents across Austin.

Jameila “Meme” Styles, MPA | Doctoral Candidate, Doctor of Business Administration



“As Austin becomes a global technology city, Responsible AI ensures innovation does not outpace equity, democracy, or human dignity.”

ABSTRACT

As artificial intelligence (AI) rapidly expands into public systems, municipalities across the United States face increasing pressure to establish governance frameworks that ensure transparency, accountability, equity, and public trust. While many cities are developing internal AI policies and procurement standards, few have meaningfully engaged residents, particularly communities most impacted by inequitable systems, in shaping municipal AI governance.

This paper documents Austin, Texas’s “Community in the Loop” initiative, a multi-session, community-led engagement series facilitated by Measure in partnership with the City of Austin and the Austin AI Alliance. The initiative was further strengthened through collaboration with neighborhood organizations, digital equity leaders, libraries, and community-based partners who helped engage participants, host conversations, and ensure that historically underserved residents were represented throughout the engagement process. Using Measure’s CARE Model framework, the initiative gathered longitudinal qualitative data from residents across multiple convenings between 2025 and 2026. Residents responded to four core governance questions centered on community participation, safeguards, trust, and accountability in AI systems.



Across these engagements, 117 participants were engaged through three Community in the Loop sessions that the City helped lead and support. Additionally, Measure’s independent community outreach and engagement efforts reached 292 individuals who responded to the same prompt, for a total of 409 residents who contributed lived-experience insights that informed emerging recommendations for municipal AI governance. Findings suggest that residents consistently prioritized transparency, human oversight, participatory governance, environmental justice, privacy protections, and equitable access to AI literacy.

This paper argues that Austin’s approach represents one of the nation’s emerging models for community-led AI governance and offers a scalable blueprint for democratic municipal engagement around artificial intelligence.

Keywords: artificial intelligence, AI governance, municipal governance, participatory governance, community engagement, ethical AI, CARE Model, public accountability, digital equity

INTRODUCTION

Artificial intelligence, in its simplest form, refers to computer systems that identify patterns in data to make predictions, recommendations, or decisions that typically require human intelligence. And as natural the definition may seem, the complexity has resulted in AI rapidly reshaping how governments operate, make decisions, allocate resources, and interact with residents. Municipal governments increasingly use AI systems in areas such as permitting, public safety, emergency response, customer service, budgeting, transportation, and infrastructure management (National Institute of Standards and Technology [NIST], 2023).

However, as cities, counties, and states accelerate AI adoption, significant concerns remain regarding transparency, bias, accountability, privacy, environmental impact, and democratic oversight. Even more, the core ethical question remains: who benefits from AI and who is closest to the potential harm the tool creates? Despite growing national conversations about responsible AI, many governance efforts remain heavily technical, compliance-driven, or institutionally centralized. Public participation is often limited, inconsistent, or absent entirely. As a result, communities most likely to be impacted by AI systems, particularly historically excluded populations, frequently have little influence over how these technologies are designed, implemented, or governed. This connects to the Theory of Interrogative Reasoning (TIR), which argues that communities closest to potential harms should actively interrogate AI systems as part of governance (Styles et al., 2026). This lack of inclusion is a familiar challenge for cities that have grappled with community engagement strategies over the years. While the intentions of resident outreach are often present, the execution remains difficult, frequently resulting in a disconnect between technical validation and real-world outcomes.

In response to this gap, the City of Austin partnered with Measure and the Austin AI Alliance to launch “Community in the Loop,” a community-led engagement initiative designed to gather resident perspectives on municipal AI governance. The initiative aligns with Austin’s Open Government Partnership (OGP) commitment to co-create an AI Accountability Framework grounded in transparency, fairness, and public trust. This framework seeks to address the “black box” nature of many algorithmic systems by ensuring that lived experience informs the standards of municipal technology use.

By leveraging the Measure CARE Model, Community, Advocacy, Resilience, and Evidence, the initiative translates community insights into actionable accountability metrics. Unlike traditional public engagement efforts, Community in the Loop used a longitudinal, participatory governance approach. Residents were not treated as passive stakeholders but as co-governors whose insights help translate qualitative feedback into measurable accountability metrics. Through multiple convenings, nearly 400 residents engaged in structured interrogation of AI systems, prioritizing human oversight, environmental justice, and equitable access to AI literacy. This process utilized the Theory of Interrogative Reasoning (TIR), which emphasizes that reducing algorithmic unfairness requires a circle of "humans in the loop" who possess proximity to the problem and its impact.

The initiative also benefited from collaboration with community partners involved in the City's Digital Equity Community Listening Sessions (City of Austin, 2026), the Building Bridges Summit hosted by Joyce James and neighborhood engagement efforts. These organizations helped extend outreach into communities that are often underrepresented in technology policy conversations and reinforced the initiative's commitment to participatory governance.

This Measure Black Paper documents findings from these engagements and positions Austin's work as an emerging national model for democratic, community-centered AI governance. It argues that for innovation to be truly responsible, it must be co-created with the communities it serves, ensuring that technological advancement does not outpace equity or human dignity.

THE RISE OF MUNICIPAL AI GOVERNANCE

Cities across the United States are increasingly experimenting with AI technologies to improve efficiency, automate administrative processes, and modernize public services (Crawford, 2021). Municipal governments have adopted AI-powered systems in transportation management, emergency response, public communications, predictive analytics, and permitting operations (World Economic Forum, 2024). Austin itself has expanded the use of AI pilots and enterprise AI tools across city departments, including wildfire detection systems, smart safety tools, permitting support systems, and generative AI productivity tools (City of Austin, 2025).

However, scholars and practitioners warn that AI systems can reproduce existing inequalities if implemented without accountability mechanisms, transparency standards, or community oversight (Benjamin, 2019; Noble, 2018).

WHAT IS ETHICAL AND RESPONSIBLE AI?

Ethical and Responsible AI is the practice of designing, building, deploying, and governing artificial intelligence in ways that protect people, reduce harm, increase fairness, and strengthen human dignity, especially for communities historically excluded from technology decision-making. At its core, it asks a simple but powerful question: Who benefits from AI, who could be harmed by it, and who gets to shape it? Responsible AI is not just about the technology itself. It is about the values, power structures, data, policies, and human decisions behind the technology.

In a city like Austin, where AI is rapidly shaping transportation, housing, policing, public services, hiring, education, and digital access, Ethical and Responsible AI becomes a civic issue, not just a tech issue.

Here are examples grounded in what Responsible AI can look like at the local government and community level:



FAIRNESS

AI systems used by the city should not disproportionately harm Black, Brown, immigrant, disabled, unhoused, or low-income communities. (Benjamin, 2019; Noble, 2018)

Austin Example:

If AI is used to prioritize code enforcement complaints, predict infrastructure repairs, or identify “high-risk” neighborhoods, the system could unintentionally reinforce historical inequities if it relies on biased historical data.

For example:

- Neighborhoods that were historically over-policed may appear “higher risk.”
- Wealthier neighborhoods may receive faster city responses because they historically submit more digital service requests.

Responsible AI asks:

- Are all neighborhoods being served equitably?
- Is historical bias being automated?



TRANSPARENCY

Transparency builds public trust. Residents should know when AI is being used in city systems and how decisions are being made (NIST, 2023; White House, 2022).

Austin Example:

If the city uses AI chatbots for public services, automated traffic systems, permit processing, or emergency response triage, residents should understand:

- When they are interacting with AI
- What data is being collected
- How decisions are generated
- How to appeal or speak to a human



ACCOUNTABILITY

Government agencies and vendors should remain responsible for the outcomes of AI systems (NIST, 2023).

Austin Example:

If an AI-powered housing screening system incorrectly flags families for eviction risk or incorrectly deprioritizes someone for housing support, accountability cannot stop at:

“The algorithm made the decision.”

Responsible AI means:

- Human review processes exist
- Harm can be corrected
- Residents have pathways to challenge decisions



PRIVACY & DATA PROTECTION

Cities collect enormous amounts of public data. Responsible AI requires careful protection of community information (Executive Office of the President, 2023).

Austin Example:

Smart city technologies, including traffic cameras, public Wi-Fi systems, sensors, and predictive analytics, can improve services, but they also raise concerns about surveillance and data misuse.

Questions include:

- Who owns the data?
- How long is it stored?
- Could it be shared with third parties?
- Are immigrant or activist communities being disproportionately monitored?



INCLUSION & COMMUNITY GOVERNANCE

The communities most impacted by AI should help shape the rules around it (Costanza-Chock, 2020; Styles et al., 2026).

Austin Example:

Before implementing AI in areas like:

- public safety
- education
- housing
- workforce development
- transportation
- healthcare access

NOTE: The city could continue to host community convening (such as our Community In The Loop town halls) where residents, nonprofit leaders, youth, artists, technologists, and neighborhood organizations help define:

- what “harm” looks like
- what protections are needed
- what success should actually mean

This shifts AI governance from:

“experts deciding for communities” to: “communities governing alongside experts.”



ENVIRONMENTAL RESPONSIBILITY

AI infrastructure has a real environmental footprint that has caused real harm (World Economic Forum, 2024).

Austin Example:

As Austin grows into a major AI and data center hub, Responsible AI also means asking:

- How much water are data centers consuming during drought conditions?
- How much energy is powering AI systems?
- Who bears the environmental burden?
- Are historically underserved communities closest to industrial expansion?



HUMANS IN THE LOOP / COMMUNITY OVERSIGHT

AI should support human decision-making, not replace human judgment (Styles et al., 2026).

Austin Example:

If AI tools are used in:

- homelessness outreach
- mental health response
- healthcare triage
- workforce eligibility
- education pathways

...there must still be humans empowered to question, override, and interrogate the system when needed. This is especially important because communities closest to harm often see problems long before systems do.

Although the current federal administration has shifted its AI policy priorities, the White House Blueprint for an AI Bill of Rights (2022) continues to provide an influential framework for responsible AI governance. Similarly, the NIST AI Risk Management Framework (2023) emphasizes that trustworthy AI systems must be governed through ongoing monitoring, transparency, accountability, and stakeholder engagement. Importantly, many scholars argue that governance cannot rely solely on technical experts. Instead, AI governance must include communities closest to potential harms (Costanza-Chock, 2020).

PARTICIPATORY GOVERNANCE AND COMMUNITY-LED MODELS

Participatory governance models emphasize shared decision-making between institutions and residents. In the context of AI, participatory governance seeks to move communities from consultation to co-governance (Delacroix & Lawrence, 2019).

While several cities, including Amsterdam, San Jose, Barcelona, Boston, and New York City, have developed AI governance initiatives, few have established recurring community-led engagements specifically designed to generate longitudinal public governance data (UCLG, 2026).

Austin's Community in the Loop initiative is notable because it combines:

- recurring engagement,
- structured qualitative data collection,
- multilingual access,
- public education,
- and policy-oriented synthesis.

This positions the initiative as one of the more developed examples of community-centered municipal AI governance currently documented in the United States.

METHODOLOGY

Research Framework

This study employed a community-based participatory qualitative research design using Measure's Community, Advocacy, Resilience, and Evidence (CARE) Model as the primary engagement framework. The study also draws upon the Theory of Interrogative Reasoning (TIR), which positions communities most proximate to potential technological harms as essential contributors to the governance of artificial intelligence (Styles et al., 2026).

Participant Recruitment and Sampling

Participants were recruited through purposive and convenience sampling across multiple community engagement events held between November 2025 and April 2026. Recruitment intentionally prioritized participation from residents historically underrepresented in technology governance discussions, including Black, Brown, immigrant, low-income, disability, youth, nonprofit, and neighborhood communities. Although the sample is not statistically representative of Austin's population, the objective of this research was not population estimation but rather the identification of recurring governance themes grounded in lived experience.

Data Collection Procedures

Data were collected during five facilitated community convenings utilizing a standardized discussion protocol centered on four recurring governance questions aligned with the CARE Model domains:

- Community
- Advocacy
- Resilience
- Evidence

Maintaining identical core prompts across all convenings enabled longitudinal thematic comparison while allowing participants to elaborate on emerging concerns unique to their lived experiences. Facilitators documented qualitative responses through written artifacts, facilitator notes, participant worksheets, digital submissions, and post-session synthesis. Data triangulation was achieved by collecting perspectives across multiple community events, venues, and participant groups.

Trustworthiness was strengthened through investigator triangulation, repeated questioning across multiple convening, longitudinal data collection, and thematic comparison across participant groups.



DATA COLLECTION

Data were collected across five major convenings between 2025 and 2026:

Event	Approximate Attendance
Community in the Loop - November 19, 2025	~100
Building Bridges Event - December 12, 2025	~200
Virtual Community in the Loop - March 19, 2026	~40
Community in the Loop - Carver Library - April 7, 2026	~9
Digital Equity Community Listening Session - March 7, 2026	~60

TOTAL ESTIMATED PARTICIPANTS:

Approximately 409 total residents (117 engaged through City-supported Community in the Loop sessions; 292 engaged through Measure’s independent outreach efforts).

Participants responded to four recurring qualitative prompts:

- Community: Who should participate in AI decisions?
- Advocacy: What rights or safeguards do residents need?
- Resilience: How can AI strengthen trust?
- Evidence: What does accountability look like?

Limitations:

Because this study employed a community-based participatory research approach with purposive and convenience sampling, the findings are intended to illuminate community perspectives rather than provide statistically representative estimates. Future research should build upon these findings by examining how these themes emerge across broader populations and additional municipal contexts.

FINDINGS

1. Community Participation Must Be Broad and Inclusive

Residents consistently emphasized that AI governance should not be limited to technical experts or government officials.

Participants identified:

- underserved communities
- educators
- youth
- faith leaders
- business owners
- residents
- advocacy groups
- and individuals without financial incentives

as essential participants in governance conversations.

Many residents stressed that:

“The people who don’t yet realize how profoundly they will be impacted” must also be included.

Participants framed AI governance as a democratic issue rather than solely a technical issue.



2. Transparency and Accountability Were Central Concerns

Residents repeatedly identified transparency as foundational to public trust.

Participants wanted:

- public disclosure of AI systems
- explainability
- traceability
- public audits
- clear ownership structures

Questions raised included:

- Who built the system?
- What data was it trained on?
- Who benefits?
- Who could be harmed?
- Who is accountable when systems fail?

Residents emphasized that accountability must be:

- visible
- ongoing
- iterative
- understandable to the public

3. Human Oversight Remains Essential

Across all convenings, participants repeatedly emphasized the importance of maintaining “humans in the loop.”

Residents expressed concern about:

- over-automation
- AI hallucinations
- discriminatory outcomes
- and the removal of human judgment from public systems

Participants consistently advocated for:

- human review structures
- community feedback loops
- and public intervention mechanisms

4. Privacy and Safeguards Were Significant Public Concerns

Residents voiced strong concerns around:

- data privacy
- surveillance
- cybersecurity
- youth protections
- environmental impacts
- data ownership

Several participants specifically connected AI governance to:

- environmental justice
- water usage
- energy infrastructure
- data center accountability

Participants also raised concerns about AI's impact on youth mental health and misinformation exposure.



5. AI Literacy and Public Education Are Critical

Residents repeatedly emphasized the need for accessible AI literacy opportunities.

Participants stated that communities cannot participate meaningfully in governance without:

- understandable information
- accessible language
- public learning opportunities

This finding aligns with broader national conversations regarding digital equity and democratic participation in emerging technologies.

DISCUSSION

Austin's Community in the Loop initiative demonstrates that residents are not only willing to engage in AI governance discussions, they are capable of producing nuanced, actionable governance insights when provided with accessible information and meaningful participation structures.

Unlike many municipal AI governance initiatives that center institutional compliance or technical management, this initiative centered:

- lived experience
- participatory democracy
- public accountability

The recurring use of the same four governance questions across multiple convenings created an emerging longitudinal qualitative dataset capable of informing both local policy and broader governance models.

This work suggests that effective AI governance requires more than procurement standards or ethics statements. It requires:

ongoing democratic infrastructure.

The initiative also demonstrates the importance of community-centered facilitation models that:

- reduce technical barriers,
- support multilingual participation,
- and create psychologically safe engagement spaces.

Austin's work may therefore serve as a replicable blueprint for other municipalities seeking to operationalize ethical AI governance through public participation.

RECOMMENDATIONS FOR THE CITY OF AUSTIN

BASED ON RESIDENT FEEDBACK, OUR COMMUNITY RECOMMENDS THAT THE CITY OF AUSTIN

1. Establish a Permanent Community AI Advisory Structure

Create a recurring resident governance body inclusive of impacted communities.

2. Publish an AI Transparency Registry

Publicly disclose:

- AI systems in use
- data sources
- vendors
- intended uses
- risks
- and oversight mechanisms

3. Maintain Human Oversight Requirements

Ensure all high-impact municipal AI systems retain human review and intervention protocols.

4. Expand Community AI Literacy

Invest in:

- public workshops
- library partnerships
- youth education
- multilingual resources
- and neighborhood-based AI literacy initiatives



5. Develop Public Accountability Mechanisms


Implement:

- public audits
- resident complaint systems
- community review opportunities
- and iterative public reporting

6. Include Environmental Justice in AI Governance

Evaluate the environmental impacts of:

- data centers
- water consumption
- energy use
- and technological infrastructure expansion



CONCLUSION

As municipalities across the nation grapple with the rapid expansion of artificial intelligence, Austin's Community in the Loop initiative offers an important emerging model for democratic AI governance.

This work demonstrates that communities are not barriers to innovation; they are essential governance partners.

By centering lived experience, transparency, accessibility, and shared accountability, Austin has begun building one of the nation's more comprehensive community-led municipal AI governance efforts. The future of ethical AI governance will depend on whether cities are willing to build systems that treat residents as co-authors of the public technologies shaping their lives.

Austin's work suggests that meaningful AI accountability begins not with algorithms, but with people.



Transparency: Artificial intelligence was used as an editorial support tool during the preparation of this report. Human authors maintained oversight of all content, analysis, and final recommendations.



REFERENCES

Ada Lovelace Institute. (2023). Participatory AI governance.

Benjamin, R. (2019). Race after technology: Abolitionist tools for the new Jim code. Polity Press.

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101.

City of Austin. (2025). Artificial intelligence initiatives and Open Government Partnership AI accountability framework.

City of Austin. (2026). Digital Equity Community Listening Session & Resource Fair.

Costanza-Chock, S. (2020). Design justice: Community-led practices to build the worlds we need. MIT Press.

Crawford, K. (2021). Atlas of AI: Power, politics, and the planetary costs of artificial intelligence. Yale University Press.

Delacroix, S., & Lawrence, N. D. (2019). Bottom-up data trusts: Disturbing the "one size fits all" approach to data governance. *International Data Privacy Law*, 9(4), 236–252.

Executive Office of the President. (2023). Executive Order 14110: Safe, secure, and trustworthy development and use of artificial intelligence. The White House.

Gilgun, J. F. (2005). "Grab" and good science: Writing up the results of qualitative research. *Qualitative Health Research*, 15(2), 256–262.

National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). U.S. Department of Commerce.

Noble, S. U. (2018). Algorithms of oppression: How search engines reinforce racism. New York University Press.

Styles, J. M., Williams, Z., Teran, J., & Johnson, S. (2026, June 10). Interrogative reasoning and the problem with the "human in the loop." Just Tech. Social Science Research Council.

The White House, Office of Science and Technology Policy. (2022). Blueprint for an AI Bill of Rights: Making automated systems work for the American people.

United Cities and Local Governments. (2026). AI governance and participatory local government.

World Economic Forum. (2024). Artificial Intelligence Governance Alliance: Presidio recommendations on responsible generative AI