

AI Localism: Counter Imaginaries from City AI Governance

Bhavana Bheem¹, Samuel Gerstein², Karrie Karahalios¹, Madelyn Rose Sanfilippo²

¹Massachusetts Institute of Technology

²University of Illinois Urbana-Champaign

bbheem@mit.edu, samuel37@illinois.edu, kkarahal@mit.edu, madelyns@illinois.edu

Abstract

Local governments across the United States have quietly become frontline regulators of artificial intelligence (AI), advancing governance efforts that both fill gaps in federal oversight and contest dominant national and industry imaginaries. Drawing on 77 policies from 55 cities and counties, this study uses qualitative content analysis to analyze policy documents for their targeted AI domains and the values and priorities they articulate. Results identify three overarching counter-imaginaries—democratic oversight, rights-oriented governance, and place-based alignment—that structure how local governments define problems, justify intervention, and envision desirable AI-enabled futures. These findings highlight local AI governance as a critical site of policy innovation and democratic negotiation, with implications for how federal frameworks, preemption debates, and industry practices might better account for situated harms, community priorities, and the plural values that shape AI’s integration into public life.

Datasets — https://osf.io/82b6h/overview?view_only=5a52427b8b544d48a37480554b37030b

Introduction

Over the past decade, city and county governments across the United States have quietly become frontline regulators of artificial intelligence (Marcucci, Kalkar, and Verhulst 2022). Their proximity and insight to community needs and concerns have enabled them to craft AI policies that harness the technology’s benefits, while also exercising informed caution regarding potential harms. In parallel, state governments have also moved quickly; in 2025, all 50 US States, Puerto Rico, the Virgin Islands, and Washington, D.C. introduced AI legislation, with over 100 pieces of legislation adopted across 38 states (National Conference of State Legislatures 2025). In contrast, federal action around AI has lagged and been largely non-existent. During the 118th Congress, more than 150 AI-related bills were introduced into federal government, but none were passed into law (Brennan Center for Justice 2025), and as of early 2026, the 119th Congress has introduced 72 additional AI bills without passing any comprehensive, stand-alone AI legislation (American Action Forum 2026).

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In this absence of robust, unified federal regulation in the US, local governments have taken it upon themselves to develop ordinances and policies around how AI systems will be acquired, monitored, and deployed. Many have created internal AI-use guidelines for their employees, adopted bans on facial recognition in policing, and experimented with procurement rules, registries, and oversight bodies. AI localism is defined as “actions taken by local decision-makers to address the governance of AI within a city or community” (Marcucci et al., 2022). This concept builds on Katz and Nowak’s theory of “new localism,” which describes a broader twenty-first-century shift in which local municipalities, jurisdictions, and multisectoral networks step in to address social, economic, and environmental challenges amid federal inaction (Katz and Nowak 2018). Applied to AI governance, localism highlights how cities are not merely waiting for and adopting policies set by higher levels of government, but rather they are actively creating their own regulatory approaches in response to locally specific risks, values, and institutional capacities.

Despite the rapid expansion of local AI governance, only a small number of public repositories have aggregated these policies (ail 2026; Arnold et al. 2024), and there have been limited efforts to systematically analyze their substantive content. Thus, this study builds upon and extends these efforts, as maintaining up-to-date analyses of local AI policies is both analytically significant and theoretically generative.

AI policy documents can serve as meaningful sites for examining how city governments articulate priorities, define risks, and imagine desirable futures shaped by AI systems. Moreover, because local governments operate in a more immediate democratic context (Katz and Bradley 2013; Stoker 1996), these policies support analysis into how direct stakeholders interpret, negotiate, and institutionalize the role of AI in everyday socio-political life.

This paper treats local AI policy documents as expressions of desirable AI-enabled futures and the social orders that accompany them via a sociotechnical lens. The approach is grounded in the concept of *counter-imaginaries*: alternative visions that emerge outside dominant national or corporate narratives and reflect localized concerns, constraints, and values (Kazansky and Milan 2021). Existing scholarship uncovering sociotechnical imaginaries within AI governance has largely focused on evaluating national

policy and strategies, industry perspectives, or other institutionally mediated governance documents (Bakiner 2023; Bareis and Katzenbach 2022; Hjaltalin and Sigurdarson 2024). While this work has yielded important insights, it tends to overlook other less formalized, locally-situated, and often non-elite perspectives around AI development and integration (Mager and Katzenbach 2021). Local governance embodies imaginaries that differ, diverge, resist and subvert national and industry technology narratives which have thus far dominated both policy discourse and academic research on AI governance.

The present study attends to the plurality of counter-imaginaries embedded in local AI governance. By systematically coding local AI policy documents for both the AI domains they target and the values and priorities they articulate, this analysis seeks to surface how municipalities define problems, justify intervention, and imagine appropriate roles for AI in public life. Results illuminate both the kinds of AI systems local governments seek to regulate and also the normative commitments that underpin these regulatory efforts in the absence of comprehensive federal oversight. Moreover, the values found may directly contest the interests and priorities of industry actors and national government. Accordingly, this paper addresses the following research questions:

- **RQ1:** What domains and applications of artificial intelligence are targeted by local government policies?
- **RQ2:** What values, priorities, and motivations are reflected in local government policies regulating artificial intelligence?

Analysis identifies three overarching counter-imaginaries embedded in local AI governance. The first counter-imaginary foregrounds democratic oversight, envisioning AI systems as legible, accountable, and ultimately governable by the public institutions and communities they affect. This challenges the widespread assumption that since AI systems are inherently opaque “black boxes,” they are resistant to meaningful public understanding or control (Von Eschenbach 2021; Gryz and Rojszczak 2021). The second counter-imaginary advances a rights-oriented framework, grounding AI governance in established legal and ethical traditions, particularly civil, digital, and privacy rights, rather than treating AI as an unprecedented rupture that requires entirely new normative foundations. The third counter-imaginary reflects a commitment to place-based alignment, wherein AI systems are expected to conform to locally defined needs, values, and priorities, and are justified primarily through their concrete, tangible benefits to the communities in which they are deployed. This counter-imaginary pushes back against the idea that AI adoption is inevitable and must proceed without hesitation or friction. It also resists notions of generalized systems, such as artificial general intelligence, in favor of localized, context-specific AI tailored to particular municipal problems.

These AI governance approaches vary from, diverge from, and, at times, contest dominant national and industry imaginaries that often prioritize unhindered innovation, scale, and economic competitiveness. In addition to these findings, this

paper contributes a database of 77 local AI governance documents alongside a structured codebook designed to facilitate future comparative analyses of AI policy.

Background

AI Localism

Broadly speaking, local governments in the US have been sites of innovation across multiple policy domains, and AI governance is one prominent example of this shift. This pattern can be understood through what Katz and Nowak describe as “new localism” a 21st-century development in which cities increasingly take it upon themselves to address pressing social, economic, and technological challenges in the context of federal inaction (Katz and Nowak 2018). In this context, AI governance emerges as a particularly clear illustration of new localism in practice. While the federal government continues to debate comprehensive AI regulatory frameworks, cities and counties already confront immediate questions related to surveillance, procurement, public service delivery, and civil rights protections (Marcucci, Kalkar, and Verhulst 2022; Verhulst and Sloane 2020).

AI localism captures the ways in which local actors actively shape how AI is understood, regulated, and contested within their communities. Marcucci, Kalkar, and Verhulst (2022) define it as “actions taken by local decision-makers to address the governance of AI within a city or community.” Importantly, AI localism is not confined to formal government or legal regulation. It also includes public institutions, such as school boards and libraries. This study, however, focuses specifically on AI policies and ordinances enacted by local governments, enabling systematic comparison across jurisdictions and treating these formal policy documents as artifacts that reveal local priorities, risks, and values surrounding AI. Throughout the paper, we use “city” as shorthand for the range of local government types represented in our corpus which includes cities, counties, and boroughs.

While AI localism highlights the growing role of local actors, it operates at a sub-national and sub-state level and has bounded autonomy. Evans, Marsh, and Stoker (2013) refers to localism as: “the devolution of power and/or functions and/or resources away from central control and towards front-line managers, local democratic structures, local institutions and local communities, within an agreed framework of minimum standards” (p. 405). Local governance functions within shared normative and institutional constraints that establish baseline expectations. And within these expectations, local actors retain discretion to design policies that respond to local conditions. Thus, local governance is always relevant to state and federal government because it operates within shared normative and institutional constraints. It is therefore a valuable test case to understand how contextual needs and values can be incorporated into broader policy frameworks, and how federal policymaking can better account for on-the-ground variation.

Existing work on AI localism has often drawn on adjacent literatures, particularly privacy localism and smart city governance. Rubinstein (2018), through studying city-level privacy policies, has shown how cities adopt different reg-

ulatory approaches to data collection and surveillance in response to local political pressures. Similarly, smart city scholarship has examined how municipalities experiment with digital technologies to reshape urban services and infrastructure (Angelidou 2014). More recent discussions of “urban AI” focus specifically on the integration of AI systems into city processes, such as transportation, housing, and public safety (Luusua et al. 2023). While these literatures are closely related, this study departs from them in two critical ways. First, the analysis is not limited to large urban cities, instead it examines local governance across a broader range of municipality sizes. Second, rather than focusing on the operational deployment of AI in city services, the analysis centers on governance documents that shape a wide range of AI uses, including informal and ad hoc applications.

Studying Imaginaries

Local governments’ AI policies, ordinances, and executive orders can be approached as both regulatory documents and a form of discourse. On the surface, these governance documents define the acceptable and restricted uses of AI technologies, and control how these technologies are integrated into city services and civic life (Marcucci, Kalkar, and Verhulst 2022). At the same time that they establish what one might call “rules,” these documents also articulate the city’s values, priorities, and visions with respect to AI technologies. From this perspective, governance and government applications of AI and associated technologies form a complex and reciprocal socio-technical system, in the co-constructive social informatics sense (Meyer 2014; Meyer et al. 2019). To foreground these orientations and frame content analysis of documentation around these relationships, this paper draws on the framework of counter-imaginaries.

Imaginaries are broadly understood as collectively held visions and background understandings that shape how society perceives what is possible, desirable, and legitimate. Taylor (2004) constructs social imaginaries as shared ways of imagining social life, including the norms, expectations, and practices through which social order is made sense of. Imaginaries are tangible and inherently futuristic. They shape how visions of the future are translated into governance and social arrangements today.

Sociotechnical imaginaries extend this broader framework of imaginaries into the study of science, technology, and governance, via a science and technology studies lens. They are defined as: “collectively held, institutionally stabilized, and publicly performed visions of desirable futures, animated by shared understandings of forms of social life and social order attainable through, and supportive of, advances in science and technology” (Jasanoff et al. 2015, p. 4). The framework has been widely used in policy and governance research to examine how future-oriented visions shape, and are shaped by, technological trajectories and institutional decision-making (Hendriks, Karhunen, and Delvenne 2025; Kuchler and Stigson 2024; Arboledas-Lérida 2025; Cave, Dihal, and Dillon 2020).

However, applications of sociotechnical imaginaries to AI governance have largely focused on national policy and institutional texts, foregrounding dominant state and industry

visions while giving less attention to non-state and localized actors. For example, Bareis and Katzenbach (2022) analyze national AI policy documents from four countries and show how these strategies consistently frame AI as both inevitable and disruptive. Similarly, Kao (2024) through an analysis of Australian AI governance efforts, demonstrates how economic priorities often supersede concerns surrounding safety and responsibility. While Oldenburg and Papishev (2025) examine a broader range of sociotechnical imaginaries, their analysis similarly relied on formally published texts produced within institutional and corporate actors.

Critique argues that the sociotechnical imaginaries framework tends to foreground dominant or elite visions. Much of the literature has examined imaginaries in highly institutionalized settings, expert discourse, or pre-selected social groups, thereby implicitly treating these perspectives as representative of society as a whole (Kuchler and Stigson 2024; Mager and Katzenbach 2021; Smallman 2020; Tidwell and Tidwell 2018). Even studies that explicitly seek to account for plural or non-dominant imaginaries remain constrained by the settings in which these imaginaries are accessed. For example, while Bakiner (2023) seeks to capture pluralistic imaginaries through civil society responses to the European Union’s AI Act, the analysis remained situated within a institutionalized setting, where imaginaries are already filtered through dominant governance frameworks. This is reinforced by Smallman (2020) whose interview study of UK policymakers shows that perspectives outside dominant elite sociotechnical imaginaries are either excluded from policy outcomes or rearticulated to fit prevailing frameworks.

In an effort to begin to address these limitations around power, consistency, and constructivism, a related strand of literature on counter-imaginaries in data-driven, surveillance, and AI contexts has primarily focused on practices of resistance and activism among civil society actors and groups. Kazansky and Milan (2021) conceptualize counter-imaginaries in their analysis through the open-source software projects engineered by civil society and grassroots activists resisting processes of datafication. Other scholarship, while not explicitly using the term counter-imaginaries, has similarly identified alternative or oppositional imaginaries within data activism and open data initiatives. This work argues that open data, participation, and democratic control of data function to challenge and subvert the power of both state institutions and corporate actors (Baack 2015; Lehtiniemi and Ruckenstein 2018; Milan 2024). Importantly, counter-imaginaries are typically located outside dominant government and industry frameworks, emerging instead within smaller, more participatory collectives (Gutierrez, Nguyen, and van Es 2025; Laba 2025; Mager and Katzenbach 2021; van Es, van der Weijden, and Bakker 2023). This literature conceptualizes counter-imaginaries around technology as practices that: (1) resist dominant imaginaries; (2) constitute forms of participatory and democratic engagement; and (3) are oriented toward addressing the needs of those most affected.

While counter-imaginaries have most often been examined through practices of activism and resistance, this framework is also well suited to the study of AI localism. Lo-

cal governance constitutes a particularly important site for counter-imaginaries because it operates through comparatively proximate and participatory democratic processes, in which residents, advocacy groups, and community organizations can meaningfully influence policy outcomes (Katz and Bradley 2013; Stoker 1996). Local AI policymaking frequently emerges in direct response to concrete harms, risks, and inequities experienced by communities, positioning residents as both participants in policymaking and those most directly affected by its consequences. As a result, local AI policies can serve as sites where situated priorities, concerns, and values are articulated, often diverging from dominant national or industry imaginaries.

For this reason, this paper adopts counter-imaginaries as its primary analytical framework. Rather than assuming convergence around a dominant sociotechnical vision, it foregrounds how subnational governance sites articulate alternative, resistant, and community-oriented understandings of AI. In the US context, local AI policy does not emerge in a vacuum. Federal inaction does not signal the absence of a dominant sociotechnical imaginary around AI; rather, such an imaginary persists and continues to structure technological development and governance priorities. Local imaginaries therefore do not constitute dominant alternatives but instead function as subnational articulations that may contest, supplement, or push back against prevailing national trajectories (Verhulst and Sloane 2020). Through this frame, we understand that local AI governance is an active site of negotiation over technological futures rather than a derivative reflection of higher-level policy regimes.

Method

Corpus Construction

Relevant AI policies were identified through a systematic review of publicly available policies, directives, and executive orders enacted by local governments before October 1, 2025. The initial corpus was compiled from GovLab’s AI Localism repository (ail 2026), Local Progress’ Algorithmic Rental Price Fixing report (rea 2025), and the Emerging Technology Observatory AI Governance and Regulatory Archive (AGORA) (Arnold et al. 2024).

Since these repositories were incomplete, Google searches using Boolean operators were conducted to identify additional local AI policies.¹ For New York and Los Angeles, more targeted search queries were used, as these large municipal regions produced disproportionately high volumes of irrelevant results under general search terms.² These searches were conducted semi-weekly from mid-September 2025 to mid-October 2025, with the first ten pages of results systematically screened during each iteration. Links were

¹The primary Google boolean search query was: “(‘City’ OR ‘County’) (‘just passed’ OR ‘adopted’ OR ‘approved’ OR ‘enacted’) (‘AI’ OR ‘artificial intelligence’ OR ‘facial recognition’ OR ‘biometric surveillance’) (‘policy’ OR ‘guideline’ OR ‘ordinance’ OR ‘ban’ OR ‘regulation’ OR ‘law’).”

²The following targeted searches were conducted: “New York City technology ordinance,” “New York City AI ordinance,” “New York City borough AI,” and “Los Angeles County AI policy.”

Domain	Count
General AI	24
Generative AI	18
Facial Recognition/Surveillance Technologies	18
Automated Decision Making	16
Data Privacy	1

Table 1: Total Counts of Domains

followed through government websites, news articles, and blog posts to locate the policy documents. The query was repeated over multiple weeks because Google results can vary by date, time, location, user profile, and various other algorithmic factors (Hannak et al. 2013). The collection process concluded when subsequent searches were no longer yielding new ordinances or policies.

The resulting dataset includes local governments’ AI ordinances, administrative guidelines, executive orders, departmental policies, directives, roadmaps, and laws. While AI localism extends beyond policies enacted by government offices or officials (Evans, Marsh, and Stoker 2013; Marcucci, Kalkar, and Verhulst 2022), other types of decision-making entities, such as school boards, were excluded in order to keep the qualitative analysis manageable and to establish a clear analytical boundary.

The final corpus comprised 77 documents representing 55 local governments across 25 US states, Washington, D.C., and Guam. These sources were categorized according to the primary object of regulation in the following domains: AI (General), Generative AI (including Deepfakes), Facial Recognition Technology (FRT) / Surveillance Technologies, Automatic Decision Making, and Data Privacy.³ A repository of the documents is provided with publication to enable further analyses and research.

Qualitative Content Analysis

Qualitative content analysis with inductive coding was conducted on the policy documents (Zhang, Wildemuth et al. 2009). An initial subsample of 20 documents was coded by the first author at the sentence level, allowing codes to emerge from the data rather than being defined a priori. Guided by the research questions, the coding focused on identifying definitions, priorities, and values articulated within the policies. This sentence-level approach enabled each statement to be interpreted in terms of its functional role. Definitions appearing in policy documents were also coded to capture how technologies and related concepts were defined and characterized, helping delineate the regulatory scope of each policy. Identifying which technologies are explicitly named is central to the first research question.

Once a preliminary code list and accompanying code definitions had been developed, a second coder independently coded six documents (approximately 8% of the corpus) using this codebook. Through approximately six calibration meetings, the coders reorganized codes, refined

³We included only city data privacy regulations that explicitly address AI applications.

Theme	Explanation	Sub-Codes
Responsible and Ethical Usage	Promotion of responsible, ethical, and safe use of AI within city operations.	Practice: Encouraging ethical, responsible, and diligent AI use. Includes expectations for sound judgment, appropriate use, and avoiding high-stakes decisions. Risk: Requirements to recognize and mitigate the limitations, biases, or potential harms of AI. Includes duties to avoid harms, document risks, and report concerns. Betterment: Places importance on AI being leveraged to tangibly improve government services, societal well-being, or community outcomes. Education: Requirements or programming for training and educational resources to ensure employees understand how to use AI responsibly and effectively.
Privacy and Security	Protection of the privacy and security of individuals and systems.	Practice: Protecting the privacy and security of individuals and systems. Includes not entering sensitive information into AI models and safeguarding sensitive data. Unauthorized Access: Restricts AI use to gain unauthorized access to devices, systems, or information, and prohibits sharing data with unauthorized parties. Data should only be shared on a need-to-know and right-to-know basis.
Public Rights	Preserving or allocating individual, public, or community rights in the face of city AI use.	Consent: Guarantees the public's right to know when, how, and why AI is used; the right to explanations; the ability to opt out; and the ability to contest or audit AI systems. Equality: Affirms equal treatment and ensures AI does not harm or disadvantage marginalized or historically disenfranchised groups. Privacy: Protection from unwanted surveillance, biometric tracking, data exploitation, and other privacy intrusions.
Contextual Values	Situating AI policy within broader governance and local community contexts.	Alignment: Aligns policy goals with community or public values, encouraging stakeholder participation in policy development. Cross-Jurisdictional: References higher-level government rules or allows sub-departments to establish additional rules.
Accountability	Assigning responsibility for AI outputs and establishing remediation processes.	Output Review: Assigns responsibility for AI outputs, mandates human review, limits verbatim use, or requires human-in-the-loop oversight. Intervention: Establishes processes to address harmful or incorrect AI outputs, including mechanisms to contest, reverse, or rectify AI-generated decisions.
Technological Determinism	Reflecting assumptions about AI's role and trajectory.	Industry: Language promoting private sector values and aligning policy with industry, capital, or third-party AI experts rather than community needs. Adoption: References promoting or advancing AI development, adoption, and innovation within the city. Frames AI as beneficial or essential to city services and as a technology that must be harnessed.

Table 2: Themes with Multiple Sub-Codes

code definitions, and resolved discrepancies until consensus was reached. The final codebook consisted of 29 codes, and intercoder reliability, assessed using Krippendorff's alpha, ranged from the minimum threshold of 0.80, indicating excellent agreement, to 1.0, indicating perfect agreement (Krippendorff 2018).

Using the validated codebook, the first author coded the remaining documents. The resulting coded segments were subsequently grouped into higher-level themes to both organize the presentation of results and support the interpretation and delineation of codes. These themes are presented in the next section.

Results

Given that the coding was structured around the research questions, the resulting codes map directly onto them. In total, 29 validated codes were organized into 12 broader themes. Note that "Definition" functions more as a structural one, capturing the explicit boundaries governments draw around regulated technologies.

Table 2 presents the themes that contain multiple sub-codes, while Table 3 presents themes consisting of a single code. Table 4 describes the sub-codes included under the structural theme of "Definitions." These tables constitute the complete codebook, which can be used to guide the coding of additional AI governance documents.

Theme	Explanation
Oversight	Establishes or delegates an oversight body or review process to approve AI usage, monitor compliance, and audit AI tools.
Disclosure	Requirements to document, report, or publicly disclose AI systems, including where AI is used, how it functions, its capabilities, and its limitations.
Compliance	Emphasizes adherence to policy, including requirements to follow rules, as well as the legal or disciplinary consequences for non-compliance.
Voluntary Usage	Explicitly states that AI adoption is not required, or discourages use in certain contexts.
Scope	Defines who and what the policy covers, including acceptable uses, approved tools, exceptions, and provisions for periodic review.

Table 3: Themes with a Single Code

Sub-Code	Explanation
Human	Defines AI in terms of human intelligence, its ability to mimic human behavior, or its performance of tasks typically done by humans.
Produce Content	Defines AI in terms of generating or producing content.
Automated Decision	Defines automated decision making as a process in which decisions are generated using algorithms, machine learning, predictive models, or other forms of data analysis.
Data	Defines that AI models use, rely on, or process data.
Sensitive Information and Privacy	Defines sensitive or personally identifiable information and/or scopes privacy practices.
Equity	Defines fairness, equity, equality, discrimination, and unequal treatment.
Smart City	Defines the use of technology or AI to manage city operations or public services.
Stewardship	Defines the management, oversight, or handling of data.
Surveillance	Defines monitoring, tracking, or surveillance activities. Includes facial recognition, biometric surveillance, social scoring, emotion analysis, etc.

Table 4: Sub-Codes Under the Structural Theme “Definitions”

City Tier Categorization

Using the 77 coded documents, quantitative methods were applied to visualize how codes are distributed across the dataset. For each city, the population size, sourced from US Census data, is defined as the number of residents governed by the authority that produced the policy document.

Governments were categorized into tiers using an approximate scheme loosely informed by United Nations classifications for lower population ranges (United Nations Statistical Commission 2020), with higher tiers defined through study-specific judgment. This classification is intentionally approximate because the dataset includes cities, counties, and boroughs with differing administrative boundaries and population definitions. Although no formal statistical analyses were conducted, these tiers provide descriptive context for differences in code application across governments of varying population sizes.

Rural areas are defined as those with fewer than 5,000 residents and towns as those with populations between 5,000 and 50,000. Because only one document, from Haines Borough, Alaska, fell into the rural category, the rural and town categories were collapsed into a single “Town” tier. The remaining tiers are defined as follows: Large Town (50,000–300,000), Small City (300,000–500,000), Medium City (500,000–1,000,000), Large City (1,000,000–3,000,000), and Megacity (3,000,000+). In total, the dataset comprises 7 documents in the Town tier, 20 in the Large Town tier, 7 in the

Small City tier, 20 in the Medium City tier, 12 in the Large City tier, and 11 in the Megacity tier.

Thematic Differences Across Tiers

Thematic differences across these tiers are visualized using a heatmap which shows the percentage share of codes associated with each theme in each tier category. The “Definition” theme is presented in a separate heatmap because it serves as a structural category rather than reflecting policy motivations or values. Figure 1 presents the resulting heatmap for the non-definition themes, illustrating how the relative emphasis on different policy themes varies across government size tiers. Similarly, Figure 2 illustrates the distribution of the definition sub-codes across different government sizes.

The heatmap of non-definition codes reveals several patterns across city sizes. Overall, the responsible and ethical usage theme is the most popular for five out of the six city tiers, highlighting the widespread emphasis on ethical AI practices across all sizes of municipalities. Common statements coded under this theme included language explicitly framing the purpose of the policy as guiding the responsible and ethical use of AI. Other examples included prohibitions on employees relying on AI in high-stakes decisions and expectations that AI systems demonstrate tangible benefits to the community. Policies also set requirements for employee AI training and included language encouraging awareness of AI-related risks, such as copyright issues, hallucinations, bias, misinformation, and privacy harms.

Additionally, the contextual alignment and privacy and se-




Figure 1: Heatmap showing the proportional distribution of non-definition themes across municipality size tiers. Values represent the percentage share of coded segments within each city tier associated with a given theme.




Figure 2: Heatmap showing the proportional distribution of definition sub-codes across municipality size tiers. Values represent the percentage share of coded segments within each city tier associated with a sub-code.

curity themes are also consistently present across the board, however to a lesser extent. Contextual alignment is typically expressed through language aligning AI policy with the needs and values of local communities, as well as framing the policy as a continuation of the city’s responsibil-

ity to its residents. For example, the Generative AI policy from Hutto, Texas states: “This policy reflects the Texas AI Code of Ethics as required under H.B. 3512 and S.B. 1964, adopted by the Texas Department of Information Resources (DIR), and aligns with our responsibilities to residents.” This

frames the generative AI policy as a continuation of past commitments, and an alignment with the agreed upon standards of the community. Similarly, privacy and security provisions emphasize protection of resident data, including restrictions on sensitive inputs and unauthorized access. For example, the San Jose AI Policy states that “Privacy is preserved in all AI systems by safeguarding personally identifiable information (PII) and sensitive data from unauthorized access, disclosure, and manipulation in accordance with the City Council’s Digital Privacy Policy.” This framing embeds AI use within existing data protection frameworks.

The disclosure theme is particularly prominent in four of the middle tiers: Large Town, Small City, Medium City, and Large City. Sentences coded under this theme typically include transparency and reporting requirements regarding AI use. For example, Alameda County’s Generative AI Policy requires employees to “disclose when and how GenAI was used in your output.” Other policies extend these expectations to documentation of review processes; for instance, Woodburn’s AI policy states that “all attributions should include the name of the AI system used plus a human-in-the-loop assertion (which should include the department or group who reviewed/edited the content).” Furthermore, these provisions for transparency are often paired with audit and review processes, and in some cases, mechanisms for community input. In Berkeley’s Surveillance Policy, for example, a “Surveillance Acquisition Report” is defined as a “publicly-released written report produced prior to acquisition or to proposed permanent use after use in Exigent Circumstances pursuant to Section 2.99.040.” While this report is intended to support transparency, it also enables formal review prior to deployment. The policy further specifies that, through this process, “decisions regarding whether and how Surveillance Technologies should be funded, acquired, or used should be governed by the City Council as the elected representatives of the City.” This embeds disclosure within a broader structure of democratic authorization, linking transparency directly to institutional oversight and decision-making authority.

Concerns over individual rights are reflected in the public rights theme, which is emphasized in the Town to Medium City tiers as well as the Megacity tier. Codes in this category include language on the right to know about AI use and explicit consent provisions, as well as references to digital, civil, and human rights frameworks that ground AI governance in established rights-based traditions. For example, New York City’s Local Law on automated decision systems requires the “development and implementation of a process for making information publicly available that, for each agency automated decision system, will allow the public to meaningfully assess how such system functions and is used by the city, including making technical information about such system publicly available where appropriate.” Similarly, San Francisco’s Surveillance Policy invokes civil rights and civil liberties frameworks, stating that “when-ever possible, decisions relating to surveillance technology should occur with strong consideration given to the impact such technologies may have on civil rights and civil liberties, including those rights guaranteed by the First, Fourth,

and Fourteenth Amendments to the United States Constitution as well as Sections 1, 2, and 13 of Article I of the California Constitution.” These provisions situate AI governance within constitutional and rights-based legal traditions, framing public oversight and individual protections as central to technological decision-making.

A different pattern emerges in the technological determinism theme, which appears primarily in Large City and Megacity policies. This suggests greater alignment with industry imaginaries in larger urban centers, likely reflecting the higher concentration of technology firms and industry stakeholders in cities such as New York City, Los Angeles, and King County. In these urban environments, AI is more often framed in terms of system-wide transformation and inevitability, consistent with industry-oriented imaginaries. In contrast, smaller municipalities place less emphasis on technological determinism, reflecting more cautious and context-dependent approaches to AI governance. This pattern points to a paradigmatic divergence across city sizes, with larger cities exhibiting greater proximity to dominant innovation-driven imaginaries, and smaller cities more strongly foregrounding situated governance concerns. This likely reflects differences in administrative responsibilities and size, as well as exposure to technology-sector ecosystems across urban scales.

The heatmap of definitions reveals systematic differences by city size, as well as areas of shared concern. Larger cities tend to define AI in terms of models that learn from data, whereas smaller cities describe them as algorithms performing human-like tasks, reflecting variation in conceptual framing. Surveillance appears more prominently in the Large Town, Small City, and Medium City tiers, indicating heightened concern about monitoring and oversight in these contexts. Generative AI—defined as AI that produces content—is a concern across five out of the six tiers, suggesting widespread attention to emerging AI capabilities. Automated decision-making is most concentrated in the Megacity tier, likely reflecting the greater need for automated processes to manage complex and high-volume decisions. These patterns suggest that while larger cities emphasize technological sophistication and operational scale, smaller and medium cities are particularly attentive to oversight and the social implications of AI, with common values of transparency, ethical responsibility, public rights, privacy and security cutting across all municipality sizes.

Discussion

Addressing RQ1, the analysis shows that local AI governance targets a consistent set of technological domains, most prominently general AI, generative AI, surveillance technologies, and automated decision-making systems. Across city tiers, the heatmap of definitions also reveals variation in how AI is conceptualized and delimited as a regulatory object, particularly in relation to surveillance, automation, and generative systems. These patterns suggest that while certain domains are widely shared, the framing of AI differs across municipal contexts and scales.

These patterns establish both the substantive scope of AI governance across municipalities and the ways in which dif-

ferent jurisdictions frame and delineate AI as a regulatory object. This descriptive foundation provides the basis for the interpretive analysis that follows, in which RQ2 is addressed through an examination of the counter-imaginaries embedded within these policy frameworks.

Democratic Oversight

Many local policies foreground democratic oversight as a key mechanism within AI governance. Across the corpus, formalized approval processes, procedural requirements for review prior to adoption, and the establishment of oversight bodies were identified. These mechanisms are designed to ensure institutional accountability and public visibility into how AI systems are procured, deployed, and monitored. This orientation also reflects an implicit rejection of technological determinism. Rather than treating AI adoption as inevitable or governance as secondary to innovation, these policies position AI systems as subject to deliberate authorization, evaluation, and constraint.

The disclosure theme recurred across AI policies in all city tiers, emerging as the most prevalent theme in the Small City tier and the second most prevalent theme in the Large Town tier. Local policies frequently specified a requirement for detailed reporting on the purpose, scope, and functionality of AI systems, including surveillance applications, automated decision-making tools, and usage of generative AI. In cases involving surveillance technologies, disclosure typically extended to specifying deployment locations, operational practices, financial expenditures, and vendor relationships. Some localities mandated public-facing reports, while others required internal audits or regular updates to city councils or ethics committees.

Alongside disclosure requirements, many policies emphasize accountability by specifying that harms arising from AI use are attributable to responsible individuals or departments rather than the technology itself. This framing prevents harm from being obscured onto the AI systems and reinforces identifiable lines of responsibility. This is reinforced through oversight provisions, including designated review bodies, compliance monitoring structures, and defined consequences for non-compliance. In addition, many policies establish avenues for contestation, including mechanisms for residents to challenge algorithmic decisions or request review of specific AI applications. These provisions suggest a broader commitment to participatory governance within local AI policymaking.

With disclosure, accountability, and contestation, these policies do not simply accommodate the acceleration of AI without friction; rather, they establish conditions under which AI can be contested and must be contestable through democratic processes. This orientation contrasts with dominant sociotechnical imaginaries, where deregulatory approaches often frame AI development in terms of acceleration and innovation (Oldenburg and Papsyshev 2025). Moreover, by explicitly incorporating mechanisms of contestation, these policies establish AI as inherently challengeable rather than something to be adopted uncritically (Campolo and Crawford 2020). Overall, municipalities position democratic oversight as both a procedural safeguard and an ethical

commitment that mediates power among government institutions, technology vendors, and affected communities.

Rights-Oriented Framework

Many local AI policies begin by asserting a commitment to ethical, responsible, and safe AI use. These statements framed the purpose of policy intervention as guiding adoption and deployment of AI in ways that protect public rights, ensure equity, maintain institutional integrity and respect privacy. In practice, such policies translated these goals into concrete expectations, including guidelines and training for employees on responsible conduct, emphasis on data stewardship, and guardrails to ensure ethical decision-making.

Beyond procedural ethics, these commitments were often rooted in broader rights-based frameworks. Many policies invoked digital, civil, human, and privacy rights as foundational principles. They frame responsible AI use through the lens of equality, fairness, and nondiscrimination. This rights-oriented stance moved ethical governance from abstract aspiration toward enforceable obligation, linking local policy decisions to long-standing democratic ideals. This approach diverges from national AI policy, where, as Bareis and Katzenbach (2022) note, AI is often framed as a historical rupture—a technology unlike anything seen before. Such framing can distance governance from precedent, positioning AI as exceptional. In contrast, local AI governance treats AI as a technology with measurable risks, governed through existing ethical and rights-based frameworks. AI is thus framed not as exceptional, but as a familiar governance challenge addressed through institutionalized values.

A recurring emphasis in these documents is that responsible AI use should demonstrably improve, rather than degrade, the quality of civic life. AI is framed as a tool to enhance city services, strengthen public trust, and support residents' well-being, while avoiding harm or exacerbating inequality. This establishes a standard of conditional deployment, where technologies are adopted only when they clearly produce public benefit. This differs from dominant sociotechnical imaginaries, in which AI is often uncritically accepted and innovation is prioritized as inherently beneficial, with limited space for contestation (Oldenburg and Papsyshev 2025; Campolo and Crawford 2020).

Local policies, which frequently reflect this rights-based orientation, require that benefits be demonstrated prior to adoption, reflecting a precautionary, values-driven stance that ties technology use directly to the public good. This finding is consistent with Wernick (2024), who, through analysis of GDPR and EU regulatory frameworks, shows that while existing regimes often treat innovation as self-justifying, rights-based approaches more effectively embed accountability, oversight, and fundamental rights protections into technological design and deployment.

These rights-oriented approaches also carry particular weight in addressing the disproportionate risks faced by marginalized communities. Some municipalities explicitly identify potential biases and unequal impacts as central governance concerns and articulate protective measures to mitigate and redress these harms. By linking ethical and responsible AI use to both civic benefit and civil rights, local gov-

ernments position themselves as stewards of justice, equity, and accountability in the digital age.

These commitments ground AI governance in established legal and democratic frameworks, framing AI not as an exceptional or self-regulating technology, but as one that must be governed through existing principles of rights, responsibility, and public accountability.

Place-Based Alignment

Local policies often emphasized the importance of ensuring that AI governance aligns with each city's distinct values, priorities, and identity. These municipalities framed their AI regulation as a continuation of their broader civic commitments. Policies frequently referenced community values and development goals, seeking to integrate emerging technologies in ways that reflect local needs and aspirations. This contrasts with dominant imaginaries that have been criticized for framing AI as a general-purpose, solutionist technology applicable across diverse social and institutional contexts (Mager and Katzenbach 2021; Oldenburg and Papyshv 2025).

Several municipalities describe consulting residents, advocacy groups, or advisory boards in shaping AI-related rules and guidelines. This public input played an important role in shaping the policies themselves, aligning governance approaches with the contextual values and concerns of the communities affected by these technologies. This dynamic contrasts with findings from studies of national policymaking. For example, Smallman (2020) notes that in national-level policymaking, public input that does not align with dominant national priorities is often ignored or reframed to fit prevailing imaginaries. In contrast, much local AI policymaking incorporates community needs and local priorities from the outset, with policies continuing to evolve in response to local concerns.

Industry perspectives and motivations sometimes appeared as well, particularly in policies from large metropolitan areas like Los Angeles, New York, and Miami. This likely reflects both the presence of technology companies in those regions and the scale of their potential applications for urban management and service delivery. In many cases, especially with smaller cities, policy documents typically balanced optimism about AI's potential with caution regarding its risks. This cautious optimism reflected a broader hesitancy toward uncritical or large-scale adoption and emphasized that technological innovation should not override the fundamental values or social priorities of the cities and its residents.

Finally, many policies situate themselves within a broader regulatory ecosystem, referencing existing municipal, state, or federal frameworks that inform or constrain their approach. This positioning underscores that while AI governance is locally articulated, it remains embedded within multi-level institutional structures that shape its scope and implementation.

Limitations

The corpus is limited to policy documents that were publicly available during the study period—a common limitation in

document analysis (Bowen 2009). As a result, it does not represent an exhaustive inventory of local AI governance. Additionally, the corpus may be biased toward jurisdictions with greater digital visibility and administrative capacity, as document discoverability depends on factors such as search engine indexing and search terminology. Jurisdictions using alternative terminology for AI technologies or maintaining policies that were not publicly accessible may therefore be underrepresented. Thus, the geographic distribution of the corpus reflects where eligible policies were identified rather than the absence of AI governance in other jurisdictions.

Conclusion

Local AI governance is a site where distinct, localized counter-imaginaries of AI are articulated and institutionalized. Drawing on 77 policies from 55 local governments, the analysis identifies how municipalities target specific AI domains while advancing counter-imaginaries centered on democratic oversight, rights-oriented governance, and place-based alignment. Thus, it contributes to AI localism scholarship by clarifying how local governments define the scope of AI governance and the values motivating regulatory action.

These local imaginaries stand in contrast to dominant national and industry narratives that emphasize innovation, scale, accelerated deployment and competitiveness. Instead, local governments frame AI as a technology that must remain accountable to affected communities, constrained by established rights frameworks, and justified through demonstrable public benefit and context-specific purpose. This distinction is particularly important within the United States' federal system, where state and local governments have historically functioned as sites of regulatory experimentation and policy learning (Oates 1999).

Recent federal developments in AI policy, however, suggest a shift away from this traditional dynamic. The Trump administration's 2026 national AI framework places strong emphasis on economic growth and global technological leadership, while also signaling support for forms of federal preemption (White House 2026). Although national governments inevitably operate with broader priorities, including military and economic concerns that shape their governance approaches differently from local governments, this orientation risks narrowing the institutional space available for subnational experimentation and localized governance. The findings of this study suggest that local efforts offer valuable insight into how communities understand the acceptable integration of AI into civic life, particularly in relation to accountability, public participation, rights protections, and social impact.

As debates over federal AI regulation continue, preserving space for local variation and experimentation remains critical. It is within these local governance settings that situated harms, community priorities, and context-specific concerns are most clearly surfaced. Local AI governance therefore provides not only an alternative regulatory approach, but also a distinct vision of how technological systems should be governed in relation to democratic accountability and the public good.

Acknowledgments

We thank Inyoung Jang for serving as the second coder in the codebook validation and inter-rater reliability assessment. We also thank the anonymous reviewers for their thoughtful feedback, which helped strengthen the paper.

References

2025. RealPage & Algorithmic Rental Price Fixing. <https://localprogress.org/issues/housing/algorithmic-rental-price-fixing/>.
2026. AI Localism Project List. <https://list.ailocalism.org/>. Accessed: 2026-04-25.
- American Action Forum. 2026. AI Legislation Tracker. <https://www.americanactionforum.org/list-of-proposed-ai-bills-table/>. Retrieved March 11, 2026.
- Angelidou, M. 2014. Smart city policies: A spatial approach. *Cities*, 41: S3–S11.
- Arboledas-Lérida, L. 2025. Imagining Technology-Mediated Social Change? A Critique of the Theory of Sociotechnical Imaginaries. *Minerva*, 1–22.
- Arnold, Z.; Schiff, D. S.; Schiff, K. J.; Love, B.; Melot, J.; Singh, N.; Jenkins, L.; Lin, A.; Pilz, K.; Enweareazu, O.; et al. 2024. Introducing the ai governance and regulatory archive (agora): An analytic infrastructure for navigating the emerging ai governance landscape. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 7, 39–48.
- Baack, S. 2015. Datafication and empowerment: How the open data movement re-articulates notions of democracy, participation, and journalism. *Big Data & Society*, 2(2): 2053951715594634.
- 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.
- Bowen, G. A. 2009. Document analysis as a qualitative research method. *Qualitative research journal*, 9(2): 27–40.
- Brennan Center for Justice. 2025. Artificial Intelligence Legislation Tracker. <https://www.brennancenter.org/our-work/research-reports/artificial-intelligence-legislation-tracker>.
- Campolo, A.; and Crawford, K. 2020. Enchanted determinism: Power without responsibility in artificial intelligence. *Engaging Science, Technology, and Society*.
- Cave, S.; Dihal, K.; and Dillon, S. 2020. *AI narratives: A history of imaginative thinking about intelligent machines*. Oxford University Press.
- Evans, M.; Marsh, D.; and Stoker, G. 2013. Understanding localism. *Policy studies*, 34(4): 401–407.
- Gryz, J.; and Rojszczak, M. 2021. Black box algorithms and the rights of individuals: No easy solution to the “explainability” problem. *Internet Policy Review*, 10(2): 1–24.
- Gutierrez, S.; Nguyen, D.; and van Es, K. 2025. Tool, companion or a catalyst force? Exploring sociotechnical imaginaries Within AI livestreams’ communities of practice. *Big Data & Society*, 12(4): 20539517251381663.
- Hannak, A.; Sapiezynski, P.; Molavi Kakhki, A.; Krishnamurthy, B.; Lazer, D.; Mislove, A.; and Wilson, C. 2013. Measuring personalization of web search. In *Proceedings of the 22nd international conference on World Wide Web*, 527–538.
- Hendriks, A.; Karhunmaa, K.; and Delvenne, P. 2025. Shaping the future: A conceptual review of sociotechnical imaginaries. *Futures*, 170: 103607.
- 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.
- Jasanoff, S.; et al. 2015. Future imperfect: Science, technology, and the imaginations of modernity. *Dreamscapes of modernity: Sociotechnical imaginaries and the fabrication of power*, 1–33.
- Kao, K.-T. 2024. From Robodebt to responsible AI: sociotechnical imaginaries of AI in Australia. *Communication Research and Practice*, 10(3): 387–397.
- Katz, B.; and Bradley, J. 2013. *The metropolitan revolution: How cities and metros are fixing our broken politics and fragile economy*. Bloomsbury Publishing USA.
- Katz, B.; and Nowak, J. 2018. *The new localism: How cities can thrive in the age of populism*. Bloomsbury Publishing USA.
- Kazansky, B.; and Milan, S. 2021. “Bodies not templates”: Contesting dominant algorithmic imaginaries. *New Media & Society*, 23(2): 363–381.
- Krippendorff, K. 2018. *Content analysis: An introduction to its methodology*. Sage publications.
- Kuchler, M.; and Stigson, G. M. 2024. Unravelling the ‘collective’ in sociotechnical imaginaries: A literature review. *Energy Research & Social Science*, 110: 103422.
- Laba, N. 2025. Whose imagination? Conflicting narratives and sociotechnical futures of visual generative AI. *AI & SOCIETY*, 1–17.
- Lehtiniemi, T.; and Ruckenstein, M. 2018. The social imaginaries of data activism. *Big Data & Society*, 6(1): 2053951718821146.
- Luusua, A.; Ylipulli, J.; Foth, M.; and Aurigi, A. 2023. Urban AI: Understanding the emerging role of artificial intelligence in smart cities. *AI & society*, 38(3): 1039–1044.
- Mager, A.; and Katzenbach, C. 2021. Future imaginaries in the making and governing of digital technology: Multiple, contested, commodified.
- Marcucci, S.; Kalkar, U.; and Verhulst, S. 2022. AI localism in practice: Examining how cities govern AI. *Available at SSRN 4284013*.
- Meyer, E. T. 2014. CHAPTER THREE EXAMINING THE HYPHEN: THE VALUE OF SOCIAL INFORMATICS FOR RESEARCH AND TEACHING. *Social informatics: Past, present and future*, 56.

- Meyer, E. T.; Shankar, K.; Willis, M.; Sharma, S.; and Sawyer, S. 2019. The social informatics of knowledge.
- Milan, S. 2024. Resistance in the data-driven society. *Internet policy review*, 13.
- National Conference of State Legislatures. 2025. Artificial Intelligence 2025 Legislation. <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2025-legislation>.
- Oates, W. E. 1999. An essay on fiscal federalism. *Journal of economic literature*, 37(3): 1120–1149.
- Oldenburg, N.; and Papyshchev, G. 2025. The Stories We Govern By: AI, Risk, and the Power of Imaginaries. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, volume 8, 1939–1950.
- Rubinstein, I. S. 2018. Privacy localism. *Wash. L. Rev.*, 93: 1961.
- Smallman, M. 2020. ‘Nothing to do with the science’: How an elite sociotechnical imaginary cements policy resistance to public perspectives on science and technology through the machinery of government. *Social studies of science*, 50(4): 589–608.
- Stoker, G. 1996. Redefining local democracy. In *Local democracy and local government*, 188–209. Springer.
- Taylor, C. 2004. *Modern social imaginaries*. Duke University Press.
- Tidwell, J. H.; and Tidwell, A. S. 2018. Energy ideals, visions, narratives, and rhetoric: Examining sociotechnical imaginaries theory and methodology in energy research. *Energy Research & Social Science*, 39: 103–107.
- United Nations Statistical Commission. 2020. Recommendation on the Method to Delineate Cities, Urban and Rural Areas for International Statistical Comparisons (BG-Item3j). <https://unstats.un.org/unsd/statcom/51st-session/documents/BG-Item3j-Recommendation-E.pdf>.
- van Es, K.; van der Weijden, D.; and Bakker, J. 2023. The multifaceted and situated data center imaginary of Dutch Twitter. *Big Data & Society*, 10(1): 20539517231155064.
- Verhulst, S.; and Sloane, M. 2020. Realizing the potential of AI localism. *Project Syndicate*, 7.
- Von Eschenbach, W. J. 2021. Transparency and the black box problem: Why we do not trust AI. *Philosophy & technology*, 34(4): 1607–1622.
- Wernick, A. 2024. Impact Assessment as a Legal Design Pattern—A “Timeless Way” of Managing Future Risks? *Digital Society*, 3(2): 29.
- White House. 2026. National Policy Framework for Artificial Intelligence: Legislative Recommendations. <https://www.whitehouse.gov/wp-content/uploads/2026/03/03.20.26-National-Policy-Framework-for-Artificial-Intelligence-Legislative-Recommendations.pdf>.
- 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.