



October 2025

Making AI Work for the Public: An ALT Perspective

Neil Kleiman, Eric Gordon, & Mai-Ling Garcia

RethinkAI

Last edited on October 08, 2025 at 4:31 p.m. EDT

Acknowledgments

We would like to thank the Kresge Foundation and the Chan Zuckerberg Initiative for their generous support of this work. Special thanks also to all the subject matter expert interviewees who participated in our research.

Editorial disclosure: The views expressed in this report are solely those of the authors and do not reflect the views of New America, its staff, fellows, funders, or board of directors.

About the Authors

Neil Kleiman is an urban policy professor at Northeastern University and a senior fellow at the Burnes Center for Social Change and the GovLab.

Eric Gordon is a professor and director of the Center for Media Innovation & Social Impact at Boston University.

Mai-Ling Garcia is the Director of Emerging Technology and AI at the Bloomberg Centers for Government Excellence and Public Innovation at Johns Hopkins University.

About New America

We are dedicated to renewing the promise of America by continuing the quest to realize our nation's highest ideals, honestly confronting the challenges caused by rapid technological and social change, and seizing the opportunities those changes create.

About RethinkAI

RethinkAI builds ambitious pilots with city and community leaders and then translates learnings from those efforts into concrete guidance, policy recommendations, and an AI blueprint for everyone.

About Technology and Democracy

New America's Technology and Democracy (T&D) programs work to foster a sustainable digital future that advances equitable opportunity, innovation, fundamental rights, and participatory governance—where democracy, human rights, and the planet flourish.

Contents

Introduction	6
About This Report and RethinkAI	6
How the Lessons of Civic Tech Can Inform Public Use of AI	7
Key Findings	9
What’s Happening On the Ground	9
The ALT Framework	10
AI in Government: A Field-Level Review	12
What’s Happening in States	12
What’s Happening in Cities	16
The Experience of AI Adoption	17
Supporting Institutions	20
Emerging Themes	25
Recommendations: An ALT Framework	26
ALT in Action	27
The Leadership We Need	34

Contents Cont'd

Appendix	36
Credits	36
Supporting Institutions	20
RethinkAI	36
Methods	36
Interviewees	37

Introduction

In Boston's Dorchester, a mixed-income and culturally diverse neighborhood, a community group was fed up with the city government's misperceptions of the neighborhood. They wanted to represent themselves and what was happening in their own backyards. So they built an artificial intelligence (AI) tool to make it happen.

The result, *On the Porch*, highlights the opportunity before all of us to change the way the public sector serves our communities—to rethink the public sector enterprise.

The problem is that the public sector may not know this work in communities is happening. It's been focused on building its own capacity. Just in the first half of 2025, attention to AI has dramatically increased: In the United States, 700 state bills have advanced, many states are requiring on-the-job tech training and data literacy for all staff, and city-level pilots have mushroomed across the country.¹ There is indeed a lot happening, but mostly it's focused on improving what government is already doing, not changing what it does.

AI is certainly billed as boosting efficiency, but also as a revolutionary technology. We believe that while the field has been focused on reforming, the real opportunity is in transforming.

We have some big ideas for how we can do it.

About This Report and RethinkAI

This work was conducted by **RethinkAI**. We are a coalition of researchers, activists, and practitioners long focused on the role of technology in the civic sector, brought together by the collective hunch that both the hype and the fear around AI are misplaced. Starting in the summer of 2023, we launched two parallel efforts to make sense of what is happening—an ongoing survey of trends and developments and a series of pilots with civic sector partners in three cities.

The ideas in this report were formed over the course of two years when we interviewed 40 officials; ran pilots in Boston, New York City, and San Jose; and reviewed the state of play throughout the country. We conducted this work with two goals: (1) provide a sober evaluation of the current landscape and a clear understanding of what AI implementation looks like on the ground, and (2) propose an alternative, results-based framework for civic and government actors to incorporate AI into their work.

How the Lessons of Civic Tech Can Inform Public Use of AI

To harness AI effectively, civic sector leaders need to build on reform efforts of the past. To genuinely rebuild trust in our democracy, we need to put results and people first. On the city and state levels, we have an obligation to protect our institutions, and an opportunity to reshape them to better serve residents by being more adaptable and responsive to community needs.

Our current moment is not entirely new. We have ridden tech waves in the past with varying results. Take the civic tech movement, of which we consider ourselves a part. This broad-based coalition includes a combination of practitioners from inside and outside of government. Over nearly two decades, civic technologists have launched thousands of projects, created and supported government apps, written books and white papers, and ultimately, created a community of practice committed to making public services better. This big tent of civic tech has adopted a kaleidoscope of different terms and names, including reinventing government, e-gov, smart cities, public interest technology, and on and on. Similarly, open data advocates shined a light on making the activities of government more transparent to the public, shifting attention from the mere delivery of services to the visibility and accountability of the institutions behind them.

Past efforts have led to significant change, but often change centered on improving process, leaving genuine transformation for another day. Rather than just making services better, faster, and cheaper, we should have focused more directly on what residents needed. We often fell into an efficiency trap. For those of us who have worked in public service, we found that bold ideas were met with a centrifugal force towards incrementally improving existing programs. Past choices lock institutions into familiar patterns, making it hard to break away from incremental fixes. The public administration term for this is “path dependency,” which essentially means that within governmental bureaucracies, it is almost impossible to break free from traditional policies and programs and embark on new paths. For someone who just lost their job and needs to feed their family, a simple, expedited enrollment process to their state’s food assistance program is critical. But over the long run, what that individual needs is access to resources to turn their employment and economic conditions around. We can’t just stop at access to public benefits.

Civic tech, despite good intentions, sometimes reinforced these paths instead of disrupting them. We weren’t doing enough to focus on outcomes or reverse the course of residents’ steadily diminishing trust in democratic institutions. Armed with tech tools and methods, our field celebrated basic government improvements on the merits of saving time and money. And while many of these measures led to improvements, a lot of people—especially in underserved communities—remain unconvinced by declarations of success.

This marginally-better-with-technology approach has run its course. Our public institutions are under attack. And many of our attempts at reform didn't fully address residents' needs, discontent, or apathy. We are at an inflection point, and we need to rethink the role of civic technology as institutions change to meet the advent of AI and a new federal landscape.

What follows is a high-level analysis of policies, programs, and practices in local governments in the United States and their partnering organizations. We share insights about general trends and novel approaches. And we identify open doors that we might collectively be able to walk through to shape AI to meet the needs of our democracy, as opposed to shaping our democracy to meet the needs of AI.

Key Findings

The hard truth is that AI is already shaping government, mostly to optimize old processes. If we keep chasing efficiency, we'll entrench the status quo. If we instead aim for **Adaptation, Listening, and Trust (ALT)**, we can reset how public institutions deliver value and rebuild legitimacy.

What's Happening On the Ground

- **Legislatures are sprinting—defensively.** Since 2019, states have proposed or passed more than 1,600 AI bills, 735 in the first half of 2025 alone (~45 percent of the total).² Seventy-seven percent are controlling (guardrails, bans, audits), split roughly evenly between constraining government use and regulating private providers.
- **Policy focus is fragmented** across education, safety, health, elections, etc.—a sign of no shared focus or paradigm.
- **Workforce training matters.** States are pairing adoption with mandatory training. About 100,000 public employees are enrolled in AI workforce development; some states require completion before staff touch AI tools.³
- **State and city approaches differ.**
 - States are advancing a whole-of-government approach by establishing safe environments to experiment with AI, evaluate results, and then scale.
 - Cities are more oriented towards stand-alone pilots addressing urgent problems.
- **Chief information officers (CIOs) are driving adoption.** Unlike in prior tech eras, information technology (IT) leads are the main actors writing the rules, conducting outreach, and publishing what works.
- **Ecosystem readiness is uneven but promising.**
 - Philanthropy is beginning to shift from internal productivity to field-building.

- Higher education is investing heavily in AI, and the best programs now deliver production-grade civic projects.
- Peer coalitions are moving the field from expert-led to operator-driven practice.

Diagnosis: By optimizing machinery that residents already distrust, we are building momentum without vision or a framework.

The ALT Framework

We recommend a governance—not government—approach organized around a model to **adapt, listen, and trust**.

A — Adapt: Plan for the Demand AI Unleashes

- Forecast demand before product or system launch; stress-test capacity; reallocate budgets and roles in advance.
- Use AI agents for routine steps; reserve discretion for humans.
- Redeploy people to work requiring high-level judgment; train staff for compassionate decision-making and bias awareness.
- Pass enabling policies (not just controls) so services can iterate as demand patterns change.

L — Listen: Clarify Needs, Don't Just Translate Text

- Create a public AI sandbox where staff and communities co-translate budgets, ordinances, and services into plain language with low or no-code tools.
- Train models for context engineering—institutional memory plus data architecture—so that they understand problems, not clever prompts.
- Combine structured signals (such as reports from 311, benefits information) with unstructured inputs (e.g., meeting transcripts, proposals) for richer insight.

T — Trust: Offer Two-Way Accountability, Not One-Way Transparency

- Stand up civic data trusts and community-controlled data; explore sovereign local models.
- Evaluate for trustworthiness (based on fairness, responsiveness, and usefulness) alongside efficiency.
- Tie resident sentiment to impact metrics, including environmental stress, service usage, and safety.
- Create public compacts with shared goals, timelines, and measures.

If we aim AI at efficiency, we'll get faster bureaucracy. If we aim it at ALT, we'll get a civic sector that residents actually believe in.

AI in Government: A Field-Level Review

AI use in the civic sector is at a very early stage. We evaluated local and state governments' use of AI to provide a thorough review of current policies and regulations, alongside indicators of willingness to experiment with new tools and practices. This includes an analysis of emerging strategies being put in place at supporting institutions such as universities, philanthropies, professional organizations, and a new crop of coalitions—including ours—that have sprouted up to support AI use in the civic sector. This is a fast-moving field, and we did our best to make sense of the moment amidst a dizzying pace of change.

Our interviews with over 40 practitioners and experts, pilot efforts, and thorough literature review yielded a representative sample that is by no means exhaustive, but provides a coherent picture of where the field is—and where we believe it can go from here.

What's Happening in States

Legislatures Are Moving to Control AI

Compared to past tech waves, this one has seen a truly unprecedented amount of legislation, regulation, and executive orders. This all happened quite fast and most of it is at the state level with jurisdictions below playing catch-up.

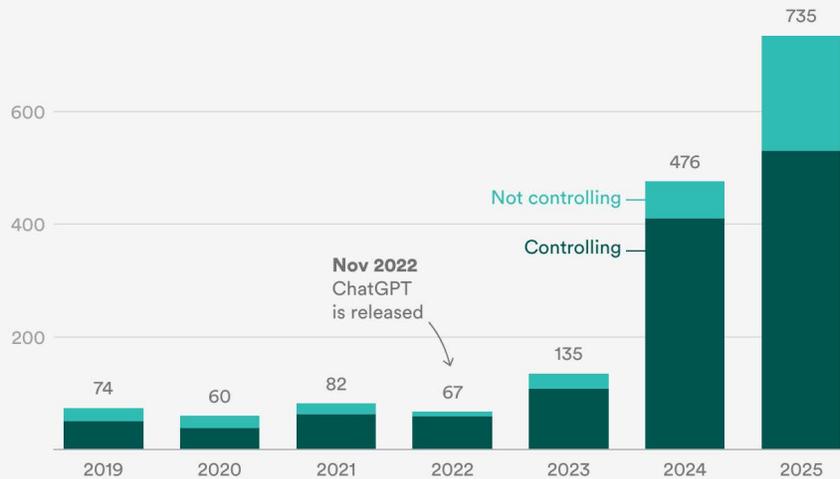
In 2025 alone, over 735 AI bills have been proposed—bringing the total since 2019 to over 1,600 bills, according to our analysis of data generated by the Beeck Center.⁴

There are three notable aspects of these bills that jump out: the sheer volume, the defensive nature, and the variety of policy areas covered.

First is the considerably large amount of legislation advanced in a short period of time. Based on our analysis of National Conference of State Legislatures data, there has never been a period in state history since the introduction of the internet that saw as much legislation proposed or passed around one technology.⁵ In fact, when we ran the numbers by field leaders, many commented, “That can’t be right. Is there really that much legislation being debated and passed?” There is. State governments are not sitting on their hands legislatively.

Figure 1 | Yearly AI Bill Proposals Have Grown 10x Since 2019, Most of Which Seek to Control the Technology’s Usage

State-level AI bill proposals by year. **Controlling legislation**—laws that set guardrails around how AI can be used—makes up 77% of the 1,600+ bills proposed since 2019.



Bills are counted as “controlling” if they include at least one controlling measure, such as guardrails, bans, or audits.

Source: AI Legislative Database (Beeck Center for Social Impact + Innovation, Digital Service Network).

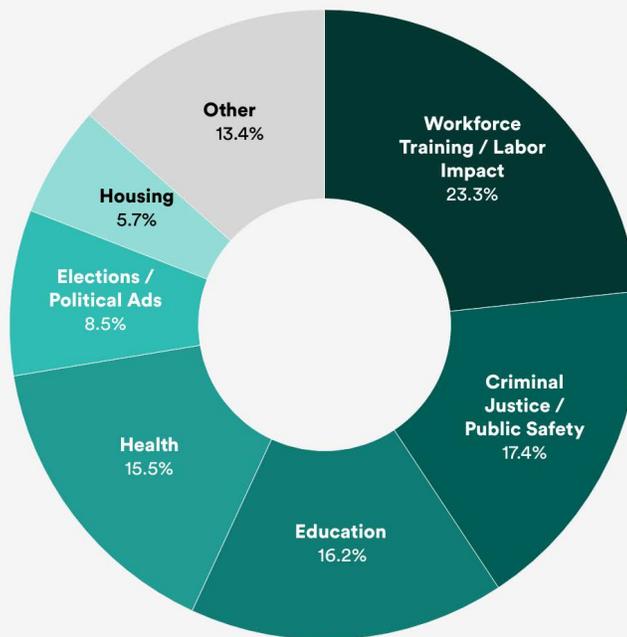
NEW AMERICA

Fully 77 percent of these bills can be classified as controlling legislation (see Figure 1). Rather than a focus on how to advance AI or enable its use, controlling legislation refers to laws that restrict or set guardrails on how AI can be used. Public sector legislation constrains how government agencies adopt AI, while private sector rules set accountability standards for firms that build and deploy it. Of the controlling pieces of legislation, roughly half focus on the public sector and include inventories, procurement audits, or ongoing impact assessments. For example, Maryland’s Artificial Intelligence Governance Act of 2024 (SB818) mandates agency inventories, bans deployment of high-risk systems without safeguards, and creates an executive subcabinet to oversee compliance. The other half of bills focus on controlling the private sector by placing regulations on companies themselves: New York’s pending 2025 AI Act (S1169) would require independent audits of high-risk systems that may impact the public’s health or legal rights, anti-discrimination safeguards, and enforcement by both the state attorney general, as well as allowing private lawsuits. One state administrator summed up this orientation by saying, “For the majority of bills introduced, legislators had some sense of risks to avoid and harms to punish, but very little of direction to push the state towards.”⁶

For legislation that was not controlling in nature, we assessed the breakdown within specific policy areas. Figure 2 below shows the major areas that the National Conference of State Legislatures tracks. As can be seen, the categories are quite varied, with no one policy area capturing the full attention of state lawmakers.

Figure 2 | Despite Increasing in Volume, AI Bill Proposals Remain Fragmented in Focus

Percentage of state-level AI bill proposals by policy domain, 2023–2025



Source: Artificial Intelligence Legislation Summary (National Conference of State Legislatures; 2023, 2024, 2025).

NEW AMERICA

Workforce Has Become a Priority

Government, especially state government, rarely puts workforce development or job training high on its priority list. Job training is generally thought of as a federally funded area, and in terms of internal training of government employees, it rarely rises to the policy agenda level, even though public administration research has clearly demonstrated the need.

The rapid distribution of AI has made that case better than researchers ever could. A number of the new legislative bills address workforce concerns. Take Texas, which earlier this year enacted two laws (HB3512 and HB2818) that

paired internal AI deployment with government staff training, requiring workforce upskilling alongside broader adoption efforts.⁷

We also observed the workforce commitment in the uptake of InnovateUS courses focused on AI. This is a relatively new philanthropically funded initiative offering free courses in AI basics and administration. The response has been positive. Many state governments are recommending or requiring the courses for all staff.

InnovateUS is now working in close partnership with a dozen states: Maryland, Ohio, New Jersey, Colorado, Arizona, Oregon, Minnesota, Connecticut, California, Pennsylvania, Maine, and Georgia, and at the time of this report's publication, was about to launch in New York, Tennessee, Illinois, and Massachusetts.⁸ As of September 2025, more than 110,000 state employees were enrolled in courses, workshops, and certificate programs. In fact, Colorado and a few other states have required completion of InnovateUS's "Responsible AI for Public Professionals" training in order to access AI tools.⁹ As J.R. Sloan, Arizona's chief information officer (CIO), noted, "As AI rapidly develops, it is essential we prepare our workforce with the skills they need to use this technology both safely and effectively. The State of Arizona prioritizes privacy, security, and responsible experimentation with AI technology in its government operations. [The InnovateUS] training aligns with these values, providing proper guidance and guardrails that enable the responsible use of AI."¹⁰

Sandboxes and Walled Gardens

We identified a growing trend among states: developing and executing on an organization-wide approach to AI adoption. This followed a well-defined and sequential approach of establishing a secure environment for experimenting with AI, including evaluating trials; establishing a formal governance structure, set of guidelines, or both; and moving to enterprise-wide adoption.

These state efforts create a safe environment where agencies test generative AI tools under limited risk conditions before moving toward broader adoption. The initiatives fall under different names, such as "sandboxing" or "walled gardens," with a number of overlapping categories and approaches. Some states use regulatory sandboxes—granting temporary waivers or exemptions so developers can pilot new AI systems under oversight (such as in Utah, Texas, and Delaware).¹¹ Others use internal walled gardens inside the government, where employees can experiment with commercial AI tools in a secure enterprise environment (California, New Jersey, Colorado, and Pennsylvania).¹² And a third model relies on executive orders or state innovation labs to coordinate ad-hoc pilots (Washington, Maryland, and Georgia).¹³

At least seven states have created sandbox or pilot structures: Utah, Texas, Delaware, California, New Jersey, Colorado, and Pennsylvania.¹⁴ A handful of others (Washington, Maryland, Georgia, and New York) are experimenting through innovation labs or agency pilots—though these don't fit the formal definition of sandbox programs.¹⁵ Several states have also attempted to enact sandbox programs, but the proposals either stalled or failed to pass (Oklahoma, Connecticut, Mississippi, Missouri).¹⁶

So far, the actual usage of AI in these programs has been pragmatic and modest. Most pilots have involved off-the-shelf tools deployed to improve government operations rather than cutting-edge AI applications. The most common uses are helping employees draft memos, generate translations, or distill dense rulebooks and human resources handbooks.

What's Happening in Cities

Unlike states, cities have focused on stand-alone pilots rather than enterprise-wide pilots. Cities tend to identify priority areas (e.g., wildfires in Los Angeles) and build AI tools to solve immediate problems, whereas states are establishing platforms for more general experimentation.

In order to understand what was happening on the city level, we conducted a field scan of AI-related projects. The most notable aspect of our scan was the variety of cases to choose from. When we began this work in the summer of 2023, we struggled to find many advanced AI use cases at the municipal level; that was no longer the case in the first half of 2025. Our selection criteria included projects that specifically mentioned generative AI, aimed for scale, and received some level of local or national news coverage. We identified 12 use cases and pilots and placed them into one of six application areas: permitting and automation, employee productivity, public safety, resident services, public engagement, and infrastructure analytics.

Table 1 offers a snapshot of the range of pilots taking shape locally. Taken as a whole, it reveals growing ambition in the past couple of years. Cities are engaging a wide set of partners on a disparate range of challenges. They represent the kinds of projects that are achievable at this moment, even with little in-house AI talent.

Table 1 | City-Level AI Applications and Use Cases Vary Widely

Application	City	Use Case	Task & Capability	Partner
Employee Productivity	San Francisco, CA	Microsoft CoPilot deployed to 30,000 city employees	NLP (LLM) Summarize, generate text, workflow automation	Microsoft (built on OpenAI)
Employee Productivity	Lebanon, NH	Emergency medical services (EMS) reports automated with AI	NLP (LLM) Summarize and generate text	EMSReports.ai (startup)
Infrastructure Analytics	San Antonio, TX	AI-powered heat resilience mapping for climate planning	Regression + Predictive Analytics Pattern finding, forecasting	Google (big tech)
Infrastructure Analytics	Detroit, MI	Annual street mapping and change detection, foundations for digital twin	Computer Vision + Classification Geospatial analysis, pattern detection	Mapillary (Meta), University of Michigan MIDAS
Permitting & Automation	Los Angeles, CA	AI-assisted wildfire plan review for permits	Computer Vision + Optimization Automate workflows, detect patterns	Archistar AI (startup)
Permitting & Automation	Santa Clara, CA	Racial deed scanning project	Classification (pattern detection) Archival record analysis	Stanford Reglab
Public Engagement	Washington, DC	Public input pilot for policymaking	NLP (LLM + survey analysis) Text summary, engagement insights	MIT Gov Lab, Stanford Digital Economy Lab
Public Engagement	Bowling Green, KY	“BG 2050” participatory planning tool	NLP (LLM) Analyze resident input, summarize feedback	Jigsaw (Google), Polis (startup)
Public Safety	New York, NY	AI-Powered noise Cameras for traffic enforcement	Computer Vision + Classification Detect violations, pattern recognition	Intelligent Instruments (startup)
Public Safety	New Orleans, LA	Real-time AI translation in 911 calls	NLP (machine translation/speech recognition) Real-time translation & pattern recognition	Carbyne (startup)
Resident Services	San Jose, CA	SJ311 provides real-time translation for resident requests	NLP (machine translation) Retrieve/search & translation	Google AutoML
Resident Services	Boston, MA	City Council roll call summaries for public access	NLP (LLM) Summarize & generate text	Google Vertex AI

NLP = Natural Language Processing. LLM = Large Language Model. Employee Productivity includes applications supporting internal staff efficiency. Infrastructure Analytics includes applications strengthening asset monitoring and planning. Permitting & Automation includes applications streamlining record-keeping and approvals. Public Engagement includes applications strengthening citizen input and planning. Public Safety includes applications improving emergency response and enforcement. Resident Services includes applications improving efficiency of services that directly touch residents.

Source: Original New America analysis based on publicly available data.

NEW AMERICA

The Experience of AI Adoption

To complement the data analysis above, we conducted a qualitative analysis based on over 40 interviews with government actors (current and former officials), academics who write about government, and related practitioners in business and philanthropy fields. Here is what we discovered.

Locales Are Still Making Sense of AI

States are proposing legislation on an order never before seen, and AI experimentation is emerging in virtually every corner of the country, but people in the civic sector are very cautious. Most states we spoke with are still in the sandbox/walled garden phase of experimentation, and while expanding quickly, it is still only a handful of cities developing major AI pilots. And, most of this activity is internal to government. There is certainly enthusiasm, but generally, this is a field that is still making sense of the technology.

Lack of trust and vision for leveraging AI: Much of this is attributable to the fact that AI is not as intuitive as other tech waves. As Aimee Sprung of Microsoft noted, unlike during the smart cities era when many were dazzled by the possibilities, “Now we fundamentally have an environment in which there is a challenge of understanding the capabilities of the technology. Yes, the public needs to trust AI, but so does government.”¹⁷ As one local government professional association leader said, “Cities are tiptoeing forward. If anything, we have gone backwards in the past year—there was a lot of interest in 2023 and then it leveled off.”¹⁸

Ongoing state legislative activity is not helping to advance local-level AI. Many officials characterized the 1,600 bills as background noise that doesn’t add up to a strategy and lacks clarity in terms of how to proceed. Some officials noted that they are eager for the coherent guidance that President Joseph Biden’s October 2023 AI Executive Order began to provide. Since the termination of that order, the path forward has felt murky. One state leader noted, “There is a narrative vacuum filled by different voices—voices from the legislature wanting a strategy, agencies trying their own thing, and public comments [urging us to do a better job].”¹⁹

Need for technical capacity and infrastructure: A challenge that a number of interviewees brought up was the persistent lack of foundational tech capacity to equip them for successful AI experimentation, especially at the city level. A professional association leader said, “I hear all the time from our city members that we lack the capacity and the tech talent.”²⁰ Another local official stated flatly, “We aren’t doing much with AI because we don’t have the infrastructure in place to execute.”²¹ We also heard from a few university professors wanting to provide pro bono assistance who couldn’t even get their calls returned. One professor said, “Trying to work with city hall [on AI] is hard because they are so woefully understaffed. Even project scoping takes time.”²²

Part of the reason for local government lag is that the value and return-on-investment for AI isn’t clear yet. A typical tech innovation life cycle has a large number of adopters at the outset that try it out, then attrition, followed by much more practical uses at more affordable prices. But at the municipal level, there

is a dynamic of waiting for the AI cycle to play out. One professional organization official said, “Cities want to see a full proof of concept for these tools before experimenting. Right now [AI] is still vague and ambiguous.”²³ Santi Garces, Boston’s CIO, noted, “Enterprise software can be [exceedingly expensive]. Costs are utterly misaligned with public sector budgeting realities. We are still [waiting] for the private sector to refine and ‘productize’ the technology.”²⁴

Early stages of greater experimentation: Despite these challenges, there are indeed locales advancing homegrown AI efforts. Where this has happened we have seen a dynamic of employing what one field leader called “minimum viable governance”—quickly cobbling together a permission structure for administrators to safely use AI. Indeed, that is what states are doing with their walled garden approach: establishing a secure environment to try things out and see what works.

At the city level, a few locales are taking this approach. As Kat Hartman, chief data officer of Detroit, put it: “We cannot regulate something we do not understand. Regarding AI, we must roll up our sleeves and learn to do it ourselves, internally. This is the only way to build up true knowledge and literacy at the level of local government—whether we build or we buy.”²⁵

The cases of experimentation are notable, but it is apparent that most local AI activity is still around basic automation and productivity. As Stephen Goldsmith, Harvard professor and former mayor of Indianapolis, noted, “One of the greatest challenges with local government is a lack of imagination.”²⁶ And Lane Dilg, former Santa Monica city manager who led policy partnerships for OpenAI, said, “Basic productivity gains are far easier to achieve than ambitious transformations, like rethinking an entire permitting process.”²⁷

As local-level AI is more focused on productivity and making sense of the technology, we found it is CIOs who are often leading at both the state and city levels. Although in more recent tech-driven waves, a few CIOs have been in the forefront, typically big ideas and vision emanated from mayors (think Michael Bloomberg and Pete Buttigieg) or highly visible strategy chiefs like innovation or chief data officers. This time, most leadership is coming from more backend, system-wide administrators.

Albert Gehami, one of the founders of an international peer group called the GovAI Coalition and the city privacy and AI officer in San Jose, explains that the City has published its own analysis and conducted its own resident outreach. “We are issuing reports with honesty about what worked and what didn’t. And we are talking to residents. Resident outreach is not a typical IT [information technology] function, but it is showing results. This is [all] new and scary, but technical staff need to get out there.”²⁸

This new role is being recognized by many field leaders. Erin McKinney from Amazon Web Services noted, “CIOs have historically been charged with implementing new state policies, and are now writing the rules for how state and local government agencies adopt AI. CIOs are stepping into leadership roles that are increasingly innovative and strategic, not just operational.”²⁹

Taken together, local governments are taking their time to get to know AI, begin the process of incorporating it into daily routines, and start to think about more ambitious, potentially transformative uses. One state administrator summed it up by saying, “Most of us are primarily—and necessarily—crossing the river by feeling the stones, but would appreciate pragmatic and positive visions to work towards.”³⁰

Supporting Institutions

Local governments, particularly cities, have always been adept at forming partnerships. Working in concert with the business community and chambers of commerce has driven local growth and development efforts for decades. More recently, university and philanthropic collaboratives have become helpful. Local governments at all levels are eager for new and stronger partnerships around AI. In fact, we decided to draft this section about such institutions because it came up in virtually every interview we conducted. Interviewees across the board noted that partners—particularly in higher education—could provide meaningful help with AI guidance and evaluation of use. Below is our assessment of various institutions that could potentially partner with local governments.

Federal Government

Few at the local level ever thought of the federal government as an AI problem solver. But there was fairly wide recognition that the Biden administration’s 2023 AI Executive Order provided helpful guidance and strongly worded guardrails for those at the sub-national level. While our interviewees did not discuss the Trump administration’s position on AI (or lack thereof), many expressed disappointment toward the absence of national-level tone setting and framing. The current administration may ultimately release a stance that preempts many state AI bills that have been passed and are being debated. But in our scan, that did not raise many flags, as local-level officials were more focused on setting their own plan and strategy without federal support.

Professional Government Organizations

There are over half a dozen major professional organizations serving state, county, and municipal governments in the United States. These entities provide a range of services, including peer networking, training, and policy advocacy agendas. They tend to follow their members' lead and shift programming accordingly. As local governments are still making sense of AI, professional organizations, not surprisingly, have yet to take strong positions or provide much technical or operational guidance. A 2024 International City/County Management Association (ICMA) member survey found nearly 50 percent of the respondents reported AI utilization as a low priority, and only about 5 percent placed a high priority on AI use.³¹

In fact, a couple of the professional organizations have ratcheted down their efforts over the past two years. At the time of publication, we found that this tide is turning, as evidenced by the recent spike in city use cases noted in the previous section. Tad McGalliard, managing director for research at ICMA, noted that in a recently updated internal survey, AI was the number one issue members want addressed. He said, "AI policy and practice diffusion hasn't happened, but in just the past year, interest has skyrocketed. Everyone now knows they need to do something."³²

Philanthropy

Private philanthropy since the 1950s has mostly steered clear of directly supporting public sector efforts in favor of targeted nonprofit funding, but when it does wade into new areas it can have an enormous impact, as its funds are flexible and seen as creative capital in the civic sector. In that regard, philanthropy could significantly influence where the field goes next.

Our scan of philanthropy revealed that, as for other sectors, its approach to AI is still taking shape. Organizations are still wrapping their heads around the new technology and trying to assess how it might impact them internally and how they can support their grantees. Jean Westrick, executive director of the Technology Association of Grantmakers, noted that "AI adoption in philanthropy is marked by a paradox: individual use is high, but enterprise-wide strategies remain limited. Funders are exploring how AI benefits their operations, yet too few are investing in their nonprofit partners. Without intentional support, this imbalance risks widening the technology equity gap—leaving the very organizations closest to communities behind."³³

There are a few notable exceptions, as philanthropy around AI has grown both in scale and urgency since the 2022 release of ChatGPT. A few large foundations and mega-donors are moving billions toward AI safety, governance, and public interest applications. Similarly, new multiyear

coalitions like the \$1 billion NextLadder Ventures initiative for frontline workers, and major corporate commitments such as Microsoft's \$4 billion for AI education, are aiming to diffuse the use of AI across society.³⁴ Bloomberg Philanthropies is supporting a few different city AI initiatives, and Robin Hood has created an AI Poverty Challenge.³⁵ The result is an ecosystem where funders are seeking ambitious, scalable ideas that don't just react to the risks of AI but shape the infrastructure, workforce, and governance frameworks that will define its societal impact.

These shifts open up space for initiatives that move from responding to present needs toward forecasting and preparing for future risks and opportunities. At the same time, trends toward trust-based philanthropy and capacity-building align well with the idea of positioning communities to lead. Many funders are shifting to longer-term, flexible support with reduced reporting burdens—an opening for community-led initiatives to be models for field-building, training, and public interest AI.

Higher Education

Colleges and universities have an uneven record of working with local governments and communities. In many respects they have a different mission; as Kate Burns from MetroLab noted, “Fundamentally, research does not have a client.”³⁶

Our assessment of higher education revealed considerably more activity than with other supporting institutions. Virtually every top-rated research institution has some combination of a presidential task force, a newly created institute, AI faculty hires, and additional academic programs focused on AI use. A few institutions are home to multimillion dollar new investments including \$500 million for the Kempner Institute for AI at Harvard, \$1.2 billion from public and national lab resources for an AI research facility at the University of Michigan, \$20 million from New York State and IBM into AI efforts in the State University of New York System, and University of California, Berkeley's brand new College of Computing, Data Science, and Society.³⁷

While there is tremendous investment in AI at many universities, there hasn't been a particular uptick in civic AI commitments at the institutional level. But we did find a few notable, if isolated, examples of higher education playing an outsized role in advancing new AI efforts with local government and communities. These were high-impact yet low-fanfare initiatives that came together quickly over the past couple of years. Each one used AI as a tool to

quickly surface value and change the traditional relationship between universities and surrounding communities:

- Two different experiential learning programs have rapidly become more focused on policy impact with AI. At Tulane University, the mandatory undergraduate public service requirement has been known to focus on basic community outreach, such as tutoring and park beautification. But Aron Culotta and Nicholas Mattei, two enterprising computer science professors and directors at the Center for Community-Engaged AI, work with local nonprofits and undergraduate students to advance significant criminal justice reforms.³⁸
- At Northeastern University, the co-op program providing students with work experience is widely recognized as one of the best administered in the country, but to date, it has focused almost exclusively on the private sector. Now, through the new AI for Impact program, there is a shift to public partnership that has quickly yielded concrete results. In six-month sprints, the program has built 18 generative AI production-ready solutions, co-designed with community partners for the governor of Massachusetts and other civic leaders in the state.³⁹
- Georgia Tech houses the Partnership for Innovation, a public-private partnership aimed at boosting economic growth. One of its programs aligns regional universities with local challenges. Its founding executive director, Debra Lam, noted that AI was an excellent vehicle that proved useful. She said, “We don’t lead with tech, but with problems. We have a number of projects, including working with rural farmers, in which AI was able to meet the need and get to results in short order.”⁴⁰
- At the University of Michigan, working on a National Science Foundation grant, a group of scholars who had never engaged with the public sector found that within days, they were able to develop new AI-powered solutions for the Detroit city government around urban planning and climate change.

Emerging Coalitions

As with higher ed, in coalition-building we saw a similar dynamic of new endeavors and pairings coming together to quickly advance AI efforts. New enterprises like InnovateUS, Bloomberg’s CityAI Connect, the GovAI Coalition, and the Digital Services Network have shifted the model from expert-led networks to more peer-driven movements (see Table 2).

The GovAI Coalition, for example, brings together over 800 local, state, and national government agencies around a shared set of principles to create common ground in how they engage with vendors that are approaching local governments with new intensity.⁴¹ This network provides members with the leverage needed to negotiate in the best interest of constituents.

Now that a field around civic change-making has been established, its movement has shifted from one of big announcements to showcasing more practical steps—and these new associations are leading the way.

Table 2 | Emerging AI Coalitions

Coalition / Network	Purpose	Participants	Examples of Activity
GovAI Coalition	Convened by the City of San Jose to advance shared AI principles and procurement practices	Over 800 local, state, and national government agencies	Published tools and templates for cities, particularly around common RFP language; hosted convenings; developed shared ethics commitments
CityAI Connect	Offered as a partnership through Bloomberg Philanthropies and Johns Hopkins University's Bloomberg Center for Government Excellence as a closed platform for city staff to exchange AI learnings	City government staff interested in AI-related initiatives from nearly 500 cities and 1,000 participants	Facilitated peer-to-peer live discussions, genAI use case sharing, a managed discussion forum, extensive resource library, weekly news highlights, and event and conference postings
Digital Service Network	Civic tech network serviced by the Georgetown Beeck Center, focused on digital services practices, providing actionable, practitioner-focused research, case studies, and resources to accelerate digital transformation, including AI capacity	8,000 digital service and innovation practitioners at all levels of government	Convened and facilitated active practitioner-based communities of practice based on role, identity, and topic area. Publish original research, host events like job fairs and webinars. Curate digital government resources and artifacts in the Digital Government Hub for replication
Collective Intelligence Project	Global initiative exploring how society can govern transformative technologies like AI in democratic, participatory ways	Policymakers, researchers, technologists, civil society organizations	Designed large-scale deliberation pilots (e.g., global citizens' assemblies on AI), produced research on governance models, and advised governments on participatory approaches to AI policy
RethinkAI	National network in the U.S. focused on governance reform through scaffolded experimentation, original research, and knowledge sharing	Researchers, policymakers, technologists, civil society organizations	Published ALT framework and national field assessment, held national convening, and co-designed pilots in Boston, New York, and San Jose

NEW AMERICA

Emerging Themes

AI is indeed taking shape in the civic sector, but little has been settled. While there has been significant progress, AI implementation and integration are still very much in a Wild West phase.

Some themes that emerged from this field review include:

- **There is a vision vacuum:** Policy, training, and experimentation are now advancing quickly. But there is a palpable narrative void. There are no set paradigms, example strategies, or governance frameworks like the ones that quickly took root in the smart cities and big data eras.
 - **Implication:** There is demand for a clear AI operating framework.

- **Many governments are employing an organization-wide orientation:** There is a greater openness to an all-of-government approach, not just one that is contained to IT offices. States in particular are getting strategic and enterprise-oriented.
 - **Implication:** Major reform can take root throughout government and at the departmental level, not just through specific tech applications.

- **Partnerships are in high demand:** The true impact of this technology will not be contained to any one organization. Everyone is eager to learn and to reduce the risk of AI. There is a necessary orientation to collaborate with a range of civic actors.
 - **Implication:** There is a greater chance of establishing a collaborative ecosystem of civic life beyond one dominated by or centered around government.

These are dynamics that point to openness to—and in some cases demand for—a new governance framework. A number of new resources and partners are waiting in the wings to assist.

The big question is: Where do we go from here—how do we use AI for (significant) good, and guard against (significant) harm?

Recommendations: An ALT Framework

The public sector is at a crossroads. Data and information have never been more contested or more valuable. A primary driver of this is AI: a high-reward, high-risk technology.

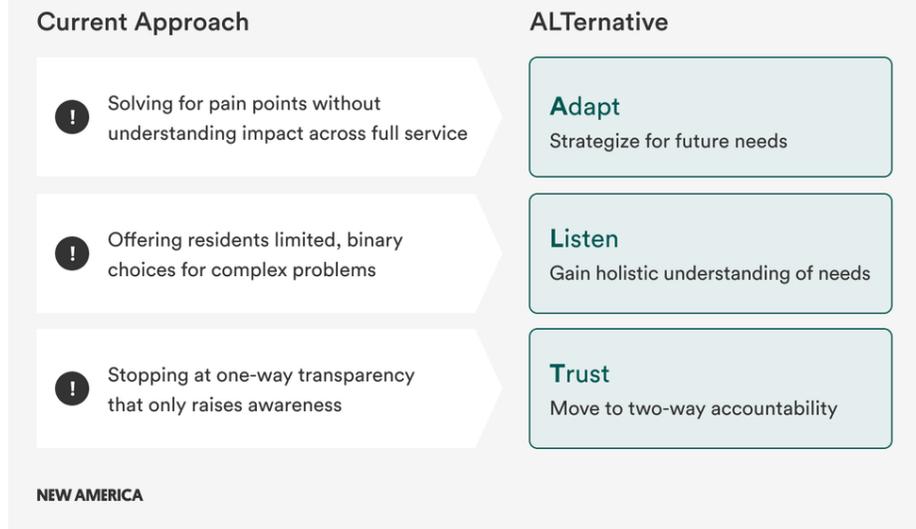
As we found in our review, AI has indeed shaken up the field. Governments everywhere are preparing for a big shift: There is a significant move to build computational infrastructure, a serious commitment to on-the-job training and data literacy, and a near-universal demand for good partners.

But there is also a void in imagination and vision. The civic sector needs a framework.

We propose an ALT approach to governance, not government. This is an alternative method of AI implementation and assessment in the public sector that focuses on how decisions are made, not just what decisions are made. It includes three imperatives for institutional reform: adapt, listen, and trust (see Figure 3).

Adaptation means anticipating future needs and repositioning institutions to meet them. **Listening** is a commitment to a holistic understanding of what people need. **Trust** is accountability to residents by demonstrating results.

Figure 3 | ALT Framework



Source: Alex Briñas/New America

What follows is a more in-depth look at each of these three components, providing hypothetical use cases and recommendations for implementation. We also recount the pilots we facilitated in Boston, New York, and San Jose and explain how they contributed to developing the ALT framework. Through worked examples, we developed a practical understanding of what works and what’s missing in AI development and implementation, and how to design a structure for guiding the field into the future.⁴² The goal is to guide the civic sector, including the public sector and its partners, towards the responsible and transformative adoption of AI.

ALT in Action

Adapt

AI is about anticipating future needs with more meaningful work, not just addressing current crises.

Adapt in a nutshell: How governments design and execute programs can be significantly altered by AI. For years, civic technology promised efficiency so that it could free up time for higher-touch services. In reality and in practice, better technology doesn’t just free up administrative time; it removes friction and makes it easier for people to access services. The result of well-placed AI isn’t efficiency, but more demand revealed. As a result, while AI agents don’t

necessarily solve service surges, they will ultimately shift what people will spend time on.

Traditionally, governments and civic organizations have been able to act on data collected and analyzed through standard channels: existing government data, surveys, and participation metrics. But AI gives practitioners the opportunity to understand not only what people need now, but what they will need in the future. Fundamentally, this is a planning function and necessitates reallocation of resources and staff time. Done well, it can help invent jobs of the future, while offloading routine work to AI. As a result, governments must not design for job elimination; they should design for human redeployment and organization adaptation.

This shift is not only about how governments manage their people, but also about how policies set the conditions for adaptation. Most recent policymaking has leaned controlling and defensive, to include restrictions, audits, bans, and measures that control risk but don't help agencies absorb new demand. Enabling policies, in contrast, create the flexibility to reallocate funds, redeploy staff, and iterate service models as needs surge. Without them, efficiency gains can backfire, turning frictionless access into overwhelming backlogs, as the next example illustrates.

Example: A city deploys a 311 chatbot to make it easier for residents to request bulky waste pickup. Service requests surge overnight. But staffing and routing processes remain unchanged. Instead of freeing up capacity, the system overwhelms crews with backlogs and longer delays. Efficiency gains on paper turn into operational strain in practice.

Alternative: The city prepares for amplified demand before rollout.

- AI models forecast how much service use will increase once friction is removed.
- Staffing is adjusted and funding is reallocated to handle the surge and address the better-understood environment.
- Digital requests flow into redesigned triage and case management systems so pickups are routed efficiently.
- Routine scheduling is automated, while staff time is explicitly reallocated to higher-value engagement with residents (e.g., resolving complex or repeated requests).
- Managers are trained to exercise discretion and compassion in how scarce resources are deployed.

To achieve the ALternative, we recommend the following:

- Realign budgets and staffing so saved time is explicitly reallocated to high-value human work, not lost to administrative creep.
- Forecast demand surges before pilot launch—model how reduced friction will increase requests, then stress-test systems for bottlenecks. This will require nimble, responsive strategies to AI grounded in service delivery, but adaptive approaches that respond to the changes in organizational need.
- Train staff in compassionate decision-making and bias awareness as humans ‘in the loop.’ Staff must be empowered to make decisions and be clear on the decisions they can and should make. These are cases that require human judgment, vulnerability, and fairness, which require more attention or human intervention outside of AI processing. This also presents a good opportunity to engage organized labor and unions to address issues of AI in the workplace.
- Leverage AI to support adaptation, not replace people. AI agents are autonomous software systems that can take in data and instructions, make decisions, and take authorized actions to achieve specific goals without constant human supervision. Used carefully, agents can automate processes, and with planning and forecasting, deploy humans for higher-value work.

In **San Jose**, the city government is addressing the challenge of homelessness by developing foundational applications that will support broader AI initiatives in the future. The city’s IT and housing departments have targeted the issue of evictions with an in-house predictive analytics tool. This tool helps them better understand when evictions are likely to occur and allows for preventative actions. While evictions represent only one aspect of the overall homelessness issue, the city’s innovative use of data is a crucial first step to **adapting** its programs, anticipating constituent needs, and proactively designing solutions to address them.

Since homelessness is a complex problem that extends beyond city borders, Rethink AI is collaborating with San Jose to explore the role a regional coalition—comprising government, nonprofit, and academic organizations—can play in adopting a new approach to homelessness. Building on the eviction application as a foundation, this coalition will leverage more extensive datasets and AI to develop regional solutions to the homelessness crisis. By applying AI regionally, all coalition members can adopt and **adapt** innovative strategies designed to anticipate future needs and improve program effectiveness.

Listen

AI is about clarifying information, not muddying it.

Listen in a nutshell: The ability for institutions to listen at scale has significant implications for governance.⁴³ It promises that governments can make sense of existing structured data (like from 311) and solicit new unstructured data (like community conversations), and use both to make decisions. For all the talk of AI fueling disinformation, it can also clarify meaning by helping governments make sense of, and generate novel insights from, new and existing data sources.

Civic tech efforts have focused on accessibility. Official policy documents, planner presentations, and even meeting notes are available online. In some cases, policies are synthesized into comprehensible summaries. Residents are able to see what is being offered, but they don't receive support in connecting that information to their needs or the needs of their communities.

Decisions are presented as a binary: Do you like or dislike this new zoning requirement? Residents have rarely been offered an opportunity to present what they want, just the option of responding “yay” or “nay.”

Example: A city website uses AI to post housing policies in plain language, and, maybe, a chatbot answers FAQs.

ALternative: A resident begins their interaction with their government with an inquiry: “I need affordable housing in Cincinnati before December. What are my options?” AI analyzes such queries and recognizes a pattern of

community needs, like heating complaints rising in a neighborhood in the winter or elderly residents who are having challenges accessing housing.

- AI interprets across zoning code, housing programs, wait lists, and deadlines.
- The system returns not just an answer, but a pathway: steps, timelines, forms, contacts.
- The system lets residents ask follow-up questions in their own words —“What if I’m a single parent?” or “What if I make \$35K/year?”—and it adapts to help an individual.
- Information that was used to generate an individual response now motivates a coordinated strategy.

To achieve the ALTERNative, we recommend the following:

- Build shared civic infrastructure through standing up a public AI sandbox where residents and staff use low-code/no-code tools to translate policies, budgets, and services into plain language. Low-code and no-code AI tools give people the ability to design AI workflows through graphical interfaces. These tools are proliferating, driving quality and usability up and costs down. This will make it easier for individuals and teams in smaller municipalities and community organizations to implement AI solutions without extensive technical resources. This trend particularly benefits resource-constrained local governments that previously couldn’t access enterprise AI solutions.
- Strengthen civic capacity by training civic workers in context engineering to ensure translations are accurate, relevant, and accessible across languages and literacy levels. This context engineering is a shift away from prompt engineering to involve systematically building comprehensive knowledge architectures that give AI systems deep understanding of the problem they are being asked to solve. Instead of crafting clever questions, civic organizations will need to build institutional memory systems and data architectures that give AI the same situational awareness that civic workers rely on.

In the future, platforms offering pre-built models for common civic use cases—like permit processing, service request routing, and community feedback analysis—will lower barriers to entry.

In **Boston**, we wanted to explore if generative AI could change the nature of interactions between government decision makers and a local community. The RethinkAI team, including Chris Le Dantec from Northeastern University, partnered with Talbot Norfolk Triangle Neighborhood United, a community organization in the Dorchester neighborhood of Boston. The leadership of that group was concerned that they had become increasingly dependent on the city for access to and analysis of data about them. They wondered what it would be like to own their own data resource so that they might level the playing field in local decision-making.

We ended up creating a local large language model, or LLM, called “On the Porch.” It is trained on structured data such as 311, traffic violations, 911 calls, and unstructured data such as planning documents and transcripts of local meetings. The result is a conversational tool that anyone in the community can use to inquire about what’s going on in their neighborhood and share things they might be concerned about. After a conversation, the bot provides a high-level summary, complete with any resources mentioned for follow-up.

The tool demonstrates the possibility of renewed **listening**. In this case, the community has access to the same data and analysis tools as the city. This level playing field can clear the way for the city government to truly listen to what the community has to say.

Trust

AI is about increasing accountability, not skirting it.

Trust in a nutshell: In order for people to trust government, they need to trust the tools that government uses to deliver services and communicate. The use of AI in the civic sector needs to be oriented around using trustworthy data and being held accountable to delivering trustworthy insights from that data. This involves making sure that users are involved, or at least aware, of the methods of collection, analysis, and use of data.

Today, much of the activity in government organizations is invisible to residents. This is especially true for digital activity and increasingly true for AI. This creates a paradox: More data does not equal more trust. AI can either reinforce the gap or open up new forms of accountability. More data than ever

before is being used to inform government decision-making. Using predictive AI tools, governments are engaged in sophisticated analysis of often diverse datasets. This can lead to better decisions, but it doesn't lead to residents' trust in those decisions.

A key distinction between this and what we had in the past is building a civic tech movement on transparency. These ongoing progress checks built necessary awareness. Moving into this next era of trust relies on community members experiencing how their voices, values, and vulnerabilities influence collective decisions, alongside open data.

Example: The City of Boston's "Council Roll Call Summaries" pilot translates complex legislative votes into plain language summaries. With the help of AI, residents can see how their council members voted without reviewing transcripts or listening to meetings.

ALternative: Residents can see how the council voted, but what if they could also see resident reactions from social media or an online comment platform, alongside official votes? Government officials could see these patterns synthesized in real time and be able to prioritize their agenda accordingly. This could mark a shift from an awareness of votes to a reshaping of decision-making structure. Trust is cultivated not just when residents observe what's happening, but when there is a move from one-way transparency to two-way accountability.

To achieve the ALternative, we recommend the following:

- Develop enabling policies that facilitate pilots and transparency, while also creating opportunities for experimentation to enable rapid resident response.
- Incubate civic data trusts or community-controlled data infrastructure (e.g., local data trusts and sovereign AI models trained on local data).
- Evaluate for trustworthiness, not just efficiency. The earliest iterations of the digital era focused on efficiency measures and productivity gains. But the real currency in the public sector is trust—and we need to develop strong measures around its success.
- Integrate broad signals of public impact, like environmental stress, service usage, or safety outcomes, alongside individualized resident input. Enter into public compacts that are visible agreements between government, residents, and partners that spell out shared goals, timelines, and accountability measures.

In **New York**, we saw an opportunity to leverage generative AI to expand the visibility and resolution of funding needs for community-based initiatives. We partnered with CitizensNYC, one of the nation’s first micro-funding organizations, whose leaders wanted to analyze a large corpus of funding proposals collected over the last decade. Could an LLM generate thematic insights about shared needs across proposals, how issues are framed, and how that framing evolves over time? For CitizensNYC, this use of AI presented an opportunity to expand their capabilities and increase the **trust** between the organization, communities across the city, and the city government.

This pilot involves building a proposal analysis dashboard for CitizensNYC that will harness more than 5,000 past proposals and New York City open datasets to analyze and visualize the content and context of new proposals. Working with a broad cross section of the CitizensNYC team and leadership, the tool will support new use cases for both proposal review and communication. For instance, city data layers on flood risk can be used to understand the context of a proposal related to climate hazard risk in a particular neighborhood—and this validation in turn can reveal past proposals related to this need locally and citywide. There are several applications of this approach that will improve CitizensNYC grantmaking, which will increase the **trust** communities and the city government have in the process.

The Leadership We Need

In order to realize an ALternative perspective, we need good leadership. As former Louisville Mayor Greg Fischer put it, “These AI tools can really help build social muscle. And social muscle is the kind of strength that you need to get through adversity.”⁴⁴ That bold approach should guide this important moment and help us use AI for its best and highest applications—not just regulating its risks, but also enabling its potential to deliver real public value.

Government leaders must establish enabling conditions. We need a vision that allows us to use AI for its best and highest applications, not just regulating its risks. This means designing accountability structures that measure trustworthiness, not just data; realigning budgets; anticipating new resident needs and amplified demand and rapidly responding to those shifts in service; and translating AI’s promise into real public value.

Philanthropy could continue to shape and fund what governments can't risk: field-defining networks, bold prototypes, and practical tools that can be scaled across cities and states. This creative capital can support experiments and connect technical tools with civic priorities so that the field can learn faster, space for vulnerabilities can be acknowledged, and progress isn't defined solely by the private sector.

Universities can anchor this work with research, training, data capacity, and long-term evaluation. While basic training is available, what's needed for cities is beyond the basic vocabulary of LLMs and prompt engineering. Instead, universities can lead in practical application and measurement as well as evaluations that extend beyond political cycles, playing the crucial role as a trusted evaluator of civic data and outcomes.

Community organizations and nonprofits must ground AI in lived experience and the tangible realities of public service. They not only mobilize residents; they are also the key agents responsible for localizing data, keeping culture, and ensuring that the use of AI reflects the needs and priorities of their communities. They hold governments accountable to real resident experience, not just machine-recorded records.

Coalitions, like RethinkAI, must continue to play the role of convenor, chronicler, and practical advisor to civic actors. They need to be committed to field-building alongside communities and other sectors.

The ALT leadership ecosystem is one in which government serves as an enabler, philanthropies buffer risk, universities evolve our understanding, and nonprofits ground us in real-world experience. Together, with these four sectors in the civic space playing their leadership role, AI can move beyond efficiency and build a model that residents truly need: one that fosters adaptation, listening, and trust.

Appendix

Credits

This report was written by Neil Kleiman, Eric Gordon, and Mai-Ling Garcia. Ms. Garcia’s contributions to this work were made in her capacity as a Member of RethinkAI’s Civic AI Advisory Trust. Additional research and writing were provided by Lilian Coral, Anthony Townsend, Matt Steinberg, Jess Silverman, and Brenna Berman. This report was managed by Rebecca Ierardo.

Supporting Institutions

Grant funding was generously provided by the Kresge Foundation and the Chan Zuckerberg Initiative. New America houses and is one of the founding partners of the RethinkAI initiative and provides editorial and communications support.

RethinkAI

RethinkAI is a dynamic, cross-sector coalition of scholars, technologists, and civic practitioners united by a bold mission: to harness the power of AI to transform local governance and community development across the United States. Through original research and hands-on pilot programs with city and community leaders, we uncover what works—and then turn those insights into actionable policy recommendations and practical AI guidance for civic leaders nationwide. What began as an informal virtual working group in the spring of 2023 has since evolved into a formal collaborative committed to shepherding a transformation in governance strategy and tactics with AI.

Methods

This report draws on a series of semi-structured interviews and supporting desk research. We conducted more than 40 interviews with mayors, former local government executives, academics, funders, and civic technology practitioners. Each interview followed a structured guide but allowed flexibility to explore themes in depth. Core questions focused on:

- interviewees’ background and experience with digital services, emerging technologies, and AI;

- opportunities and risks of AI in public decision-making and service delivery;
- institutional challenges such as staffing, funding, procurement, and capacity;
- the role of partnerships with universities, philanthropy, and vendors; and
- what resources, training, or frameworks would help local governments make better use of AI.

All interviews were conducted by Neil Kleiman and Mai-Ling Garcia. Interviews lasted 45–60 minutes and were captured in detailed notes rather than recordings to encourage candid responses. We supplemented these interviews with a review of policy documents, pilots, and secondary research on AI adoption trends. Insights were synthesized across sources to identify common patterns, gaps, and opportunities, which informed the findings and recommendations in this report.

Interviewees

Special thanks to subject matter expert interviewees, including those below. The views in the report are solely of the authors and not of interviewees or their affiliates:

- **Brendan Babb**, Chief Innovation Officer, Municipality of Anchorage, Alaska
- **Grace Simrall**, former Chief of Innovation and Technology, Louisville Metro Government, Kentucky
- **Libby Schaaf**, former Mayor of Oakland, California (2015–2022); Professional Faculty, University of California, Berkeley, and Northeastern University
- **Greg Fischer**, former Mayor of Louisville, Kentucky (2011–2022); former United States Conference of Mayors President (2020–2021)
- **Santiago Garces**, Chief Innovation Officer, City of Boston, Massachusetts

- **Erin McKinney**, Principal Policy Counsel for State, Local, and Education Public Policy, Amazon Web Services
- **Albert Gehami**, AI and Privacy Officer, City of San Jose, California
- **Kirsten Wyatt**, Senior Director, Beeck Center for Social Impact + Innovation, Georgetown University
- **Deirdre K. Mulligan**, Professor, School of Information, University of California, Berkeley
- **Henry E. Brady**, Monroe Deutsch Professor of Political Science and Public Policy, Travers Department of Political Science and Goldman School of Public Policy, University of California, Berkeley
- **Ramayya Krishnan**, W. W. Cooper and Ruth F. Cooper Professor of Management Science and Information Systems at Heinz College and the Department of Engineering and Public Policy, Carnegie Mellon University
- **Michael J. Holland**, Vice Chancellor for Science Policy and Research Strategies, University of Pittsburgh
- **Jing Liu**, Executive Director, Michigan Institute for Data and AI in Society
- **Chris Mentzel**, Executive Director, Stanford Data Science
- **Kate Weber**, Executive Director, Public Exchange, Dornsife College of Letters, Arts and Sciences, University of Southern California
- **Nicholas Mattei**, Associate Professor of Computer Science, Tulane University; Co-Director, Tulane Center of Excellence for Community Engaged Artificial Intelligence
- **Eunice E. Santos**, Professor, School of Information Sciences, University of Illinois at Urbana-Champaign
- **Ahu Yildirmaz**, President and Chief Executive Officer, Coleridge Initiative
- **Dongxiao Zhu**, Professor of Computer Science and Founding Director, Wayne AI Research Initiative, Wayne State University

- **Elizabeth Langdon-Gray**, Executive Director, Harvard Data Science Initiative
- **Supratik Mukhopadhyay**, Full Professor, Louisiana State University Center for Computation and Technology
- **William Peduto**, former Mayor of Pittsburgh, Pennsylvania (2014-2022)
- **Kat Hartman**, Chief Data Officer, City of Detroit, Michigan
- **Tad McGalliard**, Managing Director for Research, Development, and Technical Assistance, International City/County Management Association
- **Harrison MacRae**, Director of Emerging Technologies, Commonwealth of Pennsylvania
- **Erin Dalton**, Director, Allegheny County Department of Human Services, Pennsylvania
- **Sarah Williams**, Director, Leventhal Center for Advanced Urbanism and the Civic Data Design Lab; Associate Professor of Technology and Urban Planning at MIT
- **Stephen Goldsmith**, Derek Bok Professor of the Practice of Urban Policy and Director of the Data-Smart Cities Program, Bloomberg Center for Cities, at Harvard University
- **Lane Dilg**, Founder and Principal, Apeiron LLC
- **Jean Westrick**, Executive Director, Technology Association of Grantmakers
- **Kyle Cranmer**, Professor, University of Wisconsin–Madison
- **Dan O’ Brien**, Professor of Public Policy and Urban Affairs and Criminology and Criminal Justice, Northeastern University; Director, Boston Area Research Initiative
- **Aimee Sprung**, Director, Government Affairs, Microsoft
- **David W. Burns**, Assistant Executive Director, United States Conference of Mayors

- **Christopher Jordan**, Senior Specialist, Urban Innovation, National League of Cities
- **Debra Lam**, Founding Executive Director, Partnership for Inclusive Innovation
- **Meredith Lee**, Head of Strategic Partnerships and Chief Technical Advisor to the Dean, College of Computing, Data Science, and Society, University of California, Berkeley
- **Azer Bestavros**, Warren Distinguished Professor of Computer Science and Associate Provost for Computing and Data Sciences, Boston University
- **Aron Culotta**, Professor of Computer Science, Tulane University; Director, Tulane Center of Excellence for Community Engaged Artificial Intelligence
- **Alex Jutca**, Director of Analytics, Technology, and Planning, Allegheny County Department of Human Services, Pennsylvania
- **Sujeet Rao**, Health and Wellbeing Practice Director, Public Exchange, Dornsife College of Letters, Arts and Sciences, University of Southern California
- **Kate Burns**, Executive Director, MetroLab
- **Steven Wray**, Executive Director, Block Center for Technology and Society, Heinz College of Information Systems and Public Policy, Carnegie Mellon University
- **Nishant Shah**, Senior Advisor for Responsible AI, State of Maryland
- **Lauren Lockwood**, CEO & Co-Founder of Bloom Works Public Benefit LLC

Notes

- 1 Beeck Center for Social Impact + Innovation, “AI Legislation Database,” Digital Government Hub, Georgetown University, 2025, <https://digitalgovernmenthub.org/publications/digital-transformation-policy-scan/#ai-legislation-database>.
- 2 Beeck Center, “AI Legislation Database,” <https://digitalgovernmenthub.org/publications/digital-transformation-policy-scan/#ai-legislation-database>; National Conference of State Legislatures, “Artificial Intelligence 2025 Legislation,” July 10, 2025, <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2025-legislation>.
- 3 InnovateUS, <https://innovate-us.org/>; “Implementing Responsible AI in Government: Colorado’s Pilot Program,” InnovateUS, February 5, 2025, <https://innovate-us.org/blog/implementing-responsible-ai-in-government-colorados-pilot-program>.
- 4 Beeck Center, “AI Legislation Database,” <https://digitalgovernmenthub.org/publications/digital-transformation-policy-scan/#ai-legislation-database>.
- 5 The few times in which anything close to as much tech legislation was advanced include 200 bills focused on data privacy in 2020, and the brief legislative spurt following President Joseph Biden’s 2021 Infrastructure Investment and Jobs Act, which ensured investments in local broadband were effectively utilized at the state level.
- 6 Interview with state administrator, September 22, 2025.
- 7 “From Red Tape to Results: What Virginia’s AI-Powered Reform Could Mean for Texas,” Texas Policy Research, July 28, 2025, <https://www.texaspolicyresearch.com/from-red-tape-to-results-what-virginias-ai-powered-reform-could-mean-for-texas/>.
- 8 InnovateUS, <https://innovate-us.org/>.
- 9 “Implementing Responsible AI in Government,” <https://innovate-us.org/blog/implementing-responsible-ai-in-government-colorados-pilot-program>.
- 10 InnovateUS, <https://innovate-us.org/>; “State of Arizona Implements Employee Gen AI Training,” Arizona Department of Administration, February 24, 2025, <https://doa.az.gov/news/state-arizona-implements-employee-gen-ai-training>.
- 11 Stuart D. Levi et al., “Utah Becomes First State To Enact AI-Centric Consumer Protection Law,” Skadden, April 5, 2024, <https://www.skadden.com/insights/publications/2024/04/utah-becomes-first-state>; “Texas Signs Responsible AI Governance Act Into Law,” Latham & Watkins, June 23, 2025, <https://www.lw.com/en/insights/texas-signs-responsible-ai-governance-act-into-law>; “Delaware Launches Bold AI Sandbox Initiative,” Delaware Prosperity Partnership, July 23, 2025, <https://www.choosedelaware.com/in-the-news/delaware-unveils-ai-sandbox-initiative/>.
- 12 Governor Gavin Newsom, “Governor Newsom Deploys First-in-the-Nation GenAI Technologies to Improve Efficiency in State Government,” State of California, <https://www.gov.ca.gov/2025/04/29/governor-newsom-deploys-first-in-the-nation-genai-technologies-to-improve-efficiency-in-state-government>; Julia Edinger, “New Jersey Advances AI Through an Economic Development Lens,” *Government Technology*, April 21, 2025, <https://www.govtech.com/artificial-intelligence/new-jersey-advances-ai-through-an-economic-development-lens>; Jordan Anderson, “The Colorado AI Act: What You Need to Know,” Baker Tilly, September 10, 2024, <https://www.bakertilly.com/insights/the-colorado-ai-act-what-you-need-to-know>; Justin Sweitzer, “Shapiro Says AI Has ‘Real Promise’ as He Unveils Pilot Program Results,” *City & State Pennsylvania*, March 21, 2025, <https://www.cityandstatepa.com/policy/2025/03/shapiro-says-ai-has-real-promise-he-unveils-pilot-program-results/403965/>.
- 13 Sophia Fox-Sowell, “New Washington Bill Would Let State Workers Influence How Agencies Use AI,”

StateScoop, February 13, 2025, <https://statescoop.com/washington-bill-labor-union-ai-2025/>; News Staff, “Maryland Submits AI Strategy, Guide to General Assembly,” *Government Technology*, January 14, 2025, <https://www.govtech.com/artificial-intelligence/maryland-submits-ai-strategy-guide-to-general-assembly/>; Nikhil Deshpande, “State of Georgia: AI Roadmap and Governance Framework,” Georgia Office of Artificial Intelligence, <https://ai.georgia.gov/blog/2025-02-25/state-georgia-ai-roadmap-and-governance-framework>.

14 One-U Responsible AI Initiative, “Responsible AI Community Consortium,” University of Utah, <https://rai.utah.edu/opportunities/community-consortium/>; Matthew Ferraro and Anna Z. Saber, “Texas Just Created A New Model for State AI Regulation,” *Tech Policy Press*, July 17, 2025, <https://www.techpolicy.press/texas-just-created-a-new-model-for-state-ai-regulation/>; Skip Descant, “Delaware AI Commission Approves Creating Agentic AI Sandbox,” *Government Technology*, July 28, 2025, <https://www.govtech.com/artificial-intelligence/delaware-ai-commission-approves-creating-agentic-ai-sandbox/>; Khari Johnson, “New Washington Bill Would Let State Workers Influence How Agencies Use AI,” *Cal Matters*, March 21, 2024, <https://calmatters.org/economy/technology/2024/03/california-ai-purchasing-guidelines/>; Chris Teale, “States Vie for Leadership Role on AI,” *Route Fifty*, January 17, 2024, <https://www.route-fifty.com/emerging-tech/2024/01/states-vie-leadership-role-ai-393406/>; Keely Quinlan, “Trust, Innovation Central to Colorado’s New AI Guidelines, Says State Data Chief,” *StateScoop*, September 9, 2024, <https://statescoop.com/ai-policy-colorado-amy-bhikha-2024/>; Sweitzer, “Shapiro Says AI Has ‘Real Promise,’” <https://www.cityandstatepa.com/policy/2025/03/shapiro-says-ai-has-real-promise-he-unveils-pilot-program-results/403965/>.

15 Colin Wood, “Washington Governor Signs AI Order Plotting Yearlong Policy Path,” *StateScoop*, January 30, 2024, <https://statescoop.com/washington-governor-signs-ai-order-plotting->

[yearlong-policy-path/](https://statescoop.com/washington-governor-signs-ai-order-plotting-yearlong-policy-path/); Department of Information Technology, “Office of AI Enablement,” State of Maryland, <https://doit.maryland.gov/About-DoIT/Offices/Pages/officeofAIenablement.aspx>; “State of Georgia: AI Roadmap and Governance Framework,” <https://ai.georgia.gov/blog/2025-02-25/state-georgia-ai-roadmap-and-governance-framework/>; Office of the New York State Comptroller, *New York State Artificial Intelligence Governance* (New York Division of State Government Accountability, 2025), <https://www.osc.ny.gov/files/state-agencies/audits/pdf/sga-2025-23s50.pdf>.

16 “OK HB1916 | 2025 | Regular Session,” LegiScan, <https://legiscan.com/OK/bill/HB1916/2025/>; Angela Eichhorst, “What Are the New AI Laws in Connecticut?,” *CT Mirror*, July 14, 2025, <https://ctmirror.org/2025/07/14/ct-ai-laws-session/>; “MS HB1535 | 2025 | Regular Session,” LegiScan, <https://legiscan.com/MS/bill/HB1535/2025/>; “MO HB1606,” Bill Track 50, <https://www.billtrack50.com/billdetail/1850814>.

17 Interview with Aimee Sprung, July 28, 2025.

18 Interview with local government professional association leader, August 7, 2025.

19 Interview with state leader, September 19, 2025.

20 Interview with local government professional association leader, January 27, 2025.

21 Interview with local official, June 5, 2025.

22 Interview with university professor, May 29, 2025.

23 Interview with local government professional association leader, August 7, 2025.

24 Interview with Santi Garces, February 5, 2025.

25 Interview with Kat Hartman, August 8, 2025.

26 Interview with Stephen Goldsmith, March 26, 2025.

- 27 Interview with Lane Dilg, August 7, 2025.
- 28 Interview with Albert Gehami, January 22, 2025.
- 29 Interview with Erin McKinney, August 12, 2025.
- 30 Interview with state administrator, September 22, 2025
- 31 International City/County Management Association (ICMA), *Artificial Intelligence in Local Government* (ICMA, 2024), <https://icma.org/documents/icma-survey-research-artificial-intelligence-local-government>.
- 32 Interview with Tad McGalliard, September 19, 2025.
- 33 Interview with Jean Westrick, August 21, 2025.
- 34 Thalia Beaty, “Funders Commit \$1 Billion Toward Developing AI Tools for Frontline Workers,” *Chronicle of Philanthropy*, July 17, 2025, <https://www.philanthropy.com/article/funders-commit-1-billion-toward-developing-ai-tools-for-frontline-workers>; Natasha Singer, “Microsoft Pledges \$4 Billion Toward A.I. Education,” *New York Times*, July 9, 2025, <https://www.nytimes.com/2025/07/09/business/microsoft-ai-education.html>.
- 35 Bloomberg Center for Government Excellence, “Welcome to CityAI Connect,” John Hopkins University, <https://cityaiconnect.jhu.edu/#welcome>; Bloomberg Center for Public Innovation, “Bloomberg Philanthropies City Data Alliance,” John Hopkins, <https://bloombergcities.jhu.edu/program/data-alliance>; “AI Poverty Challenge,” Robin Hood, <https://www.aipovertychallenge.org/>.
- 36 Interview with Kate Burns, October 28, 2024.
- 37 Jonathan Shaw, “Chan Zuckerberg Commits \$500 Million to Harvard Neuroscience and AI Institute,” *Harvard Magazine*, December 7, 2021, <https://www.harvardmagazine.com/2021/12/chan-zuckerberg-natural-and-artificial-intelligence>;
- Michigan Engineering, “U-Michigan Announces Most Advanced AI Research Complex with Historic Los Alamos Alliance,” *Michigan Engineering News*, February 3, 2025, <https://news.engin.umich.edu/2025/02/u-michigan-announces-most-advanced-ai-research-complex-with-historic-los-alamos-alliance/>;
- Governor Kathy Hochul, “Governor Hochul Announces \$20 Million Public-Private Investment to Advance Artificial Intelligence Goals,” New York State, October 16, 2023, <https://www.governor.ny.gov/news/governor-hochul-announces-20-million-public-private-investment-advance-artificial-intelligence>;
- Rachel Leven, “UC Berkeley College of Computing, Data Science, and Society Established,” University of California, Berkeley, May 18, 2023, <https://cdss.berkeley.edu/news/uc-berkeley-college-computing-data-science-and-society-established>.
- 38 “Center for Community-Engaged Artificial Intelligence,” Tulane University, <https://sse.tulane.edu/cs/ceai>.
- 39 Burnes Center for Social Change, “AI for Impact Program,” Northeastern University, <https://burnes.northeastern.edu/ai-for-impact-coop/>.
- 40 Interview with Deb Lam, July 25, 2025.
- 41 “Government AI Coalition,” City of San Jose, <https://www.sanjoseca.gov/your-government/departments-offices/information-technology/ai-reviews-algorithm-register/govai-coalition>.
- 42 Teaching + Learning Lab, “Worked Examples,” Massachusetts Institute of Technology, <https://tll.mit.edu/teaching-resources/how-people-learn/worked-examples/>.
- 43 Eric Gordon, *How Institutions Listen: AI, Civic Data, and the Path to Public Trust* (MIT Press, forthcoming, 2026).
- 44 Interview with Greg Fischer, September 25, 2024.



This report carries a Creative Commons Attribution 4.0 International license, which permits re-use of New America content when proper attribution is provided. This means you are free to share and adapt New America’s work, or include our content in derivative works, under the following conditions:

- **Attribution.** You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

For the full legal code of this Creative Commons license, please visit creativecommons.org.

If you have any questions about citing or reusing New America content, please visit www.newamerica.org.

All photos in this report are supplied by, and licensed to, [shutterstock.com](https://www.shutterstock.com) unless otherwise stated. Photos from federal government sources are used under section 105 of the Copyright Act.