



Shaping the AI Revolution: Key Policy Options and Principles for State and Local Leaders in 2024

Dean Ball

October 2024

When Thomas Edison first demonstrated his electric light bulb in 1879, the bulbs lasted as little as a few hours or as long as a few days; gas lamps, which Edison was seeking to upend, lasted months. No one understood how electricity worked with any depth; even something as basic as the electron would not be discovered for nearly two decades. Neither electricity generation nor distribution infrastructure existed anywhere save a few research facilities.

Even two decades after that initial demonstration, electricity was useful for very little beyond urban infrastructure like streetcars and streetlamps. Fewer than 5 percent of Americans had electricity in their homes. Electricity accounted for a similarly minor share of America's industrial power generation.

By 1930, though, most American factories, offices, and homes were electrified.¹ This revolutionary new energy undergirded much of American life, powering the industry that would ultimately help the Allies triumph in World War II. Nobody could have foreseen what would unfold during those decades. Nobody in 1879 could have guessed that electric streetcars would transform the layout of American cities. Nobody could have guessed that electric motors would lead to assembly lines, with the radical changes to the labor market that entailed. Nobody could have guessed that electric appliances would save women and girls countless hours in daily household work, ultimately freeing them to participate more actively in the labor market and political life.

Today society stands on the threshold of a similarly potent technological revolution: artificial intelligence (AI). Navigating safely and swiftly from the novelty of AI's early practical uses to its seamless entrenchment in daily life—as unremarkable as turning on a light switch—will be the principal task of AI policymakers and civil society actors in the years to come.

The state played an important role in the second industrial revolution. That role did not, by and large, involve regulating light bulbs, capping the maximum wattage of power-generating facilities, or making electricity producers legally liable for consumer misuse of electric appliances—a rough analogy to many of the AI-related laws that are proposed today. Indeed, much of government’s role was *helping* the diffusion of electricity by investing in electric power generation and infrastructure.

The journey from the first demonstration of the light bulb to the humble power outlet in every room of every home is not as glamorous as the story of electricity’s invention, yet it is just as much of a miracle—a triumph of American industry, engineering, government, and science. The fundamental question AI poses is: Can the United States do it again?

This policy brief argues that it can. What follows is a set of principles, informed by technological history and present-day technological realities, to help policymakers navigate the AI revolution. After that, there is a brief guide for initial steps state and local government leaders can take to adopt and regulate AI.

Principles

AI will create unexpected challenges that necessitate action across many domains. As policymakers consider legislation and other actions, the following four principles will help guide prudent decisions and increase institutional capacity.

1. Build State Agility. In the near term, policymakers should focus on reconfiguring state agencies to adapt rapidly and effectively to the many changes AI could bring. The state should avoid placing too much importance on any one risk without ample evidence. Instead, it should prepare to respond nimbly to variable emergent risks while shoring up its enforcement of existing laws. State agility is a function of the following:

- **Resources:** State actors and regulators must have the budgetary, labor, and physical capital needed to implement programs and enforce and implement the law.
- **Bureaucratic regulation:** Excessive, unclear, complicated, or contradictory rules can needlessly bind decisions and slow action. Managing AI uncertainty will require flexible rules that enable agencies to act and pivot when confronted with unexpected challenges.

2. Promote Engineering-First Solutions. No matter how well-crafted it is, a policy will fail if it cannot be practically implemented. Before considering new regulation, decision-makers must first consider whether private sector or state-sponsored engineering fixes can resolve challenges. When regulation is needed, decision-makers must also consider whether regulatory agencies have the necessary technical tools, such as forensic tools and techniques, to implement and enforce the law.

3. Center AI Diffusion. A technological invention, no matter how promising, matters little if no one uses it. The 2007 iPhone was a marvel when it debuted, but it did not truly change society until millions of developers could experiment with it and create third-party applications. Diffusion is how a new technology transitions from being a novelty to being an engine of economic productivity.

Through widespread use, society discovers a technology's limitations, trade-offs, and dangers. Policymakers cannot mitigate AI's downsides without knowing what they are, and they can't acquire that knowledge without allowing extensive practical use of the technology.

Diffusion is an inherently decentralized process, and there are limits to what public policy can encourage. However, the wrong forms of regulation can stymie diffusion. For example, some have proposed that AI developers ensure models are certified free of risk before deployment—an impossibility given the limits of laboratory testing.² Regulations of this kind are likely to deter the investment and experimentation needed to arrive at the balance of productivity and safety.

4. Prioritize Regulating Conduct, Not Models. Because AI models and the means of producing them can change so quickly, policies that focus overmuch on current technology are likely to be outdated soon. Society's collective preferences about what should constitute illegal behavior, however, evolve far more slowly. Fraud, assault, theft, and murder have all been considered illicit for centuries. While the means may change, the desire to police such conduct does not.

To remain flexible in the face of rapid change, policymakers should default to regulating conduct, not AI models. Policies should target illicit conduct in a technology-neutral way. Policymakers should also invest in measures that make society less vulnerable to potential threats.

State and Local AI Policy Menu

State and local governments are at the heart of the American project. They perform the bulk of public services that affect Americans' day-to-day lives, from education to road maintenance to public safety. The COVID-19 pandemic was a reminder of the importance of state and local governments: most of the strategies for combatting the virus fell to them rather than to the federal government. State and local governments have also played a vital role in the diffusion of technology throughout American history, enabling the construction of vast telegraph, telephone, and internet infrastructures, the development of road systems for the automobile, and the creation of the electric grid.³

It should come as no surprise, then, that state and local governments will play a leading role in the development and diffusion of AI throughout the economy and the broader world. To do this, state legislators and agencies should

1. Update and harmonize existing law to counter AI risks and **ensure regulatory clarity** for residents and businesses,
2. **Create clear standards** for government use of AI, and
3. **Lead the way in AI deployment** by using AI to improve internal processes and operations.

The policy options outlined below, organized under the three categories above, are steps that policymakers can consider in pursuit of these ends.

Ensure regulatory clarity

Many of the imagined risks from AI relate to conduct that is already unlawful. Cybercrime, identity theft, and defamation, for example, are covered by state criminal statutes. There may, however, be ways in which existing law would benefit from modifications or clarifications in light of AI.

Option 1: Review and modify existing state law to accommodate AI use cases

States should direct their attorneys general or other relevant officials to conduct reviews of state law and make appropriate recommendations. The goals of this review process should be to

1. Ensure that illicit use cases of AI are covered by existing state law,
2. Ensure that existing law does not create unnecessary or perverse barriers to AI adoption by individuals, businesses, and government, and
3. Make recommendations for updating existing law to better incorporate AI.

Some AI risks, such as the ability to spread disinformation at scale, are difficult to counter with law. This is because it can be hard to tell what constitutes disinformation, particularly regarding an ongoing development (for example, a terrorist attack). Furthermore, the First Amendment limits the government's ability to police speech, including, often, speech that the speaker knows to be false.⁴

Thus, on both constitutional and practical grounds, policing AI-enabled disinformation on social media will likely be best left to the social media platforms. However, recently passed state laws that forbid "viewpoint discrimination" by social media platforms may have a perverse effect. By forbidding companies to remove political speech, it may create an incentive for bad actors to attempt to sway political discussions on social media using AI. This is an example of the unexpected ways in which existing laws, passed with good intentions, might have unintended consequences as AI systems become more capable.

There are limited instances of truly novel AI-enabled conduct that state and federal policymakers may seek to stop. For example, laws requiring that AI systems proactively identify themselves as such are reasonable and do not impose substantial burdens on developers, users, or

businesses. Indeed, such laws likely *help* AI diffusion because they clear up confusion. However, requiring people who use AI model outputs for their own communication to disclose that they used AI is likely overreach. For example, the law does not require people to disclose if they used Photoshop to edit an image they are sharing, because doing so would be an onerous burden on free expression.

Create clear standards for government AI use

AI is a general-purpose technology. Though large language models (LLMs) have emerged as the most prominent use case for AI—and they are indeed powerful—policymakers and agency practitioners should think beyond LLMs when evaluating ways that AI can improve government services. For example, multimodal AI models (incorporating language, images, and sometimes audio and video) are increasingly common and could have applications for everything from infrastructure maintenance to public safety.⁵ Nor are generative AI models the only kind that may be of use to government agencies. AI can, for example, be used to increase the energy efficiency of public facilities by harmonizing heating, air conditioning, and other energy-intensive aspects of a large building.

Because AI is, at its core, an information technology, guidelines and regulations for its use are likely to be covered by existing statewide and agency-specific IT and cybersecurity policies. These policies should be evaluated to ensure that reasonable uses of AI can comply with them.

Beyond these existing policies, agencies will benefit from reasonable guidelines, rules, and best practices specific to AI applications. State and local government leaders should consider the following steps.

Option 2: Hire chief AI officers within each agency

The Biden Executive Order on AI (Executive Order 14110) and the subsequent AI implementation memo published by the Office of Management and Budget in 2024 direct each federal agency to appoint its own chief AI officer.⁶ State governments should take this approach as well, ensuring that the officials with this title advocate for and oversee the adoption of AI throughout each agency. Chief AI officers should not become internal “soft regulators” of AI uses (status quo bias within agencies against novel uses of technology will likely serve as its own form of regulation). Instead, chief AI officers should primarily be responsible for educating agency staff on the basics of AI, identifying potential use cases for AI, and driving the application of the technology to agency needs and workflows.

Option 3: Use AI to augment the work of government employees, not replace them

Employees of state agencies will likely be concerned about losing their jobs to AI. However, most state agencies are understaffed.⁷ If AI is implemented well, existing employees and the public will benefit from the added low-cost cognitive labor AI delivers. One of the foremost priorities

for chief AI officers, then, should be to identify the areas where state agencies have the largest backlogs or the severest understaffing. These may be the areas where AI can deliver the greatest benefit in a short time.

Option 4: Create a statewide AI implementation board

Chief AI officers should serve on an interagency AI implementation board, chaired by the governor (or a statewide officer reporting to the governor), designed to share ideas, best practices, and common pitfalls. This task force should publish regular public reports on how the government is using AI to improve services, the realized or expected improvements AI is delivering, and any insights that might be useful to a general audience.

Option 5: Share ideas, best practices, problems, and code

The AI community has historically reaped the benefits of widespread collaboration; indeed, a robust culture of information-sharing has driven the rapid progress in software development more broadly. Because of that, agencies—led by interagency task forces—should aim to share as much of their code, models, and other AI-related technical resources as possible with one another and, where feasible, with the public.⁸

Option 6: Create a culture of experimentation and encourage an initial focus on low-hanging fruit

Because of the wide range of potential AI use cases, policies should permit agencies to experiment, rather than being overly prescriptive or requiring time-consuming impact assessments. Many commentators and policy analysts rightfully fear the use of AI-enabled automated decision-making by governments. While the benefits and cost of such use cases can and should be debated, most agencies likely have lower-hanging fruit they can pick more readily. LLMs can accelerate the review and production of routine paperwork, reporting, analysis, and similar tasks. By using AI systems for these lower-risk applications, agencies will develop internal competency in assessing the weaknesses, strengths, and nuances of LLMs. This knowledge will be useful in considering higher-risk, citizen-facing AI uses in the future.

Option 7: Experiment with cost-prohibitive frontier AI when smaller models fail

Because AI capabilities improve rapidly while costs decline, agencies can experiment with frontier language and multimodal models for use cases when cheaper and less capable models have failed. While a frontier model may be too expensive for agency-wide or other large-scale use at a given task, knowing what models at the frontier can do is valuable. Today's frontier models will be significantly cheaper in a short period, often under a year. Similarly, use cases that do not work with existing frontier models may well begin to work when the next generation of frontier models is available. Having a flexible, experimental approach will allow state agencies to keep pace with improvements as they come online.

Option 8: Avoid locking in a single model or vendor to take advantage of AI’s continuous, rapid improvements

Agencies should avoid locking in one model or vendor whenever possible. Implementations of AI systems should aim to be as modular as is feasible so that new models from different vendors can be easily “dropped” into existing AI applications and workflows.

Option 9: Start by prioritizing internal use cases over external ones

Agency leaders should prioritize internal uses over external use cases in the early stages of their AI deployment. AI “hallucinations” of false information have declined in recent models and with techniques like retrieval-augmented generation (RAG) and extended context windows, both of which allow the model to be grounded in a specific document or set of documents.⁹ However, hallucinations and other problems persist, meaning that agencies should proceed carefully when opening AI-enabled services to the public—particularly essential public services.¹⁰

Leading the way in AI deployment

The above policy options will create a robust yet flexible framework for government staff to deploy AI into many different areas of public service. It is hard, however, for individuals, businesses, and governments alike to imagine the full range of possibilities with such a broadly applicable technology. Here are additional potential uses of AI within state government entities.

Option 10: Maintain legacy code bases

During the COVID-19 pandemic, state government welfare services were crippled by outages in the software used to administer benefits.¹¹ This was due in large part to reliance on legacy code, particularly Common Business Oriented Language (COBOL), a programming language first invented in 1959. Very few programmers today still know COBOL, and it is challenging for them to learn, invented as it was decades prior to modern software development paradigms.¹²

LLMs can help, but not all of them are up to the task. COBOL code is not well-represented in the training data used to create current language models. Fortunately, there is an open-source evaluation tool that grades LLMs on their ability to solve coding problems using COBOL.¹³ There are also some language models designed explicitly for this task.¹⁴ These resources can be a starting point. It is also possible to fine-tune an LLM, at low cost and with only modest technical skill, on an agency’s existing legacy code base to give it more examples of functioning code in an uncommon programming language such as COBOL. This can, in turn, be used to maintain legacy codebases or, ideally, to convert it to more modern programming languages.

Option 11: Improve grid reliability and efficiency

An electric grid is a complex network of interconnected components, each with its own operating requirements. Each is also affected by other parts of the system, and external factors like weather, in different ways. Managing the power flow and broad configuration of an electric grid is partially

aided by a process known as topology optimization. While topology optimization software has existed for some time, it is now possible to apply AI-based methods, including the same methods that allowed Google DeepMind’s AlphaGo system to achieve superhuman performance at the board game Go.¹⁵

Like a game, grid management is a complex optimization problem with many dimensions, hard rules (such as state laws, or the laws of physics), and unpredictable external factors. Systems based on a wide variety of AI approaches have already been developed, potentially resulting in substantial energy efficiency gains, cost improvements, and increased reliability. While state policymakers do not usually directly control electric utilities, optimizations of this kind can be encouraged via public utility commissions, over which state legislators generally have direct oversight and authority.

Option 12: Customize internal chatbots for employee support

Every state agency has rules and regulations for its employees, whether set internally or by state and federal laws. Modern LLMs, with the ability to keep the equivalent of a novel in their active “attention,” combined with techniques like RAG, can allow all these agency-specific rules, guidelines, and laws to be kept in a single chatbot. Agency employees can describe a situation they are facing and ask the language model how the rules might apply to it. If deployed well, these applications can significantly speed up internal processes. A language model could even be used to highlight potential discrepancies or subtle contradictions in various rules, allowing agencies to harmonize their operations.

By additionally collecting employee feedback about the model’s responses (for example, whether it accurately reflected a nuance of a particular rule or procedure), the model can be further fine-tuned after deployment. Employees can also experiment with and train future, external-facing chatbots that could help state residents understand agency rules.

Conclusion

As with any technological transformation, there are many unanswered questions about how AI will be used, what its downsides will be, and much else. Commentators can ask as many questions as they like—and many do—but answers will not come purely by *thinking* about them. Society must *use* AI to understand it. Risks must be observed before they can be mitigated.

Thus, perhaps more than anything else, this paper suggests that state and local policymakers adopt a combination of urgency and patience. States and cities will profit from gathering insight and information with alacrity, yet they will also benefit from understanding that many of their biggest questions about AI will take time to answer. AI is likely to change the way that people interact with their government in ways far more profound than a single law. It will almost certainly change the tools of statecraft, and it may even change the nature of statecraft itself. As state and

local governments consider new laws and other prescriptive actions, they would be wise to remember that AI is still in its early stages.

About the Author

Dean Woodley Ball is a research fellow at the Mercatus Center at George Mason University and author of the Substack *Hyperdimensional*. His work focuses on AI, emerging technologies, and the future of governance. Previously, he was senior program manager for the Hoover Institution's State and Local Governance Initiative.

Notes

This policy brief is the product of a collaboration between Mercatus scholars Matthew Mittelsteadt and Dean W. Ball.

1. Jill Jonnes, *Empires of Light: Edison, Tesla, Westinghouse, and the Race to Electrify the World* (Random House Trade, 2004).
2. Dean W. Ball, "California's Effort to Strangle AI," *Hyperdimensional*, February 9, 2024, <https://www.hyperdimensional.co/p/californias-effort-to-strangle-ai>.
3. David E. Nye, *Electrifying America: Social Meanings of a New Technology, 1880-1940* (MIT Press, 1992).
4. *United States v. Alvarez*, 567 U.S. 709 (2012).
5. Gemini Team, Google, "Gemini: A Family of Highly Capable Multimodal Models," preprint, arXiv, December 19, 2023, <https://doi.org/10.48550/arXiv.2312.11805>.
6. Shalanda D. Young, Executive Office of the President, Office of Management and Budget), March 28, 2024.
7. Sadie Bogard and Nikhita Ari, "State and Local Government Jobs Still Haven't Recovered from the Pandemic," *TaxVox* (blog), Tax Policy Center, August 9, 2023.
8. Decisions about whether and what to share are complex and should consider the potential sensitivity and security of any associated data. The Office of Management and Budget's draft guidance provides a basic framework for making these decisions. See section 4d, "AI Sharing and Collaboration," in Young, "Advancing Governance, Innovation, and Risk Management," 12.
9. Yunfan Gao et al. "Retrieval Augmented Generation in Large Language Models: A Survey," preprint, arXiv, December 18, 2023, <https://doi.org/10.48550/arXiv.2312.10997>.
10. Jake Offenhartz, "NYC's AI Chatbot Was Caught Telling Businesses to Break the Law. The City Isn't Taking It Down," *Associated Press*, April 3, 2024.
11. Minyvonne Burke, "Coronavirus: State Unemployment Websites Crash As Applications Surge," *NBC News*, March 18, 2020.
12. Alicia Lee, "Wanted Urgently: People Who Know a Half Century-Old Computer Language So States Can Process Unemployment Claims," *CNN*, April 8, 2020.
13. BloopAI, "COBOLEval" (dataset), <https://github.com/zorse-project/COBOLEval?tab=MIT-1-ov-file>.
14. See Kim Martineau, "COBOL Programmers Are Getting Harder to Find. IBM's Code-Writing AI Can Help," *Future Forward* (blog), IBM, October 26, 2023, and BloopAI, <https://huggingface.co/bloopai>.
15. Erica van der Sar, Alessandro Zocca, and Sandjai Bhulai, "Multi-Agent Reinforcement Learning for Power Grid Topology Optimization," preprint, arXiv, October 4, 2023, <https://doi.org/10.48550/arXiv.2310.02605>.