



Subnational Practices in AI Policy: A Working Guide

Ian Klaus and Ben Polsky

Artificial intelligence (AI) policies and frameworks are developing rapidly at the national, international, and supranational levels, as well as at the subnational level. AI policy is a developing field, and this “working guide” seeks as much to establish key areas and concepts to watch as to draw conclusions. But even now, some notable trends at the subnational level have emerged from the ongoing work.

First, while the past decade has seen an explosion in diplomatic engagement by subnational officials, a refinement in subnational diplomatic practices, and increased visibility for cities and states on the global stage, those practices have not translated to the AI policy space. Subnational officials are turning to national and international frameworks for guidance in developing their respective policies toward AI, but they do not yet have influence upon these frameworks. Here, these officials may learn important lessons from the experiences of experts and policymakers who have focused on climate change, democracy, and sustainability issues and have integrated local perspectives and solutions into international fora and agreements.

Second, subnational jurisdictions are themselves employing a wide array of approaches toward AI, even though subnational, national, and international actors have yet to build connective tissue on AI policy issues. Though subnational jurisdictions often share similar goals and are building on existing policies, their level of engagement with the technology differs dramatically. Some cities are experimenting with AI for traffic management, chatbots for service delivery, and analysis of public comment and participation. Others, such as Los Angeles, are engaging in extensive internal stakeholder consultation. Still others, including Seattle, are refining existing policies through extensive engagement with city residents. And naturally, there are those who remain in wait-and-see mode.

Third, across the spectrum of engagement, most subnational jurisdictions are involved in intense knowledge-gathering exercises that seek to develop better understanding of both the technology and its possible implications for service delivery and policy priorities. Such efforts include creating inventories of use-cases, developing sandboxes and new risk



frameworks, and building partnerships with outside institutions. These efforts will shape policy for years to come, but for the most part they focus on the use of the technologies by governments themselves. For the larger, potentially seismic changes that will occur in societies and economies, a larger set of questions still remains.

Fourth, and finally, there are important lessons to be learned from previous subnational policymaking frameworks related to other issues, including climate change and housing. In particular, climate change conversations often have focused around risks and attendant options to mitigate them, while housing increasingly has considered a rights-based approach. As knowledge around AI risks remains nascent and is being built presently, and rights regimes related to data and technology differ across jurisdictions, cities, states, and provinces are toggling back and forth between risk- and rights-based frameworks in trying to anchor their emerging approaches.

Though subnational policymaking mechanisms and authorities differ across national contexts, policy practices are beginning to emerge at the state/provincial and city levels, as is a spectrum of engagement with the technology itself. This “working guide” seeks to capture some of those practices, as well as the process and philosophies that inform their development. It represents, in part, learnings from an ongoing series of workshops co-hosted by Carnegie California and the Barcelona Centre for International Affairs (CIDOB). These workshops have included participation from industry, civil society, as well as senior officials from the states of California and Utah; the region of Catalonia; and the cities of Los Angeles, Carlsbad, Long Beach, Seattle, Boston, and Barcelona, as well as Eurocities and the United States Conference of Mayors.

Overview of Goals and Practices

The explosion of public attention to new capabilities enabled by Large Language Models (LLMs), such as Generative AI, in 2022 hastened the need for and quickened the pace of policy innovation. Cities, states, provinces, and regions have been engaged with AI for years, and many have well-developed policies around privacy and data use. LLMs and AI, which can be used for content creation, natural language generation, and creative tasks, have expanded the horizons of AI applications beyond traditional rule-based and analytical functions.

Subnational jurisdictions have a decade of experience in developing policies and governance around big data and artificial intelligence. Some of these policies are applicable to newer forms of AI, but political contexts and policymaking processes have evolved radically, as has the technology. With regard to policymaking, we have captured four broad goals and a series of evolving practices in pursuit of them. Though the goals are broadly shared, the level of engagement and the discrete approaches to advance them vary significantly.

Goals

- Increasing the **efficiency** of government service delivery, and public **trust** in it;
- Promoting **equity and transparency and preventing bias** in the deployment and use of AI;
- **Influencing** industry and establishing predictable engagement with model and application providers; and
- Developing or maintaining respective geographies as attractive sites for **AI-related economic opportunities**.

Practices

- Building out and adapting **existing AI use policies**, particularly around privacy and data, while recognizing that such policies will evolve with the technology;
- Instilling **explainability and accountability** in AI systems through “human in the loop” designs and stakeholder engagement;
- Using **procurement** purchasing power—including acting collectively with other jurisdictions—as well as **liability regimes** to influence industry and shape the market and deployment of AI more broadly;
- Establishing **internal expertise and external expert partnerships** to review applications or models to be used by officials; and
- **Enhancing knowledge** of the technology using “inventories” of potential AI uses and impacts and “sandboxes” to test the technology and, in certain instances, to develop new inventories of and frameworks for risks and benefits to understand and track impacts.

There is an experimental, innovative, even chaotic pluralism to the emerging approaches at the subnational level, and they are by no means captured in their entirety here. Governments are engaging in these categories at different speeds and with different sequences of priorities. As such, there exists a wide spectrum of engagement with AI that spans from active experimentation with the technology to extensive internal stakeholder consultation to small refinement of existing policies to full wait-and-see mode.

Emerging National, International, and Supranational Approaches

Subnational AI policymaking occurs in the context of developing, though often at a slower pace, national, international, and supranational frameworks and regulations. National governments are taking the lead on catastrophic risk and national security AI policy-related questions. More likely than not, they will also lead on related questions of electoral processes and integrity. These policies and frameworks matter for subnational policymakers, who seek broad guidance, standard setting, and even ethical frameworks for their own policymaking. Although such developing AI governance regimes cannot be captured in their entirety, some of the essential national, international, and supranational frameworks are referenced below.

National Policies

National policies, regulations, and uses of AI are also rapidly evolving and diverse in nature. Approaches range from informal guidelines to AI reporting requirements to outright bans. Stanford University’s Institute for Human-Centered Artificial Intelligence (HAI) [found](#) that since 2016, countries have passed 123 AI-related bills.

- China has implemented a [series](#) of binding regulations to regulate specific applications of AI, including algorithmic recommendations, synthetic content, and generative AI. These approaches were informed by not only the Chinese government’s demand for information control but also its desire to address other socioeconomic impacts of AI, such as effects on privacy, labor markets, and antitrust. In the process of implementing these regulations, China is building its bureaucratic and regulatory capacity to deal with the forthcoming AI explosion.



- The United States has advanced AI policy largely through executive orders (EO), including the 2020 [EO on Promoting the Use of Trustworthy Artificial Intelligence](#); the 2022 [Blueprint for an AI Bill of Rights](#); and the 2023 [EO on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence](#). The latest EO requires companies that are developing models that pose a risk to national security, national economic security, or national public health and safety to notify the government and share the results of all red-team safety tests. In a series of mandated reports, the EO seeks to develop standards for critical infrastructure to ensure safe, reliable, and effective AI. Yet this EO, however broad and ambitious, has not been complemented by binding regulation from Congress.
- The United Kingdom's 2021 [AI Strategy](#) and the 2023 policy paper "[A Pro-innovation Approach to AI Regulation](#)" outline the British government's efforts to retain and build on the country's position as an AI hub. In contrast to other approaches, including the European Union (EU) AI Act, the strategy does not assign rules or risk levels to entire sectors or technologies. Instead, it regulates based on the outcomes AI is likely to generate in particular applications with a focus on preventing existential risk.

International Frameworks and Policies

An increasing number of international forums are attempting to advance global frameworks for AI governance. These include the United Nations's High-Level Advisory Body on Artificial Intelligence, the US-EU Trade and Technology Council, the Global Partnership on Artificial Intelligence (GPAI), and the Organisation for Economic Co-operation and Development (OECD). These international organizations seek to create a framework for global AI

polymaking that establishes norms, mitigates risk, and inspires responsible collaboration between the private and public sectors. Numerous organizations, including the [Carnegie Endowment for International Peace](#), [Google Deepmind](#), and the [World Economic Forum](#), have also proposed global AI governance frameworks. A Carnegie Endowment proposal, for example, calls for a new organization, the International Panel on AI Safety (IP AIS) to target the most urgent AI governance challenge—safety and security. This organization, inspired by the Intergovernmental Panel on Climate Change (IPCC), would have a deep technical understanding of current AI capabilities and the relevant safety and security risks. The panoply of efforts around AI global governance and leadership continues apace.

- In May, the G7 leaders established the "Hiroshima AI process" to discuss issues around AI. The process will include project-based cooperation with the GPAI and the OECD to promote safe, secure, and trustworthy AI worldwide. In October, the process produced a [Guiding Principles and a Code of Conduct](#) for organizations developing and using the most advanced AI systems, including the most advanced foundation models and ai systems.
- In November, the United Kingdom convened international governments, leading AI companies, civil society groups, and experts for the world's first [global AI safety summit](#). In the resulting [Bletchley Declaration](#), signatories agreed to look collectively at the risks around frontier AI models as a means to curtail potential intentional misuse or unintended issues.
- In 2019, the Africa Union [established](#) a working group tasked with developing a common African stance on AI, developing a capacity-building framework, and establishing an AI think tank. In

2022, the AU's High-Level Panel on Emerging Technologies reiterated the need for a continental AI strategy that enables African countries to coordinate policymaking, harness the benefits of AI adoption, and mitigate damages.

Supranational Policies

Introduced in 2021 and adopted in 2022, the EU's AI Act regulates AI within EU member states. The act encourages the development of AI technologies that align with European values, emphasizing the importance of ethical AI deployment while fostering innovation and competitiveness in the EU AI landscape. It outlines rules for high-risk AI systems, including mandatory requirements for transparency, data quality, and human oversight. The act assigns applications of AI to three risk categories based on the potential danger these applications pose: unacceptable risk applications, high-risk applications and limited or low-risk applications. It bans AI applications that pose the most significant risks to safety and fundamental rights. Enforcement, importantly, lies with national governments, not subnational ones. Subnational players, through networks like [Eurocities](#) and platforms like the [Committee of the Regions](#), have had a largely consultative role in the policy process around the Act.

Existing Subnational Policies

Subnational governments are moving quickly to adopt or adapt established frameworks and policies around AI. Many, but not all, acted before their national counterparts. There is no uniformity of approach to how AI should be utilized or regulated, and subnationals are in different stages of policy development. Different subnational jurisdictions, idiosyncratic and diverse, approach AI with different degrees of comfort and fear, "anxiety and excitement." Along that spectrum, subnationals have pursued various actions, ranging

from interim guidelines and internal IT policies to EOs and legislation.

As demonstrated in the CIDOB [Atlas of Urban AI](#), a map and repository of city initiatives to regulate the use, development, and application of AI, cities are fertile ground for testing benefits of technology and mitigating risks through policy entrepreneurship. The atlas, which tracks 165 initiatives across 63 cities, reveals that even though many cities are innovating on AI use-cases, few have overarching strategies. Hundreds of cities, however, do have existing privacy, big data, and even machine learning policies. Many of these policies have been developed through collaboration and networks. The [Cities Coalition for Digital Rights \(CDDR\)](#), for instance, was launched by Amsterdam, Barcelona, and New York City in 2018. It now includes 50 cities worldwide with the goal to "promote and defend digital rights" to "ensure fair, inclusive, accessible and affordable non-discriminatory digital environments." The CDDR was not AI-specific, preceding the 2022 leap in AI by nearly four years, but it does focus on policy issues captured in the AI policy problem set, including data privacy, bias, and algorithmic transparency. Many cities, including U.S. technology hubs, are looking to the CDDR for guidance on AI policy.

City practitioners have been examining state and national regulations for guidance while seeking to influence the frameworks with ethical principles and lessons gleaned at the local level. As subnational jurisdictions, cities, states, provinces, and regions share many of the same policy levers and goals, and therefore are grouped together in the examples referenced below.

- Amsterdam, 2020: Amsterdam's [Digital City Agenda](#) named three goals: responsible use of data and technology, combating digital inequality, and accessibility of services. The Digital City agenda includes proposals on data minimization, open by default, privacy by design, and a ban on Wi-Fi



tracking. Though this agenda predates generative AI, it created a framework for new tools on data use that are applicable to the current emerging uses and AI technologies.

- Los Angeles, 2020: [SmartLA 2028](#) is a blueprint for civ-tech innovation. The plan outlines five components, including infrastructure data tools and practices, digital services and applications, connectivity and digital inclusion, and governance. Although it does not mention AI specifically, it does note the importance of AI more broadly as a means to enable contact-free essential government services, among other previously unthinkable opportunities related to the adoption of AI.
- New York City, 2021: The [New York City's AI Strategy](#) outlines what AI is, how it works, and what ethical considerations are inherent in its use in the city ecosystem. It identifies five areas of focus: data infrastructure; AI applications within the city; city governance and policy around AI; partnerships with external organizations; and business, education, and the workforce. The strategy was released with an [AI primer](#), an implementation guide to the strategy. The primer acts as a foundation and is intended mainly for an audience of technical, policy, or other decisionmakers, not specific to city government.

Rapidly developing policy processes occur in the context of preexisting policies, as well as nascent (or entirely absent) national and international efforts. For example, over the past decade the concept of the “Smart City,” now often conflated with commercial platforms, has introduced key concepts around data in policy processes into the public sphere. Some subnationals are using these existing policies to manage the influx of policy questions arising from the introduction of emergent technologies.

Emerging Subnational Practices

According to a [recent survey](#) by Bloomberg Philanthropies, the vast majority of mayors (96 percent) are interested in how they can use AI to improve local government. Of those cities surveyed, 69 percent report that they are currently exploring or testing the technology to increase the efficiency of government services for data analysis (58 percent); citizen service assistance (53 percent); and drafting memos, documents, and reports (47 percent). A large majority of cities reported that security and privacy (81 percent), and accountability and transparency (79 percent) are the key ethical principles that guide their exploration and use of AI. Cities are actively engaged in policymaking to ensure that AI, when used, is employed in a manner that reflects the preferences of their residents.

The interest of policymakers and city and state officials in engaging AI may be well matched to the interests and concerns of their residents. In 2023, for example, Carnegie California surveyed Californians on their AI perspectives. Tracking the international efforts underway, nearly 50 percent of Californians expressed support for an international agreement on AI standard setting. Meanwhile, around 40 percent of Californians noted that local, state, and federal governments are “not doing enough” to respond to the potential benefits and risks of AI. More action on AI at not just the national level but also the state and local levels was the most common sentiment from Californians.

What might that action look like? The following subsections capture emerging practices in four broad categories: experimentation with technology and new policies; explainability and accountability; procurement policies; and efforts to enhance understanding of the technology, and potential policies, within government.

Experimentation

- Boston, Massachusetts, 2023: The “responsible experimentation approach” adopted in Boston allows for the public sector’s use of and experimentation with AI across government services and activities. In its [Interim Guidelines for Using Generative AI](#), the city outlines several scenarios in which public servants might want to use AI to improve efficiency, and provides specific how-tos for effective prompt writing. The guidelines also place responsibility on the user of the tool, and instructs officials to proof any work developed using AI.
- San José, California, 2023: The City of San José is exploring the benefits of AI in improving the delivery of services to residents, including applications for traffic management and automated license plate readers. Concurrently, the city is producing guidelines and guardrails to ensure those AI systems are used effectively and trustworthy. With the advent of AI, the city released a continually updating set of [Generative AI guidelines](#) to inform the use of the technology for public use-cases. Key characteristics include a directive not to submit any information to an AI platform that should not be available to the general public, to cite and record usage of AI, and to create an account for city use to ensure that public records are kept separate from personal records.
- Utah, 2023: The [Enterprise Generative AI Policy](#) provides guidance on the use of AI for executive branch employees in the Utah state government. The policy promotes the use of AI while seeking to protect the safety, privacy, and intellectual property rights of the State of Utah. The state is exploring the creation of AI sandboxes to explore the benefits of technology and safeguard its citizens from potential harms.

- Carlsbad, California, 2023: The City of Carlsbad has evaluated its existing policies and incorporated guidance related to AI where appropriate. Focusing on the city data policy, the guidance reiterated that employees remain responsible and accountable for all information and output regardless of the tool by which it was produced and encouraged experimentation and exploration with the tool. Houston, Texas, has deployed a somewhat similar approach, providing informal IT guidance to state employees on the use of generative AI. The guidance neither recommends nor prohibits its use but notes that the tool can be used to increase efficiency and innovation across a range for customer support, document analysis, and knowledge management.

These efforts stand in contrast to other subnational jurisdictions that are taking more reserved approaches.

- Maine, 2023: The State of Maine’s Information Technology office issued in 2023 a [directive](#) prohibiting AI for state government business or on any device connected to the state’s network for at least six months. During the moratorium, the IT office will conduct a “risk assessment” to analyze any cybersecurity and regulatory issues that the technology might raise.
- Seattle, Washington, 2023: Seattle approaches AI with the primary objective to be stewards of the public’s data. Starting with an internal policy on AI use in municipal functions, the city directs its staff on how to be thoughtful about AI. The [Generative Artificial Intelligence Policy](#) highlights several key factors to responsible use in a municipality, including attributing AI-generated work, having an employee review all AI work before going live, and limiting the use of personal information to help build the materials



AI uses to develop its product. To reduce bias and harm, employees must also apply a Racial Equity Toolkit before using an AI tool.

Explainability and Accountability

Explainability and accountability are critical themes in subnational AI policy development. By incorporating mechanisms such as public registries that hold both developers and users accountable for the outcomes of AI applications, policymakers seek to foster responsible and ethical deployment of AI technologies in local contexts.

- Helsinki, Finland, 2022: The [AI Register](#) is a window into the AI systems used by the city. Through the register, citizens can get acquainted with the quick overviews of the city's AI systems or examine their more detailed information, such as what data was used and from where it was collected. Urban data and AI are only utilized with the permission of the residents. The goal of the register is to enable access to understandable and up-to-date information about how algorithms affect citizens' lives.
 - San José, California, 2023: The City of San José has implemented an Algorithm Register. Each time the city procures an AI system, it is logged in the public register to communicate the AI systems that it uses to its residents. Each log summarizes the system objective, transparency and equity standards, and human oversight mechanisms.
 - Connecticut, 2023: The [Act Concerning AI](#) requires the Department of Administrative Services to inventory AI systems in use by any state agency. The state judiciary is required to conduct annual inventories of Connecticut's AI use to prevent against "unlawful discrimination" and other harmful outcomes.
- ### Purchasing Power and Procurement
- Across multiple subnational contexts, governments have implemented general guidelines for public sector procurement of AI in particular. The public procurement process is not only a means to acquire technology, but also a process by which the cities can vet models for accuracy and anti-bias measures before implementing the technology in government services.
- California, 2023: Home to thirty-five of the top fifty AI companies, California has both the unique standing and responsibility to promote trustworthy AI. The State of California's [EO on AI](#), released in September 2023, recognizes the influence of the state's purchasing power. The order directs the state's Operations Agency, Department of General Services, Department of Technology, and Cybersecurity Integration Center to reform public sector procurement in a manner that requires agencies to consider the uses, risks, and training needed to improve AI. California's procurement policies do not govern market conditions but influence the technology development of AI products through the buying power of the state as a high-value user.
 - Barcelona, Spain, 2021: Released in May 2021, the city's [AI municipal strategy](#) aims to regulate the use of AI in municipal services and promote ethical AI standards to be followed by private companies operating in the city. The strategy outlines procurement protocol that bans certain high-risk applications based on the EU's risk classification for city use.
 - San José, California, 2022: San José is piloting an [AI procurement process](#) whereby AI vendors complete a Vendor AI FactSheet that contains basic facts about the AI system, such as the data used to build the system and under what conditions it performs well. The city vets the

details of the factsheet and matches its needs with the system’s capabilities. The City of San José is also shepherding a [Government AI Coalition](#) with 100 government signatories to ensure the agencies can obtain critical information about AI systems from vendors in the procurement process in an effort to set an adopted industry standard that promotes responsible AI.

Enhancing Knowledge: Internal Expertise and External Expert Partnerships

Subnationals recognize the knowledge gap that exists within public bureaucracy on AI. Therefore, a number of states, cities, and municipalities seek to build and train staff and their internal expertise, as well as establish partnerships with external expert bodies.

- Barcelona, Spain, 2021: The city’s AI municipal strategy created an advisory council of experts in AI and technological humanism at Barcelona City Council, which integrates the city’s main experts in the field to review, assist, and advise the council on AI uses for the public good.
- Catalonia, Spain, 2023: As outlined in its [Artificial Intelligence Strategy](#), Catalonia supports and prioritizes a technology ecosystem that produces AI applications and research. The strategy outlines a public-private collaborative approach that balances its unique interests, such as promoting the Catalan language, with the equally strong desire to create linkages with other subnational regional entities such as Scotland and Québec.
- California, 2023: The EO on AI directs the state to enter into formal partnerships with academic institutions and knowledge partners to evaluate the impacts of AI on California, and recommend

any efforts the state should make to ensure it continues to lead the industry. It also encourages the state government to learn more about and experiment with AI’s potential uses through pilot projects and sandboxes.

As outlined in the EO, the California Government Operations Agency issued in November a report on “[The Benefits and Risks of Generative Artificial Intelligence](#).” The report offered a use-case focused comparison between “conventional AI” and “generative AI,” as well as a risk framework broken down into “shared,” “amplified,” and “new” risks, applied to issues such as labor impacts and privacy. Merging knowledge building and experimentation, the EO also directed the California Department of Technology to establish infrastructure to carry out AI pilot projects by March 2024, and set up sandboxes to test the projects to ensure that state agencies can begin to consider their implementation by July.

- Connecticut, 2023: The [Act Concerning AI](#) mandates the formation of a working group inside the state legislature, tasked with making recommendations on further AI regulation and an “artificial intelligence bill of rights.” The act requires the state judiciary to conduct annual inventories of Connecticut’s AI use to prevent against “unlawful discrimination” and other harmful outcomes.
- Utah, 2023: Utah has developed a working group composed of legislators, academics, executive branch members, and other external actors. The emerging strategy prioritizes the protection of the public and its data, enabling and encouraging AI-driven economic growth, and observing and learning on how the technology can be used

and its impacts. The state is also developing an approach for an AI Lab that would bring together industry and policymakers for joint learning processes around annual areas of focus.

Looking Forward

Just as the City of Boston appended the prefix “interim” onto its AI policy, so too did the State of California in its recent report on benefits and risks note the “preliminary” nature of its findings and the “rapidly developing” nature of the technology itself. Subnational governments are learning quickly, connecting, if informally, and attempting to deliver for their residents.

Looking forward, the ability of subnational governments to develop policy locally, exchange best practices regionally and globally, and influence policy at all level levels, will be determined by a number of issues that bear watching: which transitional platforms, such as the Frontier Models Forum or the G20, will emerge as the leaders, and how will subnational governments plug into them? How will the Global South, home to some of the faster growing urban areas and some of the leading voices in subnational diplomacy, most influentially enter into the AI global governance conversation? And, ultimately, which acute risks, as well as wider societal impacts, will emerge as the most pressing—and how might cities, states, provinces, and regions prepare for and organize around them?

About the Authors

Ian Klaus is the founding director of Carnegie California.

Ben Polsky is a consultant with Carnegie California.



© 2023 Carnegie Endowment for International Peace. All rights reserved.

Carnegie does not take institutional positions on public policy issues; the views represented herein are those of the author(s) and do not necessarily reflect the views of Carnegie, its staff, or its trustees.