AI LOCALISM IN PRACTICE:

EXAMINING HOW CITIES GOVERN AI

Sara Marcucci, Uma Kalkar, and Stefaan Verhulst

AI LOCALISM

THEGOVLAB

Image by Uriel SC/ Unsplash

TABLE OF CONTENTS

AUTHORS

SARA MARCUCCI | RESEARCH FELLOW, THE GOVLAB

Sara Marcucci is a data and policy researcher focusing on projects about datafied urbanism, data justice, democratic innovation, and citizen participation. She has previously worked at the Open Data Institute where she helped develop a series of policy options for governments to create an enabling environment for bottom-up data institutions and at Nesta Italia in the Tech for Good department.

UMA KALKAR | RESEARCH FELLOW, THE GOVLAB

Uma Kalkar is a mixed-methods policy researcher focused on responsible AI, data governance, and digital infrastructure strategies and applications through a socio-technical lens. Her work on digital disparities and civic deserts has been a cornerstone of youth civic engagement strategies for the non-profit 18by Vote. She has a background in technology policy, conflict studies, and quantitative methods.

STEFAAN VERHULST | CO-FOUNDER AND CHIEF RESEARCH AND DEVELOPMENT OFFICER, THE GOVLAB

Stefaan G. Verhulst is Co-Founder and Chief Research and Development Officer of the Governance Laboratory (The GovLab) where he is building an action-research foundation on how to transform governance using advances in science, data and technology. Verhulst's latest scholarship centers on how technology can improve people's lives and the creation of more effective and collaborative forms of governance. Specifically, he is interested in the perils and promise of collaborative technologies and how to harness the unprecedented volume of information to advance the public good.

ACKNOWLEDGEMENTS

We are deeply grateful to our peer reviewers and colleagues for editing early versions of this report, which have been published as blogs. Publishing this report would not have been possible without their insight throughout the process. In particular, we are grateful to Mona Sloane (New York University & University of Tübingen), Beatriz Botero (Sciences Po), Christophe Mondin (CIRANO), Ben Snaith (Open Data Institute), Hubert Beroche (Urban AI), and Jess Reia (University of Virginia). Moreover, we are thankful for the input provided by our colleagues at The GovLab, in particular, Sampriti Saxena and Salwa Mansuri.

EXECUTIVE SUMMARY

The rapid growth of cities and advances in data has led to the growing adoption of artificial intelligence (AI) and other automated systems in cities across the world. In fact, research shows that cities and states are leading the charge in developing governance frameworks and implementing policies at a quicker, more direct, and more impactful level than their national counterparts because of the "immediacy and proximity" of policy needs and responses. Yet the expanded use of AI by municipalities raises new questions about what constitutes appropriate and ethical use of the technology and their governance models.

In 2020, The GovLab developed the AI Localism project, which investigates instances of AI governance at a state- and city-level across the world. Our previous scholarship in this field developed an AI localism canvas, which focused on seven key themes in its governance:

- **Principles and Rights:** Non-binding agreements local agencies may develop and use, sometimes in collaboration with other agencies or city partners, to ensure the responsible use of AI at a local level;
- **Procurement:** Innovations regarding the acquisition of AI by a public institution from third-party private vendors;
- **Engagement:** Novel ways to engage publics into conversations and decisions regarding AI-related concerns, such as the collection and use of urban data;
- **Laws and Policies:** Efforts to regulate government use of AI as well as how certain AI applications can be used in certain sectors, such as public education or urban mobility;
- **Accountability and Oversight:** Initiatives on a local level that are aimed at enforcing accountability mechanisms about the use of AI systems;
- **Transparency:** Local efforts to develop and encourage transparency about the acquisition and application of AI systems across governmental agencies and domains; and
- **Literacy:** Avenues to educate citizens, residents, policymakers and the public as a whole about the development and use of AI, its functionings and social impacts.

This report aims to serve as a primer for policymakers and practitioners to learn about current governance practices and inspire their own work in the field. In this report, we present the fundamentals of AI governance, the value proposition of such initiatives, and their application in cities worldwide to identify themes among city- and state-led governance actions. We close with ten lessons on AI localism for policymakers, data, AI experts, and the informed public to keep in mind as cities grow increasingly 'smarter', which include:

- **Principles provide a North Star for governance.** Establishing and articulating a clear set of guiding principles is an essential starting point for AI localism.
- **Public engagement provides a social license.** Trust is essential to fostering responsible use of technology as well as broader acceptance and uptake by the public.

- **AI literacy enables meaningful engagement.** The goal of AI literacy is to encourage familiarity with the technology itself as well as with associated ethical, political, economic and cultural issues.
- **Tap into local expertise.** Policymakers should tap into cities' AI expertise by establishing or supporting research centers.
- **Innovate in how transparency is provided.** To build trust and foster engagement, AI Localism should encompass time-tested transparency principles and practices.
- **Establish new means for accountability and oversight.** One of the signal features of AI Localism is a recognition of the need for accountability and oversight to ensure that principles of responsive governance are being adhered to.
- **Signal boundaries through binding laws and policies.** Principles are only as good as they are implemented or enforced. Regulation sends a clear message to consumers that their data rights and protections are upheld and holds corporations accountable to respecting privacy privileges.
- **Use procurement to shape responsible AI markets.** Cities should use procurement policies to encourage responsible AI initiatives to justify the use of new technologies and understand if the benefits of these tools outweigh the harms prior to procurement.
- **Establish data collaboratives to tackle asymmetries.** Data collaboratives are an emerging form of intersectoral partnership, in which private data is reused and deployed toward the public good.
- **Make good governance strategic.** Too many AI strategies don't include governance and too many governance approaches are not strategic. It is thus imperative that cities have a clear vision on how they see data and AI being used to improve local wellbeing.

We hope these can help systematize our approach to AI localism. Considered together, these lessons aim to add up to an incipient framework for implementing and assessing AI localism initiatives in cities around the world. To learn more about AI localism, visit ailocalism.org. You can also explore over 100 case studies of AI localism across the world at the AI localism repository and submit your own examples here. Professionals interested in pursuing thought leadership on the subject are invited to contact Stefaan Verhulst at sverhulst@thegovlab.org.

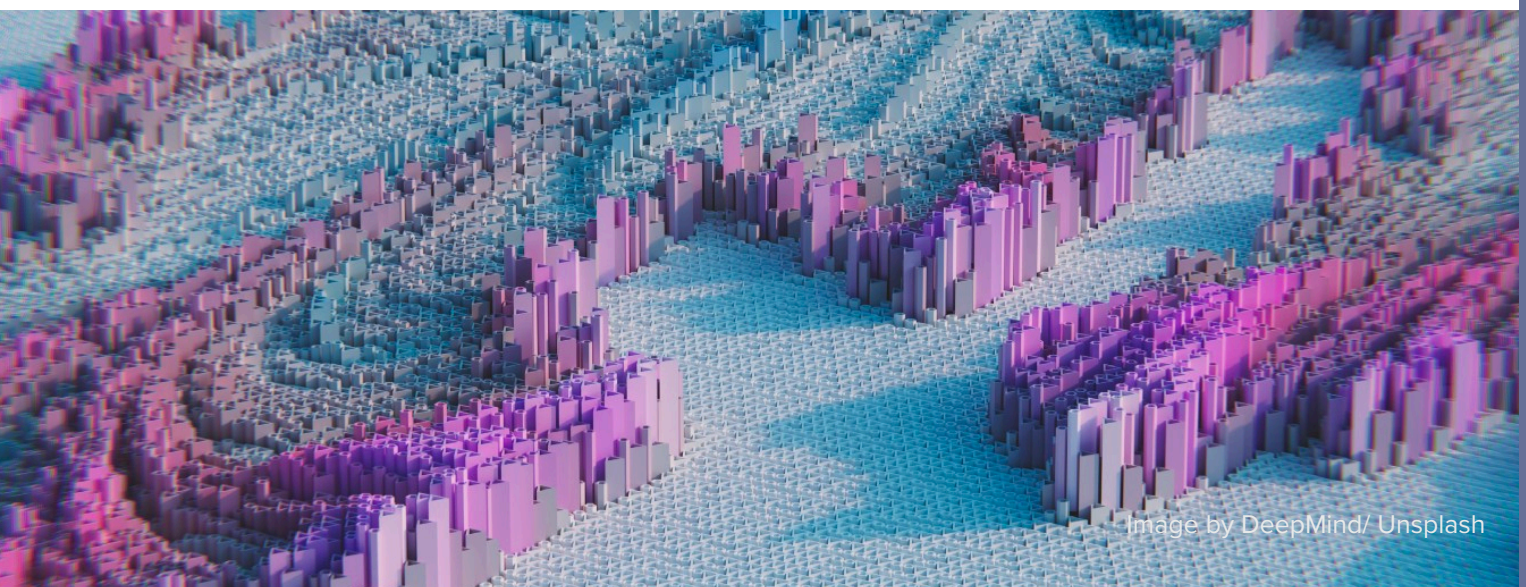

Image by DeepMind/ Unsplash

INTRODUCTION

Rapid urbanization and expansion have paved the way for increased socio-technological innovation in cities. One such socio-technical innovation is Artificial Intelligence (AI), whose definition is ever-contested, and which the European Commission describes as a digital system that "display[s] intelligent behaviour by analysing their environment and taking actions–with some degree of autonomy–to achieve specific goals."

Increasingly, nations and supranational blocs are beginning to consider and debate how to govern AI: the European Union proposed an AI Act; while the U.S. unveiled an AI Digital Bill of Rights. Similar initiatives have been conducted in Canada, Taiwan, and the United Arab Emirates, to name a few.

However, our research finds that cities and states are leading the charge of developing governance frameworks and implementing policies at a quicker, more direct, and more impactful level than their national counterparts. A number of cities have indeed proposed innovative smart urbanism visions that move away from a techno-centric approach and toward a more human-centric one.

In an article to Project Syndicate, we have called local instances of AI governance 'AI Localism'. In addition, AI Localism is an initiative that began in 2020 with our AI Localism canvas, which captures the frames under which local governance methods are developing.

WHAT IS AI?

There are numerous debates and contentions around the definition of artificial intelligence. In simple terms, AI is a system of algorithms. An algorithm uses a series of steps to automatically turn inputs (such as facial scan data from street cameras) into outputs (such as identifying people of interest for the police). Thus, a system of algorithms uses automated reasoning to organize and prioritize inputs, ultimately producing outputs upon which decisions can be made.

Oftentimes, AI-based systems in the public sector are referred to as "automated decision systems" (ADS). These are systems that, through automated computation, aid decision-making processes by standardizing, sorting, and analyzing data. In the public sector, this often involves a human that, once the AI has analyzed the data that was fed to it, will make decisions based on those analyses.

The importance of governing these systems is increasingly present, and growing together with their levels of adoption. Indeed, the more they are employed for public purposes, the more urgent it is to develop governance frameworks that can guide their employment.

WHY IS IT IMPORTANT TO GOVERN AI?

In her book *Atlas of AI*, Kate Crawford defines AI as "a registry of power". As noted by critical data studies scholars, indeed, data-driven digital technologies are the result of a process that includes contextual data generation, collection, standardization and hierarchization. As such, these technologies are "instruments of perception", and the notion that data are neutral and therefore that they report an objective version of reality is increasingly discussed and contested. In fact, as noted in *Raw Data is an Oxymoron*, edited by Lisa Gitelman, to exist and be functional data need to be "imagined as data", and such an imaginative process involves "an interpretative basis" which leaves room for an intrinsic bias in human judgment.

The relativity of AI is particularly important to note when automated tools are employed for public purposes, as these can be biased and marginalizing. For instance, the Algorithmic Justice League shows how speech and facial recognition systems can fail to recognize the voices and faces of those who are often not considered in the programming process of these models, i.e. black people and people of color. More examples include bias against people with disabilities, or teachers who got fired because their teaching style was not recognized by the algorithm that evaluated it.

Overall, it seems ever more important to find new, inclusive ways of governing AI and employing it for public purposes. AI Localism provides us with numerous tools to do so, giving us a variety of examples of how to balance the risks and take advantage of the benefits offered by automated systems.

WHAT IS AI LOCALISM?

AI Localism, a term coined by Stefaan Verhulst and Mona Sloane, builds on the concept of "New Localism"–as described by public-policy scholars Bruce Katz and Jeremy Nowak–,and refers to the actions taken by local decision-makers to address the governance of AI within a city or community. Numerous types of 'localisms' exist to address specific, local needs that national policy is not always fit to address, or to fill policy gaps in communities overlooked by national governments.

AI localism can take a variety of forms, and cities are increasingly taking it upon themselves to address technology-related issues alongside a wide range of other concerns going from the environment to social and economic equity.

For instance, broadband localism can be found in Sandy, Oregon, where the local government created 'SandyNet' to provide high-speed DSL and wireless internet connection at low costs. Examples of similar efforts can be seen in Barcelona, Spain, and Seoul, South Korea. Examples of privacy localism include New York City's local ordinance to regulate the collection and use of citizen data by governments and law enforcement agencies, as well as San Francisco's ban on police and local government from using facial recognition technology to surveil and identify residents. These kind of measures aim to address at the local level growing concerns about urban surveillance and AI-enabled racial discrimination and profiling.

It is necessary to note, however, how AI Localism does not necessarily equal "good governance" of AI at the local level. Indeed, there have been several instances where local efforts to regulate and employ AI have encroached on public freedoms and impair the public good. For example, Toronto's Harbourfront Centre neighborhood received widespread criticism from the public for its decision to commission Sidewalk Labs to collect information about locals through sensors and cameras to 'optimize' the urban environment. After several public forums between the City of Toronto, Sidewalk Labs, and residents, a lack of transparency and public trust around who all will have access to the data generated, how it will be analyzed and for what specific purposes, coupled with the uncertainty stemming from the COVID19 pandemic, led to the cancellation of the project in May 2020. As this example shows, AI Localism instances serve as important lessons not only of what national AI governance efforts may want to repurpose and further explore, but also avoid and improve.

Finally, it is important to mention that the presence of AI Localism does not mean that robust national- and state-level AI policy are not needed. Whereas local governance is fundamental in addressing local, micro-level issues, tailoring, for instance, policies for specific AI use circumstances, national AI governance serves as a key tool to complement local efforts and provide cities with a cohesive, guiding direction.

EXAMINING THE CURRENT STATE OF PLAY IN AI LOCALISM

Our AI Localism project created a knowledge base and taxonomy on the dimensions of local AI governance. To this end, we created an AI localism canvas, which identifies the seven key elements across which current examples of AI localism occur, including:

- **Principles and Rights:** Non-binding agreements local agencies may develop and use, sometimes in collaboration with other agencies or city partners, to ensure the responsible use of AI at a local level;
- **Laws and Policies:** Efforts to regulate government use of AI as well as how certain AI applications can be used in certain sectors, such as public education or urban mobility;
- **Procurement:** Innovations regarding the acquisition of AI by a public institution from third-party private vendors;

- **Engagement:** Novel ways to engage publics into conversations and decisions regarding AI-related concerns, such as the collection and use of urban data;
- **Accountability and Oversight:** Initiatives on a local level that are aimed at enforcing accountability mechanisms about the use of AI systems;
- **Transparency:** Local efforts to develop and encourage transparency about the acquisition and application of AI systems across governmental agencies and domains; and
- **Literacy:** Avenues to educate citizens, residents, policymakers and the public as a whole about the development and use of AI, its functionings and social impacts.

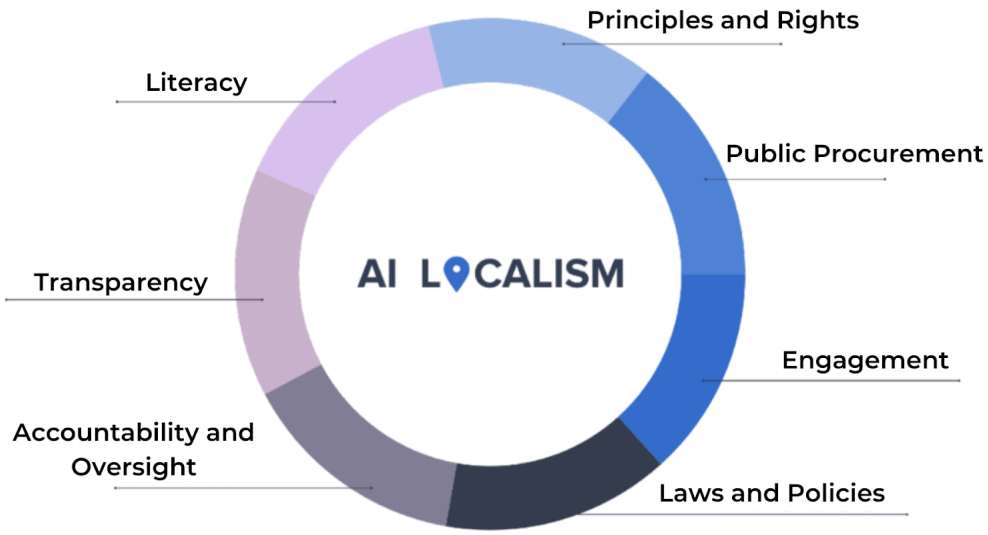

Figure 1. The AI Localism Canvas

This report serves as a primer for policymakers and practitioners to learn about current governance practices and inspire their own work. We present an overview and current examples of each frame of the canvas to identify themes among city- and state-led legislative actions. We end with ten lessons on AI localism for policymakers, data and AI experts, and the informed public to keep in mind as cities grow increasingly 'smarter'.

1. PRINCIPLES AND RIGHTS: BUILDING HUMAN-CENTRIC SOCIO-TECHNICAL SYSTEMS

Technologists and policymakers have long debated about what the foundations and tenets of AI ought to be–from a technical to a practical, to an ethical perspective. Over the past two decades, the increased adoption of AI by municipal administrations across the world have directed policy makers towards 'smart city' and 'smart urbanism' visions. Indeed, urban developers and city officials have been particularly enthusiastic about the employment of technology in cities, emphasizing the functional and economic benefits this would bring about in terms of urban management and economic boost. However, chasing after 'smart' and 'innovative' technologies can be reductive, increasing technological lock-ins and promoting tech-solutionism that centers city governance around the tools at play rather than the people impacted by them.

This section thus explores what AI principles and rights are, the reasons why they are important, and it illustrates a series of examples of how cities and local governments are implementing them and embedding them in the way they govern automated systems.

1.1. WHAT ARE AI PRINCIPLES AND RIGHTS?

Robust and enforced AI principles and rights help policymakers distance themselves from a techno-centric approach and, instead, embrace a more human-centric approach to AI localism. Principles help navigate system decisions and constrain possible solutions through responsible, practical, and ethical lenses.

AI principles and rights can be divided into two clusters, normative and dogmatic. Normative values reinforce human rights and sustainable development goals by promoting inclusion, lack of discrimination, and personal and data privacy—in other words, what AI governance at any level 'should' strive to achieve. Dogmatic principles outline 'non-negotiable' features of AI governance as decided by practitioners, policymakers, and watchdogs, such as requiring explainability of algorithms via routine audits or transparency of AI use through disclaimers.

1.2. WHY DO AI PRINCIPLES AND RIGHTS MATTER?

As Smit, Zoet, and van Meerten (2020) note, many existing 'principles' around AI are vague abstractions. They emphasize a need for design principles that scope actions and ethics for

AI around a specific problem. However, many AI principles and rights are presented divorced from each other, making it difficult to create a cohesive and coordinated approach for AI governance. Indeed, while multinational organizations such as the OECD have released AI principles and in 2021, UNESCO member states adopted the first Recommendation on the Ethics of AI, these principles remain too high-level for local governments and operations. Furthermore, reaching consensus between stakeholders to determine principles and implementing oversight to ensure they are being upheld is difficult, due to differing agendas, needs, and capacity of the myriad of players involved in AI initiatives. Thus for cities, scoping local needs and concerns around AI practices and ethics, as well as adapting top-line AI principles for their communities, can help situate and apply principles.

Multiple local administrations have taken concrete steps to regulate AI and data use in cities and protect citizens' digital rights. In 2018, Barcelona, Amsterdam, and New York City created the international Cities' Coalition for Digital Rights (CCDR) to address "algorithmic disruption" and its data and surveillance capitalism consequences. Forty-nine cities across the globe have joined the CCDR movement, implementing a diverse range of policies encompassing underlying trends of protecting citizen rights, regulating technological and social innovation, and advancing AI principles, responsibility, and justice.

1.3. HOW ARE CITIES AND LOCAL ADMINISTRATIONS REALIZING AI PRINCIPLES AND RIGHTS?

These regulatory practices usually take the form of city-wide strategies, manifestos, research labs, and citizen assemblies and center on citizen rights and needs for responsible technology use by public agencies. Montréal and Barcelona are primary examples of people-centered smart cities at the helm of AI localism. These cities have taken locally-centered steps grounded in community collaboration and personal privacy protection to develop their own set of AI principles with the aim to guide the responsible development and usage of AI and data by their local administrations.

1.3.1. THE MONTRÉAL DECLARATION FOR A RESPONSIBLE DEVELOPMENT OF ARTIFICIAL INTELLIGENCE

An early example of a municipality developing a set of guiding principles, albeit not binding, for the development and deployment of AI for public purposes is Montréal. In November 2016, the University of Montréal launched a co-construction process with citizens to design the Montréal Declaration for a Responsible Development of Artificial Intelligence. This involved a series of public consultations and citizen assemblies with 500 residents, experts and stakeholders aimed at increasing participation, transparency, and legitimacy of future AI-related policies. The Montréal Declaration promotes the following principles:

1. **Well-being:** The development and use of Artificial Intelligence Systems (AIS) must permit the growth of the well-being of all sentient beings.
2. **Privacy and Intimacy:** Privacy and intimacy must be protected from AIS intrusion and data acquisition and archiving systems (DAAS).
3. **Respect for Autonomy:** AIS must be developed and used while respecting people's autonomy and with the goal of increasing people's control over their lives and their surroundings.
4. **Responsibility:** The development and use of AIS must not contribute to lessening the responsibility of human beings when decisions must be made.
5. **Democratic Participation:** AIS must meet intelligibility, justifiability, and accessibility criteria, and must be subjected to democratic scrutiny, debate, and control.
6. **Equity:** The development and use of AIS must contribute to the creation of a just and equitable society.
7. **Solidarity:** The development of AIS must be compatible with maintaining the bonds of solidarity among people and generations.
8. **Diversity and Inclusion:** The development and use of AIS must be compatible with maintaining social and cultural diversity and must not restrict the scope of lifestyle choices or personal experiences.
9. **Prudence:** Every person involved in AIS development must exercise caution by anticipating, as far as possible, the adverse consequences of AIS use and by taking the appropriate measures to avoid them.
10. **Sustainable Development:** The development and use of AIS must be carried out so as to ensure the strong environmental sustainability of the planet.

1.3.2. BARCELONA'S AI STRATEGY

A more recent example, then, is the one of Barcelona and its AI Strategy. In April 2021, Barcelona revealed its strategy for the ethical use of algorithms and data. Building off of its previous smart city principles, such as the "Declaration for the proper development and usage of artificial intelligence" in 2017 and the Digital City initiative pioneered in October 2016, the Barcelona AI Strategy is oriented towards improving public knowledge about the use of algorithms by public services and redistributing among citizens the economic and social value produced by their data. The AI Strategy outlines the following seven governing principles required for local AI projects to uphold:

1. **Human action and supervision:** Any AI initiative that impacts residents must be overseen by humans to ensure that autonomous decision-making does not harm any one person. The level of human oversight is to be proportional to the risk that the technology used has to people.
2. **Technical robustness and security:** The City Council must be proactive in protecting and securing its technology from cyberattacks to prevent data leaks and data manipulation. Routine audits on the security of public AI are to be conducted.

3. **Data privacy and governance:** Across the entire data lifecycle, protecting and maintaining the privacy of residents must be upheld by collecting only the minimum amount of data and anonymizing data. Public data should be of high quality and published for public use, and, where possible, be accompanied by information about the data's provenance and bias mitigation strategies. Moreover, data used by public bodies has to be bias-corrected before use and be employed for "legitimate and proportional" as well as GDPR-complaint reasons.
4. **Transparency and information:** The entire process undertaken by an AI system—from how data is labeled for an algorithm to the way in which an algorithm makes its decision—must be explained and made available to the public to trace the whole decision-making process. The use and limitations of AI need to be clearly communicated and the technology and its creators must be easily identifiable to the lay public.
5. **Diversity, inclusion, and fairness:** To prevent discrimination from publicly used AI, the City Council must hold citizen consultations to understand its impact on people and ensure digital accessibility via a "universal-design approach" that allows people with disabilities to wholly interact with the technology.
6. **Social and environmental commitment:** The ways in which AI is used should contribute to attaining the 17 Sustainable Development Goals outlined by the UN. For Barcelona, this means using smart systems to uphold and bolster democratic processes, elections, and civic engagement.
7. **Responsibility, accountability, and democratic control:** When decisions are taken through AI, humans affected by these decisions should be able to get an explanation of why the decision is made in a language they can understand, and they should be able to challenge the decision with reasoned arguments. AI applications thus need to be intelligible and audited to ensure that their design and use do not harm residents.

1.4. CONCLUSION

The emergence of these guidelines is an encouraging sign that some cities are paying attention to how AI is being used and developed by local administrations in order to become laboratories of innovation in AI governance. Additional efforts to oversee the development and evolution of their AI governance include the metropolitan data charter of Nantes, France, which outlines four commitments to data sovereignty of the community, data protection, data transparency, and data innovation to increase citizen engagement and participation with the evolving use of data in public services. Similarly, in light of a wave of COVID-19 recovery plans, London released the fourth iteration of its Emerging Technology Charter, a voluntary charter that outlines guidelines for testing and deploying "data-enabled smart city technology" in a transparent and explainable manner.

Importantly, each of these examples highlight the presence and need for active collaboration and dialogue with non-governmental and grassroots organizations when designing

community principles. In his survey of principles and policies undertaken by 13 CCDR cities, Calzada (2021) finds that despite the use of different regulatory and policy practices, common objectives of furthering digital literacy, data responsibility, anti-discrimination, democracy and ethics steer city AI and data rights strategies present across AI localism efforts.

However, while these principles appear robust on paper, it is crucial to note the need for constant and critical reevaluation through a participatory approach to monitor and evaluate the success of these guidelines. To this end, we posit that re-evaluation and monitoring also be added to existing principles to ensure that other principles are current. Ultimately, as we begin our exploration of AI localism through this report, it is evident that cities across the world are taking policy steps to address the effect of technology on residents and public institutions.

2. PUBLIC PROCUREMENT OF AI: ACQUIRING AND USING AI FOR PUBLIC PURPOSES

Over the past decade, algorithms have been increasingly employed to automate or assist public decision-making processes and service delivery. Indeed, many local and national administrations have turned to automation to look for technical, "unbiased" support in a variety of areas, including urban planning, social care and welfare, education, health, housing, and public surveillance for law enforcement.

Because public administrations often have limited capacity when it comes to developing their own AI systems, they tend to contract from private companies to provide the automated systems governments wish to employ. Considering that these systems are used in the public sphere, with ramifications for constituents, it is particularly important for procurement practices to be fair and accountable, and provide clear standards and principles that account for the potential risks associated with using AI for public purposes.

2.1. WHAT IS PUBLIC PROCUREMENT OF AI?

Procurement, in general, refers to the acquisition of tools or services from third-party vendors by a given government. This is a core element of public institutions, as it allows governments to outsource efforts and resources to private companies and businesses that have the necessary means, expertise, and resources, which in turn, drives innovation.

Usually, the process of procurement involves a typical series of steps, such as (a) the definition of a specific need, (b) the development of a call for proposals which invites companies to present themselves as candidates to meet that need, (c) the bid solicitation, (d) the selection process, and finally (e) the contracting and execution.

The procurement of AI specifically refers to the process of selecting and acquiring AI systems. It involves the contracting of a business to provide automated systems to a specific government. As both national and local governments' interest in understanding and employing automated systems and information technologies increases, so does the desire to adopt such tools for public purposes. The next section will explore why there is a need for AI procurement, and what that need involves.

2.2. WHY DOES PUBLIC PROCUREMENT OF AI MATTER?

Public procurement of tools and services is a core component of government activities. In particular, the acquisition of automated systems by third-parties allows governments to rely on expertise they might not otherwise have, save time and money in the development of these tools, as well as shape the discourse around AI regulation and boost innovation and growth.

AI procurement is an important governance tool. Indeed, writing in *The Regulatory Review*, Lavi M. Ben Dor and Cary Coglianese reaffirm the need for local AI governance to protect residents and their data from use of private sector AI for public purposes without oversight. In particular, they point to a case in Texas, where the school board acquired an algorithm created by a private company to evaluate teachers, but the method by which decisions were made by the technology was concealed under trade secret protection, giving teachers no recourse to challenge the machine-made decisions. Such instances demonstrate the need for standards to assess, deploy, and monitor how governments procure AI technology and consider its risks.

2.3. HOW ARE CITIES AND LOCAL ADMINISTRATIONS REALIZING PUBLIC PROCUREMENT OF AI?

In the below, we provide a few examples of how cities have leveraged procurement and ordinances to tackle and prevent some of the harmful impacts the use of AI may have on residents, in particular, vulnerable communities, such as AI bias resulting in discrimination and marginalization. Specifically, we discuss efforts to address the threat of racial profiling by surveillance technology and the growing body of governance practices mandating disclosure and routine reporting on the acquisition and use of AI by local governments.

- In Baltimore, Maryland, the city council passed a restrictive ordinance banning public agencies and private citizens from purchasing facial recognition technology, criminalizing the technology. Balancing explainability and transparency of AI tools created or bought by governments is needed to allow for accountability to constituents of how their data is being used, and by which mechanisms.
- Berkeley, California, passed an ordinance that requires surveillance technology procured by city agencies to demonstrate the risk and benefits associated with specific technologies before getting permission to procure these tools. The Detroit City Council similarly approved an ordinance requiring transparency of all public surveillance and procurement processes for greater local awareness and transparency.
- California's state assembly introduced a bill that would require prospective contractors for high-risk applications, including automated decision systems, to submit an impact assessment in order to be considered for the contract. Similarly, the city of Amsterdam developed 'contractual terms.' including transparency and explainability of how the

algorithms function, for algorithms purchased by local governments from third-party suppliers.

- In 2022, the state of Vermont proposed a bill that requires the Secretary of Digital Services to audit all automated decision systems created or procured by state agencies, as well as outline standards for procurement development. The New York City Council also requires public agencies to annually disclose information about AI systems and the data collected and analyzed by the tool.
- In 2016, amid growing concerns about the use of surveillance technology in Oakland, California, the city formed a permanent Privacy Advisory Commission (PAC), the first of its kind in the United States. The PAC consists of nine experts who advise the city council on the "purchase and use" of surveillance technology and help craft policies and legislation related to data use and privacy with input from locals to create a smart city model that is privacy- and resident-focused.
- Washington state created an automated decision working group, which convenes representatives across sectors to develop recommendations for enhanced procurement and auditing of AI systems used by the local government.

2.4. CONCLUSION

Growingly, local governments are imposing checks and balances to oversee how AI tools are procured in order to ensure constituents' safety and avoid corporate monopolization of AI service delivery to the public sector. The New York City AI Strategy asks vendors to use "plain language" when describing their AI tools and has created a new position of Algorithms Management and Policy Officer to build a policy framework to manage the algorithmic tools used by city agencies. Recent governance practices have set in motion new rules of play that prioritize human impact over technological novelty. Indeed, AI can help streamline government practices and improve digital transformation and smart city initiatives.

Procurement of local AI tools is incredibly important for local policymakers to drive further data-driven actions and improve existing services with a fair and equitable lens. The above has provided a peek into the ways such governance is occurring, demonstrating a key lesson —that local AI use requires risk assessment, public awareness, and transparency to strengthen trust and justify decisions made with AI to residents.

3. ENGAGEMENT: INVOLVING LOCAL EXPERTISE AND THE PUBLIC

The increasing adoption of AI for public purposes requires increasing levels of engagement from the wider public. Research and impact centers, local experts, as well as citizens and local residents are essential to involve in the development and deployment of automated systems. This demonstrated engagement allows cities to create a social license when adopting AI system that will, in turn, bring about increased levels of trust and confidence that the technologies adopted are indeed serving the public interest.

3.1. WHAT IS ENGAGEMENT IN AI?

In the context of AI Localism, engagement refers to the public involvement, input, and awareness of AI use in cities or towns. Engagement can happen in multiple ways: through research between research and impact centers, which are leading public engagement and education activities; via dialogues with citizens and residents to increase participation in designing and implementing local AI; and with working groups and committees that bring stakeholders and the general public together to increase public awareness and conversation around public AI practices.

3.2. WHY DOES ENGAGEMENT MATTER?

Engagement ensures that not only experts, but also non-specialists and the broader public in general are able to participate in decision-making around AI use, and improve their knowledge about local research and investment in the use of automated decision-making systems. The involvement and participation of these stakeholders are essential to ensure AI for public use is developed and deployed in a trustworthy, inclusive way.

In particular, numerous theories and practical strategies are being developed to engage people as a way to ensure that the use of automated systems does not perpetuate and exacerbate social divisions and exclusions but, on the contrary, helps achieve a society that is more sustainable and inclusive. It is essential for decision-makers to know public preferences and lived experiences as reliably as possible, thus responding to local needs and gaps. Going further, people's engagement is valuable for the construction of a AI and data processes and practices that recognize the "particular tensions between common and ideal requirements ... and particular challenges related to AI explainability and accountability," ultimately increasing trust in and legitimacy of AI initiatives.

3.3. HOW ARE CITIES REALIZING ENGAGEMENT?

There are three main groups whose engagement in AI governance seems key: (a) research and impact centers, (b) citizens, (c) local working groups and committees. The following sections explore each of those and illustrate some distinguished examples of how they are being involved in the development of AI governance frameworks and efforts.

3.3.1. RESEARCH AND IMPACT CENTERS

Over the past few years, research around AI governance has skyrocketed in academic, policymaking, and public popularity. Increased demand to investigate AI applications and their socio-technical implications has been met with innovative, locally-centered research programs. Below, we point to examples of how research labs help ensure that cutting-edge AI scholarship is effectively implemented in practice by directly engaging experts and government officials and indirectly reaching the broader non-specialist public (see more in *Section 3.3.2: Citizen and Residents Deliberations*).

- The Urban AI think tank based in France launched a global call to uphold six key principles of smart city regulation. They advocate for smart city and urban technology to stem from a social contract, be open and accessible, decentralized, frictional, meaningful, and ecological by design. Thus far, over 100 technology and governance personalities have signed the call.
- Alongside the city of Helsinki and its partners, the Berkman Klein Center ran a three-week AI Policy Research Clinic with two teams of global scholars to turn public AI principles into tangible policy measures. One group created an oversight model for stronger collaboration and interoperability that fit with Helsinki's existing government structure, a translational matrix for ethical and regulatory requirements at a European level with use cases at a city level, a wireframe for a web-portal to increase public engagement with AI tools, and an overarching policy playbook for actions and recommendations. The second group reconfigured an existing method of multi-stakeholder engagement used in Catalonia for Helsinki's requirements and produced a playbook outlining a four-phase participatory process for introducing and implementing public AI technologies.
- The Confederation of Laboratories for Artificial Intelligence Research in Europe (CLAIRE) looks to create a pan-European network of AI research to support human-centered AI innovation. CLAIRE's main goal is to use various actors, stakeholders, and mechanisms for "citizens engagement, industry and public sector collaboration" to create a European knowledge hub that advances understanding and application of AI. Partnerships are built through knowledge sharing and the integration of stakeholders in order to boost overall European competitiveness and well-being. It also functions as a meeting place for researchers and policymakers to learn about AI and implement these lessons in their home institutions to increase the overall understanding of AI across organizations. The

group undertakes many projects, such as its collaboration with the AI, Data, and Robotics Association (ADRA).

- The Canadian Institute for Advance Research (CIFAR) conducts pan-Canadian research on AI strategies through coordination between Amii in Edmonton, Alberta, Mila in Montreal, Quebec, and the Vector Institute in Toronto, Ontario to centralize local AI research and priorities and make Canada a global leader in AI. Specifically, they focus on provincial- and national-level research for health innovation, energy and environment initiatives, and public-private collaborations.

3.3.2. CITIZEN AND RESIDENTS DELIBERATIONS

Citizens and civil society at large play an important role in bringing in new perspectives on governance and policies. A cornerstone of transparent and accountable governance is citizen deliberation, a tried and tested approach, especially for emerging and new technologies, to generate critical and democratic consideration of the risks and rewards of AI use at the local level. These practices shed light on larger gaps in public digital knowledge, highlighting which fundamental aspects of digital and data governance need to be communicated around for better general understanding.

Thus, when it comes to employing AI in public spaces and realms, it seems fundamental to encourage and initiate citizen deliberations and assemblies. For instance, in 2020, The GovLab hosted a Data Assembly in partnership with the Henry Luce Foundation. Through three 'mini-publics,' we received feedback from New Yorkers to discern what sorts of data residents did (and did not) feel comfortable sharing with city officials to address COVID-19. Similar to this endeavor, we present some examples of citizen-level engagement around local AI.

- The Laboratorio para la Ciudad de Mexico (LabCDMX), the innovation and experimental subdivision of Mexico City's local government, sought to develop a strategic plan for Mexico City to leverage the opportunities of the public use of algorithms and automated learning in the short, medium, and long terms. The first "exploratory" session defined the expectations, opportunities, and risks of AI for democracy. The second "co-creation" session brought together experts to advise on challenges, potential, and governance techniques for Mexico City's AI strategy. This session informed internal analysis to craft Mexico City's Strategic Roadmap for AI.
- Kowloon East in Hong Kong used a public participatory process to solicit local feedback in designing its smart city initiative, which includes improvements to urban infrastructure, walkability, resource management, and communication infrastructure. In addition to improving WiFi infrastructure and developing mobile apps to inform residents on updates and collect their data to improve services, the smart city initiative focuses on fostering a more sustainable and environmentally-friendly neighborhood through data-driven

initiatives for effective energy consumption monitoring, waste management, city-wide cooling, and green space additions.

- Starting in 2014, the Community Control of Police Surveillance (CCOPS) movement allows residents to voice their opinion on the use of surveillance and policing technologies in their neighborhoods. Eighteen towns across the United States have adopted CCOPS laws that require community approval before the implementation of surveillance technology and regular audits and reviews.
- In the United Kingdom, the Royal Society of the Arts organized a citizens' jury to deliberate on the ethical use of AI. The citizens' engagement and public deliberation have raised pressing concerns with regard to ways in which the public and private sectors alike must alter their mechanisms and functionalities to account for greater accountable and legitimate practices.

3.3.3. LOCAL WORKING GROUPS AND COMMITTEES

More broadly, forums for public and expert consultation around AI help open up conversations around the use, disclosure, and impact of the technology on the public. Like citizen deliberations, these consultations become spaces for education and foundational understanding of overarching technology concepts for policymakers and council members to make informed decisions.

- The Alabama State Legislature established the Alabama Council on Advanced Technology and Artificial Intelligence "to review and advise the Governor … on the use and development of advanced technology and artificial intelligence in this state." A council of policymakers and technologists will discuss and provide recommendations on the use of AI by local governments.
- Washington state established an Automated Decision Systems working group to recommend policy and regulation updates on the "development, procurement, and use" of AI by public offices. The group consists of representatives across public agencies and advocacy organizations, with a specific focus on marginalized individuals, who will debate when automated decision-making and AI systems should be banned, methods of auditing and retaining transparency in system processes, and data handling and storage processes.

3.4. CONCLUSION

Bringing together various stakeholders and sectors is only part of the work needed to enhance AI localism engagement. Including representative voices from communities, especially those who are historically underrepresented and unaware of broader digital literacy, across the design and employment of digital technologies to develop inclusive and pluralistic AI. Participation of the broader public throughout the design phase is valuable for

the construction of a datafication paradigm, a phenomenon by which social actions are transformed into quantifiable data, allowing for real-time tracking and predictive analysis, to capture experiences and visions as possible and thus improve the representative capacity of AI. Moreover, the examination and evaluation of automated systems conducted by citizens and residents are fundamental for both digital technologies to gain legitimacy and for people to be aware of participants in the increasingly digitized society they live in.

4. LAWS AND POLICIES: HOW CITIES ARE AT THE FOREFRONT OF GUIDING THE USE OF AI

Over the past few years, national and international policymakers around the world have increasingly recognized the potentially harmful outcomes of AI, especially as it relates to its use for the delivery of public services. Indeed, for instance, the U.S. Equal Employment Opportunity Commission (EEOC) and the U.S. Justice Department's Civil Rights Division (DoJ) recently warned about disability discrimination that may result from the use of AI for hiring processes. While the number of laws and policies proposed to govern AI at a national level is growing, the policies themselves largely remain lagged and ineffective in the legislative process.

More recently, then, many cities and states have stepped into the breach by taking active and future-focused steps to legislate data use and reuse by public agencies and corporations to harness AI's power in a responsible way. This section explores what laws and policies constitute as it relates to AI governance, the reasons why they are important when employing AI for public purposes, and some examples of how cities are developing and implementing local-level regulations of automated systems.

4.1. WHAT ARE AI LAWS AND POLICIES?

Within the AI Localism framework, laws and policies are key tools for local regulation of AI. Laws and policies can relate to (a) the regulations about public use of AI (see Procurement in Section 2), or (b) they can focus on how certain AI systems are to be employed for the delivery of certain public services, such as healthcare and education.

For the former, laws and policies aiming to regulate the public use of AI often pertain to procurement processes and require the local government to publicly disseminate information about the automated decision-making systems it acquires from private companies to then employ for public purposes. In addition, these measures can request for details regarding the procurement criteria, how different vendors are evaluated, and what specific AI systems are being used via procurement processes.

For the latter, laws and policies tend to focus on specific areas where local governments use automated technologies in a more risky manner that could potentially have discriminatory and human-rights-infringing outcomes in individual domains. These measures are field-

specific and can span from bans on traffic sensors that raise surveillance concerns, to suspensions of facial recognition technologies employed in schools.

4.2. WHY DO AI LAWS AND POLICIES MATTER?

As mentioned above, governments around the world are increasingly recognizing the need to regulate AI for public purposes, and local administrations are often leading the way. Yet there remains a "responsibility gap" when it comes to conceptualizing human-centric and accountable AI that is largely driven by unclear understandings of who is responsible and liable for the AI, issues in moral and public accountability of socio-technical systems, and a lack of demonstrated oversight make human-centric and explainable AI design difficult to achieve.

To this end, governments have taken to local legislative floors to decide on the limits of AI use in their communities to address concerns that automated systems can lead to unfair and undemocratic outcomes if not regulated. Laws and policies are key policy levers that decision-makers use to raise awareness to troubling consequences of AI, such as surveillance technology disproportionately targeting people color or protected characteristics influencing insurance or hiring decisions, and check the use of algorithms by governments and the private sector.

4.3. HOW ARE CITIES REALIZING AI LAWS AND POLICIES?

This section explores laws and policies specifically concentrating on who AI ignores or targets through its systems. In the following sections, we group local oversight governance into two groups, namely (a) anti-discrimination legislation and (b) surveillance and privacy regulation fields.

4.3.1. ANTI-DISCRIMINATION LEGISLATION

Discriminatory consequences of AI can be caused by a myriad of factors, including assumptions underpinning model design, the goals an algorithm is optimizing for, and unrepresentative datasets and algorithms. To address concerns about algorithmic bias and discrimination, as well as propel informed oversight of those tools, legislators have taken steps to control where and how algorithms are used by public agencies and, to some extent, industries in general.

- The state of Colorado's CO S.B. 169 bill aims to "[p]rohibit insurers from using any external consumer data and information sources, as well as any algorithms or predictive models that use external consumer data and information sources in a way that unfairly discriminates based on race, color, national or ethnic origin, religion, sex, sexual

orientation, disability, gender identity or gender expression." Working with the insurance industry, the Colorado Insurance Commissioner conducts sector-specific stakeholder meetings to design rules around big data use in a non-discriminatory manner.

- The Stop Discrimination by Algorithms Act (SDAA) in Seattle, Washington has placed bans on companies and firms that use algorithms to deliberately marginalize vulnerable individuals from accessing crucial personal and professional opportunities such as employment and housing. The legislation aims to enforce transparency and anti-discriminatory practices by mandating decision explainability for all algorithms. The bill requires insurers to disclose information about the external data sources used in the algorithm and predictive models for insurance practices. The use of external data must be assessed by a risk management framework to mitigate discrimination.

4.3.2. SURVEILLANCE REGULATION

Public pushback against increasing surveillance creep, or the ubiquity of algorithms and data collection being used to watch people, has led to targeted legislation to control the use and sharing of facial recognition and smart technology data.

- In 2019, the City of Buenos Aires implemented a 'Facial Recognition for Fugitives System' (LFRT) that installed nearly 10,000 facial recognition cameras across the city. To regulate the use of the technology, legislators mandated that the authorities who manage the facial recognition system must transfer information to the Committee of the Public Security System and the Ombudsman Office for oversight on the technical specifications and location of facial recognition technology applications. Opposition to the surveillance practice, its procurement via private contracting, and disclosure on what, how, and where the data collected will be used sparked an anti-LFRT campaign by the Observatorio de Derecho Informático Argentino (ODIA) and local watchdog organizations. In early 2022, the LFRT program was suspended.
- The cities of San Francisco and Santa Cruz, California, King County, Seattle, and Worcester, Massachusetts, have passed bans on the use of facial recognition technology by police agencies. These bans are the result of concern over civil liberty-infringement and discrimination embedded in these tools through data and algorithmic biases.
- The New York City Council proposed the KEYS (Keep Entry to Your Home Surveillance-Free) Act that requires all tenants to have traditional key entry to their homes to prevent owners from unilaterally forcing tenants to submit to facial recognition, biometric scanning, or smart key technology that comes potentially at the expense of personal and group privacy.
- In 2021, Virginia's governor signed the Consumer Data Protection Act (CDPA) which outlines a framework for controlling and processing personal data in the state. The CDPA places regulations on handling and processing personal data by both manual and/or automated decision systems, security requirements, and adherence to data minimization

standards, and grants consumers protections around accessing, correcting, and deleting their data that may focus on marginalized individuals.

4.4. CONCLUSION

The lag between local and national laws and policies around AI is reminiscent of the Red Queen's Hypothesis. This theory draws on a story from Lewis Carroll's *Alice in Wonderland*, when Alice, taking part in a race, realizes that she needs to run twice as fast as she normally would in order to move forward.

> *"Well, in our country," said Alice, still panting a little, "you'd generally get to somewhere else—if you run very fast for a long time, as we've been doing."*
>
> *"A slow sort of country!" said the Queen. "Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!"*

Managing technological innovation requires policymakers run twice as fast as AI development to 'get somewhere' in its governance. Proactive efforts are needed to curtail the dangerous effects of AI use, such as profiling or other measures that may infringe on individual freedoms of liberty, privacy, and due process. Yet, by their very nature, governance strategies tend to face a time lag and respond reactively and after the fact. The examples above point to the beginning of some legislative avenues undertaken to spur more responsible AI governance to help cities and states run the AI race more successfully.

5. ACCOUNTABILITY AND OVERSIGHT: MONITORING LOCAL AI

As the applications of AI in the public sector increase, so do the indications about the potential harm they can cause (as seen in the Introduction) and, consequently, concerns over their legitimacy, accountability, and transparency. As a result, calls for the emerging concept of "algorithmic accountability" have grown significantly, and public administrations have been more and more interested in finding new, innovative ways to respond to such calls from a policy perspective.

This section aims to illustrate the ways in which local municipalities have attempted to operationalize the concept of algorithmic accountability and oversight in cities. It particularly focuses on specific areas of urban public administration, such as law enforcement, urban planning and design, and public service delivery.

5.1. WHAT IS AI ACCOUNTABILITY AND OVERSIGHT?

To understand AI accountability and oversight, the foundations of algorithmic accountability need first to be explored. Drawing from a recently published report on "Algorithmic accountability for the public sector" by the Ada Lovelace Institute, the AI Now Institute, and the Open Government Partnerships (OGP), algorithmic accountability can be defined as "the set of policies oriented towards ensuring that those that build, procure and use algorithms are eventually answerable for their impacts."

Oversight is a fundamental element for accountability practices to be fully realized. Just because an AI-based policy or practice is accountable does not necessarily mean that they are overseen. In other words, something being verifiable does not automatically mean it is verified. However, when talking about responsible AI practices, accountability cannot be wholly achieved without oversight. Therefore, accountability and oversight are different steps of the same process.

Tools to engage in algorithmic accountability and oversight practices include the transparent publication of principles and guidelines, regulatory bans and restrictions, algorithm registers, impact assessments and audits, and rights to appeal.

Finally, the actors that can oversee AI systems and hold them accountable include citizens, tech-workers, community organizers, investigative journalists, and civil society organizations. On the other hand, actors to be held accountable mainly include policymakers employing the

given AI system for public purposes, the companies and organizations involved in designing the tool, and those chiefly participating in generating and collecting the data that feed it.

5.2. WHY DOES AI ACCOUNTABILITY AND OVERSIGHT MATTER?

Starting from the premise that automated systems can bring about benefits such as speed, efficiency, and sometimes even fairness, it is important to note that they can in fact be fallible–and, when adopted for public decision-making especially, harmful.

AI is made of a series of training data–the inputs–that are automatically organized and turned into other data–the outputs. When adopting AI to make decisions, it is important to note that the training data is always contextual and generated. Indeed, as Rob Kitchin notes in his book *Data Lives*, data is always generated before being collected; it does not exist in itself. Instead, different things are categorized and turned into numbers–i.e. datafied–according to specific systems of knowledge. Thus, the same thing can be turned into data differently, depending on different contexts and people.

Ultimately, this concept makes it essential to ensure that AI systems and those involved in their design and employment are responsible, fair, and equitable, and that there is accountable and verifiable oversight.

5.3. HOW ARE AI ACCOUNTABILITY AND OVERSIGHT REALIZED IN CITIES?

Increasingly, local administrations have been at the forefront of operationalizing the concept of AI accountability and oversight. In particular, this section investigates three specific areas where this is being practiced: law enforcement, urban planning and design, and public service delivery.

5.3.1. LAW ENFORCEMENT

Law enforcement often turns to AI for surveillance, monitoring, and identification of people. While these actions can have positive effects on public safety, they infringe on fundamental rights to privacy, and, when unchecked, can result in targeted monitoring of specific individuals.

- Shenzhen, often known as China's Silicon Valley, has become the first local government in China to regulate its artificial intelligence applications. The Regulations on the Promotion of Artificial Intelligence Industry of Shenzhen Special Economic Zone "seek to promote the use and development of AI in both the public and private sectors, establish a

framework to govern the approval of AI products and services, and regulate AI usage ethics." The regulation outlines incentives for public-private collaboration, data oversight systems, comprehensive data collection and monitoring methods, and government oversight of Big Tech.

- Following a decision by the Portland, Maine City Council to ban facial surveillance technology by its police force, legislators in the Maine House of Representatives voted in favor of a proposal that allows the use of digital technology for the investigation of serious crimes, such as rape and murder. However, the "bill would require police to have probable cause before they use facial recognition in the investigation of a crime and would limit searches to databases maintained by the Department of Motor Vehicles or the Federal Bureau of Investigation."
- The Santa Cruz City Council banned the use of predictive policing (the use of crime data and algorithms to predict where offenses are likely to occur) and facial recognition technologies by law enforcement authorities. The decision was backed by civil liberties and racial justice groups in the city, who drew attention to the racially discriminatory outcomes that the technologies often foster.

5.3.2. URBAN PLANNING AND DESIGN

Regulating the use of public AI in local agencies and shared spaces is essential to protect residents and safeguard city processes, such as waste management, from negative externalities caused by biased or uninformed algorithmic decisions.

- The City of Syracuse's Office of Accountability, Performance, and Innovation (OAPI) "develops innovative solutions to Syracuse's most pressing problems. It leverages idea generation techniques and utilizes a structured, human-centered and data-driven approach to affect change and deliver results within the city." Working on risk management of local AI implementation, "OAPI and the rest of the City government consider developing and using an evaluative framework to ensure completeness and consistency in their decision making central to this task."
- Starting in 2019, the Iowa state legislature approved the use of autonomous (self-driving) vehicles on public highways provided the vehicle met "certain conditions including that the vehicle must be capable of attaining minimal risk if the automated driving system malfunctions." Further regulation and adherence to traffic standards have been instated, including a requirement that manufacturers may not test self-driving cars without a valid permit, and the provision of oversight authority to the Transportation Commission to "restrict operation" of an autonomous vehicle on a road.
- For cities, water and sanitation form the bedrock of urban development. In particular, sewer systems act as "a form of insurance against future service disruptions. Washington, D.C.'s Water and Sewer Authority has enlisted the use of Pipe Sleuth, an AI-powered treatment system, to provide vital information about the quality and state of wastewater collection and treatment.

5.3.3. PUBLIC SERVICE DELIVERY

Similar to Urban planning and design in Section 5.3.2, management and auditing of algorithmic systems used by public agencies is required to make sure that the data and models they use are appropriate for the decisions they make, and to track these decisions in order to correct practices where needed.

- Through local law 49, New York City established an Automated Decision Systems Task Force (ADS) in 2018 to review the use of algorithms by city agencies and offices to ensure "fairer and more equitable" use of these tools. On the recommendation of the ADS, the role of Algorithms Management and Policy Officer within the Mayor's Office of Operations was developed, with the aim to create policies and guidelines to use AI systems for public service delivery in an equitable and accountable manner.
- Adopting a children's rights perspective, the New York City Council created a Special Task Force in 2017 to "investigate city agencies' use of algorithms and deliver a report with recommendations" for improving the welfare and tracking of children in protective care.
- The City of Amsterdam is developing an algorithmic register that aims to give an overview of all the artificial intelligence systems and algorithms used by the city. This allows residents to be aware of, give feedback on, and actively participate in the employment of AI systems for urban service delivery in Amsterdam.

5.4. CONCLUSION

Ultimately, accountability and oversight practices of AI systems are essential for responsible AI localism to be practiced and for equitable and accountable public policy to flourish alongside innovation. An essential element that makes it possible to do so is transparency. Without automated systems being transparent in their design and implementation, it seems ever harder to operationalize and fully realize the accountability principles and measures here illustrated. The following section will thus explore the concept of AI transparency, investigating it from a local governance perspective.

6. TRANSPARENCY: UNVEILING AND MITIGATING RISKS OF PUBLIC USE OF AI

A significant concern around AI use in the public sector is the opacity of AI-based tools. As previously described, an algorithm automatically turns inputs into outputs, making it possible to reach decisions more efficiently. However, a lack of information and traceability over how those outputs were reached creates 'black-box' AI that cannot be fully audited, overseen, or explained to neither citizens nor policymakers.

Unexplainable and opaque AI can lead to discriminatory decisions (see the Introduction) that erode public trust in the use of AI technology by cities. To this end, municipalities have attempted to operationalize the concept of algorithmic transparency. These efforts primarily attempt to regulate the use of black box algorithms and illustrate examples of informational registries as tools to implement principles of transparency.

6.1. WHAT IS AI TRANSPARENCY?

In addition to being a technical system, AI is also a social system, made of policies, laws, social contexts and norms, and cultures. In fact, AI is a socio-technical system. As a consequence, AI transparency needs to be conceptualized and operationalized both from a technical and a social perspective.

From a technical perspective, transparency requires explainability, which can be very difficult to trace, especially for advanced AI. For instance, the reasons why certain decisions were undertaken by machine learning (ML) systems are often hard to trace back because of complex codes and opacity around the data used to train the systems. To face the issue of meaningful transparency and explainability from a technical, data-focused perspective, Timnit Gebru and others advance the idea of data sheets for data sets. These data sheets aimed at providing "standard operating characteristics, test results, recommended usage" to provide an understandable 'nutrition label' of a data set's properties.

From a social and policy perspective, AI transparency is related to a series of key questions–including details about the system's purpose, the social context around its design and use and its intended users and beneficiaries–that need to be answered for the system to be deemed "transparent". For instance, some of the questions can be:

- Why is the system being developed?

- Are there any alternatives to that system? Why was that selected?
- What data is needed for the system to work?
- What are the data sources of the system?
- Who is designing the system?
- Who is using the system?
- What are the consequences of and necessary conditions to use the system in other contexts, other than the one it was initially intended for?
- How does the system impact the lives of the environment, communities, and individuals involved in its design, development, and implementation?

6.2. WHY DOES AI TRANSPARENCY MATTER?

As mentioned in *Section 5: Accountability and Oversight*, the growing concerns around the negative and discriminatory impacts automated systems can have has escalated the calls for accountability and transparency. Indeed, knowing how AI-based decisions are made requires both accountability–which, as previously detailed, refers to the ability and exercise of holding AI, its designers, and employers accountable for their impacts–and transparency to users and regulators, which is essential to understand the AI's processes and more effectively address discriminatory and intrusive effects of the technology.

Cities have started taking significant steps to address opaque, unexplainable AI in public sector use to ensure that the tools used are transparent and traceable by key stakeholders. At both national and municipal levels, governments are pioneering efforts to institutionalize algorithmic transparency as a part of their AI strategies.

6.3. HOW ARE CITIES AND LOCAL ADMINISTRATIONS REALIZING AI TRANSPARENCY?

It seems increasingly essential for public administrations to share a set of well-defined, accessible, and understandable information on the development, deployment, and evaluation of algorithmic tools used in public decision-making. This can be done in a number of ways, and this section focuses on (a) the regulation of black-box algorithms and (b) the use of informational registries.

6.3.1. REGULATING THE LOCAL USE OF BLACK BOX ALGORITHMS

To illuminate the use and inner workings of public AI, cities have taken steps to disclose and monitor where and why algorithms are applied. These measures allow citizens to monitor and

contact those in charge of using AI technologies and encourage thoughtful use or limitation of AI.

- In 2018, the municipality of Rotterdam's Data-Driven Working Program investigated an algorithm used to detect benefits fraud. The program discovered that the algorithm "ha[d] been trained with biased data," which was especially alarming because of the model's use of demographic and personal characteristics to gauge if a benefits recipient was a high fraud risk. Moreover, the ways in which decisions were made, and the reasons for certain decisions, were often "impossible to trace," making it difficult for citizens to understand how algorithms played a role in assessing them. Due to explainability and transparency issues, the algorithm has not been cleared for official use.
- In 2020, Portland, Oregon, instituted sweeping facial recognition regulation by passing two ordinances that ban the use of surveillance software by private and public agencies in public spaces. This legislation prohibits the use of facial recognition tools on video from public and private surveillance because of their "non-transparent" use of information gathered from police body cameras and hotel and pharmacy cameras to identify and target people.
- Asheville, North Carolina, created an Office of Data and Performance (ODAP) that informs Asheville residents on how local government uses data to improve its work, drive data-driven decision-making and goal measurement, and manage and govern its data stores in an equitable, secure, accurate, and accessible manner.

6.3.2. INFORMATIONAL REGISTRIES

In addition to overseeing explainability levels, cities and states across the world have begun to respond to calls for greater transparency on how they use and source AI and algorithms. Many have turned to informational registries to understand what types of algorithms are being used and how.

- The Washington state senate put forth a bill that would require public agencies to provide information about automated decision systems in plain language and make the systems and their training data publicly available.
- Amsterdam and Helsinki have created AI registry lists, which provide public records of the data used to train models, how algorithms are used, how outcomes of these models are used by human decision-makers, and what the potential biases or risks associated with the model are. The registry also includes contact information about those in charge of deploying the algorithm and allows residents to get in touch and give feedback on algorithm uses by local governments. Similarly, the Barcelona City Council's AI and data strategy creates a public register of all algorithms used by the city for open review.
- In 2019, the Nantes Metropolis in France opened its algorithms used to make decisions in public service decision-making. Currently, two public algorithms are in use: one to determine the price of public transit, which takes into consideration income and the

number of family members of an individual, and another to determine the social pricing of water for households.

- Since 2016, the city of Antibes, France has managed an inventory of algorithms used by the local government. The inventory lists what and how algorithms are fully or partially used to make decisions and are updated on a regular basis.

6.4. CONCLUSION

Overall, although AI explainability and transparency remain complex concepts that are difficult to achieve, local administrations around the world are experimenting with new ways to ensure that automated systems are employed transparently. These steps are further fueled by calls for transparency by design for AI systems and a surge of data rights protection movements. Thus, even though transparency of AI remains a difficult issue to solve, efforts at the city and local level demonstrate an opportunity to investigate new, innovative ways to provide algorithmic transparency.

7. LITERACY: IMPROVING PUBLIC UNDERSTANDING OF AI

As AI is hands-on and user-facing, meaning that knowledge about how it is created and operated is essential for informed interactions between users and technology. Digital literacy, including its subset of 'AI literacy', is paramount for a holistic understanding of emerging technologies in order to uphold and examine accountability and transparency promises by public agencies.

7.1. WHAT IS AI LITERACY?

Literacy refers to individual and community skills and understanding of a specific field. In the digital era, literacy has evolved beyond the ability to read and write to include competency of online tools, initially, to the social implications of digital systems. Indeed, being 'digitally literate' today requires a comprehensive understanding of digital tools and online platforms, sharing abilities, an ability to keep up with new technologies as they crop up, and the competence to discriminate between ethical and unethical practices. Specifically, AI literacy requires having the "essential abilities" of understanding how and why AI and AI-driven technologies are used in daily life, and what are the consequences of such use.

7.2. WHY DOES AI LITERACY MATTER?

The explosion of AI uses, research, and surrounding literature demonstrates the need for robust AI literacy for all individuals. We argue that students, residents, and policymakers all need improved AI literacy to engage meaningfully with AI, specifically by knowing what AI does, how these tools are being used by governments, what are the risks and benefits of the technology, and how decision-makers are protecting citizens' rights when using AI.

School curriculums across the world have taken steps to teach K-12 students about AI in a culturally responsive manner, enhancing computer literacy courses with an emphasis on ethics and socio-technical considerations of technology. In the below, we point to various legislation taking my cities and states to update their computer science curricula. Guided instruction has been successful at providing learners with a general understanding about AI and equipping them to identify and assess AI bias.

Similar initiatives have been taken for furthering community knowledge of AI. Community initiatives, detailed below, bring families and communities together to experience and learn

about AI. Public hearings and action groups help residents connect with decision-makers over the principles guiding the implementation of AI in their neighborhoods.

Yet surprisingly, steps to improve knowledge about AI for policymakers have been few and far between. As Michael Horowitz and Lauren Kahn note, "top policymakers—who are generally not technically trained—are at an increasing risk of being "black boxed" as technological complexity increases." They go on to warn that illiteracy comes "even at the vanguard of AI research about the "explainability" of algorithms." Thus, in addition to improving regulatory frameworks, accountability mechanisms, and legislation around public AI use, there is a need for a concerted effort by local governments to educate their people and constituents about AI in general.

7.3. HOW ARE CITIES AND LOCAL ADMINISTRATIONS REALIZING AI LITERACY?

As AI grows in popularity and ubiquity, widespread knowledge of its functionalities is required to maximize its benefits and mitigate its risks. Alongside public servant educational programs, localities have taken steps to engage students, educate families and communities, and nurture active citizenship around AI localism.

7.3.1. STUDENTS

Young people entering an AI-driven workforce and society are a prime demographic to educate about the risks and opportunities of AI. To advance awareness about AI and its ramifications, local governments have started to embed emerging technology education into curriculum and fostering extra-curricular activities around understanding AI.

- In Mississippi, the law (MS H.B. 633), encourages the State Department of Education to launch a K-12 computer science curriculum that allows students to learn about machine learning, artificial intelligence, robotics, and technology.
- Taking into account the particular needs and vulnerabilities of children and their interactions with AI, the Beijing Academy of Artificial Intelligence released cross-sectoral, child-centered values that focus on protecting children's health and privacy, preserving their dignity, reducing discrimination, and promoting education and expression of free will. The principles emphasize the need for risk-averse, explainable, and sustainable AI that enhances the development of children.
- The Machine Learning Journal Club at the University of Turin and the Polytechnic University of Turin in Italy brings together students from STEM (Science, Technology, Engineering, and Mathematics) backgrounds to engage passionately and actively in research projects about machine learning and artificial intelligence. Some of these projects involve hackathons, competitions, and published articles.

- The Elements of Artificial Intelligence (AI), run by MinnaLearn and the University of Helsinki, Finland, provides a range of online learning tools and engagement activities about AI. The course's goal is to "demystify AI" by reaching out to a diverse and broad audience that goes on to gain an in-depth understanding of AI, its prospects, and challenges. Self-paced, these courses combine both the theoretical and practical underpinnings of an ever-evolving and intriguing field of inquiry for young learners.
- Another example is Day of AI, launched by MIT Raise. It is an annual event that serves as an opportunity to introduce teachers and students to artificial intelligence. Professor Cynthia Breazeal, director of MIT RAISE, states that students "need not just knowledge of what AI is and how it works, but also the agency to use AI responsibly with confidence and creativity." The program has made close to 4 hours of module content for all age groups ranging from 3 to age 12.

7.3.2. FAMILIES AND COMMUNITIES

Families often constitute a "third space" for AI literacy and education that sits between home and school. Informational and localized toolkits offer parents and children discursive opportunities to discuss and learn about the integral role of AI in everyday life. In this context, it is worth considering several interesting examples of community-based and -led initiatives to foster and encourage family literacy around AI and its applications.

- Beta Blocks is a community-led initiative in the city of Boston that aims to "explore new approaches for community-led innovation in public spaces" to devise a bottom-up definition of 'smart cities' and galvanize civic engagement around AI in the public. Beta Blocks has taken steps to engage families around AI, such as through the 'Robot Block Party.' Organized between the city government, the MassRobotics collective, Toyota, and MIT, the event included more than 4,500 participants and featured 12 robots, ranging from self-driving cars to reboot service delivery tools.
- In partnership with the Raspberry Pi Foundation, the Alan Turing Institute hosted a series of research seminars to teach young people about data and artificial intelligence. In particular, Stefania Druga from the University of Washington reinforced the crucial role families play in fostering AI literacy in her talk. She stated that "AI literacy practices and skills led some families to consider making meaningful use of AI devices they already have in their homes and redesign their interactions with them. These findings suggest that family has the potential to act as a third space for AI learning."
- The Center for Responsible AI designed We Are AI, a five-module course that discusses the basics of AI and facilitates conversations around the social and ethical considerations of AI, as well as AI governance. This course is designed for online or in-person community learning circles to bring together individuals and engage in conversations around AI.

7.3.3. ACTIVE CITIZENSHIP

Active citizenship refers to a context in which citizens are not mere recipients of policies but actively attempt to shape civil society, dismantle discriminatory structures, and seek accountability from governing bodies. In the context of AI, this refers to a public that is tech-aware and data-literate, and that is actively involved in ensuring that AI policies promote the public good and advance the responsible use of AI. Few examples:

- Under the Barcelona City Council Open Digitisation Plan, Barcelona local representatives have developed guidelines for "Ethical Digital Standards." This toolkit, alongside the Open Digitization Plan, serves as an "open source policy toolkit for cities to develop digital policies," including using AI that "put[s] citizens, particularly families and children at the center and make governments more open, transparent and collaborative." The tool serves as an effective means to educate citizens about artificial intelligence and the digital policymaking process.
- Another example is the Ethics and Algorithms Toolkit, a collaboration between GovEx, the City and County of San Francisco, California, Harvard DataSmart, and Data Community DC. The tool kit brings together a variety of actors and stakeholders such as the media, academic institutes, and the broader public community to share and discuss everyday stories that delineate the externalities of algorithm use and their unintended repercussions. Understanding the impact machines have on human life serves as a crucial tool to ensure governments have a robust understanding of the risks associated with artificial intelligence and the best course of action to mitigate such risks.
- Created in Seattle, Washington in 2019, The Algorithmic Equity Toolkit is focused on an action-oriented approach toward political encounters, discourse, and discussion between community members, government representatives, and the broader public thus paving new pathways for AI to serve as an effective intervention. Unlike other policy toolkits, the Algorithmic Equity Toolkit is focused on equipping community members and marginalized communities with the tools necessary to foster effective community building, mobilization, and participation. While most of the other policy tools are largely focused on educating campaigners and policymakers, the toolkit is focused on "non-specialists," and provides resources to further local knowledge and activism for more productive interaction with public comment periods and citizen responses to local AI governance.

7.4. CONCLUSION

All in all, even though AI is becoming increasingly common and even colloquialized within the technology and data realms, it is a relatively novel concept in the policy and governance worlds. As a growing number of local governments and policymakers seek to use artificial intelligence to curate inclusive policies, steps to enhance awareness and literacy about AI and associated technologies like machine learning are ever more critical. In the above, we

have discussed a series of AI literacy examples, drawn from the local level, and ranging from toolkits and digitization plans to learning clubs and family engagements.

The increasing ubiquity of digital technologies—and AI specifically—makes public awareness ever more urgent, so that citizens can understand the benefits and minimize the negative consequences of AI. City-led efforts to encourage and promote education and awareness are playing an increasingly central role, offering a successful example of how AI literacy efforts in families, schools, and the broader public domain can result in greater citizen participation and action.

8. WHAT CAN WE LEARN FROM AI LOCALISM? TEN AI GOVERNANCE LESSONS FROM CITIES FOR NATIONAL AND OTHER POLICYMAKERS

As artificial intelligence (AI) grows more ubiquitous, both its potential and challenges are coming increasingly into focus. How we balance potential and challenge, risk and opportunity, is shaping up as one of the defining questions of our era. Increasingly, cities are at the forefront of exploring this question. In much the same way that cities have emerged as hubs of innovation in culture, politics, and commerce, so they are defining the frontiers of AI governance.

Some models for how cities can take the lead in AI governance include the Cities Coalition for Digital Rights, the Montreal Declaration for Responsible AI, and the Open Dialogue on AI Ethics. More specific examples can be found in San Francisco's ban on facial recognition technology, New York City's push for regulating the sale of automated hiring systems, the creation of New York City's Algorithms Management and Policy Officer, and various local AI Ethics initiatives in an assortment of urban institutes, universities and other educational centers.

In what follows, we offer ten principles to help systematize our approach to AI localism. Considered together, they add up to an incipient framework for implementing and assessing AI localism initiatives in cities around the world.

8.1. TEN PRINCIPLES TO SYSTEMATIZE AI LOCALISM APPROACHES

- **Principles provide a North Star for governance:** Establishing and clearly articulating a set of guiding principles is an essential starting point for responsive and responsible AI regulation. Principles, such as those established by the Emerging Technology Charter for London, an initiative launched by the mayoral office in 2021 to outline "practical and ethical guidelines" for research around emerging technology and smart-city technology pilots in London. Similar projects exist in Nantes, France, which rolled out a data charter to underscore the local government's commitment to data sovereignty, protection, transparency, and innovation with respect to the use of residents' data by the public sector. Such efforts can serve as a North Star for policymakers: they help chart a course

that effectively balances the potential and challenges posed by AI while affirming a commitment to openness and transparency on data use for the public.

- **Public engagement provides a social license:** Establishing trust is essential to fostering responsible use of technology as well as broader acceptance and uptake by the public. Various forms of public engagement (e.g., crowdsourcing, awareness campaigns, mini-assemblies) can help to build trust and should be part of a deliberative process undertaken by policymakers. For instance, to discuss the increasing trend of AI use in hiring and HR, the California Department of Fair Employment and Housing held its first virtual public hearing to discuss the implications of algorithms and technological bias in procurement with citizens and worker advocacy groups.
- **AI literacy enables meaningful engagement:** Fostering familiarity and awareness around AI is a critical component of building informed public engagement. The goal of AI literacy is not only to encourage familiarity with the technology itself but also with associated ethical, political, economic, and cultural issues. For example, the Montreal AI Ethics Institute in Montreal, Canada, a non-profit focused on advancing AI literacy provides free, timely, and digestible information about AI and AI-related happenings from across the world to subscribers.
- **Tap into local expertise:** As hubs of knowledge and innovation, cities are tremendous (if often untapped) repositories of AI expertise. Policymakers should tap into this expertise by, for instance, establishing or supporting research centers, such as the Confederation of Laboratories for Artificial Intelligence Research in Europe (CLAIRE), a pan-European project that takes a European focus on AI uses in cities or funding locally-developed projects such as 'How Busy is Toon', a website developed by Newcastle City Council and Newcastle University in the United Kingdom to provide real-time traffic information about the city center.
- **Innovate in how transparency is provided:** Transparency is critical to building trust and fostering engagement. AI Localism should encompass time-tested transparency principles and practices such as Amsterdam and Helsinki's open disclosure of AI use and explanations of how algorithms are used for specific purposes. In addition, AI Localism can innovate in how transparency is provided, for instance instilling awareness and developing systems to identify and overcome "AI blindspots" and other forms of unconscious bias.
- **Establish new means for accountability and oversight:** One of the signal features of AI Localism is a recognition of the need for accountability and oversight to ensure that principles of responsive governance are being adhered to. Cities are at the forefront of innovating in how to ensure oversight. Examples include New York City's Algorithms Management and Policy Officer, Singapore's Advisory Council on the Ethical Use of AI and Data, and Seattle's Surveillance Advisory Working Group.
- **Signal boundaries through binding laws and policies:** Principles are only as good as they are implemented or enforced. Ratifying legislation, such as the New York City Biometrics Privacy Law, which requires clear notices that biometric data is being collected

by businesses, limits how businesses can use biometric data collected from consumers, and prohibits selling and profiting from the data sends a clear message to consumers that their data rights and protections are upheld and holds corporations accountable to respecting privacy privileges.

- **Use procurement to shape responsible AI markets:** Government procurement is a powerful tool to shape and incentivize ethical behavior. Much as municipal (and other) governments have done in other areas of public life, cities should use procurement policies to encourage responsible AI initiatives. For instance, the Berkeley, California Council passed an ordinance requiring that public departments justify the use of new surveillance technologies and that the benefits of these tools outweigh the harms prior to procurement.
- **Establish data collaboratives to tackle asymmetries:** Data collaboratives are an emerging form of intersectoral partnership, in which private data is reused and deployed toward the public good. In addition to yielding new insights and innovations, such partnerships can also be powerful tools for breaking down the data asymmetries that both underlie and drive so many wider socio-economic inequalities. Encouraging data collaboratives, by identifying possible partnerships and matching supply and demand, is thus an important component of AI Localism. Initial efforts include the Amsterdam Data Exchange, which allows for data to be securely shared to address local issues.
- **Make good governance strategic:** Too many AI strategies don't include governance and too many governance approaches are not strategic. Similar to outlining foundational principles and charters, it is imperative that cities have a clear vision of how they see data and AI being used to improve local wellbeing. Charting an AI strategy, as was undertaken by the Barcelona City Council in 2021, can create avenues to embed smart AI use across agencies and open up AI awareness to residents to make responsible data use and considerations a common thread rather than a siloed exercise within local government.

8.2. CONCLUSION

AI Localism is an increasingly growing area, and both its practice and research remain in flux. The technology itself continues to change rapidly, offering something of a moving target for governance and regulation. That state of flux highlights the need for the type of framework we have outlined here. Rather than playing catch-up, responding reactively to successive waves of technological innovation, policymakers can respond more consistently, and responsibly, from a principled bedrock that takes into account the often competing needs of various stakeholders.

MORE ABOUT AI LOCALISM

The growing use of AI within cities has fueled a phenomenon that we have coined as "AI Localism." AI Localism refers to the initiatives and experiments taken by local decision-makers to establish innovative approaches to AI governance. It seeks to fill in gaps left by incomplete governance at the national level as well as by the private sector. The AI Localism Repository compiles international examples of policy, regulation, legislation, and projects to advance city and state level AI governance.

ABOUT THE GOVLAB

The Governance Lab's (The GovLab) mission is to improve people's lives by changing the way we govern. Our goal at The GovLab is to strengthen the ability of institutions—including but not limited to governments—and people to work more openly, collaboratively, effectively, and legitimately to make better decisions and solve public problems. We believe that increased availability and use of data, new ways to leverage the capacity, intelligence, and expertise of people in the problem-solving process, combined with new advances in technology and science, can transform governance. We approach each challenge and opportunity in an interdisciplinary, collaborative way, irrespective of the problem, sector, geography, and level of government. For more information, visit thegovlab.org.

AI LOCALISM

THE GOVLAB