

TESTIMONY AND STATEMENT FOR THE RECORD

Beth Blauer

Associate Vice Provost, Public Sector Innovation

Johns Hopkins University

“Harnessing AI to Improve Government Services and Customer Experience”

U.S. Senate Committee on Homeland Security and Government Affairs

January 10, 2024

Chairman Peters, Dr. Paul, and Committee Members, thank you for inviting me to participate in today’s hearing. It is an honor to share my views on AI, which reflect my professional expertise and are not a position of Johns Hopkins University or Medicine.

I come to the committee with more than 20 years of public sector experience. My first job in government was as a juvenile probation officer, and when I left government service more than a decade later I was the de facto chief operating officer of the State of Maryland where I had the opportunity to oversee about 85 percent of the state’s budget and a comparable amount of the workforce through an innovative data-driven program called StateStat. Together with leaders across our state we used data to create better outcomes for all Marylanders. Our success led to being named the number one school system in the country, we reduced violent crime to historic lows, dramatically reduced childhood hunger and infant mortality, got the Chesapeake Bay to a healthy tipping point and secured a triple A bond rating for the state despite a national economic crisis. This was all accomplished because we maintained fidelity to a robust data practice and used evidence-based decision-making across the enterprise of government.

In 2015, I founded the Bloomberg Center for Government Excellence (GovEx) at Johns Hopkins University with a simple yet ambitious goal: to help governments, primarily cities, use data effectively to improve the lives of residents across the United States and around the globe. Over the past eight years, GovEx has established itself as a global leader, working alongside thousands of public officials to achieve our goal. We are the go-to resource for governments as they seek to improve data proficiency and develop data-driven policy solutions drawing on the most cutting-edge technology, including AI and advanced analytics. Technological progress can present unique challenges in the public sector, but, where some see barriers, we see opportunities.

AI has the potential to greatly improve the lives of people. However, as you know, irresponsibly implemented AI has the potential to cause harm. I’ve worked to help public organizations realize the benefits of AI and mitigate the risk of harm to residents. Working with governments to shore up the most fundamental data practices has been critical to the early safe adoption of more advanced tools. When public institutions attempt to dive into the deep end of analytics, automation, or machine learning without first building a strong data practice, they run the risk of propagating disparate impacts, poorly informing decisions and creating bad outcomes. Responsible adoption of AI demands high-quality data to understand public operations in new ways. When done right, a strong data practice will lead to safer AI implementation that will anticipate resident needs and promote safer, more efficient, healthier, thriving communities.

We have learned from our work that local leaders are eager to adopt tools and practices that help them respond to constituents’ needs better, faster, and safely. However, there are multiple barriers to exploring the practical applications of AI in their government operations, including but not limited to:

insufficient technical expertise, lack of awareness of the possibilities, budgetary constraints, and reduced research and development funding.

There is already a strong body of evidence that AI in local government can have a positive effect on people's lives. Local governments are integrating AI into their operations. For example, cities are using AI to reduce traffic congestion for both routine traffic and special events and have also used AI to manage traffic to reduce air pollution. Cities are managing their fleets of vehicles by using AI to predict required maintenance and reduce breakdowns, which allows them to ensure prompt service and save money. AI also has the potential to predict the needs of residents before it's too late by identifying households that are experiencing the early onset of economic distress. Early detection allows public agencies to provide immediate interventions that reduce barriers to aid and expedite connecting those in need to the right services. For example, cities have used AI to predict residents who are at risk of becoming homeless by analyzing utility and property violation data. This analysis has led to public interventions coming sooner, keeping families in their homes and saving public resources.

Consistently, local leaders ask how they can create an environment that allows for experimentation while maintaining appropriate protections from harm. To this end here are the five recommendations we advise governments to follow as they think about their AI work:

1. **Designate a senior leader:** This should be a trusted partner who knows when to deploy AI (and when not to); someone who can influence service design; and someone with the power/authority to lead governance and technical strategy while asking the right questions.
2. **Learn with the team:** Create an evolving learning agenda for AI; find out what people in your cities' public, private and academic sectors are already doing with AI; and measure what works – and what doesn't.
3. **Share early guidance:** Be transparent about when AI is deployed; assess the risks; take precautions to avoid bias in delivery; create and communicate an acceptable use policy; and bring people to the table to discuss.
4. **Explore what you already have:** Tackle an issue that matters (but manage the risk). Ideally, this will be a quick win that will relieve the workforce and improve an outcome for the public without disrupting the delivery of key services.
5. **Create space for experimentation:** Integrate a review of your AI implementation into ongoing performance routines; be clear before implementation about what success looks like; track progress; give your team the right balance of space and support.

We are learning as we go. An iterative stance with reflection and learning periods baked into the process creates fertile ground for cautious experimentation. At the same time, some organizations may not be ready to implement AI. For example, some municipalities may learn they need to get their data houses in order before they can use data for predictive modeling. Some may realize they do not have sufficient expertise on staff, and need to hire. Others may find that their policies prevent AI implementation, and they need to address that issue first.

In order to maximize the use of AI in public organizations, governments must start to invest in increasing their capacity for technology adoption across all levels and program areas. This means that we need to attract new skills to the public sector, we need to elevate the capacity of the current workforce, and we need to start investing in the creation of public space to experiment with new AI-powered technologies. In the absence of public leadership, we run the risk of being outpaced and

undermined by the private sector a risk that will lead to more expensive and less effective interventions.

\*\*\*\*\*

The following are responses to frequently asked questions and inquiry in to the practice of AI in local government.

### **How local governments and their workforces are experimenting with and implementing AI-powered tools to improve public services**

Per the [General Services Administration](#), Artificial Intelligence (AI) “refers to the computational techniques that simulate human cognitive capabilities.”

Through our conversations with cities, the Bloomberg Center for Government Excellence (GovEx) is learning that U.S. mayors want to explore the benefits of AI to improve public services and gain operational efficiencies. However, actual implementation is not yet commonplace due to budget constraints, lack of technical expertise, and ethical and legal concerns.

For those cities with which we have relationships and which have laid the groundwork for experimentation – such as designating an AI lead and establishing AI usage guidelines for city staff – implementation starts in lower-risk areas, such as:

- 311 data analysis and optimization;
- writing (meeting minutes, reports, job descriptions, letters, general communications such as press releases, etc.)
- transportation/bus service improvement;
- planning/permitting analysis;
- customer service chatbots.

Among cities that are experimenting with AI, they are using tools that are either free of charge (such as ChatGPT) or embedded in their existing enterprise software (such as Zoom’s AI Companion). In some cases, they are developing proprietary algorithms, software, etc. Usage, however, must adhere to guidelines established by the cities’ technical leadership, and staff are expected to be transparent in their use of these tools.

Examples of cities with guidelines that allow for experimentation with AI (specifically, generative AI) are: [Boston](#); [San Jose](#); [Seattle](#)

### **On the ground: what local leaders are asking**

We learned similar lessons during the three years we spent constructing a timely federal and global Covid-19 dataset. It is necessary to embrace the fact that we do not know what we do not know, and we need to make space to learn, to be nimble, and, when it may be necessary, to pivot.

### **How we can apply these lessons at the federal level**

The [November 1, 2023 Office of Management and Budget draft memorandum](#) that is now open for public comment provides risk management and accountability guidance. This followed the [October 30,](#)

[2023 Executive Order on AI](#), and both build on [OSTP's October 2022 Blueprint for an AI Bill of Rights](#) and [NIST's January 2023 AI Risk Management Framework](#).

Together, these documents reflect the same basic principles outlined above. This work requires leadership, learning, sharing, transparency, experimentation, and accountability, and the more we foster this kind of environment, the better positioned we will be for safer implementation of AI.

### **The importance of responsible experimentation of new technology and how the federal government can best foster collaboration around AI-driven solutions**

Experimentation is essential for determining the value and evaluating the behavior of AI driven tools. The performance quality and behavior of tools that rely on AI can vary wildly depending on the input data, the environment, and the amount of effort spent on fine-tuning. Consequently, governments should be prepared to run realistic experiments with real-world data or to run pilots in order to evaluate the value proposition of AI tools. Allowing a third-party tool access to data or running new tools in government systems will, of course, incur some risk to either privacy, security, or safety depending on the use case and the data.

With that understanding, the importance of responsible experimentation of new technology cannot be overstated, and we appreciate the OMB's attention to risk management. Not all experiments carry the same risk, and should not be held to the same standard of caution. For example, experiments involving personally identifiable information (PII) should have greater protections in place than pothole data that might be used to predict which streets will need service in the winter. Since risk levels vary, creating consistent guidelines for an AI risk assessment would be useful. I am encouraged that the OMB draft includes as a requirement detail about the training data provenance and characteristics (section 5.c.iv.3). If all the agencies and governments present a united front that this is a necessary disclosure in order to consider a vendor, even under an NDA, that will encourage transparency.

To foster collaborations around AI-driven solutions, we should:

- Fund research and development through public-service oriented institutions, such as local governments, universities, hospitals, and nonprofits.
- Create a community of philosophically and demographically diverse peers and encourage them to share their experiments, their successes, and most importantly, their failures. Relatedly, we should cultivate a culture where failure is understood to be an important part of the process. Failure – especially fast failure – provides learning opportunities that lead to success.
- Create a framework for diverse peer review. Attempts by different teams to replicate experimentation lead to reliability; promotes transparency; reduces bias and harm.
- Embed ethicists into the experimentation process.
- Invite public review of the process.

Wherever possible, we must break down silos. They often result in groupthink; errors; intimidation; and duplication of effort. It is better to share and receive helpful critique than to perpetuate mistakes, bias, and harm.

### **The importance of engaging the public when developing customer-facing tools to ensure services are equitable and accessible to all communities**

Engaging with the public is our best way of understanding how to meet them where they are. The purpose of this exploratory experimentation is to deliver better, faster service to constituents, but if we don't understand the needs of the community, we can't do that. For example, if a community does not have access to broadband internet, implementing AI-driven tools like chatbots that are exclusively available online will reduce access to services.

Governments also need to be open to hearing community grievances around historic (or current) government failures to provide equitable services. A community will not take our concern over equitable treatment via new AI tools seriously if they feel that the government has been unable to fairly provide these services that the new AI tools are supposed to help them access. Our use of AI must be routinely informed and centered on the needs of the public.

We should be transparent about the tools under development, who's building them, what data is being used, what they cost, and how the work is prioritized. Without providing the opportunity to ask questions in public forums, local leaders are not building "with" their constituents. They are making best guesses and assumptions, and run the risk of shutting people out of the system, paternalizing interventions that are not a match to need, and creating outcomes that drive people further away from public trust.

It must also be stated that generative AI specifically is not a mature technology. Consequently, we must exercise care and caution, especially when safety, high stakes services, or PII are involved. We should not implement customer-facing AI in situations where the failure modes are not extremely well understood. Nor should we implement it in situations in which the scale of deployment is large enough to prevent thorough review of the interactions by humans. Expert judgment is necessary in these early-stage experimentations.

Ways that the government can prepare its data infrastructure and workforces for adoption of AI tools. In addition to the previously listed suggestions, government can prepare its data infrastructure for the adoption of AI tools by:

- Develop a data maturity model that is consistent across government entities, and encourage government agencies to advance along the continuum until they are at a point where they can responsibly consider the adoption of enterprise AI.
- Develop performance metrics and systematic ways of monitoring them that are consistent with the data maturity models. If you cannot currently measure, track, and report on your performance and work quality, you will be unprepared to evaluate an AI driven tool or to provide the level of oversight recommended by the OMB memorandum.
- Develop professional development opportunities to build foundational data literacy and capacity within the civic workforce. GovEx and Baltimore have prototyped this with the [Baltimore Data Academy](#).
- Establish real-world data sets and environments that can be used to evaluate equity, safety, and quality of AI driven tools.
- Develop a framework for connecting datasets to each other. GovEx's work with Covid-19 data showed that there are significant inconsistencies in how local governments, counties, states, and the federal government define, report, and share data. Respecting the need to maintain data for longitudinal data, we could develop universal keys that connect data from one source to another, but that starts with creating the appropriate infrastructure within each data set.
- Create an on-ramp for research, development, and experimentation to make its way into cities' hands such that they can practically apply it in government operations.

- Convene roundtables of practitioners from varying levels of government to surface any policy or practical impediments to implementation. The first step to removing impediments is knowing that they are there, and the second step is to collaborate on removing them.