

A SNAPSHOT OF ARTIFICIAL INTELLIGENCE PROCUREMENT CHALLENGES

Diagnosing perceived and actual risks impeding responsible AI acquisition in government

Chloe Autio

Kate Cummings

Brinson S. Elliott

Beth Simone Noveck

Contents

SUMMARY	2
THE AI OPPORTUNITY FOR GOVERNMENT	4
PURPOSE OF THIS PAPER	5
METHODOLOGY	6
LIMITATIONS OF THE STUDY	7
CHALLENGES AND RISKS WITH AI USE IN GOVERNMENT	8
AMBIGUOUS DEFINITIONS OF AI	8
UNCERTAINTY	9
TRANSPARENCY AND ACCOUNTABILITY	9
PERCEPTIONS OF AI IN GOVERNMENT AND THE IMPACT ON PROCUREMENT	10
RISK PERCEPTIONS ABOUT AI VARY	10
RISK OF OVER-PROMISE	12
EFFECTS OF NEGATIVE PERCEPTIONS OF AI	12
MISUNDERSTANDING AND MISTRUST CREATE COMPOUND CHALLENGES	13
AI PROCUREMENT CHALLENGES	14
TECHNICAL FAMILIARITY AND AWARENESS	15
VENDOR RELATIONSHIPS	16
DATA GOVERNANCE	17
ORGANIZATIONAL CHALLENGES	18
AI GOVERNANCE AND OVERSIGHT	19
GOVERNMENT PROCUREMENT OF AI ENCOUNTERS MANY HURDLES	21
SUMMARY OF BARRIERS	22
RECOMMENDATIONS	23
TRAINING	23
TOOLS	25
RELEASE MORE OPEN DATA AND SIMULATED DATA SETS	26
REGULATION AND GUIDANCE	27
ORGANIZATIONAL CHANGE	29
NARROW THE EXPERTISE GAP	30
SUGGESTIONS FOR FUTURE RESEARCH	32
AN EXPANDED USE OF AI IS POSSIBLE	34
ACKNOWLEDGEMENTS	35
ABOUT US	36
WORKS CITED	37



SUMMARY

Artificial intelligence (AI) promises exciting new opportunities for the government to make policy, deliver services and engage with residents. But government procurement practices need to adapt if we are to ensure that rapidly-evolving AI tools meet intended purposes, avoid bias, and minimize risks to people, organizations, and communities. The euphoria around new technology demands a closer look at the question of how the federal government—the largest purchaser of technology in the country—acquires AI, with emphasis placed on strategies that enable rapid and responsible acquisition, evaluation, and testing of these rapidly-evolving technologies.

THIS PAPER LAYS OUT FIVE DISTINCT CATEGORIES OF CHALLENGES RELATED TO PROCURING AI:

1. General and Technical Familiarity and Awareness
2. Vendor Relationships
3. Data Governance
4. Organizational Challenges
5. Governance and Oversight

DISCLAIMER

NON-ENDORSEMENT: This independently authored paper was supported by a grant from Google Cloud. We discuss technology and procurement in general terms without reference to or endorsement of specific companies or products. The interviews that form the basis for our research were conducted with public sector officials who have asked not to be named. We did not interview private sector professionals for this paper.

AUDIENCE

This paper is broadly aimed at people inside government who hope to use artificial intelligence (AI), including machine learning, and natural language processing, to improve government operations and impact. Focused on procurement, the paper will be of most interest to managers who want to explore AI procurement, but it will also be useful for technologists, policymakers, contracting professionals, and government employees who might potentially interact with AI systems. Additionally, this paper offers insights to policymakers, researchers, vendors interested in government modernization and digitization, and vendors of AI systems.

No prior technical knowledge is required to read this paper. This is a general interest snapshot, intended to contribute to the broader literature and debate about bringing government into the 21st century. The paper complements other work on the benefits and risks of machine learning and shines a light specifically on the difficulties of acquiring even the most promising tools. These impediments to acquisition include many of the same issues that make acquiring any new technology difficult, along with specific concerns about AI and machine learning technology, in particular.

1. Chloe Autio and Brinson Elliott are a Director of Policy and Client Team Leader, respectively, of the Cantellus Group.
 2. Kate Cummings and Beth Simone Noveck are a Fellow and Director, respectively, of The Governance Lab.
 3. For purposes of brevity and clarity, this paper uses the term "AI" broadly to capture machine learning and other related AI use cases. We define AI as a "machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments," using the definition from the National Artificial Intelligence Act of 2020.



Our Recommendations to Accelerate Responsible AI Acquisition in Government Include:

TRAINING

Invest in building public sector capacity to differentiate between high- and low-risk opportunities, through teaching approaches to scoping, acquiring, testing, and de-biasing well-tailored AI solutions. Especially important is the need to train public professionals in problem definition to ensure that procurement solutions are responding to a true need. Procurement processes today rarely take performance and evaluation into account, highlighting a need for public servants to be taught to assess the outcomes of an algorithm in addition to the inputs. Frequently updated training will help familiarize public servants with the opportunities and challenges of AI in a fast-changing environment.

TOOLS

Invest in and update decision frameworks, contract templates, auditing tools, and pricing models that enable procurement officers to perform AI acquisitions with confidence. Release more open data and simulated datasets. Simulated datasets alongside risk-management frameworks will help with testing algorithms and spotting discriminatory effects.

REGULATION AND GUIDANCE

Recognize that AI is a broad spectrum. Some AI use cases are more complex than others and should be treated with more consideration, while others are fairly straightforward and can be acted upon more quickly. The federal government (Federal Acquisition Service working together with the National Institute of Standards and Technology and the Office of Science and Technology Policy) needs to develop a frequently-updated typology for acquisition professionals to enable them to invest time in more problematic cases while streamlining the acquisition of tools that are less likely to impact fairness, equality, or liberty.

ORGANIZATIONAL CHANGE

Establish and support collaboration, knowledge sharing, and coordination between procurement officials, technologists, and policymakers and integrate mechanisms for public input. Solving the challenges of AI demands multidisciplinary perspectives and expertise.

NARROW THE EXPERTISE GAP

Invest in integrating those with expertise in new technologies into government beyond the offices of the Chief Innovation Officer (CIO) or the Chief Technology Officer (CTO), including into procurement, legal, and policy teams. Foster linkages with academia and expand effective fellowship programs to speed up the acquisition of relevant talent, especially those with the ability to audit the outcomes of AI use. Create these programs in state and local as well as federal government. Greater expertise is needed both to 1) support the procurement process, and 2) support responsible implementation and evaluation.



THE AI OPPORTUNITY FOR GOVERNMENT

During the COVID-19 pandemic, several United States federal government agencies turned to artificial intelligence (AI) to analyze government preparedness for the next wave of the coronavirus. Officials compiled policies related to biodefense preparedness dating back to the 2009 H1N1 flu pandemic, along with relevant White House executive orders and recent COVID-19 policies. The analysts employed Natural Language Processing (NLP), a common machine learning (ML) methodology, to process document texts in order to distill key learnings.


NLP is excellent at analyzing large quantities of human text or voice data to extract meaning, content, intention, attitude, and context. The NLP in this case linked documents and revealed their consistencies and inconsistencies. The use of NLP in this context allowed federal agencies to extract proven best practices and resolve conflicts in and between their current and new COVID-19 policy documents. The information was valuable because the evolution of the pandemic required frequent updates to policies.

What began as a policy analysis use case, spurred by a health emergency, has turned into a transferable AI-based solution that governments of all sizes can implement. Such solutions can extract key insights from large quantities of information and expedite turning insight into action, therefore improving government efficiency overall.

The significant government investment in AI is predicted to save 1.2 billion hours and \$41.1 billion.

The coronavirus experience is but one example of the compelling value that AI and machine learning offer to governments. Research finds that these technologies can automate and streamline core government processes to promote cost and time savings, expedite operations, and improve service delivery to expand access and provide critical services to individuals and the public at large. The significant government investment in AI is predicted to save 1.2 billion hours and \$41.1 billion.

The promise of AI in the public sector is two-fold: helping public servants work more efficiently, and increasing the quality and reach of government services. For example, the US Bureau of Labor Statistics (BLS) has employed AI to sort and more quickly process survey responses about workplace

 4. An AI audit generally refers to a process by which the builder of an AI model or a third party dissects the model to understand the way in which it was optimized (i.e., the weight or probability assigned to certain characteristics or data points), what data was used to train it, for what purpose it was built for, and what confidence and evaluation metrics make it fit for purpose. For a good resource on AI auditing, check out the UK Information Commissioner's Office reports on this topic and a publication by leading AI researchers.



illness and incidents submitted by employees across the country. Before introducing the AI sorting tool, BLS employees were hand coding and evaluating survey responses from hundreds of thousands of employees in the US, a rote, time-intensive task that impeded BLS public servants from doing more thoughtful and action-oriented tasks such as analyzing survey results or reaching out to survey submitters for follow-up. With the AI coding tool, the processing of survey responses was expedited and simplified, allowing them to engage with the public more quickly.

A 2020 report co-authored by faculty at Stanford University and New York University found that roughly half of the 142 federal agencies they studied are pursuing, piloting, or operationalizing AI tools and techniques

An increasing number of government agencies are acquiring AI in the hopes of realizing such benefits. A [2020 report](#) co-authored by faculty at Stanford University and New York University found that roughly half of the 142 federal agencies they studied are pursuing, piloting, or operationalizing AI tools and techniques.

Despite the promise of AI, a number of gaps in policy, knowledge, and procurement practices impede AI procurement. This paper explores the practical challenges associated with government AI procurement and AI adoption, and calls attention to the unique barriers that governments face in the acquisition of AI.

PURPOSE OF THIS PAPER

Much has already been written about AI: what it is and is not, its risks and opportunities and key considerations that must be made as the technology is developed and deployed. The present paper aims to go a layer deeper, to diagnose and examine problems with AI purchasing in government. We have analyzed both actual and perceived barriers to AI acquisition. We looked at the technical, awareness-related, logistical, institutional, market, and governance problems that slow or impede procurement, including:

- Lack of familiarity with AI and know-how to execute AI projects
- Inconsistencies in government data governance and infrastructure
- Organizational and institutional dynamics that impede knowledge sharing and collaboration
- Lack of clearly defined processes and frameworks for technology oversight and governance



METHODOLOGY

This paper is informed by bench research and long-form structured interviews. We interviewed 11 state and federal public servants between January and March 2022. We then conducted two 90-minute workshops with 15 leaders at the state and local levels and with 10 leaders at the federal level to get their reactions to the paper. Following those sessions, we have integrated the key points into this paper. A complete list of interviewees is listed in the acknowledgments section.

Individuals were selected based on their roles in expanding AI use in government or their involvement with procuring technology for government agencies. Some are directly involved in the acquisition process in Contract or Acquisition Officer roles or in providing legal advice on procurement transactions. Others are a part of teams within agencies or departments that are reviewing or deciding whether to use commercially developed AI solutions.

To create substantive consistency across the original interviews, we focused our questions on data-intensive uses of AI for automated processing applications that were intended to increase efficiency and expand insights. The use cases we considered include:

- Internal/Operations Facing
- External Customer/Constituent Facing
- Research and Policy
- Public Agency Audit/Accountability

During the subsequent workshops later in 2022 and early 2023, we asked participants:

1. Have we accurately captured how governments procure AI/ML technologies?
2. Have we adequately described the benefits and risks of procurement processes?
3. Are the recommendations for what needs to change or stay the same adequate?

We did not speak with companies selling AI-based solutions to governments, nor did we conduct interviews with other non-governmental organizations.

We conducted a review of legislative and regulatory proposals pertaining to AI procurement and reviewed relevant reports regarding government AI procurement available at the time of drafting.

LIMITATIONS OF THE STUDY

Although our discussions were rich, they were time-limited. Therefore, the findings in this paper are intended as a starting point for discussion, rather than conclusive empirical research. Some of the study's limitations include:

- Time constraints limited the number of officials we could interview.
- We primarily spoke to federal officials in the first drafting of this paper. While we did speak with 15 officials at the state and local levels to receive feedback, the results of this research may not be generalizable to the state or local level.
- We did not speak to a large enough sample of defense and civilian agency officials to enable us to draw conclusions about their different levels of risk tolerance.
- Most of the individuals we spoke with are knowledgeable in the field of technology or AI, which enabled them to contextualize their assessment of AI's risks and benefits. Ideally, we would have interviewed public servants who are not technical experts to better understand their perceptions.

Suggestions for further research appear near the end of the paper.



CHALLENGES AND RISKS WITH AI USE IN GOVERNMENT

Before unpacking the hurdles to government AI procurement, we provide a high-level overview of some core AI risks and how they apply to government in particular.

AMBIGUOUS DEFINITIONS OF AI

One confounding issue with AI is the absence of clear definitions or consensus around what these technologies are and how they are used. Several respondents noted that this gap complicates acquisition processes, such as drafting contract language that accurately captures the technical capabilities of the solution while also managing its risk. This ambiguity likewise increases the chance that procurement officers struggle to map AI products and services to appropriate

problems and to policy requirements. One of the resulting drawbacks is that AI acquisitions might not undergo proper pre-solicitation disclosures and assessments when needed.

One confounding issue with AI is the absence of clear definitions or consensus around what these technologies are and how they are used

Another risk associated with AI is that government employees lack AI experience and cannot evaluate the validity of AI-driven results as well as those trained in data science. For example, some emergent technologies bring advanced AI directly to agency staff who have minimal training in AI, such as “no code” AI.

Whereas previous AI development was performed by a trained data scientist who knew how to carry out robust checks, the solutions run by less trained personnel raise new questions regarding what “procurement” of these technologies looks like and who may evaluate whether they are suitable before acquisition. Chief Data Officers and innovation leaders need to classify and define these systems to manage risks, which may not fit in with the formal software procurement processes.

Some officials offered their own typologies for AI technologies while highlighting the challenges that arise without a broadly shared categorization. For example, they suggested that we can distinguish between simple and compound AI, embedded and stand-alone AI, and other distinctions that could help us to distinguish between more and less risky AI from the standpoint of civil liberties.

Bias and Unfairness

When machine learning models are poorly designed, based on incomplete or low-quality data, or excessively relied upon subsequently for decision-making in ways that go beyond their predictive capabilities, they can present risks of unintended bias.



If trained on poor quality or under-representative data, a system's predictions, insights, and decisions can reflect and amplify biases in society or in the institutions that collected the data.

If trained on poor quality or under-representative data, a system's predictions, insights, and decisions can reflect and amplify biases in society or in the institutions that collected the data. AI systems that preserve or replicate harmful bias can violate the principles of equitable treatment and non-discrimination, key government responsibilities that are legally owed to and expected by its citizens.

To avoid or mitigate the risks of bias, organizations need to explicitly define what fairness means in the context of their duties—for instance, defining protected classes that require equal treatment—and seek both data sets and parameters that enforce fairness prior to setting up their models. From there, civilian agencies need visibility into the system's training and test data and the logic of models that underlie the system's insights and decisions. However, even AI systems trained on fully representative data, and designed with and for a diverse user base, may have unintended bias; organizations will need to consider risks and trade-offs before proceeding.

UNCERTAINTY

Some AI is used for prediction, and uncertainty is a feature of most predictions. One illustration of this problem emerged during the pandemic when public health authorities frequently relied on algorithms that failed to predict the spread of the virus accurately.

Uncertainty and error in an AI system can arise at any point in the AI lifecycle. The errors might result from inaccuracies in the technology, limitations of the data, human factors, or the product of how data, humans, and technology interact. Given the scale and importance of some systems, even small errors can result in substantial impact, magnifying the importance of being able to understand both the prediction and the means used to arrive at it.

While uncertainty is a feature of all predictions, two core components of managing uncertainty are measuring it and deciding how much uncertainty is acceptable across different AI applications and deployment contexts. While there are no agreed-upon or common legal standards dictating what levels of uncertainty or error are tolerable, agencies can incorporate their own processes to set an acceptable level of uncertainty and manage it over time.

TRANSPARENCY AND ACCOUNTABILITY

Lastly, certain AI, such as risk recidivism algorithms used to make sentencing decisions or to power driverless cars, algorithmic hiring tools used to give or deny people a job, or predictive analytics used to deliver or deny services, can alter the



government's relationship with its residents and the fairness and due process owed to them.

Even though those algorithms may be less rather than more biased than human decision-makers, government procurement needs to include a robust risk-benefit analysis to ensure that systems and tools align with agency mandates, civil rights and liberties. In particular, procurement decisions must uphold democratic duties to promote and protect the public interest, safety, and security. In some cases, results from AI models have replaced the judgments and decisions of government employees. As a result, agencies should disclose the uses of technology tools that may impact citizens.

Although AI systems can increase efficiency, decrease risk, and eliminate human bias, they can also create new opportunities for error. By acting without human intervention, decisions made or facilitated by a so-called "black box" algorithm⁵ may not be easy or possible to properly fix, contest, or appeal. As with any new technology, the procuring entity must conduct a robust assessment to ensure that the use of the technology aligns with the agency's mandates and that the agency has the resources to dedicate to the proper oversight of the technology. These decisions could have severe consequences for human safety, a defendant's right to due process in law enforcement settings, or assessments about creditworthiness.

PERCEPTIONS OF AI IN GOVERNMENT AND THE IMPACT ON PROCUREMENT

The previous sections have shown that AI is a complex and ambiguous technology whose benefits and risks are important to assess prior to and after deployment. In addition to these challenges, our interviewees suggested that negative or limited perceptions of AI among managers and potential users have real impacts on the government's ability to leverage this new technology. Respondents all agreed, however, that recent increases in everyday and government uses of AI have positively contributed to a general understanding of what AI is and is not, and that curiosity about AI is increasing overall.

RISK PERCEPTIONS ABOUT AI VARY

Rapid increases in the development, use, and deployment of AI have put its promise and perils on full display. Public concerns about AI include workforce displacement, surveillance, malicious hacking, digital privacy, and a lack of human connection. These worries co-exist with excitement over previously

⁵ "Black box" algorithms are those in which the inputs and processes are hidden from public view, which present issues for democratic decision-making.

⁶ We asked respondents about their technology familiarity and that of their colleagues. Most respondents broadly understood AI tools, and about half reported having a high degree of technical proficiency, familiarity, and comfort with AI/ML.



discussed benefits and opportunities of AI for the government, including societal improvements and cost and time savings.

Most respondents interviewed for this paper perceive AI as only slightly riskier than other technologies, with risks stemming from a lack of clarity both in definitions of AI and its constituent methods, including ML, and in the understanding of how they are used.

Most respondents interviewed for this paper perceive AI as only slightly riskier than other technologies, with risks stemming from a lack of clarity both in definitions of AI and its constituent methods, including ML, and in the understanding of how they are used⁶. Those who felt especially confident in their understanding of AI technologies said they viewed AI as no riskier than other technology, explaining that risk largely depends on the use case, context, or application and is not inherent to the AI system or tool itself. One respondent suggested that acquiring AI is not dissimilar from the shift made in the recent past to acquiring cloud-based Software as a Service (SaaS).

Given the technological familiarity and know-how of most respondents, we asked them to discuss their own perspectives about AI risk as well as AI risk perception among their peers. We saw some anecdotal evidence of differences in risk perception across agencies. Civilian agencies like the Department of Transportation (DOT) describe a more conservative approach to AI procurement, given the agency's mandate to protect public safety; for example, risks to safety and resiliency are particularly salient in autonomous vehicles. The same is true for the Department of Labor (DOL) as they deal with human and workforce rights, disability, and discrimination.

Those we spoke to at the Department of Defense (DOD) suggested that the pressure to be competitive in AI outweighed the potential risks. "The risk of inaction is greater than the risk of action," said one respondent in the military. However, while other officials agreed, some found it hard to quantify the risks



given the quickly evolving nature of the technology and its ever-increasing number of uses in government and the economy.

RISK OF OVER-PROMISE

Importantly, all respondents emphasized the danger of an over-reliance on AI or a false belief that all systems are accurate, objective, and fair. For this reason, some respondents gave more scrutiny to acquisitions brought forward by those who were less familiar with AI, and who might be more overconfident in the abilities of these systems.

Other individuals and organizations, despite lacking technical or AI expertise, were acutely aware of the possibility for AI to discriminate or preserve bias when beginning an AI project. In one example involving AI employment and hiring applications, a federal civilian organization understood that systemic bias is inherent to hiring and had the propensity to be amplified by algorithms. This prompted the organization to engage technical leaders to provide guidance on the project and tools for performing technical checks for bias.

While this organization sought technical help to mitigate the risks and over-promise of AI, not everyone working with AI knows to consider these issues. Furthermore, the use of AI requires a continuous cycle of deployment and learning, requiring evaluation over time. According to one former federal employee working on labor-related AI applications, this makes it necessary for public servants and anyone working with AI to know that all systems, even those that are designed to remove AI bias, may present risks and biases that need to be proactively anticipated and managed.

EFFECTS OF NEGATIVE PERCEPTIONS OF AI

This lack of understanding of AI and ML, intensified by frequent coverage of AI risks in the media and academic literature, can impede the adoption of beneficial tools.

Respondents attributed colleagues' negative perceptions of AI to its ambiguity and rapid evolution, a lack of training or education, and concerns about bias or discrimination. This lack of understanding of AI and ML, intensified by frequent coverage of AI risks in the media and academic literature, can impede the adoption of beneficial tools. Concerns about the issues covered



in the media “are what can make people more nervous” when using AI compared to other software, said a former Head of Artificial Intelligence at Technology Transformation Services (TTS) within the General Services Administration (GSA). “There are real concerns with AI development and deployment, and while it’s software at the end of the day, it’s important to have all who play a role in AI development build an understanding of the implications of their choices as they build an AI—and to make sure those processes and checks are built into their workflows so that they actually can implement those needed checks.”

It is unsurprising that officials would opt to not use something they do not fully understand or that introduces new risks. As one respondent noted, senior military commanders know they will be called to testify before Congress about why they relied on a tool or system if something goes wrong. This caution serves as a significant deterrent to adopting AI in places where personal accountability is crucial.

MISUNDERSTANDING AND MISTRUST CREATE COMPOUND CHALLENGES

A lack of understanding about AI can impede the vetting and auditing of new tools, such as the use of AI to spot fraud, waste and abuse in government spending and, more importantly, in uses that have a profound impact on life and limb, such as risk recidivism algorithms⁷. Uninformed judgments can lead to the procurement and use of AI tools that sometimes do more harm than good.

At the same time, this same lack of understanding and mistrust can lead to agencies not procuring lighter-weight AI solutions that can deliver much-needed efficiencies. Officials interviewed said that the lack of a shared understanding of the various typologies of AI may have led to missed opportunities and unnecessarily slow response times.



AI PROCUREMENT CHALLENGES

Government procurement of AI technologies and services is often hampered by bureaucracy and complex procurement regimes with rigorous specifications, which limit options to a small number of approved vendors. This combination makes for a cumbersome technology acquisition process that typically takes between four and 24 months. Companies new to this landscape, or without the skills to navigate it, face organizational and logistical challenges simply entering or getting approved by the government vendor system. Given the relative novelty of AI technologies, many of which are produced by young startup companies, this paperwork-intensive process is misaligned with the speed at which all technology—including AI—evolves.

Across the board, government buying needs have expanded to include not only physical “things” (e.g. missiles, computers, cars) purchased from traditional industrial suppliers, but also software. Many IT and Software as a Service (SaaS) products are not off-the-shelf solutions that can be purchased as-is, but instead require customization, ongoing updates, security patches, and monitoring. Lack of comfort with or understanding of SaaS business models had broad impacts on the acceptance and use of AI services across government and, initially, increased hesitancy among public servants to take advantage of AI services. Similarly, while some AI technologies are deployable “out of the box,” most require additional customizations.

Technology will never be less data-intensive than it is now, and therefore never less complex to manage.

Technology will never be less data-intensive than it is now, and therefore never less complex to manage. The challenges of staying on top of frontier technologies are only accelerating.

The following sections diagnose some of the novel challenges unique to the procurement and use of AI in government, and crystallize them into the following categories:

- Technical familiarity and awareness
- Vendor relationships
- Data governance
- Organizational tendencies
- AI governance and oversight



TECHNICAL FAMILIARITY AND AWARENESS

Across those interviewed, a primary concern regarding AI procurement was the purchasers' lack of familiarity and comfort with AI technologies at both a conceptual and practical level. The poor level of understanding among government purchasers mirrors the lack of understanding among the general population described in a previous section. Particular problems include:

- **The very terms “artificial intelligence” and “machine learning” are poorly defined, and staff sometimes do not understand the need for clarity in these definitions.** Some respondents acknowledged awareness of broad principles and concepts such as “Transparency” and “Explainability,” but expressed frustration over a lack of usable tools to apply these principles in practice. The resulting vagueness hampers the work both of officials who bring products to the attention of procurement teams, and of the team members themselves engaged in reviewing AI contract terms. New or exacerbated problems also turn up when staff try to define how a product will be used. Lack of clarity can often create a negative feedback loop when acquisition teams are trying to get on the same page about the intended use of an AI tool and why it is being acquired.
- **AI talent is scarce and highly competitive, even in the private sector.** Such expertise is especially sparse in government, which does not compete with tech sector salaries and benefits. Universities are producing more graduates trained in AI, but there is still a skills gap in the public sector workforce⁸. Although fellowships and rotations help, they do not provide adequate talent needed to improve procurement, let alone the ability to build, deploy, and maintain complex AI systems at scale. One state official shared that without reliable expertise, they cannot procure these new technologies.
- **The buzz around AI has created unrealistic expectations that machine learning will solve problems that are not well understood.** Some poorly scoped AI solutions have been brought into government simply because someone on the project wanted to use the latest technologies or based on unrealistic expectations for what “black box” algorithms can do. Albeit AI tools can often help ingest and make sense of large quantities of data, if the data is of poor quality or the process is poorly designed, no amount of technology will fix the problems. “Don’t acquire AI for the sake of acquiring AI,” several respondents reiterated.
- **Knowledge gaps in the basics of commercial transactions and business incentives prove particularly problematic for government AI acquisition.** More than one federal respondent, including an Acquisition Director and a Program Manager at the Defense Innovation Unit, reported these circumstances in their organizations. Many individuals, while desirous of procuring an AI solution, had little to no familiarity with business obligations and processes such as quarterly reporting, product timelines and deliverables, and navigating licenses and

Universities are producing more graduates trained in AI

8. PIF AI track and USDS



intellectual property issues. The individuals particularly lacked insight into the impacts of these dynamics on small businesses. As an example, a Program Manager within the DOD explained that some of the acquisition leads did not understand how to license and price certain AI solutions, or why these practices were necessary.

- **Inadequate problem definition creates unique challenges to scoping, tailoring, and ultimately carrying out successful acquisitions.** Respondents reported that one of the key challenges to procuring AI effectively is getting public servants to properly define the problem they are solving with the AI solution. Often public servants simply describe the desired features of the software, which are often pulled from vendor descriptions. Eight of the 11 respondents reported that one of the biggest issues with procuring AI was a lack of fit between the AI capability or application and the problem at hand, due to poorly scoped problem definition. One federal official shared how poor problem definition can lead to procuring more technical and expensive solutions than needed to solve a problem. Some colleagues struggle to map ML to a clearly defined problem or use case, reported one respondent working in a defense agency.

VENDOR RELATIONSHIPS

Without a sophisticated understanding of AI technologies and how they can be applied to problems, procuring agencies can tend to over rely on particular vendors which may lead to vendor lock-in. One former GSA employee noted that nine out of ten times, offices and departments do not have the data and tools needed to evaluate AI solutions. Because AI causes computers to learn and evolve, there is too little understanding, sometimes even by the AI's creators, of what is going on “under the hood” and why an algorithm spots a pattern, generates content or organizes information in the way that it does. Some consequences include:

- **When agencies find that their problems are not solved by their chosen AI technologies, or they outgrow the technologies, many balk at submitting a new solicitation and exploring additional technology solutions, even if a new solution is more viable.** In such situations, procurement officers often defer to their current vendor, who may inaccurately claim that only the vendor who stores the data can access and process it, and which can lead to vendor lock-in and over-reliance on one limited solution.
- **Public professionals over-rely on the vendor.** Bringing a vendor into the picture can lead to obfuscation and inefficient use of the product. Many commercial contractors do not work directly with the agencies that use the solutions they develop, so the government end users may lack access to the vendor needed to work with the product or iterate changes throughout the system lifecycle.

**Current AI offerings
are rarely designed
specifically for the tasks
for which governments
want to use them**

⁹ The 2020 report *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies* found that commercial contractors are presently under no obligation to consult or collaborate with the public institution to whom they intend to sell their algorithm. Moreover, due to the significant costs associated with externally sourcing algorithms, government agencies often grant commercial contractors significant design autonomy.



- **Current AI offerings are rarely designed specifically for the tasks for which governments want to use them.** If the AI products are designed for commercial use instead of government solutions, the government might not understand or be able to exert an impact on the design and evolution of the tool⁹. Although many use cases are very straightforward, such as translation AI, call center AI, and chatbots, challenges lie in acquiring AI for complex solutions that were not uniquely tailored, designed, and trained for the agency's problem set and its context for deployment, including responsibility for and accountability to constituents.

The complicated environment public servants must manage with vendors, and the break in the flow from AI acquisition to implementation, puts public servants in an operational dilemma. They are to decide between two difficult and complex options: on the one hand, it is often too expensive or cumbersome for governments to develop advanced AI tools themselves, but on the other hand it is both expensive and cumbersome to oversee procured AI products.

DATA GOVERNANCE

AI algorithms are driven by data. Therefore, problems related to data management and provenance lead to several challenges:

- Although there is a relative abundance of open data now available compared to a generation ago, we need more government agencies to continue to publish their data as open data that can be used to drive the development of machine learning algorithms.
- **Open data is often rife with problems and significant investment is required to clean data and develop robust and usable training data for use in machine learning algorithms.** There is no consistent framework by which data is labeled, further impeding both the development and evaluation of ML.
- **Some government data is classified to a level to which only a select few individuals have access.** One respondent within the DoD provided several examples in which data was over-classified. Only vendors with the appropriate security clearance (in addition to technical know-how) were able to work with the data, preventing a wider pool of vendors from participating in the project.
- **There are not enough individuals—usually data engineers or architects—with expertise in translating and interpreting data sets.** These can be either sets they have built themselves or external sets that they are tasked with trying to understand and translate into actionable information.
- **The absence of clear and consistent documentation for the data sets is also a major problem, because users of some data sets may not be able to see or contact whoever compiled it.** Often, a government data set is compiled and labeled by only a handful of people, and in some cases by just one person. Without proper and consistent documentation, one respondent noted, the data set could be useless even if the data is relevant, because those trying to use it cannot interpret the decisions and labels of the owner or original creator.
- **Issues with provenance complicate the use of data sets when working with third-party vendors.** Some of these issues include trouble affirming the authenticity of the data set, determining whether harmful bias exists in the data or the ways in which it was collected, and deciding whether the data set is a proper fit for the model the agency is trying to develop. The lack of universal



data standards across governments slows down data vetting because the data needs to be cleaned and structured manually to be machine-readable for AI models. “The first 80 percent of the work is getting the data right, well before you can even start to train a model,” one respondent said.

- **Departments are worried about the loss of data stewardship.** They are often reluctant to share data and there is no clear understanding of who “owns” the data outputs from a machine learning algorithm, especially when it is trained using data from multiple agencies. One participant reported that debates about ownership of a model’s output significantly delayed the acquisition process for an AI solution. In this case, the data was controlled by one entity, the engineers who planned to deploy the model were from another entity, and the model itself was owned by the vendor. Even though the inputs (non-PII¹⁰ government data is required to be open by law), there is confusion and disagreement over ownership of the model’s output.
- **Proprietary control by the vendor can interfere with the public interest.** First, government data is open data whereas vendors often want the data that underlies and trains a model to be proprietary. Second, there are issues about control over products. A Director of Acquisitions within the federal government noted that public servants often expect the government to own an AI system outright once the acquisition process is complete in contrast to vendors who do not want to license the algorithm outright. Third, getting unlimited or general-purpose rights to data “limits the types of businesses willing to participate in the government procurement system,” said another respondent.
- **Many respondents agreed that proprietary data issues apply in unique ways to AI solutions.** Issues arise particularly when a vendor is providing a subcomponent of a system, and this subcomponent must interoperate with components provided by other vendors. In such cases, proprietary data formats hinder the government’s ability to connect different vendors on a common project.
- **Historically, many governments have not required procured software to include comprehensive, well-supported, documented Application Programming Interfaces (APIs) with clear ways of accessing them.** Without the availability of these APIs, there are significant barriers to freely and programmatically using data in the ways that the government sees fit.

ORGANIZATIONAL CHALLENGES

Organizational arrangements unique to government can make it harder to request and implement AI in a timely manner:

In government, procurement is siloed in its own office separate from policy and technology functions

- **In government, procurement is siloed in its own office separate from policy and technology functions.** Respondents highlighted the need for intra-agency coordination between procurement officers, lawyers, technologists, and compliance teams within a department as well as interagency collaboration between executive agencies and the GSA to develop and promote best practices for navigating the AI procurement process.

¹⁰ PII refers to Personal Identifiable Information, which refers to any representation of information that permits the identity of an individual to whom the information applies to be reasonably inferred by either direct or indirect means.



- **Resistance to working with new and diverse suppliers can slow procurement of the most advanced AI solutions.** One respondent from the DOD reported that some of their peers were reticent to contract with companies that are not typical defense contractors, even though the outside organizations offered the solutions that the DOD desired. This is problematic because incumbent and familiar government contractors do not necessarily offer all the AI tools the government needs. At the same time, many of the newer vendors, even if they are not startups, still do not “speak the language of government” or understand how to navigate acquisition processes.
- **The lack of agile contract processes impedes acquisition of the most cutting-edge AI applications.** The rate at which AI evolves renders some solutions irrelevant before the government can benefit from their use, causing systems to become obsolete within the window of the procurement process or program cycle.
- **The relatively long use of IT systems in government often means that rapidly-acquired AI models may not have the requisite longevity.** Another respondent, a former federal employee working on labor-related AI applications, noted that a pervasive desire for “custom AI” creates logistical issues because it is more difficult to iterate on or integrate these systems into a longer-lived, larger project. Concerns about the shelf life of AI products deter both public servants and third-party vendors from working together to develop iterative and lasting AI solutions. This problem can be exacerbated by data and vendor longevity challenges in some cases, as the ecosystem is evolving very quickly.

AI GOVERNANCE AND OVERSIGHT

When contracting officers and public servants were asked what they saw as primary concerns and impediments when procuring AI solutions, a consistent answer concerned governance and management both before and after acquisition. Problems they cited include:

- **There are no clearly defined standards for AI governance within either industry or government.** The resulting ambiguity creates confusion and additional worries for those trying to acquire and implement AI solutions.
- **Similarly, there is no approved standard to which procurement officers can refer in order to properly vet or review an AI system.** Many frameworks for AI governance exist, as well as new technical and process tools to monitor, test, and audit¹¹ AI models. But several interviewees raised questions about who should bear the responsibility of attesting and monitoring governance frameworks for AI systems that have been acquired or that the agencies plan to acquire. “This issue comes up a lot,” said one respondent in reference to AI risk management, and “without guidance, there are a lot of questions about whose job it is and where and how to work with audit/risk assessment vendors if they are available.”

There are no clearly defined standards for AI governance within either industry or government

¹¹ An AI audit generally refers to a process by which the builder of an AI model or a third party dissects the model to understand the way in which it was optimized (i.e., the weight or probability assigned to certain characteristics or data points), what data was used to train it, for what purpose it was built, and what confidence and evaluation metrics make it fit for purpose. For good resources on AI auditing, check out the UK Information Commissioner’s Office reports on this topic and a publication by leading AI researchers



- **A dearth of auditing tools slows vetting and oversight.** Although a growing number of companies are building and providing tools to audit AI models, the work of these companies is perceived as nascent and not yet mature enough for government use. Agencies do not usually have easy access to such auditing services.
- **Respondents maintained that currently no single government agency dictates standards around AI acquisition.** The National Institute of Standards and Technology (NIST) released an AI Risk Management Framework and Playbook in 2023. The Office of Science and Technology Policy (OSTP) released an AI Bill of Rights in 2022 and the Office of Management and Budget (OMB) put out guidance to agencies for regulation of AI in 2020. But the landscape of AI governance is fractured.

In short, lack of policy and associated frameworks that directly speak to acquisition hinders successful [procurement] outcomes, said one Acquisitions Director. Another added that the “absence of standards hurts solution shaping and makes the procurement process unnecessarily difficult.”



GOVERNMENT PROCUREMENT OF AI ENCOUNTERS MANY HURDLES

To summarize this section, public servants face unique challenges when procuring inherently complex AI and ML technologies. Limited technical know-how and imprecise problem definition prevent AI systems from being properly scoped for effective and long-lasting government use. An ecosystem of traditional and entrenched government contractors who possess a longstanding familiarity with navigating procurement processes limits the diversity of new AI suppliers. Governments would have access to far more advanced solutions if they could more easily work with startups and smaller companies that are unfamiliar with procurement rules.

At the same time, procuring and refining custom AI solutions is a cumbersome process. Logistical and operational issues with data organization, maintenance and ownership can lead to over-reliance or lock-in with a vendor, even if its solutions are not fit for purpose.

Finally, without standards, frameworks and clearly defined governance roles and responsibilities, risks inherent to some AI systems can go unaddressed and create or exacerbate hesitation and skepticism among public servants.



SUMMARY OF BARRIERS

We can summarize the major barriers to AI procurement in the public sector as follows:

- The public sector cannot compete with the private sector for AI talent and has lacked training to create AI-conversant professionals.
- AI procurement is hampered by the inability of public servants to evaluate AI-enabled solutions and assess the quality of vendors.
- Lack of comfort with or understanding of SaaS business models hinders acceptance and use of AI services across government and can increase hesitancy among public servants about using AI services.
- Legacy systems for data governance, conventional cultures, and traditional approaches to contracting technology services prevent efficient procurement of AI.
- Often, offices and departments lack the data infrastructure to store, integrate, and process the data needed to develop AI applications, even for well-defined and properly scoped AI use cases.
- Poor data infrastructure can lead governments to potentially become over-reliant on third-party vendors that lack transparency.
- Because data formatting and maintenance are resource-intensive and key to training AI, a lack of cohesive standards in this area slows the AI development and procurement process.
- Approved government contractors often do not sell the AI tools the government needs; at the same time, AI startups with innovative tools face challenges in learning how to navigate acquisition processes.
- While AI helps to realize efficiencies, those gains are frequently lost in the added need to provide enhanced oversight of machine learning algorithms.
- Government procurement officials and technology practitioners do not have a coherent framework by which to measure and manage risks when procuring AI technologies. It is not clear who bears responsibility for overseeing government AI systems. The market for emerging technical solutions for auditing and monitoring AI systems is not yet mature enough for widespread government adoption of these solutions.



RECOMMENDATIONS

To accelerate more informed decision-making about AI that helps procurement professionals adequately scrutinize the risks while embracing new opportunities in an agile manner, we outline below at greater length the recommendations summarized on page 2.

TRAINING

BROADLY AVAILABLE TRAINING ON AI AND DAY-TO-DAY APPLICATIONS IN GOVERNMENT

Public servants need training and professional development to understand both technical and non-technical AI concepts. Professional development should be offered through a combination of traditional and experiential learning that helps public servants differentiate between high and low-risk opportunities, and learn approaches to scoping, acquiring, testing and de-biasing well-tailored AI solutions. Interviewees suggest that these asynchronous, live and experiential learning opportunities would have broad interest for public sector learners.

There is first a need for basic “AI 101” content that explains what AI is, how it works, and what it can do. Especially given the frequent use of metaphors (e.g., the term “hallucinations” refers to an AI content generator stating something false with absolute certainty and AI is often described as a sentient being), there is a lot of misunderstanding about what AI is and can do. However, we need to invest in training that goes beyond the headlines and helps learners internalize what AI is and does on a practical level. This includes understanding and gaining fluency with other technical systems that allow AI to function, such as Application Programming Interfaces (APIs) and how data processing and various ML applications work in the cloud. Use cases need to focus on the public sector context to be more readily understood.

Hands-on demonstrations or workshops, including role-playing or tabletop exercises, could be helpful for providing this type of learning, as well as fostering intra-agency and interagency collaboration for developing AI solutions via learning cohorts, working groups and communities of practice.

Officials at both the federal and state level suggested that the materials should put the technology and its consequences in context and provide actionable and operational guidance for those procuring the technology. Many suggested that public servants would find the most value in the elevation of real-world case studies of how similar teams have gone through the procurement and implementation process.



Interviewees shared that increased exposure to and technical understanding of AI will enhance procurement officers' proficiency and ability to identify and integrate appropriate solutions to defined problems

Interviewees shared that increased exposure to and technical understanding of AI will enhance procurement officers' proficiency and ability to identify and integrate appropriate solutions to defined problems. Increasing fluency in AI will support their ability to assess the efficacy of a proposed solution to adequately address the problem.

TRAINING IN PROBLEM DEFINITION

In addition to training on new technologies and digital government, procurement officials need to know how to define the problem that their AI solution is intended to address. Many we spoke to emphasized the importance of teaching and learning problem definition to help procurement officials ensure that what they are buying will respond to a real need.¹² Before they can effectively define the system's requirements and features, they need to understand what the system is designed to accomplish.

It is important to ensure that the relevant groups in and across agencies, including policy, technology, and procurement, are brought together early in the procurement process. These discussions should cover the problem to be solved, the potential solutions, and how they respond to the problem and its root causes. These early and continued conversations can be helpful in creating alignment between solution and problem, as well as helping to ensure that relevant stakeholders and affected community members are engaged and consulted. Such deliberative co-creation processes increase confidence and clarity across the participants and can speed up successful adoption.

TRAINING PROVIDERS

Many universities and commercial training providers now offer coursework in AI, machine learning, data science and other relevant disciplines. Such courses are rarely, however, designed for those in the public sector.

InnovateUS (disclosure: the authors manage this program) is a nonprofit, nonpartisan, multi-state consortium that designs and delivers free training in digital technologies and their uses for public servants. InnovateUS currently teaches AI basics for the public sector and has plans to deliver an expanded and in-depth program in AI for the public sector later in 2023.

¹² In Chapter 4 of Solving Public Problems: A Practical Guide to Fix Our Government and Change Our World Beth Simone Noveck lays out the importance of - and provides strategies for - adequately defining a problem in the public sector.



Although neither the Office of Personnel Management (OPM) nor GSA's Digital Government University currently teach courses in AI, they do teach related concepts of digital government. The OPM courses are generally several thousand dollars each.

While the Partnership for Public Service has launched a course on AI for senior federal officials, spaces are limited and it is only open to senior executive service officials at level GS-15 and above. This is a great start but we need more high-quality and accessible learning offerings to truly meet the opportunity posed by AI.

TOOLS

Practical tools will help public servants procure and use AI solutions more effectively. Procurement officers should be given expert systems — decision aids with if-then questions and flow charts — that help them categorize, evaluate, and more accurately select among AI tools that fit their needs. For example, organizations benefit from understanding where in the AI maturity curve the solution they are considering fits and not applying a “one size fits all” approach review which could lead to greater or reduced review processes. For example, unlike more customized and predictive forms of AI, translation AI is very advanced and much more deployable out of the box, with the ability to add custom glossaries. If AI is being used to determine benefits eligibility or food safety, more care and concern needs to be taken than for uses of AI to predict rust on a cargo vessel. Having good decision support tools can help procurement officials evaluate the appropriate level of risk of AI in different use cases.

Such decision support tools will guide officials and increase their confidence in the decisions they are making at each step of the AI acquisition process, from problem scoping to assessing benefits and risks, and determining the limitations across categories of AI tools.

Such decision support tools will guide officials and increase their confidence in the decisions they are making at each step of the AI acquisition process, from problem scoping to assessing benefits and risks, and determining the limitations across categories of AI tools.

The new NIST risk management framework,¹³ playbook, and explanatory materials are very helpful in this regard. However, they are designed intentionally to be general and flexible for any purpose. They are not intended to aid in procurement specifically. Lessons can be learned from the authoring process for the NIST framework, which was co-created by a large number of experts from government,

¹³ The National Institute of Standards and Technology released an Artificial Intelligence Risk Management Framework in January, 2023



industry, and academia, on how to create similar tools for procurement professionals.

As contracting specialists gain a greater understanding of AI, they will become better equipped and can collaborate on writing guidance. Such manuals should be “wiki-fied,” enabling contracting professionals across government to contribute to them. They will also become more adept at writing contracts that are technically precise and reflect how the solution will benefit the organization. The specialists can anticipate and address legal, policy, and regulatory issues before the systems become embedded within an organization and reach a point where their removal would cause disruption.

Procurement and contracting officers should also be given tools for drafting contract language, informed perhaps by training about how to map contract terms to product features and practice. These tools will help procurement officials write contracts that reflect how AI will interact with their departmental mandate. This competency will further encourage clear problem definition and use case selection that is appropriate and aligned with governments’ intent in using AI. Disciplined practices here will result in more successful AI acquisitions and enhanced public trust.

RELEASE MORE OPEN DATA AND SIMULATED DATA SETS

Open data can play a critical role in enabling the development of more accurate and effective AI models.¹⁴ By making more diverse data available to a wider community, governments play a vital role in enabling the development of robust

By making more diverse data available to a wider community, governments play a vital role in enabling the development of robust AI

AI. For example, data on population demographics, health outcomes, and environmental factors can be used to train AI models that can help identify and address health disparities in specific communities. Similarly, data on traffic patterns, weather conditions and other factors can be used to develop AI-powered transportation systems that are more efficient and safer.

By providing a larger and more representative sample of data, open data can also help to reduce bias in machine learning and predictive analytics. Furthermore, open data can also foster greater transparency and accountability in the development and deployment of AI, as it allows for greater scrutiny and understanding of the data sets used to train the model, albeit without necessarily helping to explain why and how the model works or to ensure fair outputs. By making data available to the public, open data initiatives can help to increase public awareness and understanding of how AI is being used in different contexts. This can help to build trust in AI

¹⁴ The National Science Foundation is supporting the creation of a prototype Open Knowledge Network, citing that open access to shared information is essential for the development and evolution of artificial intelligence (AI) and AI-powered solutions needed to address the complex challenges facing the nation and the world.



systems and ensure that they are being used in ways that are consistent with societal values and goals.

Open, simulated data sets are also an essential component of training AI models. They allow developers to generate large amounts of data that mimic real-world scenarios, providing the AI model with a diverse range of examples from which to learn. Simulated data sets are particularly important in situations where obtaining real-world data is difficult, such as in fields like medicine or aerospace engineering. They can also help to address issues of bias in training data by providing a controlled environment to test and refine the model. Additionally, simulated data sets enable developers to quickly iterate and improve the model without the need for time-consuming and expensive data collection. Because government data is collected or created using taxpayer dollars, governments have an obligation to make these assets freely available to the public for reuse, which can help to further spur the development of a robust AI industry and new applications.

REGULATION AND GUIDANCE

Regulation and governance frameworks will help public servants feel more confident in acquiring AI and managing its risks. Guidance should recognize that AI is a broad spectrum, and while some AI use cases are more complex than others and should be treated with more consideration, some are fairly straightforward and can be acted upon more quickly.

The federal government (Federal Acquisition Service working together with NIST and the Office of Science and Technology Policy) needs to develop a typology to guide acquisition professionals and enable them to invest time in more problematic cases while streamlining acquisition of tools that are less likely to impact fairness, equality or liberty.

Uniform or recommended assessment frameworks for AI governance could also help public servants address ethical concerns that can arise throughout the acquisition process, as well as train them to incorporate this thinking from the outset when approaching future projects without locking in inflexible processes that cannot evolve.

The expansion and adoption of regulation, guidance, and tools will also drive a shift in culture and mindset toward organizational support for AI governance.

Guidance from NIST about acceptable procurements would eliminate much uncertainty and concern, several respondents said. Respondents called for efforts in the area of a risk management framework (RMF), including one developed by NIST. The RMF should allow flexibility in innovation across departments and contexts,

because each agency acquires and works with AI differently. Respondents also noted that a federal RMF will promote much-needed inter-agency coordination.



The expansion and adoption of regulation, guidance, and tools will also drive a shift in culture and mindset toward organizational support for AI governance. With clear rules and frameworks in hand to support those rules, there will be less finger-pointing when it comes to responsibility for AI governance or compliance, because teams will need to align with each other and collaborate to comply.

AUDITS

Procurement needs to account for regular audits and ongoing monitoring of AI systems, ensuring that purchased systems function as intended. By systematically evaluating and assessing the design, development, and deployment of AI processes, audits help to identify potential biases, errors, and security risks, while also promoting transparency and trust. Audits can help to ensure compliance with relevant laws, industry standards, and ethical guidelines. By examining the underlying data, algorithms, and decision-making processes, audits can highlight areas where more transparency is needed. Transparent AI systems make it easier for stakeholders to understand how and why decisions are made, leading to greater accountability.

Audits can help detect and mitigate biases in training data, model development, and decision-making processes

Audits can help detect and mitigate biases in training data, model development, and decision-making processes. Addressing these biases is essential for building AI systems that are fair and unbiased, which in turn enhances user confidence. Audits also can be used to evaluate the security measures in place to protect the AI system and the data it processes. Ensuring the confidentiality, integrity, and availability of data is crucial to maintaining user trust and accountability.

Audits also assess the accuracy, consistency, and robustness of AI systems to determine if they meet the intended purpose and required performance levels. A high-performing AI system instills confidence in its users and demonstrates accountability for its results.

Audits overcome system opacity; in this way they help procurement officers and government employees make an informed and defensible decision to procure and use AI-enabled tools at the outset. However, audits must be an ongoing process as well. Although we are rightfully concerned with what data feeds machine learning algorithms, we also need to assess the fairness and accuracy of the outputs throughout the lifecycle of the product.

As noted earlier, many respondents believe that third-party auditing systems are not yet advanced or comprehensive enough and that the government itself should invest in creating new auditing tools and frameworks.



ORGANIZATIONAL CHANGE

Solving the challenges of AI procurement demands multidisciplinary perspectives and expertise. Interviewees shared the need for increased support for collaboration, knowledge sharing, and coordination among procurement officials, technologists, and policymakers in addition to integrating mechanisms for public input.

SUPPORT OF INTERDISCIPLINARY COLLABORATION

Interdisciplinary teams that expose individuals to different functions and perspectives allow them to become proficient in other roles, expanding their experience and improving problem-solving throughout an organization.

Two of our respondents, one of whom is a trained technologist and the other an Acquisitions Director in the federal government, told us that their roles allowed them to work together with relative agility. They reviewed and advanced AI contracts and came away with expertise in each other's field.

This collaboration produced a dual benefit. First, contracting specialists could exploit the technologists' knowledge of AI to procure strategic and responsible AI solutions and draft contractual terms that the industry can understand and reasonably achieve. At the same time, by understanding the procurement process, technologists could design systems that align with governments' needs and comply with their processes and requirements.

This pair reminded us that their experience is not the norm. Only because they were both in a program office that facilitated this partnership and collaboration did they get the opportunity to meet and work closely together. Hence, they recommended that acquisition teams be built into program management offices inside agencies and that the teams work jointly throughout an AI project.

Inter-organizational collaboration is crucial in order to fit properly-scoped AI needs to properly-scoped AI solutions

Inter-organizational collaboration is crucial in order to fit properly-scoped AI needs to properly-scoped AI solutions. Many agencies bring in a contract specialist only when it is time to write a contract or a vendor is trying to sell a product or service. Instead, organizational structures should pair technical experts with procurement officers from the point of problem definition throughout the procurement process in order to navigate acquisitions in a more agile manner.



INTEGRATING TECHNICAL AND NON-TECHNICAL EXPERTISE

Rather than hiring a few select technical experts or siloing technical talent into one pool within a department or interagency group, “organizations should balance technical and non-technical skills among and across individuals, so that technically-oriented skills are more diffuse across a department or team and at the mission and/or product level,” said the former Head of AI at TTS within the GSA. At the same time, technical non-human resources (like tools, etc.) should be more centralized so that teams can access them in a streamlined way.

DATA OWNERSHIP

Governments should make technology and procurement decisions that allow them to retain agency over how and when to use data that they own. To do so, governments should require companies to provide well-documented, fully supported APIs as part of the standard terms of their procurement. This provides numerous benefits as governments can then use the data in new and unpredictable ways even when outside the scope of the vendor that is hosting the data, reducing vendor lock-in, breaking down data silos within and across organizations, helping to streamline data sharing, and allowing government, rather than the vendor, to remain in control of how to best provide services. Ultimately, this approach is important because governments cannot fully predict how they might be able to use the data to provide better services in the future, but they can take steps to retain the flexibility needed to best adapt. Additionally, some AI tools can also help integrate and automate large volumes of data ingestion, which can enable organizations to have access to richer data sets and faster.

NARROW THE EXPERTISE GAP

Those in government need a connection to those with expertise in AI and other emerging technologies to accurately assess a potential solution’s capabilities and support the implementation and ongoing responsible use of this technology once procured. The gap in access to this expertise directly limits the procurement and implementation of AI in government. We recommend large investments to integrate those with expertise in new technologies into government beyond the offices of the CIO or CTO, including into procurement, legal, and policy teams..

Interviewees at all levels of government shared the need for greater expertise to support both the procurement process and responsible implementation and evaluation

Officials at the local, state, and federal levels encouraged the expansion of fellowship programs and an investment in creating more accessible pathways to engage academic experts in the procurement and auditing processes. One official suggested investing in new models such as creating offices within government that could provide data science as a service to agency partners and procurement teams.



Interviewees at all levels of government shared the need for greater expertise to support both the procurement process and responsible implementation and evaluation. Without qualified experts, many officials voiced concerns about reliance on expertise from vendors in the procurement process.



SUGGESTIONS FOR FUTURE RESEARCH

While this report has highlighted numerous problems in government AI procurement and suggestions for improvement, an even broader swath of topics in this area is worthy of research as new insights can produce further improvements in AI procurement. We've laid out a few areas for potential research below:

DIFFERENCES IN RISK PERCEPTION BETWEEN IN-HOUSE AND THIRD- PARTY AI SYSTEMS.

Some argue that “In-house expertise promotes AI tools that are better tailored to complex governance tasks and more likely to be designed and implemented in lawful, policy-compliant and accountable ways.” Given recent calls for and advancements in the responsible use and management of AI, research should identify the similarities and differences in risk between systems designed by industry and government. These studies can help inform efforts to govern these systems and improve the procurement process.

GOVERNMENT CULTURE AND CHANGE MANAGEMENT.

Nearly every respondent noted barriers to government innovation posed by established processes, operations, and agency cultures. Research could explore the roles that informal and formal government culture play in AI acquisition and the change management it requires.

AI ACQUISITION AND GOVERNMENT JOB GROWTH.

Given that AI systems, especially those with direct impact on the public, require continuous oversight and monitoring, their use by the government is likely to create new jobs. Researchers should examine how the government should respond to such growth and efforts taken to attract and retain top talent.

AI IMPACT ASSESSMENTS AND CONTINUOUS MONITORING.

Due to recent legislative interest in algorithmic impact assessments in California, Washington, and other states, it is worth exploring the approaches taken by agencies and departments across the state and federal governments to draft these assessments. Research should likewise assess whether AI impact assessments and other due diligence mechanisms arising from procurement contracts can help mitigate oversight and governance issues. Oversight is currently hampered by agencies' lack of visibility into the commercial development of AI solutions they later procure and deploy, and a lack of collaboration with the developers.



PUBLIC TRUST IN GOVERNMENT USE OF AI.

Per the Pew Research Center’s findings that many Americans are concerned about the increased use of AI, researchers should explore whether perceptions differ between industry and government use and the contexts in which perceptions differ. Which government uses of AI spark the most excitement and the most worry?

MISSED AI OPPORTUNITIES AND RELATED COSTS.

Does a broad, fundamental lack of understanding of AI by people without technical background, as well as by technically knowledgeable staff, prevent AI tools from being procured? What can we learn about these “missed opportunities” and the impact of not using AI?

THE ROLE OF LEADERSHIP AND ORGANIZATIONAL STRUCTURE IN AI ACQUISITION.

Some respondents identified that learning about AI and pursuing a project related to it is not in their colleagues’ career interests because they lack time, incentives, or both. A top-down culture that rewards responsible innovation could encourage the exploration of AI solutions that benefit both government and the public.

ROLE OF NEWS MEDIA COVERAGE ON AI AND RISK PERCEPTION.

Several studies have examined the news media’s coverage of AI’s abilities and how this portrayal influences public trust, comfort, and risk perception. Future research should try to determine whether and to what extent government officials—especially procurement officers—are influenced by this coverage and its impact on their willingness to acquire and use AI.

AI AND THE FIELD OF PUBLIC POLICY.

How might the use of AI’s predictive capabilities create new opportunities for government that did not exist before, such as improved scenario planning or the ability to realize risks associated with inaction. One of the main benefits of AI is its ability to analyze vast amounts of data and identify patterns and trends that humans may not be able to detect on their own. By using machine learning algorithms, governments can create predictive models that can forecast future outcomes and help policymakers make better decisions, identify potential risks and trends, and provide real-time monitoring of conditions that can accelerate and improve policymaking.



AN EXPANDED USE OF AI IS POSSIBLE

AI solutions offer promising opportunities to eliminate repetitive tasks and produce better results in government agencies, yet many government employees do not know how to address the risks of failure or bad outcomes and shy away from these new technologies. This paper has addressed skill-based and organizational requirements that can lead to government use of AI that is productive, robust, and efficient.



ACKNOWLEDGEMENTS

We want to thank the many policy, procurement and technology professionals who contributed their time and expertise to the drafting of this paper, most of whom have asked to remain anonymous. Those we can name are listed below.

STEVEN BABITCH

- *Product and Strategy Leader*
Konux GmbH
- *Former Head of*
Artificial Intelligence
GSA Technology
Transformation Services

MICHAEL BASKIN

- *Chief Innovation Officer*
Montgomery County,
Maryland

JOY BONAGURO

- *Former Statewide Chief Data*
Officer, State of California

ALEXIS BONNELL

- *Strategic Business Executive*
Google

JASON CLARKE

- *Chief Information Officer*
State of Delaware

ADDIE COOKE

- *Global AI Policy Lead*
Google

SANTIAGO GARCES

- *Chief Information Officer*
City of Boston

OMID GHAFARI-TABRIZI

- *Head of U.S. Federal*
Civilian Policy
Google

CHRIS REIN

- *Chief Technology Officer*
State of New Jersey

GRACE SIMRALL

- *Chief of Civic Innovation*
and Technology
Louisville Metro Government

RICHARD TODD

- *Chief Data Officer*
City of Boulder, Colorado

EDWARD HARTWIG

- *Partner and Chairman*
Service Design Collective



ABOUT US



The Cantellus Group is a boutique services firm advising leaders on the strategy, oversight, and governance of artificial intelligence and other frontier technologies. Leveraging the expertise of advisors across professional and industry domains, the Cantellus Group also helps organizations understand the technology policy landscape and how technology governance impacts corporate governance.



The GovLab's mission is to improve people's lives by changing the way we govern. The group's goal 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. The GovLab believes 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. Housed at New York University Tandon School of Engineering, The GovLab is funded by various donors and partner organizations.



WORKS CITED

- ⇒ Using Artificial Intelligence to Transform Government
- ⇒ Artificial Intelligence for Citizen Services and Government
- ⇒ National Artificial Intelligence Initiative Act of 2020
- ⇒ AI-Augmented Government
- ⇒ Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies
- ⇒ Best Practices for Government Procurement of Data-Driven Technologies
- ⇒ No-Code' Brings the Power of A.I. to the Masses:
- ⇒ Michigan's MiDAS Unemployment System: Algorithm Alchemy Created Lead, Not Gold - A case study into how to automate false accusations of fraud for more than 34,000 unemployed people
- ⇒ Wrongfully Accused by Algorithm
- ⇒ How Americans Think about Artificial Intelligence
- ⇒ The Seven Deadly Sins of AI Predictions
- ⇒ Challenges for Mitigating Bias in Algorithmic Hiring
- ⇒ Gov Procurement Online Marketplace CoProcure Raises \$22M
- ⇒ UK Information Commissioner's Office AI Auditing Framework
- ⇒ Closing the AI Accountability Gap: Defining an End-To-End Framework For Internal Algorithmic Auditing
- ⇒ The New York City Artificial Intelligence Strategy
- ⇒ NIST Risk Management Framework (March 2022 draft)
- ⇒ CA AB-13 Public Contracts: Automated Decision Systems
- ⇒ WA SB 5116 - Establishing guidelines for government procurement and use of automated decision systems in order to protect consumers, improve transparency, and create more market predictability.

