

Integrating AI in Government: Lessons, Applications, and Innovations

Michael J. Keegan, Forum Editor

Artificial and augmented intelligence (AI) have proliferated across all sectors of society. With the adoption of AI in government expanding as these technologies evolve and their applications broaden, this is no futuristic concept. It has become an integral part of modern governance, reshaping how the U.S. federal government operates and delivers services. From safeguarding national security to streamlining routine processes, AI technologies are transforming the public sector at an unprecedented pace. This forum is dedicated to exploring the dynamic intersection of AI and government operations, focusing on its current applications, potential, and the critical questions it raises about trust, ethics, and effective implementation.

The U.S. federal government has embraced AI across a spectrum of mission-critical activities. Tax agencies use AI to enhance compliance and detect fraud, while defense agencies rely on it for mission planning and insider threat detection. Generative AI tools are beginning to revolutionize communication, augment decision making, and unlock new efficiencies in public procurement and service delivery. The adaptability of AI offers enormous promise—but it also demands careful oversight and strategic governance.

In recent years, significant strides have been made to institutionalize AI within government frameworks. The American AI Initiative, launched in 2019, underscores a national commitment to AI innovation. Since 2020, the U.S. federal government has taken several steps to address and regulate the use and application of artificial intelligence as well as the burgeoning impact of generative AI (genAI) from the advent of the National Artificial Intelligence Initiative Act of 2020 to the E.O. 14176 on the Safe, Secure, and Trustworthy Development and Use of AI. This executive order builds on prior work to support the development of responsible AI technologies and policies, including the Office of Science and Technology Policy's (OSTP) Blueprint for an AI Bill of Rights and the National Institute of Standards and Technology's (NIST) AI Risk Management. Taken together, these efforts focus on promoting innovation, ensuring ethical standards, managing risks, and safeguarding national security while addressing issues like privacy, bias, and accountability.

The 2020 *The Business of Government* magazine dedicated a forum on the evolving use of artificial intelligence in government, which highlighted the insights, findings, and recommendations derived from three collaborative report on AI and its early evolution published by the IBM Center and Partnership for Public Service. It aimed to spark a conversation on the use of AI, help prepare federal leaders to assess the inevitable changes coming and provide government leaders with insights to navigate this transformative time. This forum continues that conversation bringing together perspectives from a sampling of recent IBM Center reports further exploring the operation and use of emerging technologies like AI and documenting lessons learned, recommendations, and next steps.

- ***AI and the Modern Tax Agency: Adopting and Deploying AI to Improve Tax Administration.*** This report provides analysis and makes recommendations for the IRS and tax agencies around the world to leverage AI to improve customer service and education, compliance and enforcement, and risk management—and to do so while mitigating risk and building trust using AI.
- ***Pathways to Trusted Progress with Artificial Intelligence.*** This report examines the governance and applications of AI and how governments need to develop and communicate a framework for the public to understand why AI is being used; what has been done to ensure that the AI is fair, transparent, and accurate; what experiments were done to ensure that the output is reliable; and how public value from AI is being measured and created.
- ***Navigating Generative AI in Government.*** This report highlights a practical set of considerations and potential actions that can help government agencies to capture benefits and minimize risks from the use and application of generative AI.
- ***Government Procurement and Acquisition: Opportunities and Challenges Presented by Artificial Intelligence and Machine Learning.*** This report addresses using AI to transform public procurement, analyzing challenges and recommending ways to capitalize on opportunities. It makes recommendations about how to use AI to transform public procurement, which can increase speed, efficiency, and effectiveness in acquiring goods and services that serve public needs by fostering facilitating intelligent automation across the federal acquisition system.

- ***Artificial Intelligence in the Public Sector: A Maturity Model.*** This report offers public sector leaders a view into the “art of the possible” by emphasizing how AI programs can accelerate the transformation of government programs to better serve the public and do this by offering them framework for establishing a successful AI program. The challenge has always been to design and implement an AI program that has all the critical elements in place to successfully achieve the goal of improved mission delivery and citizen services.

These reports highlight the opportunities and challenges AI presents to federal agencies and offer frameworks for its responsible use. By synthesizing insights from these resources, this forum aims to foster a deeper understanding of how AI can enable smarter, more efficient governance while maintaining the public’s trust.

The conversation around AI in government is not just about technology—it is about people, processes, and the policies that shape its use. As we navigate the evolving role of AI, this forum provides a platform to examine how federal agencies can balance innovation with accountability, adapt to emerging trends, and deliver on their mission in a rapidly changing world. Together, we will explore how AI can empower government to serve its citizens better while upholding the principles of transparency, equity, and ethical stewardship.

This technology will remain front and center. The curated insights offered in this forum can help those leading the upcoming presidential transition as they begin to think about what’s next and how best to develop safeguards and guardrails for this technology as it continues to expand and evolve.



AI and the Modern Tax Agency

Adopting and Deploying AI to Improve Tax Administration

Edited by Michael J. Keegan

The first contribution to this forum focuses on how artificial intelligence could improve tax administration while minimizing some of the risks. Driven by common access to AI and the potential benefits of generative AI, the U.S. Internal Revenue Service (IRS) and tax agencies around the world are now in the mist of calibrating the right balance for the use and application of this technology. To better understand opportunities and considerations, the IBM Center, in collaboration with the Kogod Tax Policy Center at American University, hosted a global discussion on *AI and the Modern Tax Agency*.

What follows are insights, analysis, and recommendations for the IRS and tax agencies around the world to leverage AI to improve customer service and education, compliance and enforcement, and risk management—and to do so while mitigating risk and building trust using AI. These are excerpted from the IBM Center report *AI and the Modern Tax Agency: Adopting and Deploying AI to Improve Tax Administration* by Caroline Bruckner and Collin Coil.

Opportunities for Tax Administration in the U.S.

The U.S. Internal Revenue Service is embarking on a transformative IT overhaul to modernize the taxpayer experience, made possible by a substantial funding boost from the Inflation Reduction Act (IRA) in August 2022. This funding enables the IRS to invest in new technology and integrate AI across operations, targeting areas like taxpayer filing, enforcement, and internal processes. Despite collecting over \$4.9 trillion and processing about 260 million tax returns in FY 2022, the IRS faces challenges due to outdated technology and a budget reduction since 2010. Its responsibilities have also expanded, including administering programs from significant legislative acts.

To manage the growing demands, AI adoption is crucial for the IRS. The IRA funding allocated \$45.6 billion for enforcement, \$25.3 billion for operational support, \$4.8 billion for systems modernization, and \$3.2 billion for taxpayer services. AI is already used in some IRS functions, like audit selection and call redirection, but more widespread deployment faces challenges. These include ensuring security, privacy, and compliance with a complex tax code, which AI cannot fully simplify due to statutory intricacies.

International tax agencies, such as those in Australia and Singapore, have successfully implemented AI-driven tools like virtual assistants, which streamline user support. The IRS, however, has yet to deploy a similar accessible AI system, though potential exists. The IRS's extensive digital framework and tax law complexity present unique hurdles to AI deployment, which differs from the seamless user experiences seen in the private sector. For the IRS, these challenges stem from both the extensive existing IRS digital footprint and the overwhelming (and increasing) complexity of U.S. tax laws.



AI in Customer Service, Outreach, and Education to Support Tax Agencies

Taxpayers in the U.S. often struggle with filing compliance, particularly those with nontraditional income like gig work, which lacks tax withholding, contributing to the tax gap (estimated at \$688 billion in 2021). Compliance challenges also stem from low tax literacy; many taxpayers lack the knowledge to navigate complex tax rules. Despite the importance of refunds, especially for low-income individuals, education around tax obligations is limited. Research shows that only a small portion of U.S. taxpayers receive tax education, contributing to errors in filings and unintentional non-compliance.

AI could play a significant role in assisting tax agencies by improving customer service, education, and outreach. AI chatbots and virtual assistants can provide 24/7 help with routine tax queries, allowing human agents to focus on complex cases. AI tools can scan forms, offer guidance in real-time, and personalize assistance based on taxpayer data, reducing filing errors and improving the taxpayer experience. Personalization, facilitated by AI, could make tax agencies more responsive and encourage voluntary compliance, ultimately leading to better outcomes for both taxpayers and tax authorities. Additionally, targeted information campaigns using AI can educate taxpayers on benefits and free assistance programs, further enhancing compliance and satisfaction.

AI in Tax Compliance and Enforcement

When deciding which AI tools to use to improve tax compliance activities, tax agencies should consider which AI implementations have the highest return on investment, weighing factors such as increasing accuracy of tax filings, enhancing the taxpayer experience, and improving compliance with the ever-evolving complex tax rules. One area where AI is likely to have an immediate impact is efforts to combat tax scams.

Disrupting Emerging Tax Scams: AI could provide the “night vision goggles” that enable the IRS to detect tax scams and tax cheats shielding income. Recent research advancements on AI anomaly detection can facilitate development of tools that identify filings involving fake W-2 forms, Employee Retention Credit schemes, fraudulent claims for unemployment compensation, or detection of ghost preparers. Using automated AI detection tools can greatly enhance the speed of detecting these scams. AI tools can also help to disrupt scams impacting taxpayers that do not directly appear in annual tax filings. For example, telephone or mail scams use generative AI (e.g., ChatGPT) to design fake IRS correspondence and target elderly populations with “notices” of fines or penalties. Tax agencies will need to deploy AI systems to counter and mitigate the effects of the growing use of generative AI in tax scams.



Audit Selection and Process: Audits represent one of the most common reasons taxpayers interact with tax agencies and are often stressful experiences for taxpayers. The IRS recognizes the utility of incorporating AI to help IRS compliance teams better detect tax cheating, identify emerging compliance threats and improve case selection tools to avoid burdening taxpayers with needless ‘no-change’ audits. Recent AI advances in task-agnostic anomaly detection may also help with managing audit selection processes by reducing the need to train new detection models from scratch every filing season. During the audit process, AI tools may provide invaluable assistance tackling intentional tax evasion schemes, which grow more complex every year.

Going forward, AI will enable tax authorities to identify these schemes and assist auditors during the review process. The IRS has already initiated the rollout of AI models to help identify risk of noncompliance in large partnerships. Overall, AI tools have the potential to enhance the speed and accuracy of audits. This can increase efficiency and enforcement effectiveness. Deploying AI can allow for more audits of highly complex evasion schemes at a lower burden to the public. This will—with compliance and reducing the tax gap—help to restore the public’s trust in the IRS’ ability to conduct audits expeditiously and fairly.

AI in Governance, Risk, and Authentication for Tax Agencies

Even with many benefits from incorporating AI across tax agency functions, the risks are real. In the U.S., leaders recognize the issues and opportunities widespread adoption of AI across agencies present and have endeavored to lead on AI governance issues and risk mitigation.

Following the Artificial Intelligence in Government Act of 2020 and Advancing American AI Act, President Biden’s Executive Order 14110 (EO 14110), issued on October 30, 2023, outlines over 100 actions for safe and responsible AI development, with eight main focus areas:

1. **Safety and Security:** Ensuring AI does not compromise biosecurity, cybersecurity, or infrastructure
2. **Innovation and Competition:** Attracting AI talent, protecting intellectual property, and supporting small business innovation
3. **Worker Support:** Addressing potential workforce disruptions from AI
4. **AI Bias and Civil Rights:** Mitigating biases in AI, especially in criminal justice and federal programs



5. **Consumer Protection:** Enforcing regulations to protect consumers
6. **Privacy:** Safeguarding personal data against AI-related privacy risks
7. **Federal Use of AI:** Creating an interagency AI council to coordinate federal AI use and provide risk management guidance, with a focus on secure generative AI (GenAI) adoption
8. **International Leadership:** Collaborating internationally to establish responsible AI standards

Subsequent guidance from the Office of Management and Budget (OMB) includes strategies for Treasury and IRS implementation of EO 14110, such as designating a Chief AI Officer, establishing AI governance boards, and performing periodic AI risk assessments. This governance framework aims to foster transparency, accountability, and a balanced approach to adopting AI within tax agencies and other federal entities.

How Can Tax Agencies Work to Mitigate AI Risks?

To mitigate AI risks in tax agencies, several strategies have been proposed, focusing on data quality, model fine-tuning, architecture, and human oversight:

1. **Robust Training Data:** Quality data is essential, as biased or erroneous data can lead to unreliable models. Tax agencies should ensure data accuracy, completeness, and representation. Using labeled data from audits and judicial cases can help improve model reliability. Additionally, diversity in the data is critical to avoid perpetuating biases, especially given past concerns with IRS algorithms.

2. **Fine-Tuning Foundation Models:** Using general foundation models, which are initially trained on massive datasets, tax agencies can fine-tune these models with tax-specific data to boost accuracy and relevance, potentially reducing errors in AI outputs.
3. **Model Architectures:** Choosing appropriate AI architectures, such as adversarial AI, can help mitigate biases and enhance performance. An “adversary” model that critiques the main model’s outputs can be effective for bias reduction, among other risks.
4. **Human Role in AI Systems:** Human oversight is crucial, particularly in complex or novel tax situations where AI may struggle. Strategies like red-teaming (testing for system weaknesses) and deploying AI as an advisory tool rather than a replacement allow humans to use critical reasoning to verify and interpret AI outputs, which is essential for minimizing risks in tax administration.


These approaches, with a strong emphasis on human involvement, aim to improve AI reliability, prevent errors, and safeguard against unintended biases or privacy violations.

How Can AI Be Used to Build Trust in the Modern Tax Agency?

To build trust in AI, tax agencies need to improve the taxpayer experience and prioritize accuracy, transparency, inclusivity, education, and continuous monitoring:

1. **Transparency and Interpretability:** Tax agencies can increase trust by being open about where and how AI is used, especially in data handling. Clear, interpretable AI models help users understand AI’s decisions, making it easier to interact with and trust these systems.
2. **Inclusive Development:** Creating AI tools based on taxpayer needs and ensuring accessibility for all, including those with disabilities or limited digital literacy, enhances trust. Building diverse development teams fosters inclusivity and ensures tools address a broad range of perspectives.
3. **AI Education:** Educating taxpayers on using AI tools and understanding their limitations increases confidence and uptake. Walkthroughs, guides, and transparency about AI’s capabilities and limitations can prevent misunderstandings that may erode trust.
4. **Continuous Monitoring and Evaluation:** Tax agencies should regularly update AI models to reflect changes in tax laws, track and address errors, and respond to taxpayer feedback. Publishing reports on AI performance, including benefits and error rates, demonstrates accountability and commitment to accuracy.

These strategies help tax agencies ensure AI tools are reliable, fair, and accessible, reinforcing taxpayer trust.





Pathways to Trusted Progress with Artificial Intelligence

Edited by Michael J. Keegan

The second contribution in this forum examines the governance and applications of AI and how governments need to develop and communicate a framework for the public to understand why AI is being used; what has been done to ensure that the AI is fair, transparent, and accurate; what experiments were done to ensure that the output is reliable; and how public value from AI is being measured and created.

What follows is an excerpt from the IBM Center report *Pathways to Trusted Progress with Artificial Intelligence* by Professors Kevin DeSouza and Greg Dawson that distills findings and recommendations from an expert roundtable of leaders in Australia that address the needs, security, and progress of delivering AI services that benefit citizens and industry. Though this report may offer perspectives from leaders in Australia, the insights present in the report and summarized here are applicable across the world. By addressing the growth and management of AI, and the governance of data aligned to AI strategies, government can take full advantage of the power of AI.

Digital Transformation Initiatives are Revolutionizing All Aspects of the Public Sector.

AI systems will play an important role in transforming government as well as the national economy. Realizing AI's potential will only occur if there is a concerted effort to ensure that citizens trust AI systems, the government, and the government use of AI.

There is a wide assortment of AI systems, and each class of AI systems has their own characteristics. However, at their core, these systems ingest vast swaths of data, employ either supervised or unsupervised learning techniques or both, and can be deployed autonomously, semiautonomously, or in an advisory capacity to augment human decision makers. Consider the following three examples of AI systems successfully deployed in the public sector:



- Fully Autonomous:** In North Carolina, AI-powered chatbots independently manage simple, repetitive customer service tasks, such as password resets. By automating these interactions without human oversight, the system frees up human agents to handle more complex inquiries.
- Semiautonomous:** Semiautonomous robots deliver food on college campuses, where students initiate orders via an app. Humans set initial parameters (restaurant and delivery location), but the robots independently navigate routes and avoid obstacles. This technology may expand to roles in search and rescue.
- Augmented Decision Making:** In African wildlife conservation, AI assists rangers by predicting poaching locations and recommending patrol routes. Here, the AI provides data-driven insights, while humans make final decisions on the patrol plan.

Each of these examples illustrates how AI operates across a spectrum of autonomy, from fully independent systems to those that support human decision making.

While there have been plenty of successful deployments of AI systems, there have also been challenges. For example, an error in the AI for Britain’s Universal Credit program caused underpayments for individuals paid multiple times a month, a common situation for lower-wage earners. This oversight placed affected recipients at financial risk, highlighting the need for rigorous testing and oversight in AI systems handling sensitive data. Given such challenges, how can government generate and maintain public trust when it comes to the design, development, and deployment of AI systems?

The Problem of Trust

Trust is a multidimensional concept that can be broken down into three components—ability, integrity, and benevolence. See Table. Trust Elements

Trust Elements		
Trust Element Definition	Example in Government	Example in AI
Ability—Belief in the competency of the trust target	Belief that government can provide national security	Belief that the AI can correctly and consistently give the correct answer
Integrity—Belief in trustee’s ability to adhere to a set of ethical principles	Belief that government will treat all people equally regardless of their gender or ethnicity	Belief that the AI will mirror society’s view of ethical principles
Benevolence—Belief that the trustee wants to do good to the trustor	Belief that government will act in the best interests of the citizen	Belief that the AI has good intentions (or not negative intentions) in its functioning and outcomes

Trust in government in general has seen a steady decline over the last few years, including in Australia. According to the 2022 Edelman Trust Barometer, only 52 percent of Australians trust government to do the right thing (down 9 points from the previous year). Interestingly, 55 percent of Australians say that their default tendency is to distrust something until they see evidence that it is trustworthy. Factors attributable to the declining trust mirror trends around the world, and include decreasing interpersonal trust, perceptions of corruption, and

deeply seated economic worries stemming from COVID-era policies. However, few of these distrust factors appear to directly involve the design or use of information systems, including AI systems.

Trust in AI, particularly in government applications, depends on whether citizens believe the government has the integrity and capability to deploy AI that serves the public interest. Examples illustrate successful AI implementations that foster trust:

- 1. **Dubai Electricity and Water Authority (DEWA):** The chatbot “Ramma,” which responds to residents’ queries in English and Arabic, reduced physical visits by 80 percent, showing effectiveness and responsiveness.
- 2. **Australian Tax Office:** This AI-enabled tax filing tool helps users verify work-related expense claims, ensuring transparency and accuracy.

Despite declining trust in government globally, studies show people regularly use AI in their private lives, with global AI adoption rising to 35 percent, according to IBM. Noteworthy private-sector AI uses include:

- **Zzapp Malaria:** Using AI to identify malaria risk areas to prevent outbreaks
- **Vistra (U.S. power producer):** An AI-powered tool improved monitoring of power plant indicators, optimizing efficiency and reducing emissions
- **Wayfair:** AI-supported logistics changes reduced inbound costs during the pandemic

These examples highlight how AI can boost efficiency and trustworthiness, especially when governments demonstrate responsible and beneficial AI use, potentially reversing trust decline trends.



Identifying Major Themes of AI in Government— Findings from the Workshop

The decline in trust in government directly impacts how much citizens are willing to trust it in the implementation of any powerful new technology. How can public sector leaders create trust in AI, given declining trust overall in government? To address this challenge, the IBM Center for The Business of Government hosted a forum of senior Australian government officials. This meeting provided a first-hand perspective from these officials on the status of AI, issues associated with AI, and the roadblocks and accelerators to implementing, culminating in the identification of five major themes of AI in government:

Theme 1—Government is in the business of providing services, and AI is simply a tool to facilitate that.

Government should remain focused on providing government services, and not get “techno dazzled” by AI.

Theme 2—Government is held to a higher standard of performance regarding AI versus private companies, making explainability and transparency of utmost importance.

Citizens expect government to get things right, and the services facilitated by AI should be sufficiently transparent and fully explained to the citizen.

Theme 3—Government needs to work holistically in terms of defining AI standard practices, operating models, etc. There is too much work and too many risks in implementing AI for standards development to happen only at the departmental level. Rather, this work needs to be coordinated at the highest level of government.

Theme 4—Adequate governance is necessary not only for AI technology, but also for the people who build AI systems and the processes used to build them. Issues emerge not only from the technology itself but also from the people and processes that implement AI.

Theme 5—There is a need to distinguish between different types of AI (fully autonomous, semiautonomous, and augmented) in establishing guidelines and approaches. Not all AI is the same, and costs, benefits, and risks differ for each type of AI. Discussing AI at a more granular can ensure optimal uses.

Recommendations Recommendations for Building Trusted AI in the Public Sector

These themes, coupled with background work done by the authors, gave rise to a list of recommendations to support building trusted AI in the public sector:

Recommendation 1—Promote AI-human collaboration when appropriate. Different kinds of AI call for different levels of human involvement, and citizens are generally more comfortable with a human being involved in providing direct services.

Recommendation 2—Focus on justifiability. Justifiability can be thought of as an outwards-facing business case, and with citizens as a primary audience. The government needs to articulate why an AI system needs to be developed, the amount of human involvement, and execution strategies.


Recommendation 3—Insist on explainability. Government must be able to explain why the AI came to a proposed decision, including the data that was used for the decision. This becomes increasingly important with decision making for high-stakes outcomes.


Recommendation 4—Build in contestability. Just as citizens can appeal to a person in government about the fairness of a decision, they also need to be able to contest the decisions made with AI. This feedback loop helps ensure that decisions are reasonable and not prone to bias.

Recommendation 5—Build in safety. While AI is deployed, risks can arise that make a safety feedback loop important. Government needs to either create or join an incidents-tracking database to capture and act upon feedback.

Recommendation 6—Ensure stability. The machine learning function in AI means that supporting algorithms will be constantly tweaked in response to new information. Not only does the AI system need auditing prior to implementation; regular examinations will ensure that AI provides stable results.

The emergence of AI in the world, and specifically in the public sector, makes this an exciting era. Given the frantic pace of AI development, government has a responsibility to be more proactive around the design, development, and deployment of AI systems to advance national goals. Government leaders can act now to implement fundamental recommendations to ensure successful AI delivery and that adequate guardrails are in place to protect their citizens.





Navigating Generative AI in Government

Edited by Michael J. Keegan

The third contribution to this forum highlights a practical set of considerations and potential actions that can help government agencies to capture benefits and minimize risks from the use and application of generative AI. Generative AI refers to algorithms that can create realistic content such as images, text, music, and videos by learning from existing data patterns. Gen AI does more than just create content, it also serves as a user-friendly interface for other AI tools, making complex results easy to understand and use. Generative AI transforms analysis and prediction results into personalized formats, improving explainability by converting complicated data into understandable content.

It has the potential to revolutionize government agencies by enhancing efficiency, improving decision making, and delivering better services to citizens, while maintaining agility and scalability. However, in order to implement generative AI solutions effectively, government agencies must address key questions—such as what problems AI can solve, data governance frameworks, and scaling strategies, to ensure a thoughtful and effective AI strategy.


What follows is an excerpt from the IBM Center report, *Navigating Generative AI in Government* by Dr. Alexander Richter that captures perspectives from two expert roundtables of leaders in Australia and presents 11 strategies for integrating generative AI in government. Though this report is based on the insights from leaders in Australia, the perspectives shared and summarized here are applicable to leaders around the globe.

Navigating Generative AI in Government— Nine Key Themes

Nine key themes essential for navigating generative AI in government are outlined in this report. These themes are based on insights from two roundtable discussions conducted in May and July 2024. These sessions convened leaders and

experts from government agencies alongside generative AI professionals, whose contributions helped identify the critical themes for successful generative AI adoption in government contexts, providing a whole-of-government perspective.

1. **Digital Transformation:** Generative AI supports digital transformation by optimizing workflows and resources, driving efficiency while encouraging innovation and learning—rather than focusing solely on new technology adoption. Successful AI adoption in government requires strong leadership, a clear strategic vision, and an environment supportive of experimentation. Agencies must identify specific use cases where AI can add value and effective AI implementation depends on both technological readiness and rigorous data governance.



2. **Use Cases and ROI:** Demonstrating tangible returns on investment through use cases such as automated IT support can justify AI investments and guide future strategies. Communicating the benefits and risks effectively is essential, supported by real-world examples of success. Learning from past successes and failures can guide future strategies.
3. **Data Foundation:** The effectiveness of generative AI relies on the quality and volume of data including dealing with legacy systems. Robust data management strategies are necessary to ensure data accuracy, relevance, and compliance with regulations. Leveraging high-quality data enables AI models to produce accurate and valuable outputs.
4. **Ethical Considerations:** Ensuring fairness, transparency, and accountability in AI practices is vital for maintaining public trust and avoiding biases. Recognizing AI as a collaborator rather than just a tool requires governance that aligns AI's actions with human values and societal norms.
5. **Balancing Experimentation with Risk Management:** Government agencies must balance the need for innovation with robust risk management, updating policies to allow safe experimentation while protecting against real risks.
6. **Shifting the Cultural Mindset:** Overcoming risk aversion is key to AI adoption. Leadership should foster a culture that encourages safe experimentation and views failure as a learning opportunity. A culture that actively encourages innovation and calculated risk-taking, rather than simply tolerating it, can help overcome risk aversion.
7. **Skills Development:** Continuous education and training programs are essential to equip the workforce with the necessary expertise to implement and manage AI technologies.
8. **Diversity of AI Tools:** Leveraging a variety of AI tools tailored to specific government needs ensures effective and secure deployments. The presence of multiple AI tools allows agencies to address a broader spectrum of challenges by selecting the most suitable technology for each specific task.
9. **Human-AI Collaboration:** Designing flexible AI systems that complement human roles enhances collaboration and decision making. These dualities highlight the need for a nuanced understanding of how AI is integrated—ensuring that the collaboration between AI and humans is both productive and contextually appropriate.



Obstacles to Adopting Generative AI

The adoption of generative AI is hindered by obstacles related to knowledge, skills, and attitudes.

- **Knowledge:** Many organizations lack a clear AI strategy, leading to confusion in defining roles for humans and AI, along with inadequate AI literacy. This results in misunderstandings, unrealistic expectations, and reluctance to collaborate with AI.
- **Skills:** Effective AI adoption requires new communication skills and role adaptations. Challenges in natural language processing and insufficient digital infrastructure make integration difficult, fueling mistrust and fear.
- **Attitude:** Cultural and ethical concerns, fears about job security, and inadequate leadership support create resistance to AI. Trust in AI is fragile, especially regarding ethical concerns, biases, and complex decision-making areas.

These obstacles collectively slow generative AI adoption and integration within organizations.


Strategic Pathways for Integrating Generative AI in Government

Along with identifying these themes, key strategic actions for responsibly integrating generative AI into government operations were curated from the roundtable discussion. To successfully integrate generative AI, government agencies should consider establishing a AI governance office to oversee initiatives and ensure ethical standards and set clear guidelines for data governance. In addition, empowering solution owners with governance capabilities will enhance

model transparency and ensure agility, while maintaining coherence across government AI strategies.

Developing adaptive governance models, investing in robust data infrastructure, promoting a culture of innovation, and implementing comprehensive training programs are critical steps. Additionally, expanding AI-driven citizen services and enhancing public engagement and transparency will build trust and ensure that AI initiatives align with public values.

Figure: Strategic Pathways for Integrating Generative AI in Government summarizes the preceding recommendations.




As government agencies increasingly embrace digital transformation, the integration of AI is not just an opportunity but a necessity for staying ahead in a rapidly evolving technological landscape.

This report has identified key themes and obstacles that government agencies must navigate to fully harness the benefits of generative AI. From ensuring robust data governance and ethical AI practices to fostering a culture of innovation and continuous learning, the path to successful AI integration is complex but achievable. By addressing these challenges with a strategic and thoughtful approach, government agencies can leverage AI to deliver public value in ways that were previously unimaginable.

Moreover, the adoption of AI in government must be underpinned by a commitment to transparency, public engagement, and ethical accountability. As AI systems take on more significant roles in public administration, it is crucial that they are designed and implemented with the public's trust and confidence in mind. This includes not only mitigating biases and ensuring fairness but also actively involving citizens in the AI journey through open communication and opportunities for feedback.

The insights provided in this report offer a roadmap for government agencies to navigate the complexities of AI integration. Establishing an AI Governance Office, investing in data infrastructure, promoting a culture of experimentation, and enhancing public engagement are all critical steps toward realizing the full potential of AI in government. As government agencies move forward with AI adoption, it is important to remember that the success of these initiatives hinges not just on the technology itself but on the people, processes, and principles that guide its use. AI should be seen as a tool that, when combined with human ingenuity, can drive meaningful improvements in public service delivery and policymaking.





Government Procurement and Acquisition

Opportunities and Challenges Presented by Artificial Intelligence and Machine Learning

Edited by Michael J. Keegan

The fourth contribution in this forum addresses using AI to transform public procurement, analyzing challenges and recommending ways to capitalize on opportunities. It makes recommendations about how to use AI to transform public procurement, which can increase speed, efficiency, and effectiveness in acquiring goods and services that serve public needs by fostering facilitating intelligent automation across the federal acquisition system.

What follows is an excerpt from the IBM Center report *Government Procurement and Acquisition: Opportunities and Challenges Presented by Artificial Intelligence and Machine Learning* by Mohammad Ahmadi and Dr. Justin B. Bullock, which recommends about how to use AI to transform public procurement. This can increase speed, efficiency, and effectiveness in acquiring goods and services that serve public needs by fostering facilitating intelligent automation across the federal acquisition system.

Challenges Facing Modern Government Procurement

To serve the public, governments procure a wide variety of goods and services. The acquisition process lies at the heart of agencies' abilities to provide goods and services in an effective, efficient, and equitable manner. Given high agency investments—annual procurement spend exceeds \$500 billion—and the complexity of government processes, procurement is an extremely challenging task to perform well. This impacts governments worldwide, and particularly large federal governments such as that of the United States.

Indeed, congressional panels and executive reports often highlight the challenges of procurement and acquisition, including prominent GAO reports addressed in this report in more detail. A 2018 GAO report in particular highlights a number of important challenges for improving federal procurement practice:



1. Requirements definition
2. Competition and pricing
3. Contractor oversight
4. Federal procurement data
5. Acquisition workforce

While these challenges do not represent the only hurdles facing procurement reform, they are challenges that AI and machine learning tools can help to address. AI applications can help improve effectiveness and efficiency across the U.S. federal procurement process. Recent studies in public administration and public management reflect a growing consensus that AI and ML tools, if carefully constructed and deployed, can help agencies complete certain types of tasks for governments. Opportunities exist for AI and ML tools to improve the functioning of procurement, but these opportunities will come with their own challenges that need to be carefully managed.

Connecting Procurement Challenges with AI Opportunities

This section highlights the AI use cases outlined in this report connecting them to a specific procurement related challenge facing U.S. federal government agencies.



Requirements definition challenge. Government procurement requires well-defined, achievable, and verifiable requirements to ensure successful contract outcomes. However, defining these requirements is often challenging due to workforce shortages, frequent turnover, and the complexity of the process, which can lead to vague or unrealistic specifications. This problem is evident in certain Department of Defense contracts, where unclear definitions have caused delays and increased costs, despite years of reforms aimed at improving procurement practices. Artificial intelligence offers promising tools to aid in specifying requirements.

AI Procurement Opportunity 1: Specification Definition. AI offers a valuable solution to improve specification accuracy. By analyzing data from past contracts, machine learning algorithms, such as recommender engines and ranking systems, can help procurement teams define requirements that are clearer, feasible, and verifiable. These tools assist by suggesting similar specifications from historical data, enabling more accurate predictions of required goods and services, especially for repeat procurements.

AI Use Case 1: Recommendation Algorithms. Additionally, deep learning models trained on historical data can help assess if specifications are realistic, providing agency staff with an added layer of verification. However, AI should complement rather than replace human expertise, enhancing efficiency in defining needs while relying on human judgment for final decisions. This approach aligns with the recommendation algorithms used by platforms like Amazon or Netflix, which effectively anticipate user preferences based on past behavior, and can similarly support procurement agents in aligning contract specifications with agency requirements.



Competition and pricing challenge. Federal procurement regulations mandate fair and reasonable pricing through competitive bidding, though agencies sometimes use sole-source contracts with only one supplier. Ensuring fair prices remains challenging, despite efforts like broadening competition, conducting market research, adopting private sector practices, and using interagency contracts. GAO reports little change in competitive versus non-competitive contract rates, with reasons including vendor preference, restrictive policies, and inadequate planning.

AI Procurement Opportunity 2: Improving Market Knowledge. AI offers promising solutions to support market research, acquisition planning, and price assessment. Machine learning can improve market knowledge by aiding acquisition teams in performing market research, planning acquisitions, and checking price fairness. For instance, virtual assistants powered by natural language processing (NLP) could search datasets and archives to provide answers and references for market questions, potentially saving time and resources. Deep learning can help identify patterns in acquisition planning issues, ensuring timely and cost-effective procurement.

AI Use Case 2: Natural Language Processing. NLP technologies, such as those in virtual assistants like Siri and Alexa, could be trained to assist with specific acquisition tasks, making market research faster and more efficient. Language models like ChatGPT can also summarize feedback, draft reports, answer complex questions, and conduct document searches. However, while AI tools offer efficiency, they should enhance human decision making rather than replace it, and ethical, privacy, and security concerns must be managed carefully as AI adoption grows in federal procurement.



Contractor oversight challenge. Contractor oversight ensures vendors deliver goods or services according to contract terms, but this process poses challenges. Uncertainty about the extent of oversight needed, combined with the technical expertise required to inspect deliveries, creates oversight gaps—exacerbated during COVID by staffing shortages and remote work environments.

AI Procurement Opportunity 3: Oversight. AI presents solutions for enhancing oversight through image and video analysis. Machine learning algorithms can detect defects, assess risks, and improve quality control efficiently and at



lower cost. Computer vision, paired with drones or robotics, could support quality inspections in diverse areas such as construction, vehicle maintenance, and product evaluation. However, sufficient training data is essential to make these systems effective.

AI Use Case 3: Some Market Offerings for Computer Vision and Machine Learning. Several AI tools showcase how image processing aids oversight. Amazon Lookout for Vision detects defects by analyzing images and generating a dashboard for process monitoring. Crowd AI uses visual inspection to reduce manual inspection time, boost defect detection accuracy, and increase equipment life across various industries. Google's Visual Inspection AI simplifies defect detection in manufacturing with minimal training data.



Federal procurement data challenge. The Federal Procurement Data System-Next Generation (FPDS-NG), the main repository for federal procurement data, faces challenges with data reliability and validation. Despite initiatives to improve data quality, GAO found FPDS-NG data to be unreliable at detailed levels and lacking in adequate validation processes.

AI Procurement Opportunity 4: Data Acquisition, Capture, and Storage. AI offers an opportunity to enhance data acquisition, capture, and storage in procurement. Automated data capturing can reduce manual data entry time, cut costs, and minimize errors. Techniques used by social media and e-businesses to track platform activity could be adapted to streamline federal acquisition data collection.

AI Use Application 4: Automated Data Capturing. Tools such as Optical Character Recognition (OCR) and Intelligent Document Recognition (IDR) leverage machine learning to digitize and extract data from physical documents, such as invoices and reports. Software like IBM Datacap combines natural language processing, text analytics, and machine learning to automate data capture from printed sources, supporting accurate, efficient data handling for procurement processes.



Acquisition workforce challenge. The federal acquisition workforce, responsible for managing billions in contracts, faces persistent challenges like understaffing, high workloads, and insufficient training. Although efforts like the Defense Acquisition Workforce Development Fund (DAWDF) support recruitment and training within the Department of Defense, ongoing gaps highlight the need for improved training and early planning.

AI Procurement Opportunity 5: Improving Task Performance. AI can play a significant role in enhancing task efficiency for federal acquisition employees by automating repetitive and time-consuming tasks, allowing personnel to focus on more complex and engaging work. Intelligent automation can speed up document processing and data interpretation, which are essential for efficient planning and decision-making.

AI Use Application 5: HHS and Grant Funding. For instance, the Department of Health and Human Services (HHS), which oversees about 70 percent of federal grant funding, has applied AI to streamline its grant administration processes. By using AI and blockchain technology in its Grant Recipient Digital Dossier, HHS consolidates and summarizes data about prospective grantees from various sources, cutting down risk assessment time from hours to minutes. This approach not only saves time but also improves risk identification, helping administrators focus on high-priority issues. With this AI-driven efficiency, HHS could save approximately \$142 million annually.

Managing the Use of AI for Overcoming Procurement and Acquisition Challenges

Using AI in federal acquisitions requires long-term investment and constant improvement, best tackled through an artificial intelligence strategy. Indeed, creating value by using AI for acquisition requires a well-managed and coordinated long-term strategy. The strategy may begin with executing small AI projects, for example by using recommendation engines or knowledge graphs for training on existing acquisition data, to help the acquisition workforce define specifications for new needs. Appropriate training will allow the acquisition workforce to know the latest AI breakthroughs and their potential applications in the acquisition process.

- **Building a Maturity Model:** The adoption of AI technologies in procurement agencies requires a structured maturity model to manage technical requirements, long-term functionality, and continuous improvement. A maturity model offers a roadmap for enhancing organizational performance, meeting objectives, and assessing capacity to avoid issues such as AI decay and bias. The model's levels vary, with frameworks like the AI Maturity Playbook and the IBM Center's maturity model for public agencies suggesting stages from entry-level to mastery, alongside six key dimensions for success: data, computational systems, analytical capacity, innovation, ethics, and strategic vision.


- **Leveraging Power of Data:** Data is foundational, with structured, unstructured, and semi-structured forms all necessary for building reliable AI solutions, although unstructured data remains underutilized. Automating data collection is also recommended for efficiency and accuracy, enabling insights throughout the acquisition process.
- **Capacity Building:** Building in-house AI capabilities is a competitive advantage, enabling agencies to collaborate across departments, develop specialized solutions, test algorithms, and maintain an effective feedback loop for continuous improvement. However, this requires financial support, as resources are essential for developing the workforce, infrastructure, and tools for effective AI integration. Managers should also establish performance indicators to measure AI adoption progress.
- **Getting Financial Support:** Financial support is essential for AI adoption. Adequate resources and performance indicators for measuring AI integration can help transition agencies toward an AI-driven future.
- **Task management:** It is critical, requiring agencies to determine which procurement tasks AI should

automate, augment, or leave to humans based on task complexity and uncertainty. For example, the Social Security Administration (SSA) uses AI to streamline data organization for judges in disability benefit decisions, preserving human judgment where complexity and nuance are needed. This approach aligns with the literature, which advocates using AI for straightforward tasks and reserving complex decisions for human intervention.

- **Avoiding Automation Bias:** Finally, risks like automation bias and “administrative evil” must be mitigated. Automation bias, the over-reliance on AI, can impair judgment, while “administrative evil” refers to potential ethical harms if AI systems lack oversight, leading to dehumanization or misalignment with public values. Overall, maturity models, ethical guidelines, and ongoing research provide frameworks to balance AI’s benefits and risks in government procurement.

This report reflects an optimism that, if properly deployed and managed, AI tools could significantly improve the effectiveness and efficiency of government procurement. However, significant risks are involved with these tools and their deployment, so management needs to deploy them carefully.





Artificial Intelligence in the Public Sector

A Maturity Model

Edited by Michael J. Keegan

The fifth and final contribution to this forum provides public sector leaders a view into the “art of the possible” by emphasizing how AI programs can accelerate the transformation of government programs to better serve the public and do this by offering them framework for establishing a successful AI program. The challenge has always been to design and implement an AI program that has all the critical elements in place to successfully achieve the goal of improved mission delivery and citizen services.

Recognizing the need to address this challenge, the IBM Center commissioned a report in 2018 by Professor Kevin DeSouza, *Delivering Artificial Intelligence in Government: Challenges and Opportunities*, which proposed an initial maturity model that gave public agencies a starting point for developing an AI capability.

Since that time much has changed in the use and application of this technology, and as such, an opportunity arose to fine tune this model, based on extensive research on how the public sector was deploying AI, documenting successful use

cases and highlighting pitfalls and lessons learned. Professor DeSouza offers this revised maturity model with significant input from frontline practitioners and academic in a follow up report, *Artificial Intelligence in the Public Sector: A Maturity Model*. What follows is an excerpt from that updated report highlighting aspects of the maturity model and insights that may help government agencies get the highest value from their efforts and investments in AI.

A Maturity Model for Designing, Developing, And Deploying AI

Maturity models are popular in a wide assortment of fields from quality management to software engineering, education and learning, organizational design, and even information systems. While each maturity model has its own peculiarities, they all provide an evolutionary framework to guide improvements and/or advancements on one or more domains.



The domain of interest in this report is AI design, development, and deployment efforts in the public sector. The maturity model has two dimensions. The first dimension represents the critical elements that need to be managed as AI projects are designed, developed, and deployed in the public sector. The model indicates that agencies must show proficiency on six core elements and those elements can be divided into two domains: technical and organizational.


- Technical domain includes big data, AI systems, and analytical capacity.
- Organizational domain includes innovation climate, governance and ethical frameworks, and strategic visioning.

The second dimension outlines the maturity levels that begin with ad hoc, followed by experimentation, planning, and deployment, scaling and learning, and finally, enterprisewide transformation.

Exploring Dimensions of the Maturity Model

The proposed maturity model outlines six core elements and five maturity levels for guiding AI initiatives in the public sector. Public agencies need to start small and be aware of the required upfront financial and time investment for data governance, computational systems, and analytical capacity.

- **Big Data:** AI relies on big data for its design and development, but once deployed, enables organizations to make sense of large data reservoirs through the application of machine learning algorithms. In an ideal world, public agencies should be able to access, integrate, and leverage data of interest in an effective and efficient manner. The COVID-19 pandemic is cited as a case study where AI-enabled data tools were quickly mobilized for public health, specifically in contact tracing and risk assessment.
- **AI Systems:** Computational systems are the engines that transform data into actionable insights and outcomes. As discussed earlier, AI applications leverage a range of computational techniques to ingest, analyze, visualize, and even act on data. AI can fully or partially automate tasks through the power of predictive data analytics, fed by multiple sources of historical and real-time data; learn from previous interactions and self-decide through the power of machine learning; and in some cases, such as chatbots, interact with users through natural language processing.
- **Analytical Capacity:** AI systems are only as good as the human analytical capacity that support them and refers to the human element related to designing, developing, and deploying AI. Organizations need a well-trained workforce that is analytically aware and has the aptitude to leverage data to derive evidence-driven insights. Public agencies face numerous challenges when it comes to recruiting, developing, and retaining analytical talent, including the general lack of analytic-savvy people in both government and in the general recruiting pool. Moreover, regardless of existing analytical talent, the need for the presence of deliberate mechanisms to leverage that talent to create organizational value is pivotal.
- **Innovation Climate:** Public agencies need to innovate if they are to deliver on their objectives given the ever-evolving environmental pressures. While innovation in the public sector continues to garner interest, we still see agencies struggle when it comes to digital transformation efforts. Experimentation is critical to the ability to innovate. Yet in public agencies, experimentation is often shunned upon due to the perception of being deemed a failure and wastage of public resources. Data challenges such as incomplete and siloed datasets and a lack of



investment necessary to upgrade legacy computational systems can significantly impact the ability of public agencies to innovate.

- **Governance and Ethical Frameworks:** Governance and ethical frameworks are vital as oversight mechanisms to ensure that AI is deployed in a responsible manner and advance public value. Governance frameworks establish accountability and assign responsibility when it comes to AI design, development, and deployment. They serve as critical coordinating mechanisms to ensure that agencywide economies of scale, learning, and value can be secured. Ethical frameworks ensure that AI mitigates issues such as bias, discrimination, and harm. When AI fails or causes harm, these frameworks can assist in providing recourse mechanisms to compensate victims.
- **Strategic Visioning:** Leadership at public agencies needs to play an active role in creating environments that are supportive of the development of AI. How they are designed, developed, deployed, and regularly enhanced need to be incorporated into the long-term strategic plans of agencies. A good strategic visioning also considers the important fact that deploying AI can change the function and design of agencies given the affordances of AI for changing work processes and engaging citizens.



Assessing Levels of Maturity

The six elements described above need to be assessed both individually and collectively in terms of their maturity. The maturity levels are noted below and go in increasing order of sophistication.

- **Ad Hoc:** The public agency does not have a plan in place to design, develop, and deploy AI. Datasets remain an underutilized asset, computational systems lack

necessary capabilities, and analytical capacity is limited or unavailable. There is limited appetite to innovate with AI, and this inertia also plays out with the absence of governance and ethical frameworks for AI.

- **Experimentation:** The public agency is actively experimenting with AI. Experimental projects leverage datasets, computational systems are being designed and/or upgraded, and analytical capacity is being mobilized around these projects. There is a growing interest in learning from early experimental efforts, and there is a recognition to invest in designing ethical and governance frameworks that support responsible experimentation and innovation.
- **Planning and Deployment:** The public agency has put in a place a plan to design, develop, and deploy its first set of AI projects. Datasets for the initial set of AI projects are of sufficient quality, investments into computational systems necessary for AI are in place, and an initiative to attract, mobilize, and retain analytical talent is underway. Senior leadership is supportive of AI efforts and initial visioning efforts are underway to incorporate AI into strategic plans of the agency. An initial set of metrics are created and agreed upon to track investments and performance of the AI.
- **Scaling and Learning:** The public agency is enacting thoughtful and repeatable processes to select and implement AI and these processes encompass all aspects of AI implementation including technical, governance, and staffing. AI projects are viewed as a critical part of the agency; a concerted effort is being made to measure efforts against the metric developed in prior maturity levels.
- **Enterprisewide Transformation:** The public agency has successfully integrated AI into a routine part of the environment and agencies can move quickly to implement additional AI projects as necessary into the environment. Because the necessary technical, governance and staffing infrastructures are in-place, design and deployment can proceed rapidly across the agency and these efforts are managed using a portfolio approach.

Going Forward

Moving up a level requires a) successfully overcoming the limitations of prior level, and b) evaluating an organization's readiness for the next level. The evaluation requires knowing what limitations public agencies need to overcome at the current level and at the next level. Therefore, the elements and levels of the proposed maturity model are intertwined and inextricably linked, rather than operating in isolation.

At the ad hoc level, some individuals who have some personal interest in AI initiatives often start talking about their ideas, which can quickly grow if a suitable innovation climate exists. Showing organizational interest in establishing analytical capacity and computational systems can greatly contribute to creating the required level of competency to prepare for moving to the next level. External pressures—for example, efforts at other countries or peer public agencies deploying AI and seeing promising results—can often act as additional stimulants for agencies to move from the ad hoc level to the experimentation level. Public agencies, however, need to start with strategic plans that consider the cost and benefit of initiating an AI initiative, particularly in terms of potential risks and harm to citizens.

Managers in charge of AI experimentations often express that the ability to share learnings from experiments with peers can effectively facilitate learning and refinements to AI initiatives. Some even believe that this enables them to do rough benchmarking across different classes of AI. Using knowledge sharing networks can support sharing of lessons learned and can therefore facilitate collaboration with both internal (e.g., middle-range managers and staff in relevant departments that contribute to AI initiatives) and external stakeholders (e.g., academia, third parties, and other relevant public agencies). These efforts are paramount for public agencies to make the leap from the experimentation level to the planning and deployment level.

At both the planning and deployment level and the scaling and learning level, ongoing collaborations between program leaders and the IT department are of paramount importance. Detailed business cases need to be developed to clearly articulate how AI initiatives advance public value and engage citizens. Thoughtful medium-range plans are required to outline how efforts on AI projects are aligned to near-term priorities. While technical infrastructure and analytical capacity are low at this level, an organization that is interested in initiating an AI initiative would benefit from developing governance and ethical frameworks and assigning key personnel to plan for recruiting or upskilling analytical capacity. This allows the agency to build a solid base for moving to the highest level of maturity—i.e., the enterprisewide level.

At any level of the model, public agencies are advised to regularly reflect on and share lessons learned and the costs and benefits of moving up a level. Metrics on AI projects should be developed and used for each level. Lack of such mechanism can lead to scaling of prior ineffective practices and poor strategies.

Check out the complete report for more details on specific steps that can enable government agencies move from one level to the next.

