

AI in the Public Sector





Table of Content

| | |
|---|----------|
| Understanding the AI landscape | 3 |
| AI risks in the public sector | 4 |
| Balancing the benefits and risks of AI | 5 |
| Getting started with AI | 6 |
| How SAP can help | 8 |
| Next steps | 9 |



Understanding the AI landscape

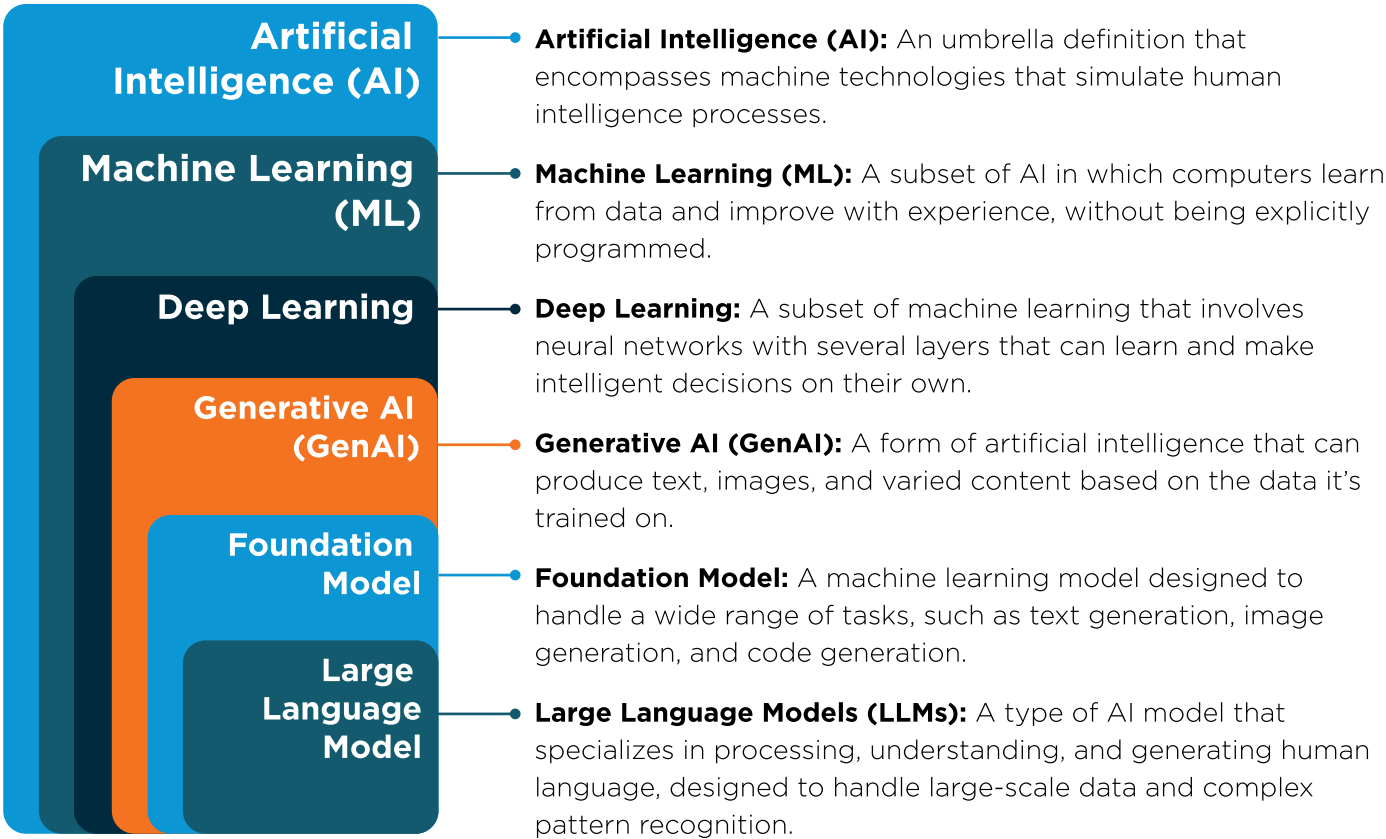
Public sector organizations are trying to find the sweet spot between innovation and risk mitigation. While they’re looking to reduce costs, increase efficiencies, and become more results-oriented, they also need to continue serving citizens — whether providing essential services, building prosperous communities, or protecting national borders.

Artificial intelligence (AI), when deployed correctly, can play a key role in that digital transformation. AI embedded into business processes and used in harnessing data insights, in particular, promises to revolutionize the efficiency, effectiveness, and agility of public services: improving citizen experiences, automating processes to optimize resources, supporting decision-making, harnessing the power of data to get insights, and achieving a stronger cybersecurity posture.

But AI isn’t a magic bullet. It comes with risks, and achieving its true potential requires strategic vision, specialist skills, and technical investment. In a survey of public sector leaders at the federal, state, and local government levels, as well as higher education, 63 percent of those surveyed by Deloitte worried that GenAI would further erode trust in public institutions, which could be why so many have taken a cautious approach to rolling out the technology.

Yet, agencies can’t lag too far behind in their adoption, since many government employees are already using AI tools to accomplish work-related tasks. Indeed, 56 percent of workers are using Generative AI on the job, according to a survey from The Conference Board. This is despite the fact that just 26 percent of those surveyed say their organization has a policy related to its use.

A glossary of AI terms



AI risks in the public sector

No matter the application, public sector organizations face a wide range of AI risks around security, privacy, ethics, and bias in data. But the public sector faces unique risks when compared to the private sector — such as the spread of misinformation that could influence elections and erode public trust.

The use of AI requires the collection and analysis of large amounts of personal data. While this raises privacy and security concerns across public and private sectors, governments in particular need to be aware of how their use of AI could impact citizen privacy. Accountability is another issue: If machines are making decisions, who is accountable for those decisions?

Whether machine learning algorithms are being used for the automation of administrative tasks or for predictive analytics, public institutions need to ensure that citizen data is safeguarded from misuse — so data collected for one purpose isn't used for another, such as AI surveillance and monitoring. Misuse could quickly erode public trust.

This issue is even more pressing with Generative AI. Training models requires the use of large datasets, which may have been collected for other purposes — so using it could violate privacy regulations, even if it's just being used for training. That's because LLMs can 'memorize' training examples and reproduce them in the production system, which could result in privacy leaks.

Generative AI use cases are typically considered high risk because of current technological limitations, such as hallucinations (when an LLM creates inaccurate, misleading, or even nonsensical outputs). This can happen for a number of reasons, such as errors in the training methodology, incomplete datasets, or inaccuracies in datasets.

LLMs are also sensitive to user prompts, so the way a prompt is phrased could lead to different responses (which is why some organizations are hiring or training prompt engineers). Other risks include biases in output, flawed algorithms, potential misuse of models, and uncertain legal or regulatory frameworks for intellectual property. Reasoning and logic issues can also be problematic for governments since the public requires transparency and accountability.

Another risk associated with prompting is that confidential data could be included in a prompt that is shared with the LLM provider. Shortly after the launch of ChatGPT, this particular risk caused many governments to publish "interim guidance" on government use of Generative AI platforms, restricting the use of GenAI tools by public servants.

Bias is another consideration. AI learns statistical patterns from training data to generate human-like language. If the training data contains stereotypes or discriminatory content, this could lead to the reinforcement of those biases in language generation.

When choosing products and services, public sector organizations should also be aware of the risk of vendor lock-in, especially in a rapidly evolving market in which LLMs are being commoditized. We're already seeing some finely-tuned models outperform more sophisticated, general-purpose models in particular domains and tasks.

There are several ways to combat hallucinations and other risks:

- Choose software vendors that prioritize responsibility. Ensure that their AI development standards and policies address security, privacy, and ethical concerns.
- Human talents should still be given final AI oversight. Using model governance tools, they should continuously review model accuracy and reliability of its outputs. LLMs need human intervention to provide context, especially for AI systems that support life-critical decision-making — and where mistakes or biases could have major consequences.
- Dataset curation (such as anonymization), advanced prompt design, and fine-tuning based on output evaluation can help to mitigate some of these risks at different stages of model development. But completely eliminating these risks may not be possible.
- Ground the models with business data and context, and provide access to multiple models so that the most appropriate model can be applied to a given domain or task.

- Since model behavior can quickly change in a short period of time, there's a need for continuous monitoring of GenAI models,
from Stanford University and University of California Berkeley.

Balancing the benefits and risks of AI

The public sector has limited resources to provide services for stakeholders. That's where AI can help.

"As public servants, it is our duty to find ways to overcome these hurdles through the use of technology. I feel strongly that AI holds the key to unlocking a plethora of services that otherwise may be impossible to implement within government," said Anthony A. Fisher, Head of Data Governance and Artificial Intelligence with Colorado's Department of Revenue.

For Fisher, the areas that excite him most are within policy interpretation through the use of AI to support stakeholder inquiries. He also sees opportunities for agencies to use AI to better understand their data and find trends that will ultimately lead to better outcomes for stakeholders.

"But, it is also our duty to ensure the privacy and security of our stakeholders' data. That must be held to the highest degree when implementing AI in the public sector," said Fisher.

Many public sector organizations are consulting with industry, academia, and the community to develop guardrails for AI. Some governments are making limited investments, although these are mostly proof-of-concepts with few production implementations. But there's good reason for taking a slow, cautious approach: they need to get it right.

"Government agencies face different risks than do private sector companies. For example, the technology can be misused to spread political propaganda or compromise national security,"
from McKinsey & Company.

"Confidential government data can be leaked or stolen if government employees inadvertently introduce that information into foundation models through prompts."

At a broader level, governments are attempting to address some of these risks through the creation of frameworks, regulations, and policies for AI:

- The European Union's *AI Act* — the first-ever legal framework for AI from a major regulator — provides clear requirements on the use of AI. This includes prohibitions against cognitive behavioral manipulation, biometric categorization, and the use of social scoring.
- California's *AI Accountability Act* aims to create guardrails for the use of AI by state agencies, including notifying citizens when they're interacting with AI.
- Canada's *Artificial Intelligence and Data Act*, part of Bill C-27, aims to put measures in place that will identify and mitigate risks of harm or biased output.

However, these efforts are still in their infancy, while Generative AI is advancing at breakneck speed. While regulations are important, individual agencies will also need to create their own guardrails to ensure responsible and ethical AI usage within their unique business context. This will include a mix of technological tools and capabilities, along with user education and training.

To ensure accountability, SAP has made 'Responsible AI' a key pillar of its strategy, with a commitment to guaranteeing the sustainable and safe development of all AI systems. The SAP Global AI Ethics Policy ensures that the development and deployment of AI systems are in line with established guiding principles, including: human agency and oversight; addressing bias and discrimination; and transparency and explainability.

Balancing benefits and risks means that many public sector leaders are “taking it slow with Generative AI.” Quite honestly, most of the many public institutions have been cautious in approaching the new technology, increasing their funding of gen AI, albeit at a lower rate than commercial counterparts.”

Are increasingly recognizing the transformative potential of AI and outlining comprehensive plans to deepen its use to boost their economies and improve societal outcomes. These plans are typically encapsulated in national AI strategies, which often include specific initiatives aimed at leveraging AI to enhance the efficiency and effectiveness of the public sector.

Getting started with AI

Getting started doesn't necessarily require a major upfront investment in technology and infrastructure, but it does require a move to the cloud. Budget-conscious public organizations can jump-start their AI development by using pre-trained solutions for specific use cases.

While some departments and agencies may have already started their AI journey with traditional AI and machine learning, others may still be in the early stages of experimentation. Moving forward means starting with low-hanging fruit and building on those successes.

“I'm optimistic about the potential of AI in the public sector. AI-enabled technology has the power to improve public services by automating routine tasks and providing personalized services to the public. This technology can help us better anticipate and respond to the needs of our communities, ensuring that government services are more effective and frictionless. As with any type of technology, it is crucial that we remain committed to ethical considerations, ensuring that AI is used responsibly and inclusively.”

- Nicholas Stowe, Chief Technology Officer, Washington Technology Solutions (WaTech)

Low-hanging fruit could include process automation, where machine learning algorithms can help automate administrative tasks such as data entry and record-keeping. This wouldn't replace public servants, but rather free them up to perform more valuable tasks. Automating repetitive admin tasks can also help to reduce errors, although it still requires human oversight.

For example, the city state of Hamburg in Germany created an aid-application platform for artists who were negatively impacted by the pandemic restrictions, which is now available in all 16 German states. Using SAP Business Technology Platform, the solution went live in three weeks, helping to quickly distribute aid to those in need.

By digitizing processes — from aid requests to payouts — processing time was shortened, even with a limited workforce. This has also provided a foundation for other funding programs run by the German government. Powered by pre-trained ML models, the platform provides AI-based decision support and automation, which reduces employee workloads and increases quality for a variety of other initiatives.

To strategically implement AI, Boston Consulting Group's **And iterative approach** that focuses on experimentation and ongoing capability development, matched to the organization's AI maturity. This can help public sector leaders “drive impactful improvements in service delivery, operational efficiency, and citizen engagement while minimizing risk and exposure and building the skills and capabilities to scale up their transformations.”

Public sector organizations can start by experimenting with AI to digitize and automate processes, side-by-side with existing business applications (an extensions approach). AI is also being embedded in next-generation business applications to boost workplace productivity, optimize often limited resources, and support better decision-making (an embedded approach).

There's no one-size-fits-all approach, and public sector organizations should consider a mix of AI capabilities to meet their diverse needs. Having access to multiple models means developers can apply the right model to the right task. For example, they may want to take a different approach for tasks that need advanced functionality or for those that require strict compliance with government regulations.

- **Assess the value of AI straight out of the box.** Evaluate how AI can create efficiencies for your core processes with pre-configured AI scenarios delivered as part of your business applications.
- **Start small and tap into AI extensions.** If you want to configure more custom AI applications, begin with small pilot projects focused on low-hanging fruit. Build on that success, and gradually scale and expand your AI footprint. Make use of available AI services to extend your core business processes.
- **Add your own AI models.** Be open to adding your own models, built for your specific use cases—whether you rely on pre-configured models or build your own from scratch. Create suitable training datasets, testing different competing algorithms on the data and assessing the models' performance against various criteria, including ethical considerations.
- **Integrate pilots and proof-of-concepts into your IT ecosystem:** Once tested and developed, AI services and models can be turned into usable scenarios. However, these models will need to be integrated into service platforms, enterprise software systems, and existing workflows. This could require changes to existing service design, as well as clarifying the role of AI in performing work.
- **Continuously monitor and review models:** With Generative AI, 'model drift' could become an issue, so it's important to continuously monitor and review outputs. For example, bias could creep into models if input data has evolved but the model hasn't been retrained on the new data.
- **Provide human oversight:** Establish a process for validating outputs — such as human-in-the-loop, human-on-the-loop, or human-in-command — where human talent reviews and validates AI-generated outputs. Human talent should be well-trained to understand how to properly review LLMs.

Over time, public institutions will likely want to overhaul their existing IT architecture with modern cloud-based platforms that can handle the sheer volume and scale of model training. They may also want to collaborate with other public agencies, as well as academic talent, to extend their capabilities and speed development.

They may also want to keep artificial general intelligence (AGI) on their radar, which is currently a hypothetical concept. AGI refers to highly autonomous systems that can outperform humans at most economically valuable work.

If realized, AGI would be able to understand, learn, adapt, and implement knowledge across a wide range of tasks. While Generative AI can be a component of such systems, it's not equivalent to AGI. Rather, GenAI focuses on creating new data instances, whereas AGI denotes a broader level of autonomy and capability.



How SAP can help

In the near future, AI will be an expected capability in enterprise applications. That's why AI is being embedded throughout the SAP portfolio, with AI tools and capabilities built into core business processes — connecting finance, human resources, IT, and procurement.

Today, SAP offers a large catalogue of AI-powered scenarios across all business functions

| | | |
|--|--|--|
| Finance <p>SAP S/4HANA Cloud Public Sector Management, Public Edition</p> <ul style="list-style-type: none"> Joule availability (planned release) <p>SAP S/4HANA Cloud for Behavioral Insights, Private Edition</p> <p>SAP S/4HANA Cloud Public Edition</p> <ul style="list-style-type: none"> Joule availability <p>SAP Sustainability Footprint Management</p> <ul style="list-style-type: none"> Intelligent emission factor mapping <p>SAP Enterprise Service Management</p> <ul style="list-style-type: none"> Next Generation Shared Services Automation | Policy & Planning <p>SAP Analytics Cloud</p> <ul style="list-style-type: none"> Just Ask Search Search to Insight Smart Discovery Smart Predict Predictive Planning Predictive Forecast Smart Grouping | Citizen Experience <p>SAP Service Cloud, Version 2</p> <ul style="list-style-type: none"> Email Draft Recommender Account Synopsis for Sales Case Summary <p>SAP Field Service Management</p> <ul style="list-style-type: none"> Intelligent Filtering Equipment Insights |
| Supply Chain and Procurement <p>SAP Ariba</p> <ul style="list-style-type: none"> Category Management market segmentation Category Management cost structure Category Management market dynamics <p>SAP Concur</p> <ul style="list-style-type: none"> Faster Itemization of Hotel Receipts Request Assistant: Cost estimates <p>SAP Fieldglass</p> <ul style="list-style-type: none"> AI-enhanced SOW description generation AI-enhanced job description generation AI-enhanced translation of job description | Human Resources <p>SAP SuccessFactors</p> <ul style="list-style-type: none"> Joule availability Job Description Generation Assisted Interview Questions with MS Teams Learning in Viva with Co-Pilot AI-assisted person insights for compensation discussions AI-assisted performance goals AI-assisted development goals AI-assisted Writing AI-assisted Talent Feedback | IT and Cross-Function <p>SAP AI Core & Launchpad</p> <ul style="list-style-type: none"> Generative AI Hub <p>Document Informative Extraction</p> <ul style="list-style-type: none"> Premium Edition <p>SAP HANA Cloud</p> <ul style="list-style-type: none"> Vector Engine <p>SAP Build Code</p> <ul style="list-style-type: none"> Code Generation (Joule) <p>SAP LeanIX</p> <ul style="list-style-type: none"> Inventory AI prompt AI-assisted text AI-generated context AI-supported translations |

<https://open.ais.sap.com/business-views/000D3ABE772D1EDCA0E6F12B2587FEC9>

SAP's Generative AI copilot, Joule, is revolutionizing the way users interact with SAP business systems in a secure and compliant way. Joule not only provides contextualized information and helps users navigate complex systems, but also improves developer productivity by incorporating code-generation capabilities for data models, application logic, and test script creation.

Applied to the context of Public Sector Management, Joule can streamline and improve the user experience for regulated industries as well as for higher educational, research and public sector customers, allowing them to navigate their financial systems more efficiently and make better strategic decisions. Financial system rollouts in the public sector and higher education typically require extensive training, which Joule can accelerate. Instead of searching through documentation for help, users can simply ask Joule for definitions, summaries or how-to guidance, saving time and reducing inflows to IT departments.

Developers can access pre-built integrations across SAP applications with SAP's Business Technology Platform (BTP), which allows them to build custom AI-enabled applications on Microsoft Azure, Google Cloud, and Amazon Web Services.

AI Foundation on SAP BTP provides a one-stop-shop for development teams to create AI- and GenAI-powered extensions and applications. AI Foundation includes everything development teams need to create business-ready AI applications, from ready-to-use AI services and access to leading LLMs, to vector database capabilities, and AI runtime and lifecycle management. With this access, developers can orchestrate multiple models.

This hub provides tooling for prompt engineering, experimentation, and other capabilities to accelerate the development of GenAI applications. Development teams can submit a prompt to multiple LLMs, compare the generated outcomes to identify the best-suited model for the task, and gain greater control and transparency with built-in prompt history.

SAP Build Process Automation, SAP Signavio, and SAP Integration Suite also feature Generative AI capabilities and deeper integration, allowing users to better analyze, integrate, automate, and monitor business processes across all SAP and third-party applications. Users can also pinpoint opportunities for automation and immediately implement process improvements.

Next Steps

Generative AI isn't meant to replace people; rather, it's meant to augment human functions and free workers from mundane, tedious tasks. It's important to address any concerns the public service workforce may have about AI, provide foundational training on GenAI, and offer upskilling or reskilling opportunities where applicable.

By developing competencies in AI, public institutions can cultivate a culture of innovation and revolutionize the efficiency and flexibility of public services. Find out how SAP can help empower your organization to be more agile, responsive, and accountable, while providing citizens with the convenience and transparency they expect in the digital age.

For more information, please contact your SAP Partner.

Use cases for AI in the public sector

AI in finance: Increase performance across a range of financial activities from funds and grants management, to invoice management, auditing, forecasting, and more.

AI in policy and planning: Tap the power of AI to forecast medium- and long-term program expenditures and to understand causal relationships between policy targets to create impact.

AI in citizen services: Increase citizen satisfaction by leveraging GenAI to offer agile, personalized, and inclusive services.

AI Copilots: Simply ask a question or frame an issue to receive intelligent insights grounded in your business data and enriched by your business context.



CONNECTING GOVERNMENT

WWW.PUBLICSECTORNETWORK.COM

NORTH AMERICA
P +1 (510) 556 0789
E contact@publicsectornetwork.com

AUSTRALIA / NEW ZEALAND
P +61 2 9057 9070
E info@publicsectornetwork.co