

A time of transformation for *the public sector*

Using generative AI to move
public services forward

Generative AI can revolutionize citizen services

The public sector must navigate a fine line; building on generative AI's potential to create value while maintaining citizen trust, controlling data, and adhering to regulations and ethics

The private sector typically leads the way in technological innovation, and the public sector follows. Private companies have the means to invest in research and development in the hope that these investments will generate higher revenue, help gain a bigger market share, and allow the business to stay ahead of its competitors.

Public sector organizations, on the other hand, are not driven by the same business goals as the private sector. Instead, they prioritize providing stable and reliable services to their communities. This creates a gap in innovation between the two sectors, with public sector organizations being unable to take the same calculated risks as their private counterparts.





But things are *changing*

When OpenAI launched ChatGPT in November 2022, the response was phenomenal. ChatGPT saw 100 million monthly active users within just two months of its release.¹

ChatGPT marked a turning point for a significant number of people who began using generative AI as a tool to enhance their work. In the initial 12 months of its release, ChatGPT was used for everything from generating articles for national newspapers to leading a church service. Almost overnight, organizations were thinking about how they could integrate generative AI into their business models.

Unlike many other emerging technologies, generative AI tools like ChatGPT, GitHub Copilot, Midjourney, and Bard have been democratized from the outset. As Guy Scriven, US technology editor for The Economist, says: “When new technologies emerge, they benefit different groups at different times.” But one of the most remarkable aspects of generative AI is the speed at which it has been adopted by millions of people daily. Due to the ease of use and access, many public sector started to almost immediately experiment with generative AI. The challenge for the public sector now is to keep the momentum going and not get left behind as organizations find innovative ways to use generative AI.

There has been a frenzy of excitement around generative AI. But as the hype begins to dissipate, the focus for both public and private organizations turns to how this technology can be best utilized and, crucially, how this can be done safely and responsibly.

Here, the public sector truly has the opportunity to lead the way. But it must tread carefully.

¹ [ChatGPT sets record for fastest-growing user base](#), Reuters, February 2023

Civil servants and citizens want to see the public sector improve



Government agencies around the world are seeing two major trends. One of these trends is a growing need for more efficiency.

The lack of efficiency is due in large part to the fact that many public sector organizations are dealing with an overwhelming increase in workload. Over half (56%) of civil servants report significant increases in their workload since 2020.² At the same time, many public sector agencies are struggling with a shrinking workforce due to demographic change and the competition for talent with the private sector.

Another trend facing governments today relates to citizen experience. Over the past years, people have grown increasingly frustrated with the services provided by government agencies. A recent study found that citizens are 21% less satisfied with government services than with private sector services.³ While private sector services have become more and more digital and quick, the same level of service is often not being provided by government agencies.

But there is also good news. While the public sector has the reputation to be technology-averse, things are different these days: according to a recent study, over two-thirds of public sector workers believe that new technologies can significantly improve their work performance and the majority agrees that investment into technology such as artificial intelligence can reduce their workload.⁴

So the question becomes: how exactly can generative AI be applied to the public sector?

² [Public servants around the world raise alarm over increasing workload](#), Global Government Forum, September 2023

³ [2023 Edelman Trust Barometer Global Report](#)

⁴ [Public servants around the world raise alarm over increasing workload](#), Global Government Forum, September 2023

Improving efficiency and enhancing user experiences

Government agencies around the world are seeing two major trends. One of these trends is a growing need for more efficiency.

Generative AI has the potential to revolutionize the way the public sector operates. One way it can do this is by taking over repetitive tasks, allowing for faster workflows. For example, a large language model (LLM) can be used as an administrative assistant to help manage a city's support desk, quickly processing and associating cases, and automating data entry to reduce processing times. This would free up human resources to focus on more complex tasks.

Generative AI-driven chatbots can also provide round-the-clock support to citizens, answering queries quickly and efficiently. This would enable citizens to access the information they need at any time of day or night, without having to wait for a human response.

In the context of tax services, generative AI could enable self-service platforms and provide personalized guidance on tax benefits and payments. This would save time and resources for both citizens and tax authorities, while also reducing the potential for errors.

The use of generative AI in the public sector can result in a decrease in operational costs due to automation of tasks and more efficient processes. This, in turn, would lead to better user experiences for citizens, as they would be able to access the services they need more quickly and easily.



The challenges of *generative AI*


While traditional AI models are designed to accurately predict an output based on input data, generative AI models go beyond prediction and create entirely new data that did not exist before.

Although this might seem intelligent, it is crucial to remember that generative AI does not have general intelligence and is completely reliant on the information it has been given. To generate new data, models are trained on a large amount of labeled or unlabeled data, which is fed into the model during the training process. The model learns patterns and relationships in the training data, and uses this knowledge to generate new data that is similar to the training data, but also unique and novel.

The performance and accuracy of a generative AI model is determined by the quality and quantity of the training data. It is essential to ensure that the training data is representative of the distribution of the data that the model will encounter in the real world.

One of the greatest weaknesses of publicly available generative AI models is that they are prone to mistakes. They may sometimes produce output that is not accurate or realistic, and may contain errors, distortions, or inconsistencies. Occurrences of false information, known as “hallucinations,” can occur due to a variety of reasons, such as insufficient training, or errors in the model architecture or implementation. Hallucinations can also result from the inherent limitations and constraints of the generative AI model, such as the inability to capture complex or nuanced relationships in the data, or the inability to generalize to new and unseen data.

Because of the possibility of hallucinations, organizations need to have a level of explainability; they should have the ability to understand and interpret the decision-making process of AI models that generate new data. The challenge with generative AI models is that they often lack transparency, making it difficult to understand how they arrived at their output. Transparency and explainability help to ensure that the output is relevant and trustworthy. Transparency also involves the AI model being open about its limitations. Rather than present incorrect information, it should be built-in that a model admits to not knowing what the correct answer is.



AI safety is about *human rights*

One of the most worrying things that people have noticed about generative AI is that it has in certain situations shown bias against specific groups of people.

As The New York Times reported: “Tech companies acknowledge machine-learning algorithms can perpetuate discrimination and need improvement.”⁵

The fact that the technology has shown unfair or prejudicial treatment of individuals based on certain characteristics like gender or race is unacceptable regardless of the organization. For the public sector, it goes against everything that they stand for. Government agencies, services, and policies must be equally accessible and provided to all entitled individuals, regardless of their race, ethnicity, gender, religion, disability, or other protected characteristics.

Failure to do so can result in legal action and penalties, as well as damage to the organization’s reputation and credibility. Public sector organizations need to have robust policies and procedures in place to prevent and address any instances of discrimination and to ensure that their actions and decisions are fair and just for all individuals.

5 Black Artists Say A.I. Shows Bias, With Algorithms Erasing Their History, The New York Times, July 4 2023



Building *unshakable trust*



A smoother experience for all citizens

The Generalitat de Catalunya, which governs Catalonia's self-rule in Spain, has introduced generative AI to respond to citizen queries more efficiently.

This initiative started in the Business Management Office, where the initial phase involved testing the usefulness of AI in anticipating and automating responses. Although the AI's responses are still monitored by an official, the model will eventually work independently. This will lighten the workload of civil servants and enable better interactions for citizens.

One of the primary concerns is whether the AI can comprehend and communicate effectively in both Catalan and Spanish, as both languages are official in Catalonia. Users must receive the same quality of experience no matter which official language they wish to communicate in. Capgemini has played a significant role in the experimentation phases of the project, and Google Cloud, also involved in the project, has already trained its Palm 2 AI language model to understand and express itself in Catalan.

Organizations can train generative AI models with their own data as a way to tackle the challenges of hallucinations, transparency, and explainability. But entering confidential information into a publicly accessible AI model would represent a serious privacy risk. If public sector organizations want to feed secure data into a generative AI model, they need to have a customized model that's for their use only.

A customized AI model also requires a trust layer around it. Establishing this trust layer means that organizations can monitor potential model usage leakages, ensure that AI tools are being used securely and privately, and that the model is not showing signs of discrimination or negative bias.

As part of the trust layer, guardrails need to be put in place. These prevent AI models from producing inaccurate information or generating responses that contradict an organization's values. With guardrails, the system can avoid hallucinations and maintain its accuracy and consistency. It is also necessary to establish boundaries and ensure that the models admit their lack of knowledge when necessary rather than presenting false data.

An aerial, high-angle photograph of a public square with a geometric, mosaic-like floor pattern. The image is heavily tinted with a vibrant blue color. A glowing, wavy blue line starts from the top left and winds its way across the square, passing through groups of people. The people are seen from above, some walking, some standing, and some looking at their phones. The shadows cast by the people and the buildings are sharp and dark, indicating bright sunlight. The overall atmosphere is one of a bustling, modern urban environment.

An ethical approach to *generative AI*

When King Charles III addressed the International AI Safety Summit in November 2023, he did not understate the potential of AI.

In a pre-recorded video, he told the summit: “The rapid rise of powerful artificial intelligence is considered by many of the greatest thinkers of our age to be no less significant, no less important than the discovery of electricity, the splitting of the atom, the creation of the World Wide Web, or even the harnessing of fire. AI holds the potential to completely transform life as we know it.”

When we hear hyperbole like this, it is easy to assume that the advent of generative AI marks a fundamental change for the public sector. But that is not the case. The role of the public sector remains fundamentally the same: to provide essential services, ensure social welfare, maintain infrastructure, promote economic growth, and provide relief in times of emergency. Generative AI itself is not the primary focus, instead it acts as a helpful resource that can assist government agencies in carrying out their duties in a more productive and impactful manner.

For generative AI to succeed, it needs to be used with a strong focus on safety and ethics. As King Charles III mentioned during his summit speech, we can only fully realize the vast benefits of AI if we address its significant risks. He urged governments, the public sector, civil society, and private companies to collaborate in protecting people’s privacy and livelihoods, securing our democracies from harm, and ensuring that the advantages of new technology are shared by everyone. And to do this in ways that adhere to our established values and laws.

At Capgemini, we strongly believe in upholding ethics and placing them at the forefront of all our actions. Since 2019, we have had our AI Code of Ethics to ensure that our advancements in AI align with our broader values.

Our commitment to ethics and technological innovation is passed on to our clients as well. We provide transparent and explainable AI, ensuring that the outcomes are easily understandable, traceable, and auditable. Our AI is controllable, meaning that humans always have the final say. We also prioritize data protection and privacy in our generative AI operations, ensuring that data remains secure and legally compliant with privacy regulations from the very beginning.

With over 30,000 data and AI consultants worldwide, we can draw from our extensive experience across various industries and technologies. As a trusted strategic partner for data-driven governments around the world, Capgemini helps public sector organizations harness the potential of generative AI.



About Capgemini

Capgemini is a global leader in partnering with companies to transform and manage their business by harnessing the power of technology. The Group is guided every day by its purpose of unleashing human energy through technology for an inclusive and sustainable future. It is a responsible and diverse organization of nearly 350,000 team members in more than 50 countries. With its strong 55-year heritage and deep industry expertise, Capgemini is trusted by its clients to address the entire breadth of their business needs, from strategy and design to operations, fueled by the fast evolving and innovative world of cloud, data, AI, connectivity, software, digital engineering, and platforms. The Group reported in 2022 global revenues of €22 billion.

Get the Future You Want | www.capgemini.com

Contact

DR. PHILIPP FUERST

VP Data-Driven Government & Offer
Leader, Global Public Sector

philipp.fuerst@capgemini.com

PIERRE-ADRIEN HANANIA

Generative AI Lead for EU Institutions

pierre-adrien.hanania@capgemini.com

Copyright © 2024 Capgemini. All rights reserved.