



Generative AI and Government Innovation

Unlocking the future



Content

04

Introduction

06

Navigating the UAE AI Strategy 2031 in the Age of Generative AI

08

What is Generative AI?

14

What is the potential of Generative AI for government innovation?

18

Future Scenarios for Generative AI in Government

30

How is Generative AI governed and regulated?

38

International Case Studies applying Generative AI in the government work

- International Case Studies applying Generative AI in the government work
- AI-Powered Civic Engagement in Singapore
- AI-Driven Healthcare Transformation in the United Kingdom

Introduction

Artificial Intelligence (AI) is a rapidly growing field that has been changing the way we live, work and interact with technology. Over the past few decades, the rise of AI has been phenomenal, and it has revolutionized every industry, from healthcare and finance to transportation and entertainment.

AI refers to the development of computer systems that can perform tasks that typically require human intelligence, such as visual perception, speech recognition, decision-making, and language translation. The field of AI has been driven by advancements in machine learning, natural language processing, and robotics, which have led to the creation of intelligent machines that can learn from experience, adapt to new situations, and perform tasks more efficiently than humans. As AI continues to evolve and become more sophisticated, it has the potential to transform the way we work, communicate, and solve complex problems.


Over the past decade, AI has become an increasingly critical technology for the public sector. The ability to sift through enormous amounts of data to create unique insights, as well as automate products and services for citizens, make AI a powerful tool for improving quality of life. AI is already used by governments around the world, allowing them to deliver more personalized citizen services, streamline bureaucracies, deploy digital assistants, enable more effective policing and judiciary systems, improve education outcomes, and stimulate collaboration across different agencies. The public sector is more adaptable and can deliver higher quality services today, thanks to AI.

AI in government also brings risks: it reinforces societal biases, it requires stronger (and more expensive) cybersecurity to safeguard citizen data, and it requires more technical expertise, both for implementation and regulation. Like all tools, AI use must be intentional and careful in order to maximize its public utility.

Yet, we are only at the beginning of harnessing AI's true potential for governments. The pace of innovation is so rapid that the advanced programs that agencies are using today will feel primitive in 10 years' time.

In November 2022, AI in government took a massive leap forward when the world entered the age of "generative AI". This jump in AI power might be likened to replacing a screwdriver with a drill, or hand saw with a power saw. Generative AI is faster, smarter, and more useful, thanks to its unique properties.

Specifically, generative AI is a subset of AI that involves creating new data or content, and has shown great potential in transforming the way governments operate, interact with citizens, and provide public services. Generative AI refers to machine learning techniques that enable computers to actually make new things that have never existed before, such as images, videos, text, or audio, based on patterns and rules learned from existing data. This technology can be used to automate and streamline various government functions, from predicting patterns in crime data to generating personalized content for citizens.




As governments face increasingly complex challenges, from managing large-scale data to improving public services in the post-pandemic era, Generative AI has emerged as a powerful new tool for addressing these challenges. The ability to generate new data and content can help governments to improve decision-making, optimize resource allocation, and enhance citizen engagement. At the same time, the use of Generative AI in the public sector raises important questions around privacy, ethics, and transparency, highlighting the need for careful consideration and governance frameworks.

This special report aims to explore the potential future applications of Generative AI in the public sector, as well as the challenges and considerations that must be taken into account when implementing this technology. We will first discuss the technology itself and what makes it unique, followed by the top emerging use cases for government agencies.

As AI rapidly transforms the government landscape, the United Arab Emirates (UAE) has positioned itself at the forefront of this revolution. This report has been co-designed with the UAE Minister's Office for Artificial Intelligence, Digital Economy and Remote Work Applications, as part of the UAE's ongoing strategy to be the world's most AI-ready nation by 2031. This ambitious mandate is a testament to the nation's commitment to integrating AI across a spectrum of societal and governmental functions. This report provides a nuanced view of the transformative potential of generative AI, and offers early insights into the intricacies of unlocking government innovation in the UAE government.

Navigating the UAE AI Strategy 2031 in the Age of Generative AI

In 2017, the UAE became the first nation with a government minister dedicated to AI, thereby signalling its commitment to this transformative technology. The UAE's AI strategy is a comprehensive roadmap that targets key sectors, including healthcare, transport, space, renewable energy, water, and education. The overarching goal is to save costs, minimize government paperwork, create new economic opportunities, and improve the quality of life.



The new 2031 strategy aims to position the UAE as a global leader in AI by 2031, driving the UAE's economic growth through a new AI-empowered digital economy. This vision is anchored on four main pillars:

- 1 Fostering a generation equipped with AI skills
- 2 Creating an enabling environment that allows businesses to thrive in the AI space
- 3 Creating the innovative regulatory landscape for AI to flourish
- 4 Nurturing the country as a testbed for AI deployment

Generative AI also plays a crucial role in adopting new technologies within the UAE's ecosystem. By integrating generative AI into sectors such as healthcare, education, and urban planning, the UAE is fostering an environment that welcomes technological progress. This proactive approach to adoption encourages businesses and institutions to explore the possibilities of AI, promoting a culture of innovation and technological advancement.

Moreover, the UAE's emphasis on generative AI is helping to create a robust digital economy. As generative AI continues to evolve, it will create new industries, job roles, and economic opportunities. It will also drive efficiency and productivity, reducing costs and enabling better decision-making. By harnessing generative AI, the UAE is positioning itself at the forefront of the digital economy, ready to capitalize on the vast opportunities it presents.

What is Generative AI?

Generative AI is a subset of artificial intelligence that focuses on creating new data from existing datasets. It is designed to learn patterns, structures, and features from input data and generate new, original output that maintains the same statistical properties as the training data. This kind of AI is often used to create content such as images, music, speech, and written text.

One of the most common methods for creating generative AI is through the use of Generative Adversarial Networks (GANs), which consist of two neural networks competing against each other: one (the generator) creates the data, and the other (the discriminator) tries to distinguish between the generated data and real data.

On the other hand, AI, or Artificial Intelligence, is a broad term that encompasses any instance where machine behavior mimics human intelligence. This includes a wide range of applications, from simple rule-based systems like a calculator, to complex deep learning systems like a self-driving car.

AI can be broadly categorized into two types: Narrow AI, which is designed to perform a specific task such as voice recognition, and General AI, which is a system that can understand, learn, adapt, and implement knowledge across a broad range of tasks at the level of a human being.

Generative AI falls under the umbrella of Narrow AI, as it is specifically tasked with generating new data based on the patterns and structures it has learned from its training data. In contrast to other forms of AI, such as discriminative models which are used to classify or predict specific outcomes from input data, generative AI's primary focus is on the creation of new, previously unseen outputs.

Generative AI is considered the most explosive technology trend of 2022. Researchers have been showcasing the industry's potential for several years already, but its first widely available public tool was ChatGPT, a chatbot which launched on November 30, 2022 (GPT stands for "Generative Pretrained Transformer"). People could simply create a free account and start asking the AI chatbot to generate any written content they wanted, such as:



Detailed essays for business school classes that convinced professors they were written by students instead of an algorithm.



Poems written in the style of legendary Arab poets like Imru' al-Qais or Labid.



Instant document translations into dozens of different languages and scripts, including Arabic, Farsi, Urdu, Hindi, and even niche ones like Assyrian.



Press releases for local government agencies following a natural disaster.




Social media posts that are statistically more likely to go viral.

Of course, The potential impact of generative AI is not defined by the amusing tasks it can perform today, but the deep well of future possibilities as the data models improve and algorithms become more sophisticated. It will be revolutionary for the public sector. For example, tech billionaire Bill Gates insists that it will fundamentally change “the way people work, learn, travel, get health care, and communicate with each other.” Other industry players call it a “once-in-a-lifetime breakthrough” that allows users to go from a “lightbulb in your head to a first draft or sketch in seconds”.

The possibilities feel endless. Other use cases currently under discussion include new architectures of search engines; explaining complex algorithms; creating personalized therapy bots; helping build mobile apps from scratch; explaining scientific concepts; and writing recipes, among others. The writing ability is important not only for inventing new content, but also editing and clarifying existing content to match specific parameters.

Behind this digital superpower is a massive trove of information. ChatGPT looks at the words that are typed in, searches through its data banks for the most probable answers, and then returns what it believes to be the most probable response. The AI model is designed to understand the patterns and relationships within the data, which it then uses to generate new examples that are similar to the training data. This magic is achieved by sampling from a probability distribution learned by the model. It should go without saying, but while tools like ChatGPT are highly accurate on certain topics, they should not be relied upon as a sole source of information, and all responses should be critically evaluated and confirmed with other sources before being used.





Of course, generative AI is not limited to text alone, but as mentioned above, can also tackle other creative mediums, including but not limited to: images, videos, architecture, biological protein modelling, video games, graphic design, financial modelling, data analysis, product design, and computer code. We are in the early stages of a massive new global ecosystem of companies, business models, regulators, applications, non-profits, researchers, and more, all trying to figure out the fullest extent of AI for public good – on par with the breadth of mobile app stores today.

In fact, pretty much every field that involves human creative endeavors can be either augmented or completed by a machine. Fundamentally, that means smaller local government agencies will have tools at their disposal only previously available to national agencies.

As generative AI becomes mainstream there are numerous risks. Will the internet be flooded with iterative artwork that is indistinguishable from reality? Will teachers continue to be able to assign take-home essays to students? Will large parts of the creative workforce, including commercial artists working in entertainment, video games, advertising, and publishing, lose their jobs because generative AI models can accomplish their roles faster and cheaper? This is where governments need to step in.

What is the potential of generative AI for government innovation?


In short, generative AI will streamline, optimize, and unlock new value from government services. Generative AI can automate repetitive tasks, freeing up time and resources for other important work. Generative AI can provide insights and recommendations based on vast amounts of data, helping decision-makers make better informed decisions and more data-driven policies. Generative AI can help government agencies better understand and respond to the needs of citizens, leading to improved services and outcomes. Generative AI can be used to detect and prevent fraud in government programs, such as welfare benefits and tax refunds. On the flip side, Generative AI algorithms can cause issues like biased and unfair outcomes, data privacy and rise of cyber security challenges. Careful attention must be paid to prevent this.

Government adoption of generative AI tools is a matter of if, not when. Search queries and chatbots are already staples of public sector services, and generative AI will take them to the next level.


For example, Google Search is found embedded in many local government websites—searching first in the city or county databases. Local governments have also found themselves using voice assistants like SIRI and Alexa licensed by Apple or Amazon. Even before the COVID-19 pandemic, search and chat enhancements were growing. But during the pandemic, the use of such technologies exploded with mostly positive results. Both in the private and public sectors, chatbot technology proved to be a worthy addition to staff who were overwhelmed with calls and having to ask the same routine questions over and over. Chatbots were often deployed as a first line of communication to augment staff resources and to fill in on nights and weekends, thus allowing staff to focus on unique challenges requiring their expertise. Generative AI is repeating this process anew, with much more processing power, speed, and phenomenal analytical capabilities.

Besides profound effects on tasks and jobs, generative AI models and associated externalities have raised alarm in the governance community. One of the problems with large language models is their ability to generate false and misleading content. Facebook parent company Meta's Galactica – a model trained on 48 million science articles with claims to summarize academic papers, solve math problems, and write scientific code – was taken down after less than three days of being online as the scientific community found it was producing incorrect results after misconstruing scientific facts and knowledge.

Nevertheless, governments shouldn't sit on the sidelines because privacy, security, and ethical questions are difficult to answer. Agencies need to understand and use these systems well to regulate them effectively and manage the changes they will bring. Governments need to use new technologies, safely and ethically, to better support and protect the communities they serve.



In fact, this is already happening as the following early-adopter examples demonstrate:



As reported in March 2023 by Semafor, In **Portugal**, the national government released a ChatGPT-powered chatbot for citizens to ask legal questions. Officials say it will give citizens information on court proceedings, answering questions about documents required for a marriage license or Portuguese citizenship, for example. The goal of the program is to help reduce the burden on court staffers usually tasked with similar requests.



As reported in March 2023 by Semafor, in **Romania**, the government unveiled a generative AI-powered chatbot that analyzes feedback from Romanian citizens on social media and online portals, and then relays that to leaders to better inform their policy decisions. Lawmakers are enthusiastic the system will help tear down walls between bureaucrats and the people, and note how the chatbot will be able to filter out machine-generated content to ensure only human voices are heard.



As reported in February 2023 by DefenseScoop (a US-based media outlet focusing on technology trends in the military and defense industries), In the **USA**, the Central Intelligence Agency (CIA) is investigating how generative AI can assist intelligence agencies. The CIA is already exploring how to use chatbots and generative AI to assist its officers in completing day-to-day job functions and their overarching spy missions. For example, the agency is creating a common platform to enable shared services — to ultimately actualize and scale AI applications — and establishing new resources and opportunities to increase its staff's familiarity with smart and automated technologies. The CIA's AI team has also started to sharpen its attention on how ChatGPT and similar capabilities might provide more advanced assistance to agency personnel as they conduct critical intelligence processes. The agency is not trying to implement these tools alone, however, and is partnering with federally funded research and development centers as well as other partners in academia and industry.



As reported in January 2023 by The Boston Globe, a **Massachusetts state** lawmaker used ChatGPT to help him author new legislation that regulates similar generative AI programs – becoming one of the first elected officials to use an AI-generated text for such a purpose – arguing that guardrails are needed on the technology. He and his staff iterated multiple times with different prompts until the finished product was a “workable draft”. The AI suggestions for regulation included a requirement that similar programs not engage in discrimination or bias against any individual or group based on federally protected characteristics, prevent plagiarism by generating all text with a distinctive watermark, and implement “reasonable” security measures to protect the data of individuals used to train the model. Any company that operates a large-scale generative AI model like ChatGPT would have to register with the state attorney general's office within 90 days of the bill's effective date.


Future Scenarios for Generative AI in Government

In this section, we offer initial ideas and insights for 25 different future ideas for generative AI in government. Generative AI itself was able to suggest some of these ideas, while others originated from leading sources and expert advice. Most of these use cases are yet available as described, but certainly will be within the coming five years.



1. Offer policy feedback

Generative AI can instantly read a proposed policy or legal document, and then generate specific points of feedback to improve the regulations. This feedback can range from improving syntax to optimizing how the policy document can advance government priorities. For example, it can suggest alternative wordings to make a regulation more binding or powerful, or close potential loopholes before they can ever be exploited, or point out sections that do not make sense. Since the algorithm can read through all existing laws instantaneously and simultaneously, it can also ensure compliance with existing policies and provide specific feedback where there are potential issues, much faster (and more accurate) than any team of lawyers.



2. Better understand policy/law/regulation/project criticism

Generative AI can produce hypothetical negative criticisms from different population demographics, allowing government leaders to better understand the project or policy's shortcomings, as well as how it will be perceived. For example, how different groups might reject a proposed road expansion, or how a new visa law change might impact various social groups. This is both a powerful and dangerous application, because moral hazard suggests that government leaders could become reliant on AI-generated criticism and ignore real criticisms that are different – this should only be considered as a thought exercise during policy drafting.



3. Answer questions about complex regulation

Many government agencies must adhere to highly complicated regulations that can be difficult to understand, even for experts. Generative AI models can help by providing clear, concise answers to questions about these regulations. Instead of submitting questions to a legal team and waiting hours (or days) for a response, government employees can do better research themselves to understand the implications of existing laws in order to stay compliant. Citizens can enjoy this feature as well: generative AI can improve customer service by allowing citizens to ask the same questions to understand complex litigation, without the need to pay for expensive lawyers.



4. Government customer service

Generative AI will become a far more effective chatbot for answering citizen questions on government websites or mobile applications. Whereas a majority of users have limited patience for existing chatbots with their stiff language, seemingly low ability to perform tasks, and desire to speak to a real human, generative AI chatbots would be indistinguishable from the latter, and be far more adept at solving customer issues. This will lead to happier customers and a higher rate of chatbot success in solving queries.



5. Automated report generation

Generative AI can instantly write reports on various government activities and initiatives, such as progress reports, impact assessments, and policy briefs. Furthermore, it can take already-written longer reports, and instantly shorten them and repackage them (by changing verb tense, keywords, or emphasis) into 1 page press releases, PowerPoint presentations for different audiences, and so forth. This makes existing documents much more accessible to a wider range of public sector and citizen audiences, as well as frees up significant time for employees when this task is automated.



6. Natural language processing for better communication

Generative AI can help government agencies analyze and understand high volumes of public feedback or correspondence, and then respond in natural language, leading to better communication. In other words, it can take surveys, feedback forms, social media posts, and other forums for citizen information, and offer faster insights than manually reading or using more discrete data analytics. It can be automated to keep track of social media trends on the topic, even sifting through millions of posts. Any communication documents / images / videos can likewise be tailored to very specific audiences in their own level of jargon, including important business partners, children, international diplomats, and so forth.



7. Language translation

Generative AI can be used to instantly translate text and speech into different languages, making it easier for people to communicate with each other regardless of their native language. This includes live translation during conferences (of course, only for lower priority conferences where a 90% accuracy is permissible – for important meetings, always use a human translator to avoid mistakes), and can extend to helping citizens access government services, or re-writing laws in different languages. For multilingual societies, this service can be a gamechanger, because it fully democratizes the language divide and makes better integration more accessible.



8. Budget scenarios

Generative AI can help government agencies develop realistic and effective budget plans by creating different frameworks from past trends or future intuition. In fact, Generative models can be trained to generate synthetic financial data that resembles real-world data. This can be used to augment existing datasets and create more comprehensive models. It can also be used to forecast future financial outcomes based on historical data. This can allow public sector leaders to discuss different budget approaches for a particular project or better understand how much money is available in general, far more quickly than traditional methods. Moreover, Generative models can be used to optimize government revenue portfolios by simulating different scenarios and identifying the optimal mix of taxation and service levies to achieve a specific goal, such as maximizing returns or minimizing risk. It can also be used to detect fraudulent financial transactions by identifying patterns in data that indicate fraud, and can be used to assess financial risk by simulating different scenarios and analyzing the potential outcomes.



9. Streamlining government processes

AI can help automate routine government processes and applications, such as licensing, permits, and regulatory compliance, leading to faster and more efficient services. Generative AI is particularly well-suited for service delivery and support because it can generate human-like responses quickly and accurately and provide 24/7 customer service and support, reducing the need for human agents and increasing efficiency. Additionally, generative AI can quickly analyze and categorize large amounts of citizen data to identify insights from customer sentiment, behavior, preferences, patterns, and trends, which can improve user experiences and identify potential issues before they become significant problems. There are at least nine use cases for generative AI in customer service and support including customer self-service, chatbots and virtual assistants, sentiment analysis, predictive support, personalized support, email responses, language translation, knowledge management, interactive tutorials, coaching and training.



10. Education augmentation – personalized learning

Generative AI can be used to create personalized learning materials for students, based on their learning style, pace, and progress. For example, it could be used to generate customized quizzes, exercises, and study guides that adapt to the student's strengths and weaknesses.



11. Education augmentation – intelligent tutoring

Generative AI could be used to create virtual tutors that provide personalized feedback and support to students. These tutors could use natural language processing to understand and respond to students' questions and provide tailored guidance and recommendations.



12. Education augmentation – content creation

Generative AI could be used to create new educational content, such as interactive simulations, videos, or visualizations that help students understand complex concepts. For example, it could be used to generate 3D models or animations that illustrate scientific or historical phenomena, providing a more engaging and immersive learning experience for students. Generative AI cannot fully replace the content created by human teachers (yet), but it can enhance the classroom experience, allowing teachers to focus their attention where it can do the most good.




13. Healthcare augmentation – medical imaging

Generative AI can be used to create high-resolution images of organs or tissues that can be difficult to capture with traditional imaging techniques. It could also be used to enhance or reconstruct images that are blurry or low-quality, enabling physicians to make more accurate diagnoses.



14. Healthcare augmentation – drug discovery

Generative AI could be used to identify new drug candidates by generating and screening a large number of virtual compounds that have the potential to treat specific diseases. This could help to accelerate the drug discovery process and reduce the time and cost associated with traditional methods.





15. Healthcare augmentation – electronic health records

Generative AI could be used to analyze large datasets of electronic health records to identify patterns and predict outcomes. For example, it could be used to identify patients at risk of developing specific conditions, such as diabetes or heart disease, and provide targeted interventions to prevent or manage these conditions. It could also be used to identify gaps in care or areas where healthcare processes could be optimized.




16. More powerful accessibility for people with visual or auditory disabilities

Generative AI can provide information and support for people with disabilities, such as those who are visually or hearing impaired. It can augment existing text-to-speech software, for example, or offer more bespoke translation and interpretation services instead of existing robotic services today. In addition, generative AI could generate computer vision algorithms that enable people with visual impairments to navigate their surroundings more effectively, or be used to better test websites and applications for accessibility, by generating simulated experiences of users with different disabilities to see how easily they can navigate it.



17. Providing questions for interviews

Generative AI can sift through relevant knowledge databases to quickly provide highly relevant and specific interview questions, whether for job applicants, or journalists, or co-workers, or peer policymakers. On the flip side, the same tool can be used for citizens to augment their questions to government leaders in town hall or other public formats, allowing them to optimize the information they receive from government leaders in return.




18. Aiding government digitization efforts by augmenting and debugging code

Generative AI can write code as well as text – in any language, whether Java, C++, Python, or something else – which can help governments scale their digitization efforts much faster. Software engineers are already using generative AI to debug their code as well as accelerate project deliverables. Government employees trying to learn code or better integrate new software features can do the same. Government agencies can even use this capability to test and debug software received from 3rd party contractors, to ensure it performs all the requested features. This feature improves on education as well as speed of delivery.




19. Tax guidance

Generative AI can act as a user's accountant, and help them prepare tax documents much faster than a live human counterpart. This saves time especially for startups, small businesses, and individuals who do not normally have the budget for hiring accountants. Furthermore, Generative AI can be used to predict the tax implications of certain financial decisions, such as investments or business transactions; it can analyze large amounts of financial data to provide accurate tax projections and help individuals and businesses make more informed decisions. For accountants and other tax professionals themselves, Generative AI can be used to be more efficient, accurate, and improve the quality of their services. It can analyze and classify financial data, generate tax forms and schedules, and even suggest tax strategies based on past data and patterns.



20. Virtual assistants

Generative AI can serve as virtual assistants for government employees, providing assistance with tasks such as scheduling, document management, and task prioritization. Users can conversationally tell their AI assistant about upcoming events, documents they need to read from emails, and checklists that need to be accomplished, and let the algorithm optimize the daily schedule by summarizing said documents, send you reminders as necessary, and deflect unnecessary calls just as adeptly as a human counterpart.



21. Automatic Transcription and Minutes Generation

Generative AI can be leveraged in official meetings to transcribe conversations and automatically generate concise meeting minutes. This includes a summary of key points discussed, action items, and assigned responsibilities. Generative AI can even go the extra mile and automatically email these assignments to each respective employee. This not only saves considerable time in manual documentation but also ensures the accuracy of record-keeping, thus improving overall productivity and operational efficiency.



22. Efficient Procurement

Generative AI can generate requests for proposals (RFPs) based on project requirements. It can also streamline the bidding analysis process by generating comparison metrics based on bids, and even draft smarter contract texts. Plus, Generative AI can identify relevant business and risk signals, such as emerging customer trends, competitors' moves, market changes, or supply shortages, and then suggest appropriate adjustments in forecasts and plans in response. Combined, this reduces the time and complexity involved in the procurement process, leading to more efficient and transparent procurement practices.




23. Crisis Communication

In emergency situations, generative AI can quickly generate comprehensive crisis communication plans and create clear, concise public announcements. This can include generating evacuation plans, public advisories, or updates on the situation, ensuring effective and timely communication during crises. In addition, aforementioned government use cases like chatbots and virtual assistants can automatically be equipped to accurately answer anxious citizen questions about the emergency at hand. Where other public sector services might be disrupted, such as healthcare or education, generative AI can automatically update records.



24. Predictive Maintenance

In public infrastructure management, generative AI can predict maintenance needs and automate repair logistics with unprecedented levels of efficiency. For example, it can suggest specific tasks lists with personalized recommendations for infrastructure managers. It can also guide less experienced technicians with AI co-pilots in a safe metaverse environment. While technicians are physically working on repairs, they can use generative AI to turn photographs into detailed schematics for managers, or turn handmade sketches into realistic equivalents. Overall, this approach can help reduce breakdowns and improve resource allocation, thereby ensuring smoother operation of public infrastructure and better citizen satisfaction.



25. Smart Urban Planning

Generative AI can generate more sophisticated urban development plans considering various factors like future population growth, environmental conditions, traffic flow, and available resources. Or, AI algorithms can consider factors such as sunlight exposure, topography, and water flow in order to determine the most suitable locations for future parks and recreation spaces. Besides improving the built environment itself, Generative AI can also be used to study how a given city would react under different scenarios, such as revamping the public transit network, or making certain existing streets pedestrian-only. This can lead to more sustainable and resilient cities for citizens to enjoy.



How is generative AI governed and regulated?

Just as drug companies cannot sell people new medicines without first subjecting their products to rigorous safety checks, nor can biotech labs cannot release new viruses into the public sphere purely to impress shareholders, generative AI systems should not be entangled with the lives of billions of people at a pace faster than cultures can safely absorb them. Language is the operating system of human culture. From language emerges myth and law, gods and money, art and science, friendships and nations and computer code. Generative AI's mastery of this space through natural language processing means it can now manipulate the operating system of civilization. This is where regulators and government agencies come into play.

Social media was the first round of contact between AI and government regulation, and governments so far have lost. In social media, primitive AI was used not to create content but to curate user-generated content. The AI systems behind social media were sufficient to create a curtain of illusions that increased societal polarization and undermined mental health. Millions of people have confused these illusions with reality. Large language models (like ChatGPT) are the second round of contact with AI systems, and this time, government regulators must do better.

In the private sector, two approaches to the governance of generative AI models are currently emerging. In one camp, companies such as OpenAI are self-governing the space through limited release strategies, monitored use of their data models, and controlled access via API's for their commercial products. In the other camp, newer organizations, such as Stability AI, believe that these models should be openly released to democratize access and create the greatest possible impact on society and the economy. Stability AI open sourced the weights of its model – as a result, developers can essentially plug it into everything to create a host of novel visual effects with little or no controls placed on the diffusion process.

In the public sector, little or no regulation governs the rapidly evolving landscape of generative AI. When policymakers sit down to develop a serious legislative response to AI, the first fundamental question they face is whether to take a more “horizontal” or “vertical” approach. In a horizontal approach, regulators create one comprehensive regulation that covers the many impacts AI can have. In a vertical strategy, policymakers take a bespoke approach, creating different regulations to target different applications or types of AI.

In a recent letter to the White House, a US Congresswoman highlighted “grave concerns about the recent unsafe release of the Stable Diffusion model by Stability AI”, including generation of violent imagery.

Other issues surround intellectual property and copyright. The datasets behind generative AI models are generally scraped from the internet without seeking consent from living artists or work still under copyright. If such models have been trained on the styles of living artists without licensing that work, there are copyright implications that are as-of-yet not legally tested. For example, in 2021, an Australian court ruled in favor of an AI system that could be named as the inventor on a patent application. However, this was later overturned by the Australian Federal Court.

The problem with copyright is also visible in the field of autocompleted code. Microsoft's GitHub Copilot is already involved in a class action lawsuit alleging the system has been built on "software piracy on an unprecedented scale." Copilot has been trained on public code repositories scraped from the web, which in many cases, are published with licenses that require crediting creators when reusing their code.

In addition, there needs to be enough time for institutions to figure out what to do. Smart regulation takes time, especially in the face of unprecedented leaps in technology. Although current-generation AI tools are not scary, future applications might be less virtuous and the public sector must be ready. Fortunately, the improvement of AI programs depends on how fast people use them, so as more people contribute, the easier it should be to understand.

The emerging discussion between centralized and controlled adoption with firm ethical boundaries on the one hand versus faster innovation and decentralized distribution on the other will be important for the generative AI community in the coming years. This is a task not only reserved for private companies, but which is equally important for civil society and for policymakers to weigh in on. This includes disruption of labor markets, legitimacy of scraped data, licensing, copyright and potential for biased or otherwise harmful content, misinformation, and so on. Only when solid checks and balances are in place can more thoughtful and beneficial expansion of generative AI technologies and products be achieved.



Since there are no best practices in AI regulation yet, it is fortunate that two major experiments are currently playing out in the European Union (EU) and China. The EU is racing to pass its draft Artificial Intelligence Act, a sweeping piece of legislation intended to govern nearly all uses of AI. Meanwhile, China is rolling out a series of regulations targeting specific types of algorithms and AI capabilities. For the host of countries starting their own AI governance initiatives, learning from the successes and failures of these two initial efforts to govern AI will be crucial.

For the EU, the AI Act groups AI applications into four risk categories, each of which is governed by a predefined set of regulatory tools. Applications deemed to pose an "unacceptable risk" (such as social scoring and certain types of biometrics) are banned. "High risk" applications that pose a threat to safety or fundamental rights (think law enforcement or hiring procedures) are subject to certain pre- and post-market requirements. Applications seen as "limited risk" (emotion detection and chatbots, for instance) face only transparency requirements. The majority of AI uses are classified as "minimal risk" and subject only to voluntary measures. Once in force, years of work by courts, national regulators, and the technical standards bodies will clarify precisely how the AI Act will apply in different contexts.

This framework strikes a balance between the dual imperatives of providing predictability and keeping pace with AI developments. Its risk-based approach allows regulators to slot new application areas into existing risk categories as AI's uses evolve, providing a balance between flexibility and regulatory certainty. Meanwhile, the AI Act's definition of relatively flexible essential requirements also alleviates the key precision challenge posed by purely horizontal frameworks, allowing compliance strategies to be flexible across sectors and as technology evolves. But the EU's broadly horizontal approach faces several risks that other countries should watch closely. Individual regulators tasked with enforcing requirements might differ in their interpretations or capacity to regulate, undermining the key capacity and harmonization benefits of a horizontal approach.

For China, the country has rolled out some of the world's first nationally binding regulations targeting algorithms and AI. It has taken a fundamentally vertical approach: picking specific algorithm applications and writing regulations that address their deployment in certain contexts. The first two regulations in this camp targeted recommendation algorithms and deep synthesis technology, also known as generative AI. The recommendation algorithm regulation focused on its use in disseminating information, as well as setting prices and dispatching workers. It required that algorithm providers "vigorously disseminate positive energy" and avoid price discrimination or overly demanding workloads for delivery drivers. The second regulation targeted deep synthesis algorithms that use training data to generate new content, such as deepfakes. The regulation again focused on concerns around "harmful information," but it also required providers to obtain consent from individuals if their images or voices are manipulated by the technology.



Yet China's regulations do contain a horizontal element: they create certain regulatory tools that can be applied across several different vertical regulations. A prime example of this is the algorithm filing system or registry. The registry was created by the recommendation algorithm regulation and reaffirmed by the deep synthesis regulation, both of which require developers to register their algorithms. It acts as a central database for Chinese officials to gather information on algorithms, such as their sources for training data and potential security risks. As such, the registry also serves as a vehicle for regulators to learn more about how AI is being built and deployed—a key goal of sectoral regulators around the globe. Looking ahead, it could be adapted to require different kinds of information depending on the algorithm's application, or it could simply provide a uniform baseline of information required on all algorithms subject to regulation.


China's approach allows it to more precisely target regulatory requirements to specific technical capabilities. In many cases, this approach might risk rules falling behind quickly evolving technology. In China's AI regulation, however, some requirements are defined so vaguely that they effectively function to shift power from technology companies to government regulators, which can wield their newfound regulatory power to force any changes on companies that they wish.

In the end, both the EU's and China's approaches both have positives and negatives. While the EU AI Act creates four risk tiers and lays out broad requirements for each one, the heavy lifting of articulating specific compliance thresholds will be done by Europe's main standardization bodies. And while China has issued vertical regulations targeted at recommendation engines and generative AI, those regulations lean on horizontal tools like the algorithm registry in order to bring coherence across them. Business groups argue that the EU's broad approach will stifle innovation, and analysts assert that China's targeted regulations will be used to tighten information controls. But by taking a step back to look at their fundamental approaches to regulating AI, policymakers can draw lessons for countries and institutions around the world. By adapting lessons learned from these initial approaches to their unique contexts, other countries can take meaningful steps toward reducing the harms produced by AI systems.

International Case Studies applying Generative AI in the government work

Estonia, a small Baltic nation known for its progressive approach to technology and governance, sought to further simplify its already efficient public services. While their digital infrastructure was robust, they aimed to reduce the burden on human service agents who handled routine inquiries from citizens.

To address this challenge, Estonia implemented generative AI chatbots in various government departments and agencies. These AI-driven chatbots were trained to understand and respond to citizen inquiries, provide information on government programs and services, and assist with basic administrative tasks.



1. Transforming Government Services with Generative AI in Estonia

The chatbots were equipped with natural language processing capabilities, allowing them to comprehend and respond to queries in Estonian, Russian, and English. They also had 24/7 availability where citizens could access these virtual assistants round the clock, reducing waiting times and improving accessibility.

The chatbots were seamlessly integrated with government databases, enabling them to fetch real-time information and perform tasks like verifying personal records or scheduling appointments.

This led to many benefits. The AI chatbots significantly reduced the workload of government employees by handling routine inquiries, allowing human agents to focus on more complex tasks. Citizens benefited from instant responses and access to government information at any time, enhancing overall satisfaction with government services. Moreover, Estonia's investment in generative AI reduced administrative costs and improved resource allocation within government agencies.

Furthermore, these chatbots are consistently accessible, highly intelligent, and proficient in understanding informal language and regional vernacular. Moreover, they stand out due to its swift response times and the inclination to provide additional information that goes beyond the initial inquiry. For example, if a user requests information on starting a new business, the chatbot actively reminds them about impending events, such as the expiration of their driver's license in four months. As a result, the chatbot serves as an excellent way to save citizens time and effort while ensuring they remain informed about critical matters.

In terms of chatbot development, the state involved the private sector in AI development to achieve the best results.

The success of the AI chatbots in one department led to their expansion across various government agencies, further optimizing service delivery.

In addition to all the above tangible benefits, the system offered Data-Driven Insights. The AI system generated valuable data on citizen inquiries and preferences, informing future policy decisions and service improvements.



2. AI-Powered Civic Engagement in Singapore

Singapore, known for its modernity and digital infrastructure, faced the challenge of ensuring that its citizens actively participated in the democratic process and had easy access to government services. Despite its many successes, there was a need to further enhance citizen engagement and create a more inclusive and responsive government.

The government of Singapore decided to leverage AI to address this challenge. They embarked on an ambitious project to create AI-powered virtual assistants known as “Civic Companions.” These companions were designed to serve as intelligent, round-the-clock guides for citizens, offering information, answering inquiries, and facilitating civic interactions.

These AI companions were versatile, capable of assisting citizens with a wide range of services, from obtaining information on public transportation schedules to providing guidance on government policies and initiatives.

The companions were designed with a user-centric approach, ensuring that they could cater to citizens with diverse needs and preferences. They could interact via text, voice, or even through mobile apps.



To provide accurate and real-time information, these companions were seamlessly integrated with various government databases and systems, allowing them to access a wealth of up-to-date data.

Benefits were delivered fairly quickly as citizens in Singapore found it easier than ever to engage with government services and obtain information on a wide array of topics, leading to increased civic participation. The government also saw significant reductions in the time and resources spent on addressing routine inquiries, allowing them to allocate these resources to more complex and impactful tasks.

The system as inclusive as well, as the AI companions were accessible to citizens across different age groups and backgrounds, making government services more inclusive. Moreover, the wealth of data generated by these AI companions provided valuable insights into citizen preferences and needs, helping policymakers make informed decisions.

Singapore’s innovative use of AI for civic engagement is not only a success story but also a glimpse into the future of responsive and inclusive governance. As AI technology continues to evolve, we can expect even more personalized and efficient government services that cater to the diverse needs of citizens.



3. AI-Driven Healthcare Transformation in the United Kingdom

The UK's National Health Service (NHS) faces the challenges of an aging population, rising healthcare costs, and increasing demand for services. There was a pressing need to enhance the efficiency and effectiveness of healthcare delivery while maintaining high-quality patient care.

The UK government and the NHS embarked on an ambitious initiative to harness AI to address these challenges. They launched a program known as "AI for Healthcare," which aimed to leverage AI technologies to improve diagnosis, treatment, and patient care. Artificial intelligence-driven machinery is the central element in 22 initiatives conducted in various UK universities and NHS trusts. These projects encompass research related to semi-automated surgical robotics for tumour excision and health predictions based on pre-existing medical conditions.

AI algorithms were deployed to analyze medical imaging, such as X-rays and MRI scans, to detect and diagnose diseases with greater accuracy and speed than human radiologists. Moreover, these AI models were developed to predict disease outbreaks, patient admissions, and resource requirements, allowing for better resource allocation and planning.

In addition, AI-powered chatbots and virtual assistants were introduced to provide patients with immediate medical information and support, reducing the burden on healthcare staff.

In the back office, AI-driven Electronic Health Record (EHR) systems were implemented to streamline patient data management, enhance data security, and improve clinical decision support.

These AI-enhanced diagnostic tools significantly reduced the time it takes to diagnose medical conditions, improving patient outcomes. Predictive analytics allowed for more efficient resource allocation, reducing healthcare costs and optimizing hospital operations. In parallel to these benefits, virtual health assistants improved patient engagement by providing immediate responses to inquiries and offering personalized health recommendations.

Last, but not least, AI-driven EHR systems improved data accuracy, enabling healthcare providers to make more informed clinical decisions.

The UK's AI-driven healthcare transformation is an ongoing process with promising prospects for the future. As AI technology continues to evolve, we can expect even more advanced applications in healthcare, such as AI-assisted surgery, drug discovery, and personalized treatment plans.





مركز محمد بن راشد
للابتكار الحكومي

MOHAMMED BIN RASHID CENTRE
FOR GOVERNMENT INNOVATION

The Mohammed Bin Rashid Centre for Government Innovation was established to cultivate a culture of innovation within the government sector through the development of an integrated framework that aims to make innovation one of the key pillars of the UAE Government in alignment with the goals of vision of “We the UAE 2031”, focusing on unlocking the individuals’ full potential and enhancing the UAE’s position as a global hub for innovation. MBRCGI also plays a pivotal role in the achievement of the UAE’s Centennial 2071 goals by building a robust innovation ecosystem and continuously developing the governmental work by deploying innovation intelligence to multiply the results of governmental effectiveness to enhance the country’s competitiveness by making the UAE Government one of the most innovative in the world and promoting a digital, knowledge-based economy.

