Deep Science Publishing
https://doi.org/10.70593/978-81-988918-1-5_10

DeepScience
Open Access Books

Chapter 10: Driving public sector innovation with generative artificial intelligence for services, compliance, and efficiency

10.1. Introduction

The digital transformation of the public sector was accelerated by the COVID-19 pandemic, which acted as a catalyst for change. Governments rapidly shifted to online service delivery as the pandemic left much of the world in lockdown, and this momentum towards digitalization is set to continue. Building on the lessons learned during the pandemic to further drive innovative change, the continued digital transformation of government can, and must, also include advances in artificial intelligence technologies (Carrasco et al., 2023; Androniceanu, 2024; Johnson & Morales, 2024).

At the close of 2022, artificial intelligence, in the form of generative AI technologies, has captured the attention of audiences everywhere. Software products leveraging artificial intelligence to generate content in various forms are suddenly all around us. The speed of adoption of generative AI technologies has stunned pundits, rivals, and AI experts alike. Generative AI tools – while still in their infancy – allow users to create content in formats both new and familiar, making use of either text prompts or some blending of prompts with user-provided images. Audiences have just begun to contemplate the positive attributes of generative AI tools or indeed the socio-technical implications, good and bad, of technology that, for better or worse, creates content that looks like it was authored by a human. And yet, at the close of 2022, the question sitting most prominently in forethought is whether generative AI tools can – and will - act as platforms on which society can build revolutionary new applications for creativity and communication, pushing the limits of innovation and new product development further than it has ever been in the past (Kulal et al., 2024; Yoon, 2024; Smith & Kumar, 2023).

10.2. Understanding Generative AI

This chapter aims to introduce the topic of Generative AI (GenAI) in detail and also highlight its relevance to innovation in public sector work. Firstly, we begin with an overview of GenAI technology itself and describe the main characteristics and features of GenAI systems. Secondly and relatedly, we explain the nature and scope of GenAI, what makes it different from prior AI technologies, and what enables GenAI technologies to be built today. Finally, while we do touch upon the necessary ethical and application considerations surrounding their use, the main focus of this chapter is to be focused on GenAI technology itself. The intention here is to familiarize the reader with the technology as they may not have had the opportunity to work with GenAI tools in detail or be cognizant of its impressive capabilities.

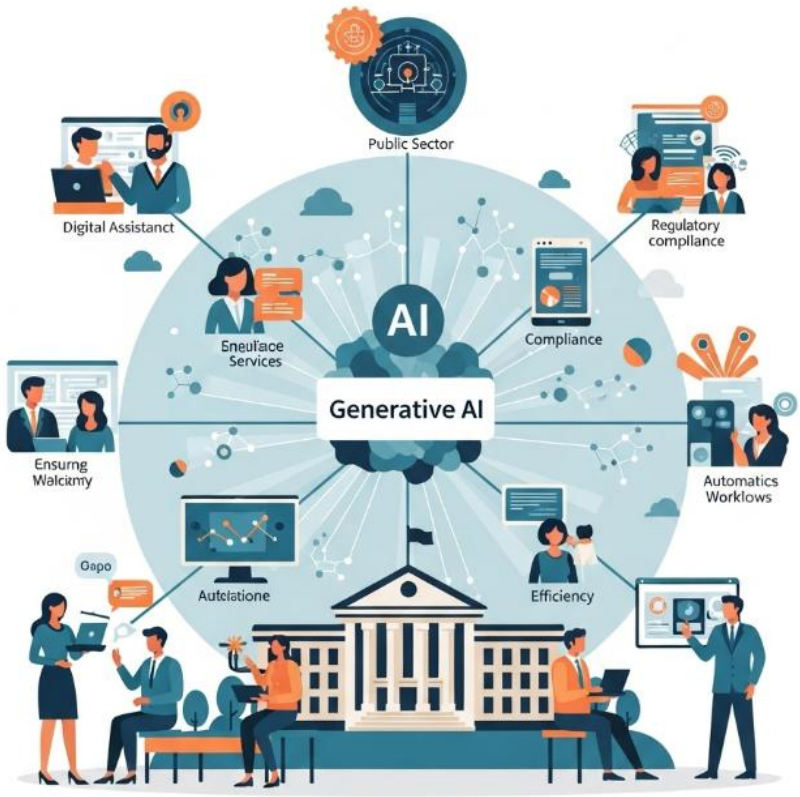

Fig 9.1: AI in Public Sector

In recent months, there has been a rapid and unprecedented rise in the use of various Generative AI (GenAI) applications in both public and private spheres. Are these technologies revolutionary? Can they create fundamentally new ideas, works and solutions? Are they reliable when they do share what appears to be confident output? These questions are not easily answered and opinions can be varied and at times, extreme

- both in terms of being overly optimistic as well as overly pessimistic. One way to become informed and knowledgeable in order to make those answers more concrete is to familiarize ourselves with GenAI technology itself. Despite all the concerns, questions, and reservations that can perhaps be validly raised about the use of GenAI tools, it is clear that they have opened up entirely new possibilities for working and collaborating with computers to generate various forms of content.

10.2.1. Definition and Scope

Generative AI has emerged as a groundbreaking class of artificial intelligence technologies capable of creating novel content such as text, code, images, video, audio, simulations, and 3D models. Generative AI systems consist of components that comprehend initial human inputs, such as broad prompts, detailed instructions, requirement specifications, sample inputs, or high-level plans. These systems produce entirely new content or build upon or modify existing content. As we delve deeper into understanding the definitions and vicinity of generative AI, it is crucial to decipher this latest addition to the expansive list of existing AI capabilities. Why is it called generative AI? How is it different from existing AI capabilities? What are the boundaries and frontiers of this concept? What is unique about generative AI? What is not unique about it? What are the implications of its capabilities? Is its utility limited to its obvious applications, or can generative AI also be used for applications at the periphery or outside the obvious definitions?

The "generative" moniker indicates the technology's distinctive ability to generate new content. Previous AI models have typically focused on prediction rather than generation. Most prior AI technologies rely on supervised learning models whose outputs are discrete decisions or predictions over a fixed and finite domain. For example, facial recognition AI predicts from a small domain of distinguishable IDs based on the input of the geometrical measurements or other distinctive features of the physical face or other facial biomarkers. Other AI technologies — such as voice recognition, speech synthesis, image recognition, and language translation — are also predictive systems at their core.

10.2.2. Historical Context

Artificial intelligence (AI) historically has been an enigma captured by its ambitions. Supposedly designed to assist humanity, it appears more as a digital genie able to conjure wondrous artifacts. AI came close to a Turing test of any sufficiently advanced technology being indistinguishable from magic. Its development has swayed wildly between naïve optimism and sanguine despair. Talented individuals, even society as a

whole, have dared to believe that the smartest machine could finally be made. The relentless pace of increasingly sophisticated hardware and data storage and computation has frequently overwhelmed uncertainty and skepticism.

The first stirrings of AI were golden years. Automated chess, vision, and verbal and almost corporeal confrontation with naive human rivals, from checkers through symbolic translation, shooting galleries, and bodiless voices, enchanted the beamish inventors, as well as hordes of onlookers. Then reality kicked in. The discovery of the importance of deep learning algorithms to AI garnered considerable attention over the past few years. Compared to the years when the field was largely experimental, the past year has seen an unprecedented number of deep learning papers accepted at major AI conferences. This disruptive technology has led to the creation of entirely new companies and products, plus significant internal investments at tech giants. AI never exited its uncanny valley.

10.2.3. Current Trends in AI

The field of artificial intelligence has passed several evolutionary milestones. Currently, a trend observable in the research community is centered on the creation of novel methods, tools and architecture to build and scale AI at lower costs, not necessarily for a better automation of narrow tasks, but for the creation of Generalist AI. From a social direction, the focus is on the building of those Generalist AI without blinders. Generalist AI which closely follows a worker path, from narrow to generalist, in a seemingly impeccable Human trajectory. The new Interphases built by AI image generators are seen as conduits to generalist AI which closely mimic humans in artistic and executive tasks, and are not empty black boxes. Generalist AI that are endowed with common sense, decency and insightfulness. For sociology, they would be Resident AI, because they live and share the space bias, culture, language, families, and social circles of the people they are helping with skills they are in a stage of learning.

The AI ecosystems encompassing Hardware, Data, Compute, the People/Community and Business/Organizational Framework are evolving and about to converge. This is happening due to several factors. Developments in superconductors, quantum computing, and photonic-like programmable accelerators may bring the hardware costs and energy consumption down significantly. Convergence of GPU Cloud Computing due to a drop in and to subsidy of costs, and the vertical integration of API based key services: Storage, and Learning/Optimization - Deployment – Inference, with shared costs, will ease the experimentation and interface transitions to vertical specialized AI Agents.

10.3. Public Sector Challenges

The public sector is different from the private sector. Definitely, leadership in the public sector can innovate and drive emergent solutions for the public good, but this is undeniably more challenging, with the following reasons.

All levels of government, despite their advanced economies and large-scale spending budgets, are continuously unable to adequately deliver services to their citizens and other stakeholders. Many critical services such as public transport systems, digital infrastructure, health and social services, and financial aid face bottlenecks, affecting both users and the organizations delivering them. Recently, the pandemic has again tested the will and capacity of public sector organizations in many ways, dealing with vaccine distribution, contracts with private providers, keeping citizens informed, and handling data privacy and security concerns. In a continually evolving world, but one that is reliant on technology more than ever, governments cannot deliver their objectives since they operate by relying on outdated means and processes, which are neither integrated nor fully digitalized.

In addition, too many countries still have poorly integrated channels to engage with their citizens. Some services are difficult to access or require lots of paperwork, meaning that many stakeholders do not use them. Citizens are judged as irresponsible for not asking for help in case of illness, or for not reporting a missing child. Why so many citizens do not claim the state pension they are entitled to? The lack of citizen engagement by the public sector is a paradox caused by the gaps in both service delivery and communication. The gain from switching to modern channels is mutual: it is not just about delivering services over the Internet, reducing costs for governments, but about better connections with the citizens.

Governments are frustrated both by the inadequacies and the complications of too many laws, regulations, and policies that they are required to comply with. Various state authorities impose many different compliance processes, and in addition, the compliance requirements are not static but are constantly upgraded. As a result, companies and citizens must deal with a burden of complexity when dealing with different governments. In the public sector, innovation occurs on the necessity to comply both with laws and with the processes set up by the government for non-compliance.

10.3.1. Service Delivery Issues

Services provided by governmental and other public sector organizations are often criticized as being difficult to navigate, unresponsive, static, and so far lacking in quality that citizens might need or want to rely on alternative private or non-profit providers. A first glimpse into the challenges related to service delivery is offered by citizen

satisfaction metrics. Citizen satisfaction metrics are used as an aggregate measure and have also been shown to correlate with other measures of good governance such as public sector transparency, regulatory quality, and the rule of law.

At the same time, generally the services being delivered are also experiencing increasing pressure to adapt to technological progress. Citizens are used to personalization and customization, a higher level of interaction, and the use of advanced technologies in other sectors, especially technology-oriented private sector services. These generic developments are also observable in specific areas of public sector service delivery. In the field of public sector human resources management, there is a growing trend to actively engage with the workforce and proactively apply branding techniques – partly enabled by big data analytics in order to address the general challenge of a labor force that is increasingly reluctant to apply for jobs in the public sector, and of talent retention in particular.

Recent advances in generative AI open up exciting new opportunities to leverage technology for the greater public good, to rethink service delivery, and to turn a number of aspirations of the past into reality. To deliver on this promise requires however new ways of prioritizing investments in AI solutions, and new radical approaches to AI design and implementation. Offering both excitement and caution, generative AI has the potential to reshape service delivery – provided that some caveats are addressed and new approaches to service delivery and service design are envisaged.

10.3.2. Compliance and Regulatory Hurdles

The legality and ethical reasoning for the use of generative AI models for public services is highly conceivable yet complicated. From determining whether GPT models' understanding is sufficient to be used as advisors in the government to deciphering laws siding with AI copyright creators, public branches such as regulatory authorities are at a risk of indecision. Laws built on human interactions have coalesced for hundreds of years, creating a thick wall against the differences that AI brings. Experts have pointed out a lack of federal AI policy would create unwanted hurdles for state and city governments, making insurance untouchable and liable.

The unique characteristics of generative AI bring a series of legal considerations in a variety of areas. Regulation is slow in general but even slower at the local level, and this delay causes paralyses within these organizations. Unless these questions are addressed, the risks of using generative AI in the public sector are enormous. The expectation on the limits and implausibilities of generative AI should be taken into center stage immediately. No doubt our society and the public sector, in particular, will create an effort to develop safe explanation frameworks.

10.3.3. Efficiency Gaps

Efficiency gaps have become a persistent challenge for governments around the world who grapple with the dual burden of increasing citizen demand for personalized, humane, high-quality services while dealing with fiscal pressure to eliminate budget deficits. A report argues that over 75 percent of government budgets consist of spending on people, investments in labor-intensive sectors, therefore, greater investments in advanced technologies have the potential to increase productivity by enabling government employees to focus on higher value-added tasks. For example, in law enforcement agencies, advanced analytics can enable police officers to focus on proactive policing rather than reactive calls for service, and predictive algorithms that determine which officers should take on what roles during a given shift can lower the lag time in response to emergency calls and improve efficiency.

Do AI and automation create speed and efficiency in the functioning of government agencies? The all-too-common answer to this question is mixed; improving efficiency has been difficult due to conflicting objectives and challenges in implementation. Efficiency improvements are complicated by old service delivery methods, outdated government sectoral processes, poor working relationships stemming from poor trust levels between agencies and citizens, lack of interagency coordination, and workforce fatigue. Cutting corners may produce short-term benefits, but long-term efficiencies can be gained only by enhancing the service delivery experience from the customer's perspective. Deep service innovation calls for true collaboration and risk-sharing between agencies and their private government contractors or providers in order to maximize the efficiency of the service delivery experience.

10.4. Generative AI Applications in Public Services

The constantly growing number of citizens digital interactions with public body services makes the latter ideally suited for being enhanced with Generative AI technology. Public service organizations are increasingly adopting it to their service provisioning activities. A survey conducted from late 2022 to early 2023 shows that 45% of responding agencies already use Generative AI in their operations and 70% expect that within some years this technology will be embedded in their core functions.

Going into detail, we can find specific public service activities that will be impacted by Generative AI technologies. Generative AI will be able to enhance citizen engagement processes, by applying Natural Language Generation to automatically compose messages to better describe governmental services to citizens, court case descriptions, and rapid responses to Frequently Asked Questions. Generative AI will also automate citizen engagement tasks, directly communicating to citizens through chatbots, emails,

and text messages. It can also provide multichannel options allowing citizens to choose which interface or interfaces they want to use. Moreover, it could facilitate contact between citizens and officials, tapping Generative AI's overall generative capabilities in producing drafts for emails and reports.

Public organizations will also improve their decision-making processes. Government organizations dealing with applications for grants, maintenance of criminal records, and credit will apply Generative AI to aid people writing applications and to help analysts compose recommendations for application evaluation. Generative AI will also use information extracted from Public Sector Knowledge Graphs to help generate suggestions and answer questions in crisis management, considering geographic and demographic data.

10.4.1. Enhancing Citizen Engagement

Organizational communication with citizens is often criticized for being technical, inaccessible, rigid, or boring. Various public sector organizations are experimenting with generative AI applications to augment their communication with citizens, using chatterbots and work assistants to support employees in responding to citizen inquiries; and speech-writing and speech-delivery tools to generate engaging content for video messages delivered to citizens. Organizations are also investigating the potential of generative AI applications to deepen citizen consultation regarding policy development, as traditional methods for citizen consultation are perceived as ineffective, uninclusive, or exclusive of marginalized groups in society. Generative AIs can be applied to conduct simulations of citizens with specific identities and pain points or to help shape policy proposals that encourage exchanges by citizens with divergent needs. By analyzing responses, boards could better gauge citizen values and respond accordingly.

In terms of specific applications, a generative AI chatbot allows customers to ask detailed questions about certain topics in a conversational manner, while a hybrid model maintains continuity in topics addressing nuances such as sarcasm, metaphors, and humor while offering more precise, human-like responses to answer customer queries in real-time on website platforms. There has been significant investment in the public application of generative AI to reshape Q&A by enabling users to ask questions in a natural and conversational tone and receive detailed answers.

10.4.2. Automating Administrative Tasks

Generative AI can automate mundane yet important administrative tasks for civil servants and support professionals in the delivery of their jobs. Routine tasks

conveniently take a lot of time and resources, fueled by bureaucracy-intensive work across creative and productive sectors in the public sector business model. AI can help with research and planning, quality assurance, verification, record-keeping, scheduling, prioritization, and monitoring. Generative AI can also be deployed to assist and improve the job of translators and court and interpreter transcribers, streamlining the process, speeding it up or improving accuracy.

New AI tools have also become available for effective website management. Bots powered by generative AI can personally advise clients about service availability and eligibility and use cases directly from government websites or through social media messaging channels. This also means lowering the cost of providing these services. Language generation tools can summarize website and complex legal information into plain language or common languages, bridging the language divide and increasing access to the rule of law. Moreover, virtual assistants are also available, while others can build specific scripts for your needs.

Job automation was traditionally identified as a risk by some anthropologists, as they characterized services that have a social significance more than a true functional one. Civil servants may offer insurance for some of their co-citizens unless jobs are performed by redundant ephemeral bots half-jokingly known as “the insert name”. Although this critique has undeniable merits, it is also clear that these jobs might offer little personal satisfaction and it is to be expected that AI would favor fewer and smarter interactions between civil servants and society in several countries according to society’s preference.

10.4.3. Improving Decision-Making Processes

One of the most relevant applications where generative AI may contribute value in public services is to improve decision-making processes. Government agencies spend extraordinarily large amounts of resources explaining the why and how of their decisions to courts, citizens and other affected parties; so the quality of these decisions is paramount. Making good decisions requires experience and technical know-how, particularly when it comes to decide on cases at the margins where little precedent exists, such as in granting or denying a special condition waiver. The problem is exacerbated in agencies in charge of making thousands of similar decisions every year and to justify each of them exhaustively, such as adjudicating immigration petitions for employment-based temporary visas.

On the one hand, large language models are capable of reproducing complex decision trees based on internal datasets, legal language and particular prompts provided by trained officials from these agencies. After some initial guidance by an official who is a subject-matter expert, such models may decrease human intervention to a minimum

while adapting the decisions to a variety of edge cases that can also be fed to the AI, provided that the output is not used as-is but is vetted by an official in charge of these matters. Similar models are also capable of translating these decisions into human language a layperson can understand, with little human intervention if not zero.

10.5. Generative AI for Compliance

To be considered sound, any framework for generative AI-enabled services must incorporate insights from regulation, inspection, audit, and compliance functions. As much as generative AI can enhance service delivery, it can also create risk and negative outcomes. Given the black box nature of generative AI, these functions will need to evolve significantly in order to be able to determine whether both generative AI and the services enhanced by generative AI are providing quality, reliability, trustworthiness, and so on. These frameworks are the subject of relatively little research but are clearly of increasing importance, which makes these tasks ideal use cases for the flexible, iterative functionality inherent in generative AI.

Regulations often create significant administrative burden on both public and private sector organizations. For example, organizations are often required to have systems in place to monitor risk areas and be able to demonstrate this to an auditor. Generative AI can significantly reduce the burden and cost of these endeavors. However, as with any tool, its effectiveness will depend on the skill of its users and the governance framework within which it is employed. Generative AI can contribute to regulatory compliance in at least the following ways.

Much of compliance work consists of conducting searches over knowledge bases. Auditors then evaluate whether risks have been properly considered and whether sufficient action has been taken to address them. For example, organizations in many regions are required to record and assess their exposure to different risks, such as bribery and corruption, money-laundering, and fraud. However, such processes are often checklist exercises, in which limited thought is applied to the actual wording within a document. Generative AI-based tools can speed up the laborious process of gathering and coding data, effectively allowing companies to code significantly larger data sets or document flows to inform assessments and to do so more frequently, enabling faster remediation of any issues that come to light. This can ultimately create greater transparency and benefit for companies if they are facing a negative issue during a compliance checklist or if it is known that they have been evaluated transparently in more detail versus competitors.

10.5.1. Monitoring Regulatory Compliance

Compliance management is a complex task; rules come from multiple jurisdictions, and keep on changing, sometimes suddenly and without warning. Rules have to be implemented throughout the organization, and then monitored. Keeping track of internal compliance in small, technology-savvy organizations can be difficult; it is a Herculean effort for multinationals. The cost of non-compliance can be substantial, with corporations facing regulatory and enforcement actions, and incurring loss of business, reputational damage, litigation, and settlements. Such penalties can be even more severe for nonprofits, which operate under intense public scrutiny.

Generative AI can help in many areas of compliance. It can track regulatory authority announcements, issue statements, and other sources of applicable rules, simplifying the task of compliance officers who have to prepare this information for the whole organization. It can extract obligations from such rules, and define in simple terms what has to be done, or avoided, in order to comply. It can help organizations figure out how to implement specific provisions, support the employees responsible for doing so, then generate the policies that define those processes. Tracking compliance can be unattractive, repetitive work that runs the risk of human error; generative AI can carry out a lot of it for everyone involved. From emails to actions, or hiring processes to due diligence on third parties, generative AI can monitor activities, aggregate the relevant data, and report them – or flag them for someone to check. Finally, generative AI can provide guidance on how to take the right actions for each activity, offer reminders, and explain the underlying rationale.

10.5.2. Risk Assessment and Mitigation

Unsupervised learning methods, such as rule mining and anomaly detection, can help organizations quickly identify potentially risky behavior in user data by detecting deviations from the normal activities that users would typically run. In many use cases, organizations have access to multiple channels of data to understand how users are using the tools provided to them, and anomalous behavior may relate to several of the data channels. As an example, an organization may have channel data available to them, including profile attributes, message synthesis, summary statistics, sharing behavior, message content, sentiment analysis, user connections, and templating. For organizations that have a reasonable expectation of what normal levels of each metric might be in the data from a particular generative AI system, and what misuse of the system might entail, anomaly detection for these multichannel statistics can provide risk assessment and mitigation at scale in near-real-time.

Organizations also may have purpose-built LLMs that are designed for a specific set of topics, for a specific industry, or for a small number of regulated functions. In this case, these organization-specific LLMs do not have the same risks of content generation that public models and that general-purpose models have. These organization-specific models are also being run on data that is subject to internal security. Thus, the risks those organizations should face are limited to those that would face all other data systems, and security measures could be tuned to those specific risks fairly easily. Public access to organization-specific models should be limited, as should clearing in the models prior to access. It would be useful to explicitly curate prompts that trigger the generation of sensitive data and monitor the use of those prompts to assess risk. Historically, organizations with similar needs for risk mitigation were encouraged to assess risk based on the dynamic monitor score of particular commands.

Fig 10.2: AI Public Sector Cost Optimization

10.5.3. Reporting and Transparency

The objective of reporting is usually to collect and summarize data from projects, products, and lines of business in order to inform, influence, and enable action for various stakeholders. For example, both internal stakeholders such as departmental auditors, program initiators, and executives, and external stakeholders such as budget scrutinizers and legislative lobbyists want to understand how IT funds are allocated, how held funds are spent, how IT resources are used, and what the impact of the funds is on the overall business of the government. Data anomalies across various datasets can suggest misappropriation of funds, false reporting of project outcomes, or hidden engagements in politically detrimental activities like lobbying. As with compliance monitoring, the challenge here is that collecting and summarizing such information from vast, heterogeneous datasets is a tedious and expensive workflow that has little impact on the bulk of the projects which are compliant. LLM models can make this process much more efficient.

Government agencies are also held accountable for their actions in measurable ways, such as impact on the environment, impacts on toll payer tax burdens, or failure to deliver goods and services in a timely manner. Missing, inaccurate, or suddenly changed data reported for these purposes can trigger suspicion of potentially unethical behavior. Further, the types of data and the degree of granularity of reported data are often supplemented by legislation or executive orders. LLMs and generative tools can assist agencies in reviewing, validating, and modifying proposed transparency reports in order to efficiently prepare a required document or find existing reports that match specific criteria.

10.6. Driving Efficiency with AI

In 2018, the UK Government published an AI Sector Deal that encouraged sectors "in desperate need of productivity improvements" to collaborate with the UK's AI research community to develop relevant solutions, which are expected to have a transformative impact in their field. The report cited a number of examples from the public sector – particularly Ministry of Defence and the National Health Service– including AI projects for military surveillance, body-tracking and motion prediction for soldiers, automated breast cancer detection, drug discovery, radiotherapy treatment planning, predictive modelling of patient risk, virtual home health assistant, AI-augmented clinical surveillance assistant, AI-powered chatbot, and video risk screening for mental illness. In many cases, local governments are at the forefront of efforts and initiatives aimed at promoting and implementing generative AI in public services. With relatively limited resources available, cities are increasingly keen to harness AI's potential in order to deliver more efficient and prompt services to their citizens. City governments are

rethinking the nature of how they approach service delivery to residents, transfer services to the digital realm, build confidence in digital services, and actively seek out innovation and new technologies to better serve constituents. It is particularly important that automated solutions are well designed with human oversight and involvement because citizens will trust services delivered by their government agencies only if the core values of accountability, justice, privacy, and transparency are observed in the process. At the same time, however, as initiatives increase, research into the lessons learned from failures remains less comprehensive. Although we may not possess detail about the relative amount of failed AI initiatives, we have enough research data to be concerned about the high stakes of failure. The increasing sophistication of purpose-built AI systems are not enough to compensate for the non-AI-specific organizational inertia, private-sector-style ruthless, profit-obsessed prioritization, and a general disregard for hard-to-measure-with-numbers crux factors associated with the public good that have long plagued public sector innovation efforts.

10.7. Global Examples of AI Implementation

Countries all over the globe have kitted out with AI tools to tackle challenges in national security, public health, and disaster recovery. The United States of America started a project to help its intelligence community address the volume of transactions on digital platforms in monitoring terror and insider threats using natural language processing, deep learning, and machine learning technologies. This project has attracted debate among developers, social media activists, civil society organizations, and numerous digital companies about ethical AI use and abuse of digital privacy.

With China's increasing military exercises in the Taiwan Strait, and protracted confrontation between Russia and Ukraine, Japan established the National Security Innovation Strategy to ensure economic security. AI will play a key role in enhancing national defense capabilities. The National Institute of Advanced Industrial Science and Technology has developed a machine-learning based general-purpose object detection system to assist in maritime security and defense.

In Canada, national leaders are leveraging the capabilities of AI to monitor traffic flow patterns, and assist in regulating transit traffic in urban areas. Anomaly detection using advanced machine learning and AI technologies can alert intelligence and police agencies about irregular movements and challenges prior to actual events. The United Kingdom set up a program to support the use of AI in the health and care sectors, exploiting recent advances in the fields of secure data access, AI, and digital technology.

Finland is an example of a publicly-minded approach to AI and has prioritized knowledge-sharing to avoid the pitfalls of data silos that limit data availability, model

reproducibility, and fair AI service delivery. The strategic agenda on AI was developed with input from over 300 stakeholders, and sets priorities on how to develop ethical AI. The agenda makes a strong point that AI produced externally should comply with national values and be approved for use by Finland's public sector.

10.7.1. Success Stories in Local Government

Although generative AI is not widely deployed in local government, cities across the globe have begun experimenting with the tool. In 2023, several U.S. cities started uncovering ways for GPT to serve their needs through pilot projects. For example, in March 2023, the city of San Diego, California, gathered staff from 10 departments to experiment with its use of ChatGPT for drafting speeches, creating social media posts, answering questions, generating summaries, and more. Similar efforts were launched by employees in Miami, Florida; El Paso, Texas; and Austin, Texas, with suggestions ranging from streamlining report summaries and mapping out city budgets to drafting job descriptions and creating grant proposals. In a recent project, city of Los Angeles staff trained a custom version of ChatGPT to perform over 200 tasks in fields ranging from city planning to communications to public safety.

Internationally, many other localities have also tested various AI applications. In the United Kingdom, for example, the city of Birmingham recently piloted a generative AI tool to help city staff in creating meeting minutes. In a specific instance, the tool saved college student interns nearly a day of work through the process of producing summaries based on videos of city council meetings. Elsewhere in the UK, the city of Liverpool recently announced its testing of a specific version of GPT as a virtual assistant for residents, especially those dealing with social welfare and other challenging life issues. Nearby, the London Borough of Croydon also announced its trial of ChatGPT as a customer care assistant for residents.

10.7.2. Lessons Learned from Failures

Recent examples of AI deployment highlight its promise as a potential catalyst of previous failures and of broader aspirations, while its various forms still raise ethical, safety and legal questions. The use of GenAI tools has seen rapid adoption, with diverse use in a wide variety of areas. Local governments are currently using GenAI tools to create forms, all types of policy and regulatory documents, answers to frequently asked questions, and speech drafts for mayors or council members. Vendors are proposing additional use cases. Those not currently using GenAI tools are pragmatically weighing costs and benefits, and exploring how governors, regulators, privacy officers and ethics boards will guide and govern AI use going forward.

Governments must learn the recent lessons learned from previous generation autonomous AI. These utilized reinforcement learning and deep learning neural network concepts around for many years, and cost billions of dollars to early investors before their recent successes. The lessons are clear, and revolve around the need for a clear purpose and a limited scope. Be clear on why you are deploying GenAI tools, make sure it creates value, and limit its use to processes that can be smart-enabled to apply to a single domain process. And institute controls to regularly evaluate and modify GenAI use. Consider the ambitious example of prospective uses of Generative AI where for many citizens, "AI will help deliver on the promise to re-engineer government services in ways that enhance efficiency, accuracy, creativity, data-driven insights, personalization, and scalability while reducing costs."

10.8. Ethical Considerations

Generative AI tools have the potential to augment many tasks that collaborators could perform. However, the implementation of these tools presents new challenges in equity and ethics. These challenges require both researchers and practitioners to identify areas in which they must adapt standard practices of responsibility, such as recognizing how these tools make use of data, avoiding generated issues in public settings generated by a tool, and avoiding surreptitious processes of using these tools. Questions around how AI might change the ethical terrain are especially pressing, as these are use cases that are poorly covered by normative and practical approaches in the ethics literature. As a radical new approach to how work gets done, the ethical-landscape shift comes layered with a series of challenges. We first elaborate the most pressing ethical challenges we've identified around generative AI and public-sector innovation, and then turn to exploring concrete strategies that policy researchers and public administrators can adopt to move from a simple and perhaps naïve implementation of AI tools to a responsive, ethical, and ultimately effective one.

The use of machine-learning technologies and especially generative AI tools elicits concerns of bias and fairness, that is, the responsibility to avoid introducing bias and generating outputs that could result in majorities or historically marginalized groups being subjected to harm. The central focus of this work, generative AI, raises further ethical questions that are particularly germane to its role in shaping policy and administrative processes. We then discuss ethical questions around accountability and transparency, the responsibility of organizations to ensure that tools are adopted responsibly and the work generated through these tools is transparent, respectively, drawing not only from common ethical research themes, but also from the unique ethical-landscape shift based on the public-policy decision area – specifically, our

interest in how generative AI tools shape the policy and administrative-discursive landscape.

10.8.1. Bias and Fairness in AI

The idea that AI should be fair and not discriminate is an intuitive one. Yet the ways in which AI can, and does, encode and propagate deep inequities in society are numerous and complex. In addition to being trained on inputs that are discriminatory, which include biased labeling and stereotyped, unbalanced datasets, AI tools are also able to confuse commonly-associated groups in society — e.g., using a search model to match faces to prison records with devastating results. The unique nature of these effects makes it difficult to ascribe blame in AI — is it the engineer and employer creating it? The firm developing the tools? In many areas of public policy, from criminal justice to labor to welfare reform, biased decision-making inherently creates inequities, but these are often masked by the assumed objectivity and neutrality of algorithmic decision-making. For instance, in the case of a commercial AI-based recidivism risk assessment tool judges were required to use, a white woman convicted of murder who had just completed a sentence served no more time before parole than did a black man convicted of attempted murder and armed robbery, one with a longer record of incarceration.

This raises a larger point: AI tools are inherently incomplete, ignoring important factors that aren't quantitatively assessed or calculated when making decisions. A recidivism risk tool examining prior arrest records, for instance, doesn't calculate the varying levels of community support or career guidance a reentering individual will have access to on release. It can therefore only be one element in a larger determination of whether or not a previously incarcerated individual is functionally rehabilitated. Because purportedly objective AI tools provide a clear road map for public officials to be held accountable — clearly showing the basis for a recidivism decision — they are, in theory, an improvement on prior systems. But just as criminal justice data has a long history of racist inequity, so too will the AI tools built using such data, especially if they operate without oversight. Despite the accountability benefits of AI decision models, ethical questions persist.

10.8.2. Data Privacy Concerns

Citizens increasingly value their privacy and worry about ChatGPT storing their private information. They also express concern about using ChatGPT for work-related meetings and documents, suggesting worries about data security and leakage. Questions have arisen about monitoring prompts submitted by users, and several companies have banned

the use of generative AI at work. Therefore, it is important to understand data privacy-related safeguards before using ChatGPT in the public sector.

Security safeguards can help protect sensitive information from unintentional sharing through prompts and undesirable leaks by using ChatGPT to generate responses. Public administration can leverage different pre- and post-security controls. Before submitting prompts, ask if the information can be anonymized or removed. Instead of actual scenario prompts, consider using general approaches. Some companies offer on-premise models that can be run locally; government agencies may consider these. Besides using public models, organizations can deploy open-source models that can be locally run and modified. Sometimes, the degree of reidentification risk can inform the decision of using a public model and whether a generalized prompt is acceptable.

Data leakage by ChatGPT also needs to be addressed. Diverse approaches can be used to control the response, including adjusting temperature to reduce diversity. The model can be given guidance about both the role of a transformer and the knowledge it should draw upon. Internal or confidential information can be added to prompts as memory, internal models can be converted into different languages, and input-output mappings are provided to keep conversations in a specific area or define when responses should be ignored. Auditing prompts and responses can help track anomalies in data loss.

10.8.3. Accountability and Transparency

Generative AI has advanced so quickly that it presents a learning opportunity about who is accountable for the results, as well as the processes that produced the results. Quality work will almost certainly involve professionals with expertise in the subject matter reviewing and improving the results. However, work produced by a generative system can easily mislead a casual user. For example, concern for privacy has made this use controversial. Laws governing personal privacy are evolving quickly, partly in response to legal precedents. Organizations are liable for privacy of decisions affected by the data they collect, including decisions made by AI systems relying on that data. In addition, obligations of transparency that come with such rules mean that people affected by those decisions have a natural right to replay decisions to understand how they were made. Consider a generative AI system producing an advertisement for an influence campaign for a political candidate. The decisions branded as misleading require caution, but right to transparency extends only to adjectives that are truly descriptive — can AI systems result in such content?

More generally, organizations that want to use generative systems must be careful to provide documentation that lists the features described in the prompt and is used by the user as the prompt, as well as the type of model, version, parameters and features used

when the model produced the output, again, in quantifiable form. Professional systems should track AI-generated advertising directly and add an AI flag as an additional disclosure. Participants in that work should also discuss the potential exploitation of loopholes in the disclosure policy, two examples being just in time use of a previously described advertisement that turns out to be from a third party and partnerships offering the particular product at discount prices.

10.9. Future Directions

Emerging Technologies in AI Future directions to drive public sector innovation have several aspects: Policy decisions on whether to support innovation in generative AI, its research agenda in the public sector, and the longer-term impacts of possibly shifting AI capabilities. In particular, although advanced before, the wide availability and low cost of generative models and extensive tool chains make it possible for practically anyone to build novel and useful services lacking only the development of complex user-facing features that would require engineer development cycles, design experiences, and possibly user evaluations. Thus, we expect the current paths of sustainable innovation and momentum in the private sector to continue. This innovation may address the reliability and safety concerns associated with current LLMs by incorporating future advanced models. While generative AI will likely not put traditional knowledge workers out of business, it will augment their capabilities.

With sufficient speed, creativity, and direction, these innovations could reshape the interactions of the public sector with its constituents by replacing existing systems. They could enable contract workers, philanthropic organizations, and others with a lower marginal cost to create better systems that more directly fulfill public missions, enabling the public sector to take advantage of the collaborative model that has been demonstrated with crowdsourcing platforms.

Policy Recommendations Early ideas for policy recommendations concerning these directions include the following. 1) Expanding and accelerating the push for more rapid discovery and development of state-of-the-art AI tool chains in the public sector; for example, sustain, augment, or replenish incentives in private sector ecosystems that build supporting commercial AI developer platforms. 2) Commit capital to funding contracts so that startups sensitive and responsive to the needs of the public sector can build the missing production-quality public options for specific use cases, and enable collaborators from the broader ecosystem to connect with decision-makers and help leverage research results more effectively. Long-Term Impacts on Public Sector

10.9.1. Emerging Technologies in AI

Observing the last decades' accelerating technological progress provides incentive to wonder what is on the horizon for the future. Computers, previously mostly used by special interest groups, are becoming more ubiquitous and more powerful. What is more, it is becoming increasingly easy to use them. In the coming years, we will see an unprecedented democratization of digital tools. This phenomenon already appears in the rise of Generative Pretrained Transformers: Just a few years ago, researchers began training large-scale neural networks. Today, non-experts can use advanced dialog systems to prompt the system to write software code or entire essays, with only little chance of failure. These advancements do not only provide interesting demos, but are paving the road to commercial applications that will become integrated in an even larger array of digital systems that make human lives easier.

As we are about to witness highly practical, agentified versions of Python and AutoGPT, there is a strong possibility for a Cambrian explosion of easy to use practical AI systems. This has implications in diverse areas: Businessmen will develop increasingly better products and services. Journalists will hopefully focus on more hair-raising stories, as they will be able to skip hours of background research. School will get less boring for students and teachers, as education will be optimized for the individual. Gamers will see the emergence of novel game genres that integrate, mix and match procedural content generation with other ideas. Data will no longer be static: The internet as well as our personal databases will provide livestreams of digital twins that respond to the real world.

10.9.2. Policy Recommendations

This chapter contributes to the growing literature base on AI for government. The adoption and use of emerging technologies in AI for public services faces many challenges. This chapter identified challenges in three main factors: access to technology; capacity to use AI in public services delivery; and a Maturity-Driven Framework for government units to integrate these technologies into service delivery. With these challenges in place, several guidelines are proposed in order to mitigate and solve those challenges, fostering the positive effects of the use of AI and its aspirational goals. These recommendations can be elaborated as follows:

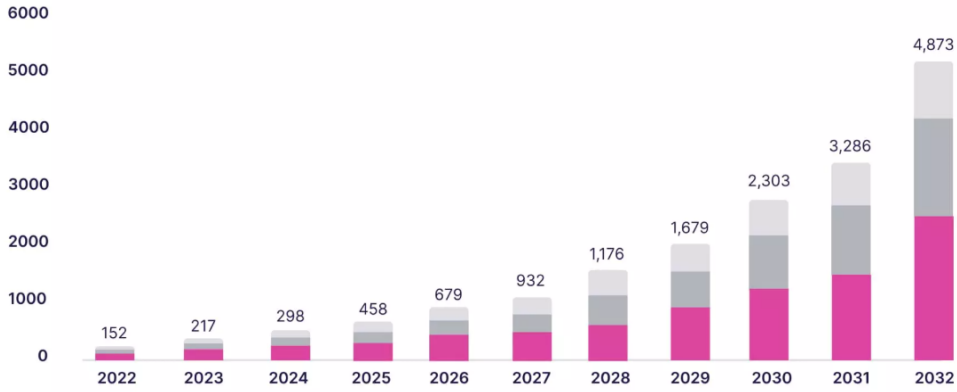

Fig : AI in Public Sector: Efficiency & Innovation

While at first sight it is easy to assume that an organization relatively big and strong as a government may have full, unrestricted access to emerging technologies such as generative AI, this is not necessarily true – the government is in a way a customer of the companies that provide those technologies, and have no control over how prices, licenses and condition of service are steered. While the market for AI products and services is still emerging, it is important to promote health and fair competition, which depends both on transparent and clear regulation on software and algorithms, and the development of a market of robust, open-source tools. These two actions are necessary to government units, especially those belonging to less developed economies or regions, which would otherwise simply not be able to access these technologies, or would face exorbitant prices, such as some generative AI tools. Another aspect that is relevant in terms of access is the existence of language barriers. With the more popular tools mainly supporting the English language at this point, less spoken languages are facing higher barriers to the access to innovation in the delivery of government services. Without complete access, governments and services will be locked into a vicious cycle of baseline services that would only reinforce non-consumption.

10.9.3. Long-Term Impacts on Public Sector

The idea that AI will have a much needed positive impact on the public sector is certainly embraced with enthusiasm and a fair amount of goodwill but largely dismissed as fantasy when a consideration of the long-term impacts of AI is undertaken. We are still in the early stages of developing the advanced technologies required even for the most productive implementations of many of the private sector uses we discuss. In much of the public sector, strategic foresight is in short supply. The many forces of opposition to making the choice that AI will be a powerful ally to ease the pressures of managing change are only just starting to respond to the challenges generated in the public sector.

It should not be forgotten that the development of these 'new' AI tools is not conducted in a vacuum, unaffected by the past. For the foreseeable future, the AI innovations that will assist the public workforce will continue to be fed from the streams of private sector implementation, reflecting the priorities and values that have animated the drive to maximize the bottom line.

Various counter-forces may move in different directions. On the one hand, the pressure for some form of governmental intervention in the process, whether through regulation or the carving out of specific functions and sectors, could cause the stream of support for private sector-derived tools to dry up. If the shorter term private sector gains are accompanied by social disruption, AI tools could be viewed as a cause to be mitigated rather than as a partner for development. On the other, a successful roll-out of public sector measures to integrate AI in a balanced and productive way could encourage its broader use as a partner in economic transformation. The priorities shaping the private toolbox could then be influenced by its success, resulting in a competitive dynamic. These flows originate in the values driving the private and public sectors. Management tools both reflect and shape the ethos.

10.10. Conclusion

In this brief chapter, we summarize our motivation for this book and its contents. We then discuss future work in this area.

As outlined in previous chapters, recent advancements in artificial intelligence (AI) – and, in particular, generative AI – have made it increasingly possible for organizations in any sector to derive great value from this technology. Some of this interest is being driven by the excitement surrounding generative AI, and the recognition that it could significantly change the nature of work. In this work, we are primarily interested in how generative AI can be used to help public sector organizations do their work more effectively and efficiently. We even work to highlight the fact that some emerging use cases have the potential to transform the way work is carried out today, especially tasks that rely heavily on expert knowledge, creativity, judgment, or providing personal services to clients. The emerging synergy between technology and people in the workplace would, in many ways, be an extension of existing technological advancements and the changes they have brought in the nature of work. Along with the potential, we also discuss the risks and challenges involved in using generative AI as a decision-making assistant in a public sector context.

We discuss four themes in depth: the characteristics of generative AI and how it could augment human capabilities, the types of use cases in the public sector in which generative AI could be applied, the risk and challenge mitigation strategies that public

sector decision-makers should consider, and the effective strategies for implementing generative AI to drive outcomes for public sector organizations and their stakeholders, both internal and external. With this foundation of knowledge, we also hope to motivate a fuller agenda for research in this area – one with deeper linkages between use cases, implementation strategies, and risk mitigation.

References:

Smith, J., & Kumar, A. (2023). Generative AI in government: Enhancing public service delivery and compliance. Government Technology Review, 12(3), 45–60.

Johnson, L., & Morales, F. (2024). Leveraging generative AI for smarter compliance and public sector workflows. Journal of Public Administration and Digital Innovation, 8(1), 13–29.

Androniceanu, A. (2024). Generative artificial intelligence, present and perspectives in public administration. Administration & Public Management Review, (43).

Kulal, A., Rahiman, H. U., Suvarna, H., Abhishek, N., & Dinesh, S. (2024). Enhancing public service delivery efficiency: Exploring the impact of AI. Journal of Open Innovation: Technology, Market, and Complexity, 10(3), 100329.

Yoon, S. (2024). Digital Innovation in Public Administration through Intelligent Public Sector Automation (IPSA): Strategies and Challenges. Journal of Multimedia Information System, 11(4), 249-260.

Carrasco, M., Habib, C., Felden, F., Sargeant, R., Mills, S., Shenton, S., ... & Dando, G. (2023). Generative AI for the public sector: From opportunities to value. Boston Consulting Group.