



AI in GovTech: what works, what fails, and why

The AI market is flooded with tools that promise everything but often deliver little. The key isn't more AI, it's the right AI built around real public sector needs.

AUTHORS:

Elena Sügis, PhD

Senior Data Scientist, Nortal

Grace Smith

Principal Consultant, GovTech and Healthcare, Nortal

Bobby Chandan

Client Principal, GovTech for Local and Regional Government, Nortal

Andres Raieste

Senior Vice President – Public Sector (Global), Nortal

This report may be cited as “Sügis, E., Smith, G., Chandan, B., Raieste, A. (2025). AI in govtech: what works, what fails, and why”

Nortal is a global strategic innovation and technology company and a valued partner for governments, healthcare institutions, leading businesses, and Fortune 500 companies.

Table of contents

| | |
|--|----|
| Introduction | 4 |
| Making AI work for government | 7 |
| From pilots to impact | 7 |
| 01. Sector-driven AI capability building | 9 |
| 02. AI-augmented government services | 10 |
| 03. AI-ready digital public infrastructure | 11 |
| Evaluating AI projects: when to pursue, when to pause | 13 |
| Diversity of AI projects | 14 |
| Why do most AI projects stall in the PoC phase? | 15 |
| Citizen adoption | 17 |
| Insights from successful AI deployment | 19 |
| The challenge | 20 |
| AI in the public sector: key questions to ask | 23 |
| Expectations & ROI | 24 |
| Budget constraints | 24 |
| Data challenges & legal restrictions | 25 |
| Citizen adoption | 25 |
| References | 26 |

Introduction

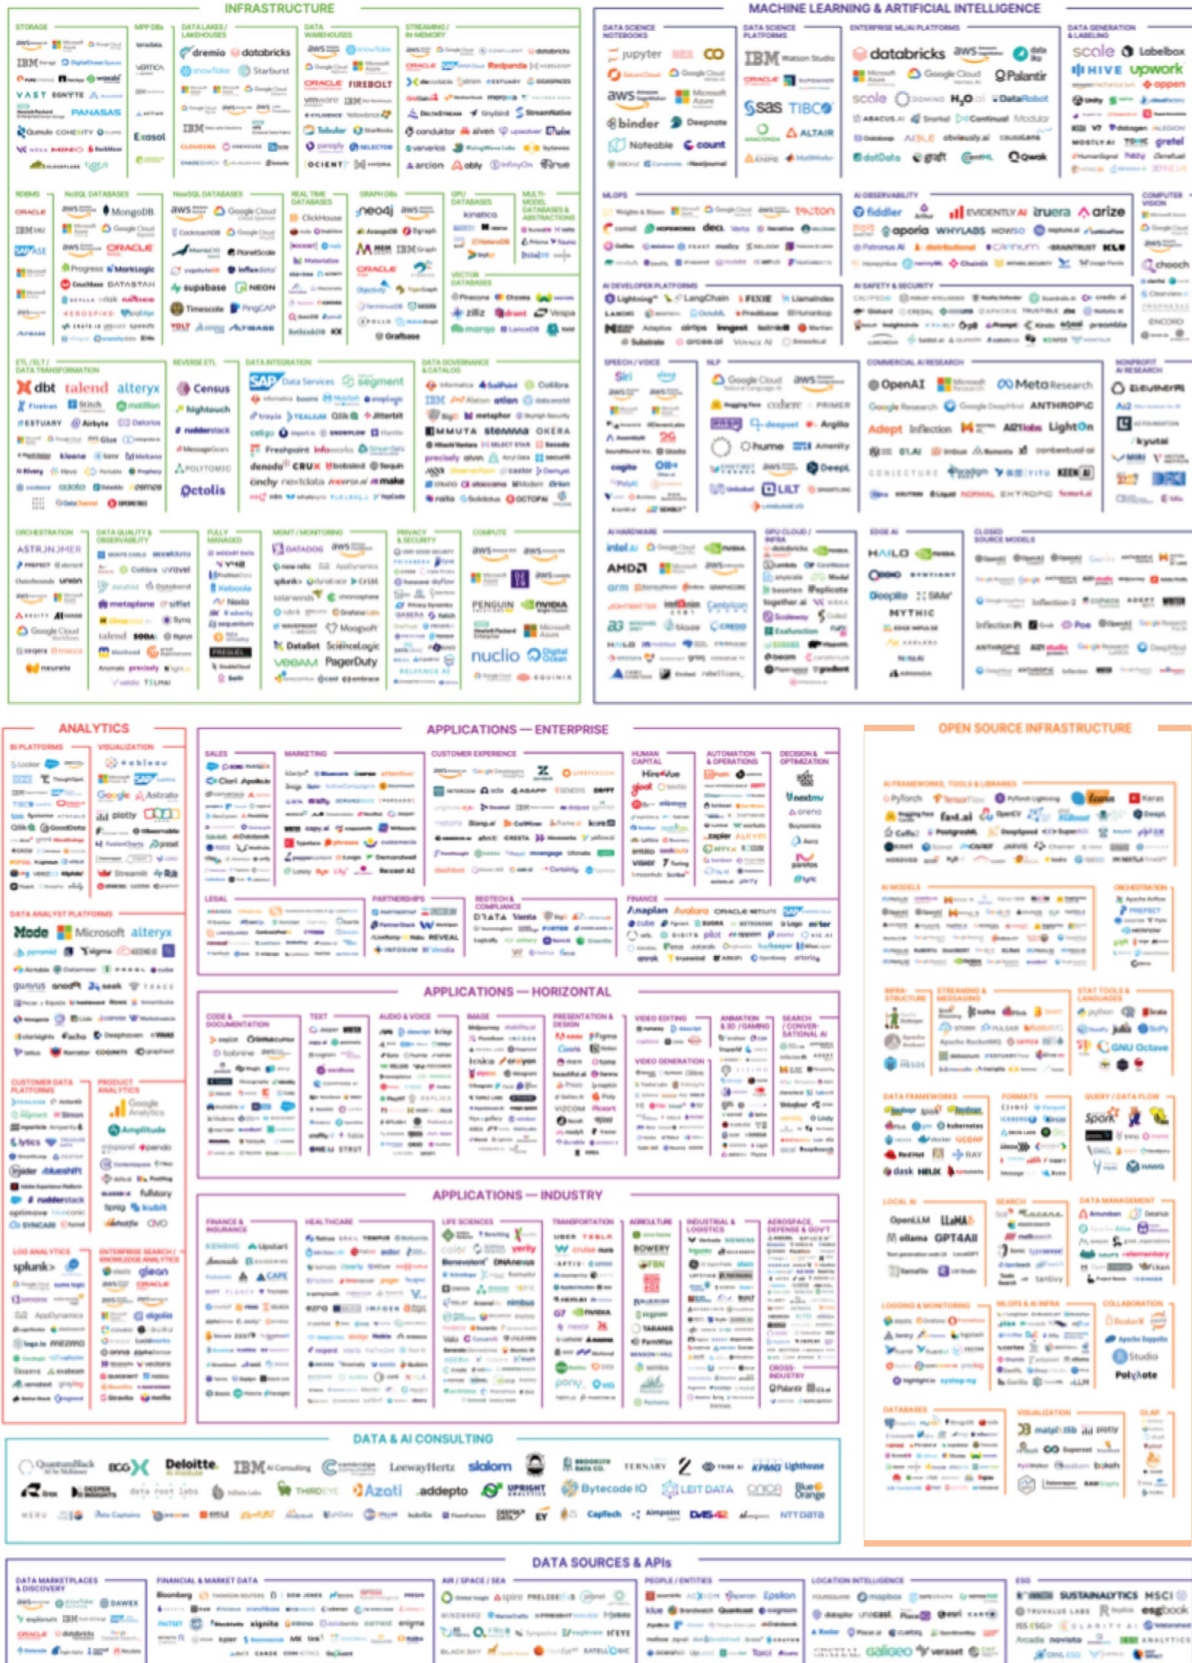
AI is advancing faster than governments can adapt. Only a few years ago, the dominant assumption was that AI would primarily automate repetitive tasks, but it is already extending into the realm of transforming public services, optimizing resources, and reshaping how governments interact with citizens. Despite its vast potential, most public sector AI initiatives never progress beyond the pilot stage. Over 70% stall at proof-of-concept—not because the technology isn’t ready, but because AI is not being embedded into the core machinery of government in a way that aligns with long-term public priorities.

AI is not like traditional IT systems that follow fixed rules – it is probabilistic, self-learning, and constantly evolving. Its output depends on real-time datasets and genuinely embedded intelligence, making it less predictable, harder to explain, and reliant on continuous oversight. In high-stakes areas like healthcare, criminal justice, and social services, public officials must balance the benefits from increased automation with accountability and long-term trust. [01]

For governments, this rapid progress presents both major opportunities and deep structural challenges. Advanced analytics, generative models, and decision-support systems offer unprecedented potential to enhance public services, optimize resource allocation, and improve citizen interactions. Yet their strengths—adaptability, learning capabilities, and probabilistic nature—also make them difficult to regulate, govern, and scale responsibly.

The challenge is clear: How can governments move beyond pilots and make AI work at scale?

THE 2024 MAD (MACHINE LEARNING, ARTIFICIAL INTELLIGENCE & DATA) LANDSCAPE



Version 1.0 - March 2024 © Matt Turk (@mattturck), Aman Kaberi (@AmanKaberi1) & FirstMark (@firstmarkcap)
 Blog post: mattturck.com/MAD2024 Interactive version: MAD.firstmarkcap.com Comments? Email MAD2024@firstmarkcap.com

FIRSTMARK
 EARLY STAGE VENTURE CAPITAL

The growing AI ecosystem - how can these tools solve my specific problem? Image source: Matt Turk, <https://mattturck.com/landscape/mad2024.pdf>



Making AI work for government

From pilots to impact

AI has the potential to make government more responsive, efficient, and effective—but only when it's designed for real-world public sector challenges. Too often, AI remains stuck in pilots, solving narrow problems but failing to scale.

The issue isn't the technology itself. It's that governments don't always have a clear process for matching AI's capabilities to sector needs, ensuring long-term adoption, and embedding AI within digital infrastructure that supports growth.

For AI to work, governments need to focus on three essential elements:

01

Sector-driven AI capability-building

AI adoption should begin with a clear view of what's possible in a given sector—where AI adds real value, how it fits into existing systems, and what risks need managing. This means leveraging industry expertise and proven use cases. The best AI deployments enhance services rather than disrupt them, and they are built by those who understand the sector from the inside out.

For practical insights, explore Nortal's bank of 750 public sector use cases [02] to understand the art of the possible and see how capability matching works in practice.

02

AI-augmented public services

AI must fit into government, not the other way around. The most effective AI is tailored to public sector realities, working within policy and operational constraints rather than forcing disruptive change.

Targeted adoption beats experimentation, and governments must focus on specific, high-impact AI use cases that improve efficiency, reduce bottlenecks, and enable smarter decision-making.

03

AI-ready digital public infrastructure (DPI)

Strong digital foundations are essential for scaling AI. Without interoperable, high-quality, and secure data infrastructure—and the governance to manage it—AI will remain a series of disconnected pilots.

By embedding AI within sector expertise, operational workflows, and a strong digital foundation, governments can move from AI as a pilot to AI as a core capability.

Sector-driven AI capability building

Many governments struggle because they attempt to apply AI without first understanding the unique challenges and opportunities within each sector. Without a clear process for matching AI capabilities to specific public sector needs, investments remain fragmented.

Sector-driven approach in practice

Define the art of the possible

- **Identify high-impact opportunities:** Use industry expertise to pinpoint where AI can add significant value—be it in healthcare, justice, public safety, or beyond.
- **Capability matching:** Leverage practical insights from Nortal's bank of 750 public sector use cases to match AI tools to sector-specific challenges.

Embed expertise within the sector

- **Bring industry experts into the process:** Collaborate with professionals who deeply understand sector operations to design AI solutions that truly support existing workflows.
- **Build in-house capability:** Ensure that not only IT teams but also policymakers, regulators, and service designers are involved in AI planning and governance.

Establish sustainable governance

- **Develop clear oversight:** Implement rules and accountability measures that evolve alongside AI deployments.
- **Focus on long-term adoption:** Move beyond pilots by ensuring that AI systems are scalable and integrated with sector-specific operational models.

Why it matters

- **Better outcomes:** When AI is matched to the specific needs of a sector, it transforms services rather than creating isolated projects.
- **Enhanced trust:** In-house capability and strong governance [03] foster accountability and public confidence.
- **Sustainable impact:** Industry expertise ensures that AI projects are practical, scalable, and aligned with long-term sector goals.

AI-augmented government services

Government services are often reactive, leaving citizens to navigate complex, cumbersome processes. Many AI projects address isolated tasks without transforming the overall service experience.

AI-augmented public services in practice

Personal government

- **Proactive engagement:** AI can anticipate citizen needs and deliver tailored information and services at the right time, reducing the need for citizens to seek out help.
- **Customized interactions:** Use data-driven insights to provide personalized services that align with each individual's circumstances.

Seamless integration into workflows

- **Invisible efficiency:** The best AI solutions work behind the scenes—streamlining processes, automating routine tasks, and freeing up public servants for more complex decision-making.
- **Support, not replacement:** Ensure AI tools enhance human expertise rather than substituting it, allowing public servants to focus on high-touch interactions.

Transparency and trust

- **Clear accountability:** AI systems should be explainable and auditable, so citizens understand how decisions are made.
- **Ethical deployment:** Maintain robust oversight to ensure fairness and minimize bias in AI-driven processes.

Why it matters

- **Improved citizen experience:** Proactive, personalized services reduce friction and make government more accessible.
- **More efficient operations:** Streamlined workflows lead to faster, more responsive public services.
- **Strengthened public trust:** Transparency and clear accountability build confidence in digital government initiatives.

03

AI-ready digital public infrastructure

Even the most promising AI solutions cannot scale without a strong digital backbone. Many government AI initiatives stall because they rely on outdated, siloed, or fragmented data systems.

AI-ready digital public infrastructure in practice

Interoperability and data standards

- **Unified data frameworks:** Develop and adopt common data standards to ensure seamless sharing across departments.
- **Real-time connectivity:** Enable secure, real-time data flows that give AI systems access to comprehensive, up-to-date information.

Robust security and privacy


- **High-quality data:** Invest in modernizing legacy systems to deliver structured, machine-readable data.
- **Strong protection measures:** Build in data privacy, security, and ethical guidelines to maintain public trust and meet regulatory requirements.

Flexible, adaptive infrastructure

- **Modular systems:** Design digital platforms that can evolve with AI advancements, avoiding the pitfalls of rigid, outdated technology.
- **Embedded governance:** Integrate clear oversight and accountability into the digital infrastructure from the start, ensuring AI remains transparent and fair.

Why it matters

- **Scalable AI deployment:** A robust digital infrastructure allows AI initiatives to move beyond isolated pilots and operate across entire agencies.
- **Enhanced collaboration:** Interoperability facilitates data sharing, enabling more comprehensive and integrated AI solutions.
- **Long-term sustainability:** Secure, adaptive infrastructure supports the continuous evolution and improvement of AI applications.



Evaluating AI projects: when to pursue, when to pause

The National Audit Office reported in March 2024 [04] that 70% of government bodies surveyed are piloting or planning the use of AI, with applications including supporting operational decision-making and improving internal processes. However, the report also highlighted risks to value for money if issues such as aging IT infrastructure, skills gaps, and data management are not addressed.

Artificial Intelligence has emerged as a transformative force in the public sector, with potential applications ranging from improving citizen services to optimizing internal operations. Yet, despite its promise, many AI projects fail to deliver tangible impact, with 70% remaining stuck at the Proof-of-Concept (PoC) stage.

Why do so many AI projects stall? Our analysis of 4,372 AI use cases (including 750 publicly available cases) reveals recurring pitfalls that hinder AI from moving beyond PoCs. To address these challenges, public sector leaders must move beyond unrealistic expectations, carefully evaluate opportunities, and take strategic steps to achieve meaningful AI adoption.

Diversity of AI projects

Based on their goal, AI projects can be broadly categorized into three groups: process automation, predictions for decision support, and generative AI for increasing efficiency. Each comes with distinct benefits, challenges, and considerations.

Automation projects

These aim to streamline repetitive tasks, such as processing documents, managing workflows, or responding to standard citizen inquiries. Automation projects are ideal when processes are well-defined and data is structured. However, they are less suitable for tasks requiring nuanced judgment or creativity.

Example:

A municipality automates building permit applications, reducing processing time from weeks to days.

Prediction for data-driven decision-support

These projects leverage machine learning to forecast outcomes, such as predicting healthcare needs or traffic patterns. These projects succeed when high-quality, unbiased data is available. Challenges arise when data is fragmented, biased, or lacks historical depth, leading to unreliable predictions.

Example:

A national health agency predicts flu outbreaks but struggles due to gaps in patient data, limiting model accuracy.

Generative AI for increasing efficiency

This category includes AI-powered content creation, such as drafting reports, summarizing documents, or simulating citizen feedback. While promising, generative AI demands cautious adoption due to risks like hallucinations, misinformation, and lack of explainability.

Example:

A government agency deploys a generative AI tool to draft policy briefs, but officials must still validate outputs due to potential inaccuracies.



The decision to pursue AI must align with strategic objectives and operational readiness. [05] AI projects that chase hype rather than solve well-defined problems often fail.

Why do most AI projects stall in the PoC phase?

Based on our case study analysis, we have identified the most common pitfalls that prevent AI projects from scaling beyond PoCs.

Lack of clear leadership & ownership

A major challenge in the public sector is the absence of clear ownership. AI projects often lack a dedicated leader responsible for driving them from PoC to production. AI projects are frequently assigned as an additional responsibility rather than a primary role. Without someone accountable for long-term success, projects stall at the PoC stage.

Example:

A government agency develops an AI-based case management tool that shows strong PoC results but fails to scale because no single leader is responsible for integrating it into daily operations.

Solution:

Assign a dedicated AI leader or task force with clear authority and budget to oversee implementation and long-term success.

Misaligned expectations & ROI

One of the biggest reasons AI projects fail to progress beyond the PoC stage is misaligned expectations. Organizations, particularly in the public sector, often expect AI to deliver instant results, underestimating the complexity of implementation, integration, and maintenance. A common issue is overly optimistic ROI projections. Many believe AI will immediately reduce costs or automate entire processes, but the reality is that AI requires significant upfront investment before delivering measurable returns. Projects struggle to demonstrate value without clear baseline metrics, such as how much time is spent on a task and how AI could realistically improve it.

Another challenge is vague or undefined success criteria. Many projects launch without a clear answer to what success looks like—should AI improve speed, accuracy, cost efficiency, or user experience? Without well-defined KPIs, stakeholders may lose confidence in the project, leading to abandonment.

Example:

A government agency deploys AI to optimize document processing but fails to define KPIs like time saved per case, leading to uncertainty about the project's impact.

Solution:

Clearly outline expected before-and-after metrics, such as time reduction, accuracy improvements, or cost savings.

Budget constraints as a hidden barrier to AI growth

Many AI projects receive funding only for the PoC stage, with no long-term financial planning for scaling and maintenance. AI budgets are front-loaded on experimentation, but operationalization costs (infrastructure, retraining, compliance) are overlooked. AI is not a standalone tool; it requires integration into workflows and processes, staff training, and ongoing monitoring. If these factors are overlooked, AI adoption struggles, especially among digitally excluded groups like the elderly. Many public sector projects rely on one-time funding grants, which are not renewed for scaling AI systems. ROI concerns lead to hesitation in approving additional funding.

Example:

A government agency launches an AI initiative to improve service delivery but allocates funding only for the initial proof-of-concept phase. Without budgeting for essential long-term costs like infrastructure upgrades, staff training, and continuous monitoring, the project struggles to scale, leaving critical integration gaps and marginalizing digitally excluded groups.

Solution:

Develop a comprehensive funding strategy that includes not only the PoC phase but also long-term operational costs. Incorporate phased rollouts, change management strategies, ongoing monitoring, and staff training to ensure sustainable and inclusive AI adoption.

Data challenges & legal restrictions

AI systems rely on high-quality, well-structured data to function effectively. Many projects fail because of

- Poor data quality (inconsistent, outdated, or biased datasets).
- Lack of access (data silos prevent AI from leveraging relevant information).
- Legal constraints (privacy laws restrict data sharing).

Example:

A healthcare AI tool for personalized patient care is developed, but cannot access critical patient data due to privacy regulations, limiting its effectiveness.

Solution:

Establish data governance frameworks to ensure quality and access. Implement privacy-preserving AI techniques (e.g., federated learning) to comply with regulations.

Citizen adoption

Citizen adoption remains uneven, with trust issues, digital literacy gaps, and accessibility barriers affecting uptake. This is especially pronounced among digitally excluded groups, such as the elderly, who may struggle to adapt to AI-driven services. Citizen adoption is a significant factor in the success of AI initiatives within the public sector. While not an insurmountable barrier, it presents challenges that require careful consideration and proactive management.

In the United Kingdom, the integration of AI into public services has been met with both enthusiasm and caution. A recent YouGov poll indicates that 87% of the population support laws requiring AI developers to prove their systems' safety before release, reflecting a strong public desire for responsible AI deployment.

The most resource-intensive AI solutions often provide the greatest benefits—such as automating complex bureaucratic processes—yet they also face the biggest adoption hurdles. Addressing issues related to public trust, digital accessibility, and user engagement is essential for the effective implementation of AI-driven public services. This requires transparent communication, AI literacy programs, and inclusive service design to ensure that both decision-makers and citizens can effectively engage with AI-powered systems.



Case Study

Insights from successful AI deployment

One of the most compelling examples of AI-driven public sector transformation is Nortal's implementation of the AI-powered decision support tool (OTT) for the Estonian Unemployment Insurance Fund (EUIF). This tool calculates the prognosis of moving into employment for every newly registered person who has been unemployed for 35 days. It also calculates the probability of becoming unemployed again within a year, and indicates the factors affecting these probabilities. Based on the received information, the EUIF councillor can offer data-driven personalized labour market services. This initiative demonstrates how AI-powered decision support and process automation can significantly enhance efficiency, accuracy, and service delivery, ensuring better outcomes for both administrators and citizens.



The challenge

With the start of the COVID-19 pandemic in 2020, 96 thousand people in Estonia, a country of 1.3 million, registered themselves as unemployed, resulting in the highest unemployment rate (6.8%) since the recovery from the 2008 global economic crisis. The 350 consultants at the Estonian Unemployment Insurance Fund (EUIF) were tasked with inspecting and analyzing vast amounts of information to design personalized action plans that enable each individual to successfully return to the labour market. This process is time-consuming, but it is crucial to act fast - the longer an individual remains unemployed, the greater the financial strain on both the person and the state, while also diminishing the likelihood of reemployment. Given that there were 96,000 applications to process, it appeared clear that 350 consultants were simply not enough to handle the workload quickly and efficiently.

Several critical challenges emerged:

- Staff capacity constraints: Caseworkers struggled to keep up, leading to delays and inefficiencies.
- Current business processes (such as manually reviewing and analyzing each application, creating individual action plans, and following up) made it impossible to scale operations effectively.
- The abundance of data made it laborious and error-prone to manually distinguish the unemployed who have more difficulties getting back into the labour market and should receive more attention and resources.

EUIF needed a scalable, efficient, and fair solution that could handle high demand without increasing staff while ensuring equitable service delivery across regions.

Nortal and Center for IT Impact Studies (CITIS) designed and developed decision support tool for counsellors [06, 07] that streamlined their workload to focus more on individuals needing extra support.

The solution allowed the Fund to manage the enormous workload and offer even more personalized services, all without increasing the staff.

The Estonian Unemployment Insurance Fund faced many of the same challenges that cause 70% of AI projects to stall at the Proof-of-Concept (PoC) phase. However, through careful planning, process redesign, and strategic AI-powered decision-support tool integration into the Employment Information System successfully transitioned it from PoC to full-scale deployment. We can extract valuable insights for other public sector AI initiatives by analyzing these challenges and solutions.

Ownership & leadership

Unlike many AI projects where responsibilities are vague, OTT had a dedicated implementation team within EUIF, ensuring clear leadership. Instead of treating OTT as an experimental initiative, it was integrated into EUIF's long-term digital strategy, making its success a priority at the executive level. EUIF ensured that IT teams, policymakers, and employment counselors worked together to align the system with organizational needs.

Expectations & ROI

Instead of aiming for full automation, OTT was designed as a decision-support tool, helping caseworkers enhance, not replace, their role. This ensured that AI complemented human expertise, increasing acceptance and adoption.

The project measured success through specific, tangible outcomes, including:

- Increased employment rates due to personalized counseling.
- Saved 3.8% of Estonia's yearly expenditures by reducing unemployment by a single day.
- Provided 95% forecast accuracy.
- Achieved 100% visibility of the unemployed, risks, and regional trends daily.

Data challenges & legal restrictions

OTT was integrated with national databases, enabling real-time data validation and reducing errors in benefit claims. AI-driven cross-checking against national registers ensured accuracy in claims while maintaining compliance with legal frameworks. The system was designed to comply with Estonia's data protection regulations, ensuring that sensitive citizen data remained secure.

Budget constraints

Unlike many AI PoCs that rely on short-term funding, OTT was developed with sustainability in mind, ensuring funding was secured beyond the pilot phase. Instead of increasing staffing costs, OTT allowed caseworkers to handle more cases effectively, demonstrating a direct return on investment to policymakers.

In 2021, the AI-powered decision support tool for unemployment services and benefits was awarded The Best Data-Based Digital Service in the Estonian public sector.

Citizen adoption & trust

Instead of fully automating services, OTT provided AI-powered recommendations to caseworkers, ensuring citizens still received personalized support. AI-powered services were available both online and via human caseworkers, ensuring that digitally excluded citizens (such as the elderly) were not left behind. AI decisions were clearly documented, and caseworkers could override recommendations, ensuring public confidence in AI-driven decision-making.

This case serves as a blueprint for other public sector organizations looking to scale AI solutions successfully. It highlights that AI's potential is only realized when it is backed by thoughtful planning, process adaptation, and continuous optimization. The lesson for governments and institutions aiming to leverage AI effectively is clear: technology alone is not enough – success depends on ownership, strategy, and trust.



AI in the public sector: key questions to ask

AI projects in government often face similar hurdles: misaligned expectations, lack of clear leadership, and difficulty in scaling beyond the pilot phase. When considering AI for public sector projects, it's crucial to ask the right questions before moving forward. Governments must take a thoughtful and strategic approach to AI adoption, ensuring that the technology aligns with their objectives and solves real challenges. To maximize the potential of AI investments, public sector organizations should evaluate proposals with a critical eye. Here are the essential questions to help determine if AI is the right solution and how it can create value.



Expectations & ROI

Identify pain points and opportunities

Define the problem statement clearly – AI can help automate repetitive tasks, improve decision-making, enhance customer service, and more.

- Are the areas where your organization faces challenges or inefficiencies right for an AI - type solution? Could it be addressed using another solution?
- Can your organization learn from innovative new business models deployed elsewhere that are transferrable?
- Are there any proven AI use cases across sectors with similar characteristics?

Focus on ROI

Choose projects that have a clear return on investment. This will help justify the initial costs and resources required for AI implementation.

- What specific problem does an AI solution address?
- Is there a clear case for efficiency, cost reduction, improvement in quality, customer retention, additional revenue, etc?
- How is success measured, and what KPIs will be used to evaluate outcomes?



Budget constraints

Identify and budget for total running costs

- What are the long-term costs for maintenance and updates?


Access to AI skills and expertise

- Can we access and collaborate with the appropriate AI experts, consultants, and stakeholders?
- What level of MLOps expertise is required to deploy and maintain this AI solution?
- Should MLOps be handled internally, or will it require external service providers for continuous monitoring and optimization?

Ease of implementation

Consider if AI can be deployed in a select area to test viability.

- Can we implement AI solutions in selected areas and monitor performance before deciding to scale up?
- Can we lower the risks of poor outcomes by testing iteratively?
- How will this solution scale across departments or regions?




Data challenges & legal restrictions

Evaluate data availability and governance

AI thrives on data. Ensure you can access to quality data in the areas you want to improve. This data will be crucial for training AI models.

- What data is required, and is it available in the needed quality and quantity?
- Are there any data privacy or security concerns to address?



Citizen adoption

Gauge public acceptance

Consider public opinion and ethical implications, especially for services directly impacting citizens.

- How does the solution ensure fairness, transparency, and accountability?
- Are there risks of bias, and how are they mitigated?
- What training and support will users need to adopt this solution?
- How will workflows and processes be adjusted to integrate AI?

In conclusion, AI holds immense promise for the public sector, but its success hinges on aligning the technology with the unique needs of government operations. Too often, projects stall because they fail to integrate AI in a way that supports long-term goals and existing systems. By focusing on sector-driven solutions, ensuring AI adoption is scalable, and embedding robust digital infrastructure, governments can unlock the true potential of AI, moving from isolated pilots to impactful, sustainable change.

References

1. Raieste, A., Rozov, J., Solvak, M. (2024). Formula for creating trust in digital government. <https://www.personalgovernment.com/trust>
2. Nortal (2024). Real-world AI use cases. Retrieved from <https://info.nortal.com/ai-use-cases-public-sector>
3. Nortal. (2024) Mastering Data Governance. Available: <https://info.nortal.com/mastering-data-governance>
4. National Audit Office (2024). Use of artificial intelligence in government (HC 612). Cabinet Office & Department for Science, Innovation & Technology. Retrieved from <https://www.nao.org.uk/reports/use-of-artificial-intelligence-in-government/>
5. Nortal (2024). Data readiness: A prerequisite for the public sector benefits of AI. Retrieved from <https://info.nortal.com/uk-ai-maturity>
6. Nortal. (2022). Estonian Unemployment Insurance Fund prevents unemployment with Artificial Intelligence. Retrieved from <https://info.nortal.com/estonian-unemployment-insurance-fund-ai>
7. European Commission: Joint Research Centre, Misuraca, G. and Noordt, C., AI watch, artificial intelligence in public services – Overview of the use and impact of AI in public services in the EU, Publications Office, 2020, <https://data.europa.eu/doi/10.2760/039619>

25 years of shaping
the future