ISEA
CENTRO STUDI

# Delivery of Public Services in the AI State

Zijian Khor, Dr. Alex Lin

Article n° 3

October 2025

*This article has been published following the agreement between ISEA – Istituto di Studi per l'Economia Applicata srls and SR21. SR21 offers selected scholars the opportunity to present their research work to the Saturday Global Calls network and to publish it on the ISEA website (courtesy of ISEA) in the section dedicated to the ISEA / SR21 partnership.*

*The articles published in this section have not undergone academic peer review and are approved by the editorial committee based on the recommendation of SR21, which endorses their publication. The author remains solely responsible for the content of the article.*

ISEA – Istituto di Studi per l'Economia Applicata srls

🌐 www.iseacentrodstudi.com

✉ info@iseacentrostudi.com

# Delivery of Public Services in the AI State [1]

Mr. Zijian Khor[2], Dr. Alex Lin

[1] This article is the sole responsibility of the author and neither ISEA nor any of its affiliates have any role in the drafting of this article.

[2] Zijian Khor (zijian.khor@gmail.com), is an ESG public leader and employed with the Singapore Government.

# Abstract

Delivery of Public Services in the AI State [3]

Artificial Intelligence (AI) is redefining how governments deliver public goods and services, moving from automation to anticipation. Unlike earlier digitization efforts that streamlined processes, AI empowers states to predict needs, personalize delivery, and allocate resources dynamically. This white paper argues that AI is now part of the governing architecture itself—shifting governments from reactive bureaucracies to proactive, data-driven systems capable of disbursing welfare automatically, forecasting public health crises, and preventing infrastructure failures before they occur.

The paper examines both opportunities and risks. While AI could reduce administrative costs by up to 30–35% (Boston Consulting Group, 2025) and enhance public sector productivity (Tech Monitor, 2025), it also reshapes the social contract, raising questions about fairness, transparency, and accountability. Drawing from global examples—including Singapore, Estonia, South Korea, Finland, and Brazil—the study explores sectoral applications in healthcare, education, welfare, taxation, infrastructure, and policymaking. A qualitative readiness index assesses governments' capacity to adopt AI based on technical maturity, data integration, and public trust.

Findings show that AI can strengthen governance legitimacy through efficiency, inclusiveness, and cross-agency coherence. Yet its success depends on strong guardrails—ethical design, transparency, accountability, and international collaboration—to ensure that AI augments, rather than undermines, democratic governance and citizen trust.


---

[3] *This is a pro bono academic initiative.*


*Keywords: Artificial Intelligence, Digital Government, Public Goods, Governance Innovation, State Capacity, Automation, Policy Design, Public Sector Efficiency, Technology Adoption, Accountability, Trust*

*JEL Classification: H11 – Structure, Scope, and Performance of Government; O33 – Technological Change: Choices and Consequences; Diffusion Processes; O38 – Government Policy; Technological Change and Innovation; D73 – Bureaucracy; Administrative Processes in Public Organizations*

# Introduction: Governing with AI

Public goods—education, healthcare, social protection, justice, infrastructure—are the foundation of the modern state. The central challenge has always been scale: how to ensure universal access while managing *limited* resources and maintaining fairness. Traditional bureaucracies relied on human-led, hierarchical processes, which often meant delays, inefficiencies, and uneven outcomes. AI offers a new operating model. By embedding machine learning, predictive analytics, and natural language processing into public systems, governments can:

- Anticipate needs instead of waiting for citizens to apply (for example, flagging a household eligible for benefits before they fall into hardship).
- Optimize resources dynamically, such as reallocating welfare funds or adjusting transport in real time based on predicted demand.
- Personalize services at scale, tailoring interventions to individual circumstances while upholding universal access.
- Monitor and respond at speed, closing the gap between decision and delivery through real-time data feeds and adaptive algorithms.

This is not simply administrative reform—it is a redefinition of the social contract. Citizens who once queued at offices or waited weeks for approvals increasingly interact with AI-driven services. Decisions that once rested on human discretion are now shaped by algorithms. This raises fundamental governance questions: Who is accountable for an automated denial of benefits? How do citizens appeal opaque AI decisions? Can trust in public institutions survive when service delivery is mediated by machines? Early evidence suggests that when done right, AI-driven responsiveness can improve public trust (Global Government Forum, 2025), but missteps can just as easily undermine legitimacy.

To explore these questions, this study proceeds in several parts. First, it examines how AI is being deployed across domains of public service—from education and healthcare to taxation and welfare—detailing the technologies involved and the impacts on efficiency and equity. Second, it reviews the international landscape, comparing different national approaches. Brownfield transformations (where existing institutions adopt AI incrementally) are contrasted with greenfield innovations (where AI-native government functions are built from the ground up). For example, *Estonia* incrementally adds AI assistants to its e-government systems,

whereas *Dubai* created a dedicated Ministry of AI to accelerate adoption. Nordic countries integrate AI into their strong social welfare systems with an emphasis on **ethical, human-centric design**, while countries like *Brazil* leverage AI to leapfrog capacity constraints (Brazil's courts have launched over 140 AI projects to cope with 70+ million lawsuits) (Rest of World, 2025). *South Korea*, consistently top-ranked in e-government, employs AI for purposes ranging from urban traffic management to detecting online fraud and even predicting demand for school programs (Inter-American Development Bank [IDB] Gobernarte, 2024).

Third, the paper explores cross-cutting governance implications of AI in the public sector. It includes new sections on how AI can augment **policymaking and regulatory design**, and how governments are beginning to collaborate across borders on AI for public service delivery. AI isn't just automating services—it's also informing how policies are made (for instance, analyzing vast datasets to draft smarter regulations (OECD, 2024)) and enabling shared solutions (such as international data standards and collaborative AI initiatives for global challenges (Brookings Institution, 2023).

The central claim is that AI in public services is not just about technology, it is about governing differently. Done well, AI can help states evolve from reactive administrators to proactive stewards of public welfare: faster, smarter, and potentially fairer. The gains include greater efficiency, predictive power to prevent problems, and mass personalization of services. But these gains are not automatic. Policymakers must ensure that AI **enhances**, rather than erodes, the principles of accountability, inclusion, and transparency that underpin public goods. This requires investing in data infrastructure, skills, and legal frameworks, as well as setting clear ethical limits (for example, banning AI-driven mass surveillance that violates rights). In the pages that follow, we map both the opportunities and the guardrails needed as governments integrate AI into the very machinery of governance.

Before turning to individual sectors, we outline the AI **technology stack** enabling this transformation.

## AI Technology Stack for Public Services

Successfully deploying AI in government requires a robust stack of technologies, from foundational infrastructure to sophisticated algorithms. Broadly, the public sector AI stack comprises:

- **Data Infrastructure and Integration:** High-quality data is the fuel for AI. Governments need interoperable data architectures (secure data exchanges, APIs, cloud platforms) to break down agency silos. Initiatives like *Estonia's X-Road* illustrate how a secure cross-agency data backbone enables AI systems to draw from multiple databases seamlessly. Privacy-preserving data techniques (encryption, federated learning) are increasingly used to balance openness with individual rights. Without solid data foundations, even the best AI algorithms will produce poor results or remain pilot projects.

- **Computing Power and Cloud Platforms:** AI workloads often demand significant computing resources. Many governments partner with cloud providers or build government clouds (as *South Korea* did with its G-Cloud in 2012 (Inter-American Development Bank [IDB] Gobernarte, 2024)) to ensure scalable, secure computing capacity. Edge computing is also used for latency-sensitive tasks (e.g. processing traffic camera data in real time on location).

- **Core AI Algorithms and Models:** A range of AI techniques are employed:

    - *Natural Language Processing (NLP):* powering chatbots, virtual assistants, and document analysis. For example, conversational AI helps answer citizens' queries in multiple languages, and large language models (LLMs) can draft summaries of legislation or provide first-line support to users.
    - *Machine Learning & Predictive Analytics:* statistical models and neural networks analyze historical data to predict future events or risks (from predicting tax fraud or disease outbreaks to forecasting infrastructure failures).
    - *Computer Vision:* image and video analysis used for everything from reading medical scans and detecting road potholes to monitoring security camera feeds or satellite imagery for environmental changes.
    - *Knowledge Graphs & Anomaly Detection:* graph analytics link data across systems (e.g. to uncover fraud networks in welfare or tax), and anomaly detection algorithms flag irregular patterns (like unusual spending in budgets or anomalies in public health data).
    - *Robotic Process Automation (RPA):* though not "intelligent" in itself, RPA bots handle repetitive digital tasks (data entry, form processing) and often work in

tandem with AI (for instance, an AI might decide an invoice is fraudulent and an RPA bot then blocks the payment).

    - *Generative AI:* a recent breakthrough, generative models (like GPT-4 or equivalents) can produce human-like text, code, or synthetic data. Governments are experimenting with these for drafting reports, translating bureaucratic jargon into plain language, and even **assisting in writing legislation** (in 2023, a U.S. Congressman introduced the first bill co-written by ChatGPT (IE University – Center for the Governance of Change, 2024)).

- **Platforms and Applications:** On top of core models, there are platforms tailored to public sector needs. Examples include integrated data dashboards for policymakers, digital assistants like *Ask Jamie* in Singapore for citizen services, or Alibaba's **City Brain** platform for smart city traffic orchestration. Some governments use open-source AI tools adapted to their context, while others procure solutions from vendors (with a trend toward requiring explainability and openness from those vendors).

- **IoT Sensors and Digital Twins:** IoT devices (sensors on infrastructure, smart meters, environmental monitors) supply real-time data that AI systems consume. **Digital twins**—virtual models of physical assets or even whole cities—are increasingly used to simulate scenarios. For instance, digital twins of power grids or water systems let officials run "what-if" analyses (like how a flood would impact the network) and test AI-driven responses in a sandbox.

- **Experimental Technologies:** Governments are also testing frontier technologies that complement AI:

    - *Autonomous systems:* from drones for inspecting remote infrastructure to self-driving buses, these rely on AI and may transform service delivery (e.g. automated emergency drones in firefighting).
    - *Quantum computing:* still nascent but being explored for optimizing complex logistics or cryptography to secure government data.
    - *Blockchain:* occasionally paired with AI for transparency (as in some pilots for tracking public finances or verifying identity in AI-managed processes).

This stack is continuously evolving. Importantly, many components are market-available and mature (NLP, computer vision, predictive analytics for specific use cases), while others are experimental (advanced autonomous agents or quantum algorithms for policy problems). Governments must choose the right tools for the job, often blending off-the-shelf platforms with custom development.

Crucially, technology alone is not enough—organizational and human factors (discussed in Cross-Cutting Enablers) determine success. But understanding the stack helps policymakers see what is possible today. With this technical foundation in mind, we now turn to how AI is being applied, domain by domain, in the delivery of public goods.

# Chapter 1: Education & Lifelong Learning

Education has always been one of the most visible expressions of a government's promise to its citizens: to prepare the next generation, to equalize opportunities, and to create pathways into the economy. Yet despite decades of reform, education systems remain burdened by the limits of scale. Teachers face large classrooms with diverse needs, resources are stretched thin, and curricula struggle to keep pace with the demands of a rapidly changing economy. Artificial Intelligence is beginning to shift this equation, offering governments tools to deliver learning that is not only more efficient, but also more adaptive, personalized, and continuous across a lifetime.

AI's impact on education lies in its ability to transform what was once a standardized, one-size-fits-all model into a system capable of responding to individual needs in real time. Adaptive learning platforms can track how a student engages with a lesson and adjust difficulty accordingly, much like a skilled tutor would. In China, platforms such as Squirrel AI already provide millions of students with AI-driven tutoring that adapts to their performance. Natural language models extend this further, functioning as virtual teaching assistants that can answer questions, explain concepts, and even generate exercises on demand in multiple languages. At the system level, learning analytics dashboards allow policymakers to see where cohorts are falling behind and target interventions before failure rates escalate.

This transformation does not stop at childhood education. For adults, AI enables governments to build lifelong learning systems tied directly to labor market demand. Training pathways can be dynamically adjusted based on job vacancies and skill shortages, giving workers a clearer route to reemployment or upskilling. In Singapore, where continuous reskilling is seen as essential for economic competitiveness, government-funded training programs are increasingly supported by AI that matches workers to courses based on their profiles and market trends.

This new model of education also reshapes the role of teachers. Rather than spending hours grading or preparing lesson plans, AI systems can automate routine tasks, freeing teachers to focus on mentoring and personal interaction. Lesson-planning engines generate material aligned with national curricula, while assessment tools provide detailed analytics on student progress. Teachers become not less important, but differently important: custodians of critical thinking, empathy, and social development, while AI handles the repetitive mechanics of instruction.

Yet these shifts are not without risk. The very data that enables personalization—fine-grained records of student performance—raises questions about surveillance, profiling, and privacy. Algorithms can misclassify students, locking them prematurely into tracks or labels that shape their futures. The reliance on proprietary platforms also risks creating dependency on commercial vendors, potentially embedding biases into national education systems. Civil society has voiced concerns that technology designed to equalize could deepen divides if connectivity and devices are not universally accessible. Indeed, even in AI-rich education environments, ensuring equitable access to digital tools remains a challenge.

The promise of AI in education is clear: better learning outcomes, reduced dropouts, stronger alignment with the labor market, and greater inclusion for rural and underserved groups. But these benefits will only be realized if risks are acknowledged and managed. Governments need to ensure transparency in how AI assigns, assesses, or streams students; guarantee that every child has access to the digital infrastructure required to benefit from AI tools; and resist overreliance on private vendors by adopting open standards and, where possible, open-source AI educational resources.

**Benefits:** Personalized learning at scale; improved teacher productivity; early intervention for at-risk students; responsive lifelong reskilling systems.

**Risks:** Bias in learner profiling; data privacy breaches; exclusion of those without access to devices or connectivity; dependence on proprietary platforms.

**Actions:** Invest in national AI teaching co-pilots that support (rather than replace) educators; guarantee universal digital access; require transparency and appeal mechanisms in AI-driven student decisions; establish public oversight of educational AI to preserve equity and trust.

# Chapter 2: Healthcare & Public Health

Healthcare sits at the heart of the social contract. Citizens expect their governments to protect them against illness, provide care in times of need, and ensure access regardless of wealth. Yet health systems everywhere face the same strain: rising costs, aging populations, workforce shortages, and persistent inequalities in access. Artificial Intelligence is increasingly being woven into the fabric of public health delivery, promising to shift systems from being reactive and episodic to being predictive, preventive, and continuous.

At its core, AI in healthcare enables governments to move from treatment to anticipation. Algorithms trained on imaging data can detect cancers or cardiovascular risks earlier and more accurately than human specialists. Predictive models fed by electronic health records can flag patients at risk of sepsis or hospital readmission, allowing interventions before crises occur. During the COVID-19 pandemic, AI-driven epidemiological models helped governments predict outbreaks, target lockdowns, and allocate resources—foreshadowing a future where health ministries operate as much on forecasts as on clinical reports.

The technology stack behind this transformation is wide and layered. Computer vision systems read X-rays and MRIs, identifying anomalies in seconds. Natural language processing extracts insights from medical notes and unstructured health records. AI chatbots and triage tools handle basic patient queries, easing pressure on call centers and clinics while providing reassurance to citizens around the clock. Wearables and remote monitoring devices feed real-time streams of patient data into predictive models, shifting care into homes rather than hospitals. At the system level, digital twins of hospitals and even entire health networks allow governments to simulate resource allocation under different scenarios, improving preparedness for pandemics and surges.

For patients, these capabilities mean faster diagnoses, shorter waiting times, and expanded access to care—especially for those in rural areas or underserved communities where telemedicine paired with AI helps fill gaps left by scarce doctors. For clinicians, AI offers support in the most time-intensive tasks: reading scans, documenting notes, drafting discharge summaries. Yet many doctors worry about liability: if an AI flags a cancer that a later biopsy disproves, or misses a diagnosis that a human might have caught, who is accountable? For policymakers, AI promises a more efficient and data-driven health system, but also raises

thorny questions about privacy, secondary use of sensitive health data, and the risk of public backlash if citizens feel their health records are being exploited.

The societal impacts are profound. A well-governed AI health system could reduce mortality through earlier detection, extend access to populations historically excluded, and lower overall costs by shifting care upstream. But a poorly governed system could produce the opposite—inequities amplified by biased training data, citizens deterred from seeking care out of fear of surveillance, and clinicians demoralized by opaque black-box tools imposed on their practice.

The promise, however, remains too significant to ignore. Governments can take practical steps to tilt outcomes toward trust and value. They can establish national model evaluation units to rigorously test and certify clinical AI tools for accuracy and bias before deployment. They can create health data trusts with clear consent frameworks, giving citizens visibility and control over how their data is used. They can adjust reimbursement systems so that AI-augmented services—from tele-triage to remote diagnostics—are recognized and incentivized rather than treated as unregulated add-ons.

**Benefits:** Earlier and more accurate diagnoses; improved access to care in underserved areas; reduced costs from unnecessary hospital admissions; stronger public health forecasting and epidemic prevention.

**Risks:** Data privacy violations; opaque algorithms influencing life-and-death decisions; unclear liability for AI errors; public mistrust if health data is misused.

**Actions:** Establish model evaluation units for clinical AI; legislate health data trusts with citizen consent dashboards; update reimbursement codes to integrate AI-driven care; guarantee appeal routes and human oversight for AI-influenced clinical decisions.

# Chapter 3: Social Protection & Welfare

Social protection is where the state meets its citizens at their most vulnerable—in times of unemployment, disability, poverty, or crisis. Welfare programs embody the promise of security and inclusion, but they are also among the most administratively complex parts of government. Eligibility criteria, fraud prevention, case management, and disbursement all require immense resources, often leading to delays, errors, and citizen frustration. AI offers a powerful tool to make welfare systems more responsive, targeted, and efficient—but it also tests the limits of fairness and accountability in government decision-making.

The appeal of AI in welfare lies in its capacity to turn reactive systems into proactive ones. Instead of waiting for citizens to apply - often through lengthy forms and bureaucratic hurdles - AI can cross-reference tax data, employment records, and social service databases to identify those who qualify for assistance automatically. Governments in Europe and Latin America have begun experimenting with predictive analytics to detect households at risk of poverty, enabling early intervention before hardship becomes entrenched. In the United States, state-level agencies use AI models to flag fraudulent claims in unemployment insurance, cutting leakage and saving billions of dollars.

The technology stack enabling this shift spans anomaly detection systems that identify suspicious claims patterns, machine learning classifiers that assess eligibility in near real time, and integrated citizen service portals that offer personalized information about benefits. Behind the scenes, graph analytics tools map relationships between claimants, employers, and dependents to uncover hidden fraud networks. Increasingly, conversational AI is deployed to guide citizens through the claims process, reducing errors and improving accessibility for those with low literacy or language barriers.

For citizens, the experience could be revolutionary: faster approval, fewer forms, and benefits delivered directly without repeated verification. For caseworkers, AI-assisted triage means lower caseloads and more time to focus on complex cases that require human judgment. For policymakers, the promise is efficiency, reduced leakage, and better targeting of limited resources. Yet the risks are equally stark. Automated systems have been known to wrongfully deny benefits, leaving vulnerable people without recourse. Eligibility models trained on historical data can embed discrimination, disproportionately flagging certain communities as

fraudulent. The opacity of algorithms can make it difficult for citizens to understand or appeal decisions, eroding trust in institutions designed to protect them.

The societal implications of this transformation cut both ways. A well-governed AI welfare system could expand inclusion, reduce poverty traps, and restore citizen confidence in the state's ability to deliver support. A poorly governed one could deepen inequality, stigmatize vulnerable groups, and undermine the legitimacy of social protection itself. To avoid these pitfalls, governments must embed accountability into every stage of welfare automation. This means mandating algorithmic impact assessments for welfare AI systems, ensuring that citizens retain a right to explanations and appeals for any AI-driven decision, and keeping analog channels open for those who cannot or will not use digital services. Independent audits of welfare AI should measure not only fraud detection rates but also false positives and demographic fairness. Civil society organizations can play a vital oversight role, ensuring that technology designed to include does not inadvertently exclude.

**Benefits:** Faster and more accessible benefit delivery; proactive outreach to at-risk households; reduced fraud and leakage; lower administrative costs.

**Risks:** Automated wrongful denials with little recourse; embedded bias against vulnerable communities; lack of transparency in eligibility decisions; erosion of trust if citizens feel the system is a "black box."

**Actions:** Require algorithmic impact assessments for welfare AI; guarantee citizens' rights to understand decisions and appeal; maintain in-person/analog service options for digital have-nots; mandate independent fairness audits of welfare algorithms.

# Chapter 4: Public Safety, Justice & Borders

The provision of safety and justice is one of the most fundamental obligations of government. Citizens expect their states to protect them from crime, uphold the rule of law, and manage borders fairly and securely. Yet these functions are also among the most resource-intensive and politically sensitive. Courts face backlogs stretching years; police forces struggle to allocate officers effectively; border agencies juggle rising flows of people and goods. Artificial Intelligence is entering this space with the promise of greater efficiency, faster justice, and more effective prevention. At the same time, it brings some of the gravest risks to civil liberties and public trust.

In law enforcement, AI is being used to shift policing from patrol-based models to intelligence-led strategies. Predictive algorithms trained on historical crime data can suggest where resources should be deployed, aiming to prevent incidents before they occur. Cities in the United States and Europe have experimented with such models, though controversy has followed: the same systems that **reduce burglaries** can also entrench discriminatory policing if trained on biased arrest records. In Dubai, AI-enhanced surveillance networks allow police to monitor crowds in real time, automatically detecting unusual behavior or traffic violations.

Courts are also beginning to use AI to ease their crushing administrative burden. In Estonia, an AI "robot judge" has been piloted for small claims disputes, issuing decisions subject to human appeal. Elsewhere, natural language processing is used to sift through vast volumes of case law, helping clerks and judges identify relevant precedents. Translation AI reduces language barriers for migrants navigating the justice system. These tools speed up proceedings, but they also challenge notions of judicial accountability: if an algorithm contributes to a ruling, who ultimately bears responsibility for the outcome?

Borders represent another frontier. AI systems already screen travelers using facial recognition, risk scoring, and advanced document verification. In the United States, Customs and Border Protection deploys machine learning models to flag high-risk cargo for inspection. In Singapore, automated immigration counters use biometric AI to handle vast flows of travelers with minimal delays. These systems are efficient, but they also provoke questions of proportionality, fairness, and surveillance creep.

The technology stack underpinning these shifts spans computer vision for facial recognition and CCTV surveillance feeds, predictive analytics for crime pattern modeling, natural language

processing for sorting case documents, and multimodal biometrics for border management. Increasingly, digital twins of cities are used to simulate crowd flows for policing and emergency planning, while AI-powered translation ensures accessibility in multicultural legal contexts.

The potential benefits are clear: faster case resolution, better allocation of police resources, safer borders with smoother processing, and greater accessibility of legal services (e.g. instant translation or AI-guided legal aid). Yet the risks are profound. A biased predictive policing model can unfairly target communities that are already over-policed. Automated border systems can wrongly flag innocent travelers with little recourse. Opaque AI involvement in court decisions risks undermining the legitimacy of justice itself. Public safety and justice require not only effectiveness but also fairness, and AI makes balancing these principles more challenging.

Governments can mitigate these risks through statutory safeguards. For example, real-time biometric identification (facial recognition surveillance) should be tightly limited to contexts explicitly authorized by law or courts. All AI-derived evidence in policing or trials should meet clear admissibility standards. Algorithms used in justice should be logged in public registries to allow defense and public scrutiny. Independent oversight boards - including civil society and community representatives - are essential to scrutinize deployment and impacts. Sunset clauses on high-risk tools can ensure that technologies are periodically reviewed rather than silently entrenched if they prove problematic.

**Benefits:** Faster resolution of court cases; better allocation of police and border patrol resources; improved border security and throughput; greater accessibility of legal information and services via AI translation and document review.

**Risks:** Bias and discrimination in predictive policing; erosion of privacy through mass surveillance; wrongful detentions or denials at borders due to algorithmic error; loss of legitimacy in judicial decision-making if AI is unaccountable.

**Actions:** Enact clear legal safeguards limiting use of facial recognition and biometric surveillance; require algorithmic transparency and public registries for AI in justice; set

evidentiary standards for AI-derived insights in investigations and trials; establish independent oversight and periodic review of all public safety AI systems.

# Chapter 5: Transport, Mobility & Infrastructure

Transport and infrastructure are the arteries of modern economies. Efficient movement of people, goods, and resources underpins productivity and quality of life. Governments have long struggled with congestion, maintenance backlogs, and uneven service coverage—challenges that traditional planning cycles and static models have failed to resolve. Artificial Intelligence is beginning to change this, enabling systems that are adaptive, predictive, and continuously optimized rather than fixed and reactive.

AI in transport transforms how cities manage flows. In Dubai, AI-powered traffic signals respond dynamically to congestion, reducing waiting times and fuel consumption. In China, Alibaba's City Brain platform analyzes data from traffic cameras, ride-hailing apps, and public transit to orchestrate vehicle movement across entire districts, cutting emergency response times by minutes that can save lives. Singapore's Land Transport Authority uses predictive models to anticipate ridership peaks and adjust bus and train frequencies in real time. These examples show how AI allows governments to move beyond merely reacting to breakdowns, and instead actively shape the flows of people and goods.

The same predictive capabilities extend to infrastructure maintenance. AI models trained on sensor data can identify signs of stress in bridges, pipes, or railways before failure occurs, prioritizing repairs where risks are highest. For utilities, AI forecasts energy and water demand, allowing grids and reservoirs to be balanced more efficiently. Some cities are building digital twins—virtual models of entire transport or utility systems—to simulate disruptions, test policy scenarios, and stress-test infrastructure against climate risks.

The technology stack behind these changes includes IoT sensors embedded in roads, vehicles, and public utilities; computer vision systems monitoring traffic and infrastructure integrity; machine learning models predicting failures or bottlenecks; and orchestration platforms that integrate multiple data streams into real-time decisions. Cloud computing and edge AI provide the computational backbone, while digital twins serve as decision-support environments for planners.

For citizens, these capabilities mean shorter commutes, fewer delays, and improved accessibility. In underserved areas, AI-guided micro-transit systems can fill service gaps, dynamically routing vehicles to where they are most needed. For operators and engineers, AI reduces outages, improves asset management, and enables smarter long-term investments. For

governments, the result is both economic—higher productivity, lower maintenance costs—and political, as visible improvements in daily mobility reinforce public trust in government services.

The risks, however, are real. Continuous monitoring through sensors and cameras raises privacy concerns if data is not properly anonymized. AI systems trained on historical traffic data may inadvertently prioritize affluent districts over poorer neighborhoods, embedding inequities into service provision. Dependence on complex vendor systems could leave governments locked into proprietary platforms with limited control or negotiating power. And as recent cyberattacks on critical infrastructure have shown, greater connectivity also increases vulnerability to hacking or systemic failure.

To maximize benefits, governments should pursue open standards for mobility and infrastructure data, and adopt privacy-preserving analytics to protect citizens. Contracts with vendors must include equity metrics—ensuring that improvements (like reduced commute times) are fairly distributed across communities, not just in wealthier areas. Digital twin pilot projects should publish their findings openly, demonstrating both the benefits and the trade-offs of AI-driven optimizations. Above all, resilience must be designed in: AI-managed transport and utility systems need manual fallback modes and robust cybersecurity, given their critical role in daily life.

**Benefits:** Reduced congestion and travel times; lower vehicle emissions due to optimized traffic flow; predictive maintenance that prevents infrastructure failures; better allocation of public transit resources; improved accessibility in underserved areas through dynamic routing.

**Risks:** Privacy violations from pervasive sensors; inequitable service distribution if AI optimizes for efficiency over equity; vendor lock-in with proprietary smart city platforms; heightened cyber vulnerabilities in connected infrastructure.

**Actions:** Adopt open data standards for mobility and infrastructure; mandate that AI systems in transport meet equity and transparency benchmarks; publish results of digital twin simulations for public accountability; embed resilience by requiring fail-safes and cybersecurity measures in all AI-enabled infrastructure systems.

# Chapter 6: Housing, Planning & Environment

The planning and regulation of housing, buildings, and environmental resources has always been a slow, paper-heavy process. Permit applications take months, building inspections lag behind reality, and environmental monitoring often relies on sparse or outdated data. Citizens and developers alike experience these bottlenecks as friction that erodes trust in government capacity. Artificial Intelligence is beginning to rewire this space by introducing automated compliance, predictive oversight, and climate-aware planning into the way states regulate the built and natural environment.

AI is first transforming permitting and approvals. In Singapore, the **CORENET X** platform uses AI to pre-check digital building models against regulatory codes before they reach human reviewers. Errors that once triggered weeks of back-and-forth can now be caught in minutes, drastically shortening approval times. In Los Angeles, e-permitting systems augmented with AI are being piloted to validate zoning compliance, allowing governments to streamline approvals while focusing human inspectors on complex cases. These systems not only accelerate development but also improve predictability for both citizens and businesses.

Beyond permits, AI is reshaping urban inspections and safety enforcement. Computer vision systems analyze drone or CCTV footage to spot illegal construction, fire hazards, or unsafe crowding in real time. Predictive models identify which buildings are most likely to face structural stress or code violations, enabling inspectors to prioritize scarce resources. In environmental management, governments are deploying AI to monitor air quality, water flows, and deforestation in near real time, often using satellite imagery. This shift makes compliance proactive: regulators can act on problems as they emerge rather than waiting for complaints or disasters.

The technology stack includes building information modeling (BIM) integrated with AI rule-checkers, geospatial analytics that combine satellite and sensor data, computer vision for automated inspection, predictive risk models for infrastructure stress, and increasingly, digital twins of cities that combine environmental, housing, and infrastructure data into unified simulation platforms. These systems allow policymakers to model the impacts of climate change—such as flooding or heat waves—on urban layouts and building stock, embedding resilience into future planning decisions.

For citizens, these capabilities mean faster approvals for home improvements or construction, safer buildings, and more responsive environmental protection (for instance, catching pollution incidents early). For developers, they bring greater clarity and consistency, reducing costly delays and uncertainty. For regulators, AI enables scarce staff to focus on complex, high-stakes issues rather than routine box-checking. For policymakers, AI turns planning into a living process—one where environmental, housing, and social data can be modeled together to support better long-term choices about land use and resource management.

The risks are tied to fairness and legitimacy. Automated permitting systems can deny applications without clear explanation, leaving individual homeowners or small developers frustrated and at a disadvantage compared to those who can navigate the algorithms. If AI is trained on historical zoning or housing patterns, it may simply embed past biases or exclusionary practices into future decisions (for example, reinforcing patterns of segregation). There is also the danger of over-centralization: when all decisions flow through opaque AI systems, the discretionary space for negotiation, appeal, or citizen input may narrow, potentially undermining the democratic aspect of planning processes.

To avoid these pitfalls, governments must embed transparency into planning AI. Every automated decision should include an explanation traceable to the regulatory logic (and that logic should be openly published). A clear right to appeal must be guaranteed for those who feel a permit or decision was unfairly denied by an algorithm. Building codes and environmental rules should be published as open digital standards so that AI validations can be independently verified by third parties. And climate and sustainability criteria—such as energy efficiency standards or flood resilience requirements—should be formally integrated into automated checks, ensuring that AI-driven planning aligns with broader environmental goals.

**Benefits:** Faster and more predictable permitting and approvals; improved compliance and building safety through targeted inspections; proactive environmental monitoring and enforcement; climate-resilient urban planning guided by simulations.

**Risks:** Automated denials without explanation or recourse; reinforcement of discriminatory historical patterns in housing; over-centralization of decision-making with reduced public input; potential for “algorithmic zoning” that lacks flexibility or empathy for unique cases.

**Actions:** Require transparency and a right to appeal for AI-made planning decisions; publish regulatory codes in machine-readable open formats to enable external review of AI rules; embed climate and equity criteria into algorithmic decision rules; maintain hybrid systems where human planners can override or adjust AI recommendations when warranted, to ensure accountability.

# Chapter 7: Taxation & Revenue Services

Taxation is the lifeblood of the state, financing every public good from schools to hospitals. Yet for most citizens, it is also one of the most frustrating interactions with government: complex forms, opaque audits, and long delays for refunds or rulings. For governments, the challenge is twofold—maximizing compliance and revenue while minimizing administrative costs and taxpayer burden. AI is proving to be a natural fit in this space, offering tools that simplify compliance for taxpayers, sharpen risk detection for authorities, and increase trust in fiscal systems when implemented with transparency.

The most visible shift has been the rise of pre-filled and AI-assisted tax filings. In Estonia, Spain, and several Latin American countries, governments now issue citizens tax returns that are already completed by the system, leaving only confirmation or minor adjustments by the taxpayer. Natural language processing and conversational AI are used to guide taxpayers through filing, answering questions in real time and reducing errors. For small businesses, AI-enabled accounting portals can automatically reconcile invoices, payroll data, and expenses, dramatically lowering compliance burdens.

For governments, AI brings precision to the other side of the ledger: fraud detection and smarter audits. Machine learning models trained on historical filings flag anomalies—such as patterns of underreporting or suspicious deductions—allowing auditors to focus on the cases most likely to yield findings. Graph analytics uncover networks of shell companies or coordinated VAT evasion. In Canada, the revenue agency's AI models have improved both detection rates and speed, saving auditors from trawling through millions of filings manually.

The technology stack enabling these advances includes anomaly detection engines, machine learning classifiers for risk scoring, NLP-powered virtual assistants, and integrated tax platforms linked to national identity systems. Increasingly, cloud-based systems allow governments to unify tax data across agencies, while blockchain-backed ledgers are being piloted for transparent VAT collection and invoice verification, adding trust to the system of record.

For citizens, the benefits are immediate: less time spent filing, fewer errors, and greater clarity about one's obligations. For small businesses, automated reconciliation reduces both compliance costs and the risk of penalties. For governments, the results include higher voluntary compliance, faster detection of fraud, and improved fiscal yield. But these systems

are not without danger. Risk-scoring models can be opaque, flagging certain taxpayers or transactions without clear explanation. If citizens feel they are targeted unfairly by "black box" algorithms, trust in the tax system can erode quickly. Small enterprises might feel that the same anomaly detection tools that catch big fraud also subject them to disproportionate scrutiny for minor mistakes, especially if appeal processes are cumbersome.

Ultimately, AI in taxation illustrates both the promise and peril of automation in public goods. When well-designed, it increases both efficiency *and* fairness, ensuring that honest taxpayers comply with ease and that evasion or fraud is rapidly detected. When poorly designed, it risks creating fear and resentment, widening the gap between state and citizen in an area (taxes) that is already sensitive.

**Benefits:** Reduced compliance burdens for citizens and SMEs; easier filing (e.g., pre-filled returns); faster refunds and responses; higher voluntary compliance and revenue collection; more effective fraud detection focusing on high-risk cases.

**Risks:** Opaque risk scoring and selection for audit; wrongful flagging of compliant taxpayers; disproportionate burden or suspicion on small firms; erosion of trust if AI-driven notices or audits feel like a "big brother" with no explanation.

**Actions:** Publish understandable summaries of factors that trigger audits or investigations (to demystify the AI scoring); guarantee taxpayers the right to appeal decisions and receive an explanation for any AI-flagged irregularity; adopt secure, interoperable data systems that protect privacy; ensure independent validation of tax AI models for fairness and accuracy (e.g., check that error rates are not higher for certain income or demographic groups).

# Chapter 8: Labor & Employment Services

Employment is both an economic necessity and a marker of dignity. Governments invest heavily in labor and employment services to connect workers with jobs, provide training, and manage unemployment benefits. Yet traditional systems are often sluggish, reactive, and poorly aligned with the fast-changing realities of the labor market. AI promises to rewire this domain, enabling governments to create dynamic, skills-based, and personalized employment systems that respond in real time to economic shifts.

The greatest shift lies in moving from static job boards to AI-driven matching platforms. Instead of merely listing vacancies and leaving workers to search, machine learning models can analyze an individual's skills, past employment, and even micro-credentials to suggest tailored job opportunities. In Singapore, for example, the national SkillsFuture platform uses AI to guide workers toward courses and jobs that align with both their existing capabilities and projected demand. In China, provincial employment platforms use natural language processing to match migrant workers to opportunities across sectors, reducing friction in labor mobility.

AI also enables predictive workforce planning. Governments can analyze trends across industries to forecast where skill shortages are likely—whether in nurses, data engineers, or teachers—and proactively adjust training subsidies or education programs to close those gaps. For individuals, AI can provide career path simulations, showing how different upskilling choices might translate into future wages and employability, thereby helping jobseekers make informed decisions.

The technology stack behind this transformation includes recommendation engines that map worker profiles to vacancies, NLP systems to parse CVs and job descriptions, graph-based models that capture relationships between skills and occupations, and predictive analytics that forecast demand for certain roles. Increasingly, these systems are integrated into digital labor exchanges linked to national ID or social security systems, ensuring continuity between welfare, training, and job placement services (for instance, someone coming off unemployment benefits can be directly nudged toward reskilling programs and new job matches).

For workers, these advances mean better matches and shorter job searches—finding the "right job" faster. It can also mean clearer guidance on what new skills to acquire for future opportunities. For caseworkers in public employment centers, AI-assisted triage can reduce caseloads, allowing staff to focus on coaching and support rather than paperwork and database

searches. For employers, the result is access to a more accurately matched talent pool, potentially reducing hiring times and improving candidate fit. And for governments, the payoff is both economic—faster re-employment reduces welfare outlays and boosts tax revenues—and political, as citizens experience more responsive support during times of job transition or economic disruption.

But the risks are serious. Matching algorithms can reproduce historical biases in hiring data, systematically disadvantaging groups (like women, minorities, or older workers) if not carefully designed. Risk scores used to prioritize who gets extra help might inadvertently exclude those who need help the most (e.g., an algorithm might decide some unemployed persons are "unemployable" and thus not worth investing training resources in—creating a self-fulfilling prophecy). For citizens, being "scored out" of opportunities by an opaque system can feel dehumanizing, especially if appeal mechanisms are weak. There's also a question of fairness in transparency: should employers have access to the same predictive scores about workers that governments use internally? If not handled properly, trust in public labor services could suffer.

To govern responsibly, states must ensure these systems are designed with explainability and fairness at their core. Matching algorithms should be subject to independent evaluation, and metrics about their outcomes should be published to show whether AI is improving equity or not. Workers should have clear control over their data and the right to opt out of certain kinds of algorithmic inference (for example, if a platform is inferring things about one's aptitude or risk level, a worker should know and consent). Training subsidies should be tied not simply to enrollment but to measurable skills gains validated by independent assessments, ensuring that AI is guiding people to programs that truly improve employment prospects, rather than just shuffling the unemployed through courses.

**Benefits:** Faster reemployment; better matches between workers and jobs; more efficient targeting of training funds to where they'll have impact; higher overall productivity as skills gaps are closed more quickly.

**Risks:** Bias in matching algorithms that reflect or reinforce discrimination; exclusion of vulnerable groups if AI deems them "low potential"; opacity in how people are ranked or scored; potential misuse or over-reliance on worker data (privacy concerns and fairness).

**Actions:** Mandate explainability in job-matching AI (workers should be able to understand why certain jobs or training were recommended or not); ensure workers retain control over their profiles and data (including correction or deletion rights); regularly evaluate and publish the equity outcomes of AI-driven employment services (e.g., are placements improving for historically disadvantaged groups?); tie government training subsidies to demonstrable skill acquisition and employment outcomes, not just course attendance, using independent assessments to validate AI recommendations.

# Chapter 9: Civic Participation, Information & Democracy

Democracy depends on informed citizens and accessible channels for participation. Yet governments often struggle to engage the public meaningfully, especially in an era of information overload, disinformation, and declining trust in institutions. Artificial Intelligence has the potential to expand civic participation by making information more accessible, government communication more responsive, and consultation processes more inclusive. At the same time, it risks tipping into manipulation, surveillance, or erosion of democratic norms if misused.

AI is already reshaping how governments communicate with citizens. Natural language models can translate laws and regulations into plain, multilingual summaries, lowering barriers for non-experts to understand complex policies. Parliamentary transcripts are increasingly processed by AI for fully searchable, real-time public access. In Taiwan, civic platforms like vTaiwan use AI-based clustering to synthesize thousands of citizen comments into coherent themes for policymakers, improving the quality of deliberation. In Europe, AI-enabled platforms have been trialed to parse feedback from large-scale public consultations, ensuring that every voice is captured rather than drowned out by sheer volume.

The technology stack for civic engagement includes natural language processing to generate accessible summaries of policy documents, sentiment analysis tools that aggregate citizen feedback, machine translation engines to break language barriers, and clustering algorithms that surface key themes in crowdsourced input. Increasingly, governments are experimenting with generative AI to draft first-pass responses to common public inquiries or comments, which human officials can then refine—speeding up response times without replacing the human judgment needed for sensitive issues.

For citizens, these tools mean clearer access to information (less bureaucratic jargon), less frustration in finding answers, and more effective channels to be heard (since AI can help ensure each comment is analyzed and not ignored). For governments, they mean a sharper signal from the noise of public discourse—being able to distill thousands of comments into actionable insights—and the ability to demonstrate responsiveness at scale. But they also come with dangers: the same generative models that can inform could also be used to *persuade* or even manipulate. For instance, micro-targeted messages crafted by AI could border on propaganda if used to shape public opinion rather than respond to it. AI-driven moderation of

online discourse might inadvertently (or deliberately) suppress dissenting voices in the name of maintaining "civility" or reducing misinformation. The very tools designed to strengthen democracy could, without safeguards, undermine it.

The societal impact therefore depends on governance choices. When implemented transparently, AI can lower the costs of participation (e.g., making it easier for anyone to understand a proposed policy), reduce the impact of misinformation by rapidly debunking or clarifying false claims, and bring more diverse voices into policy debates (since AI can translate and categorize inputs from multiple languages or regions). But if used without oversight, AI in civic tech could corrode public trust—especially if citizens suspect the government is using AI to surveil discussions or to automate astroturfing (fake grassroots support).

Governments must set clear boundaries. AI-generated communication should be *disclosed* or watermarked, so people know when they're reading content created by a machine. Consultation platforms should publish open datasets showing what inputs were received and how the AI aggregated them, enabling civil society and journalists to validate the results and check for bias. Civic literacy programs should teach citizens how to critically interpret AI-mediated information (for example, knowing that a quick summary might miss nuances). Above all, persuasion must never be allowed to replace genuine participation: AI should serve as an amplifier of public voices, not a substitute for them. The goal is to augment democratic engagement, not automate it.

**Benefits:** Greater accessibility of government information; improved synthesis of citizen feedback (policymakers can actually digest what large crowds are saying); more inclusive consultation processes that scale to millions of inputs; reduced misinformation through transparent AI-generated summaries and fact-checks.

**Risks:** Manipulative or biased use of generative AI for government propaganda; suppression of dissent under the guise of automated content moderation; erosion of trust if citizens feel AI is being used to "manage" them rather than empower them; privacy concerns if AI is surveilling public sentiment too closely.

**Actions:** Mandate disclosure of AI-generated text or chat responses in government communications; publish raw and processed consultation data for public auditing, along with

AI algorithms' methodology, to ensure accountability; invest in civic AI literacy so the public understands these tools; establish ethical rules (possibly legislation) preventing authorities from using AI primarily for one-sided persuasion or mass surveillance of legitimate democratic debate.

# Chapter 10: Cross-Cutting Enablers & Governance

While sector-specific applications of AI—in education, health, welfare, taxation, transport, etc.—generate visible outcomes for citizens, the ability of governments to harness AI sustainably depends on a set of cross-cutting enablers. These are the invisible scaffolds of an AI-powered state: **data infrastructure**, **institutional capacity**, **procurement frameworks**, **workforce readiness**, and **trust mechanisms**. Without them, even the most sophisticated AI tools risk remaining isolated pilots that fail to scale or fail to win public legitimacy.

**Data infrastructure** is the first enabler. AI systems depend on high-quality, interoperable data—yet most governments still operate in silos. *Estonia's X-Road* demonstrates how a secure, cross-agency digital backbone can allow AI to function seamlessly across domains by enabling data sharing under strict governance. By contrast, in many countries, fragmented data environments and legacy IT systems make integration slow and costly. Privacy-preserving analytics and data trusts are emerging as models to balance openness with protection of personal data, but adoption is uneven.

Second is **institutional capacity**. Governments need structures that can oversee, evaluate, and regulate AI use. This means not just appointing Chief Data or AI Officers or creating digital innovation teams, but also establishing mechanisms like model registries (public lists of algorithms in use), independent audit units to probe algorithmic decisions, and "red-teaming" practices that stress-test AI systems before deployment. Albania's experiment with "ADA" (nicknamed an AI procurement minister) highlights both the promise and risks of radical innovation: it demonstrates ambition to automate government decisions, but raises unresolved questions about legitimacy and accountability when a symbolic AI "official" takes on roles traditionally reserved for elected or appointed humans.

The third enabler is **procurement**. AI in the public sector is often delivered through private vendors, yet traditional procurement processes are ill-suited for fast-moving technologies. Governments that procure AI as if it were concrete or paper clips risk outcomes like vendor lock-in, biased solutions, or black-box systems they don't truly control. Forward-looking states are shifting toward outcome-based and performance-based procurement, requiring vendors to demonstrate attributes like fairness, explainability, and interoperability of their AI systems rather than just delivering a piece of software. This might include asking for access to training data or algorithms for auditing as part of contracts, or mandating use of open standards.

Fourth, the **workforce** must be prepared for co-existence with AI. Teachers, doctors, caseworkers, analysts—frontline public servants need not only new digital skills, but also clarity about how their roles will evolve. If AI is seen as a tool that *supports* professionals, it can strengthen morale and performance. If it's seen (or sold) as a replacement, resistance will rise and adoption will falter. Change management and continuous re-skilling programs are therefore as vital as the technology itself. For example, training programs for civil servants on AI literacy, and creating new roles (like "AI ethicist" or "automation coordinator" within agencies) can help integrate AI smoothly. Governments also need to plan for transitions: certain routine jobs may be reduced, but new jobs analyzing or maintaining AI systems will appear. A just transition for the public workforce is essential to maintain trust and motivation.

Finally, **trust mechanisms** must underpin every deployment. Citizens will only embrace AI in public services if they believe the systems are fair, transparent, and accountable. This requires measures such as algorithmic registries (so the public knows what automated systems are in use), explainability for high-stakes decisions (e.g., a person denied a service by an algorithm is given a clear reason), guaranteed appeal rights (no AI decision should be final without human review), and public reporting of AI system performance and error rates. It also requires political leadership to set clear "red lines"—for example, prohibiting certain uses of AI that are deemed too risky. Many countries are considering or enacting bans on real-time face recognition surveillance in public spaces, or prohibiting the use of AI for social scoring of citizens, to draw boundaries around acceptable use.

The benefits of getting these enablers right are systemic: faster scaling of effective AI across government, improved efficiency through shared infrastructure, greater citizen trust (because of visible safeguards), and reduced risk of vendor capture or techno-solutionism. The risks of neglecting them are equally systemic: fragmented pilot projects that never integrate, public backlash at the first serious mistake, inequities embedded into critical systems without oversight, and erosion of legitimacy if people feel AI is something done *to* them by an unaccountable state.

**Benefits:** Scalable and interoperable AI solutions across government agencies; improved public trust through transparency and accountability; more efficient and fair procurement results; a public sector workforce that is empowered by AI rather than displaced by it.

**Risks:** Vendor lock-in and dependency if governments lack bargaining power or internal skills; a proliferation of pilot projects that don't translate to enterprise solutions; citizen mistrust or even active resistance if AI implementations cause scandals; workforce resistance and low morale if automation is handled poorly; symbolic adoption (like appointing an "AI minister" with no real plan) without substance, which can waste resources and credibility.

**Actions:** Invest in cross-agency data backbones with strong privacy and security safeguards; create independent AI evaluation units and public algorithm registries; reform procurement rules toward agile, outcomes-focused approaches (including requiring auditability of vendor AI); embed continuous workforce reskilling and change management for civil servants; and **codify red lines** around high-risk AI uses (e.g. via legislation or binding policies that forbid mass surveillance, secret profiling, or other abuses).

# Chapter 11: AI-Augmented Policymaking & Regulatory Design

Beyond service delivery, AI is starting to **augment the policy and regulatory process itself**. Governments traditionally craft policies and regulations through labor-intensive analysis, stakeholder consultation, and incremental updates. AI offers tools to make policymaking more data-driven, adaptive, and responsive, potentially shifting from a *"regulate and forget"* model to an *"adapt and learn"* approach (OECD, 2024). This section explores how AI can assist in designing better rules, forecasting the effects of policies, and improving regulatory enforcement in a technology-forward manner.

On the **policy design** side, AI can help policymakers navigate complex legal and economic landscapes. Modern regulatory systems consist of thousands of pages of legislation and rules, often with overlapping or outdated provisions. AI systems can rapidly search and summarize this "stock" of regulations, identifying gaps, overlaps, or inconsistencies that humans might miss (OECD, 2024). This can enable more informed and targeted updates to laws. For example, an AI analysis might reveal that two agencies have regulations that conflict, or that a certain industry practice is unregulated due to an overlooked loophole. By spotting patterns in large datasets (legal texts, economic data, public feedback), AI can highlight areas where policy intervention could yield benefits.

AI can also **automate routine policymaking tasks**, creating efficiencies in processes like regulatory impact analysis (RIA) and even initial drafting. For instance, natural language generation tools can produce first drafts of simple regulations or policy documents based on parameters given by humans (some legislators have already tested having AI draft proclamations (IE University – Center for the Governance of Change, 2024). While human experts must refine and approve any AI-drafted text, this can accelerate the workflow. Additionally, AI-driven simulations can help predict the outcomes of a proposed policy—similar to how weather models forecast storms, policy models could forecast, say, the impact of a new tax incentive on employment or the effect of a pollution regulation on emissions over time. This *anticipatory analysis* allows for experimentation in a virtual setting, so policymakers can iterate and improve design before implementation (OECD, 2024).

AI can strengthen **stakeholder engagement and public consultation** as well. By analyzing public comments, social media, or consultation submissions, AI can help regulators understand public sentiment and key concerns. As noted in Chapter 9, tools like clustering algorithms can

sift large volumes of feedback to identify common themes (OECD, 2024). This ensures that even open-ended feedback from thousands of citizens can be distilled into insights that inform regulatory adjustments, rather than being ignored or sampled non-scientifically.

On the **regulatory delivery and enforcement** side, AI can vastly improve how rules are implemented on the ground. Regulatory agencies are often inundated with compliance data safety reports, financial filings, inspection results. AI systems can monitor these streams and detect patterns of non-compliance or emerging risks in real time. For instance:

- **Risk-based Inspections:** AI can enhance risk models used to target inspections, increasing accuracy so that regulators inspect the entities most likely to be violating rules (OECD, 2024). If data (from sensors, reports, etc.) suggests a particular factory has anomalies in emissions, an AI can flag it for immediate inspection instead of waiting for an annual audit cycle.
- **Monitoring and Detection:** AI algorithms can monitor media, social networks, or company data to detect violations—such as businesses engaging in prohibited practices or emerging safety hazards. For example, if restaurants are required to maintain hygiene standards, an AI vision system might automatically detect rodent sightings in CCTV feeds or analyze online reviews for complaints about cleanliness, prompting a quicker response.
- **Automating Compliance Checks:** Much like the permitting examples in Chapter 6, AI can automate certain compliance approvals. Companies could submit data (financial statements, environmental measurements) and an AI system could automatically check compliance with thresholds or rules, only flagging human regulators when something is amiss. This reduces burdens on both regulators and the regulated entities.

Governments are already piloting such approaches. For instance, some financial regulators use AI to spot fraudulent transactions or signs of money laundering in real time, rather than relying purely on after-the-fact reports. Environmental regulators employ satellite imagery AI to detect illegal deforestation or unreported mining activities almost as they happen. **Digital regulation sandboxes** have emerged, where AI tools simulate how new fintech products would behave under existing rules, revealing whether regulations need updating to account for innovations.

However, **AI-augmented regulation** raises important governance challenges:

- **Accountability:** If an AI system makes a regulatory decision (e.g., approving a license or flagging a violation that leads to a penalty), who is accountable for errors? Automated decisions must remain subject to human oversight, and agencies need clear protocols for review.
- **Transparency:** The logic of regulatory AI tools should be explainable to businesses and the public. Otherwise, companies might face an inscrutable "AI regulator" and not know how to comply, undermining the rule of law. As the OECD notes, maintaining accountability and transparency in algorithmic regulatory decisions is critical (OECD, 2024).
- **Bias and Fairness:** Regulatory AI must be carefully tested for bias. If the data used to train an AI enforcement system contains bias (e.g., historically focusing on certain neighborhoods or businesses due to human bias), the AI could reinforce that unfairness. For example, a risk model might unfairly target small businesses or certain ethnic communities for inspections unless checks are in place (OECD, 2024).
- **Adaptive Regulation:** AI allows for more dynamic rules (like speeding up approvals as conditions change). But laws typically change slowly for stability and predictability. There is a balance to strike between adaptability and the need for clear, stable rules. "Adaptive regulation" might mean granting regulators more discretion to update technical requirements via AI insights, but the democratic process must still set the goals and guardrails.

In sum, **AI offers policymakers a richer toolkit**: the ability to mine vast information for smarter policy design, the power to simulate outcomes for forward-looking regulation, and the means to enforce rules more effectively and efficiently. Embracing these tools could lead to regulations that better achieve economic, social, and environmental outcomes while reducing unnecessary burdens. But to do so, governments must update their policymaking processes—integrating AI in a way that upholds accountability, transparency, and the rule of law. As one international analysis put it, digital technologies (AI included) can enable innovative, data-driven rulemaking that continuously learns and improves (OECD, 2024).

**Benefits:** More evidence-based and anticipatory policymaking; faster identification of regulatory gaps or overlaps; streamlined drafting and analysis processes; smarter, risk-based

enforcement that protects the public while easing up on compliant actors; ability to adapt regulations to emerging technologies and trends quickly.

**Risks:** Lack of accountability if AI's role in policy decisions is not transparent; potential bias in AI-generated policies or enforcement targeting; reduced human judgment in nuanced legal questions; possible over-reliance on automated analysis, missing the political or ethical dimensions that data can't capture.

**Actions:** Treat AI as a policy aide, not policymaker—ensure human officials set objectives and review AI suggestions; invest in AI systems to assist with regulatory research, but pair them with legal experts to interpret results; maintain public transparency about how AI is used in policymaking (e.g., publish if an AI tool was used to draft or analyze a bill); develop clear guidelines for algorithmic regulatory enforcement (including appeal rights for regulated parties and periodic audits of the AI's decisions). By doing so, governments can leverage AI to craft smarter rules and institutions for the future while reinforcing, not weakening, the principles of good governance.

# Chapter 12: Cross-Border and Multilateral Collaboration

AI's transformative impact on governance is not confined within national borders. Many of the challenges and opportunities presented by AI in public goods delivery are global in nature—ranging from cross-border data flows and regulatory standards to ethical considerations that transcend jurisdictions. This section examines how countries are working together (or in some cases, struggling to work together) in harnessing AI for public service delivery, and why **multilateral collaboration** is crucial for maximizing benefits and managing risks.

**Global public goods and shared challenges:** No country can "go it alone" in fully addressing issues like pandemics, climate change, cybersecurity, or large-scale migration. AI is being seen as a tool to tackle these global challenges, but doing so requires cooperation. For example, the G7 and the Global Partnership on AI (GPAI) have called for harnessing AI to advance the United Nations Sustainable Development Goals (Brookings Institution, 2023). Collaborative "moonshot" projects are envisioned where countries pool resources to develop AI solutions for climate modeling, disaster prediction, or agricultural productivity—problems that cross borders. An AI system that predicts flooding or disease outbreaks is far more powerful if it draws on data from multiple countries. We saw this during COVID-19, when sharing data and models internationally helped improve pandemic response. In the climate realm, countries are discussing shared AI early warning systems for extreme weather.

**Aligning rules and standards:** If every nation develops totally different AI regulations, companies and even governments face a fragmented environment that stifles innovation and raises costs. Consider a scenario where one country's public healthcare AI must meet stringent transparency rules, but a neighboring country has no such rules—the result could be hesitancy in adopting life-saving AI across borders, or difficulty in integrating systems (like contact-tracing apps or pandemic AI tools) between countries. **International regulatory cooperation** can reduce these frictions. Aligning key aspects of AI regulation (for instance, agreeing on baseline safety requirements or banning certain high-risk uses globally) can enable a larger market for AI solutions and prevent a race to the bottom (Brookings Institution, 2023) (Brookings Institution, 2023). Specialized AI firms could thrive if they can offer their tools across many countries without having to completely re-engineer for each legal regime. For example, common standards for *AI in medical devices* or *digital health records* would let an approved AI diagnostic tool be used in multiple countries' hospitals, benefiting more patients.

That said, complete global harmonization is unlikely—nations differ in legal traditions, values, and risk tolerances. The aim is **"co-opetition"**: cooperate on foundational principles and shared interests, while still competing to develop the best solutions and maintaining sovereignty (Brookings Institution, 2023). We see early moves toward this in efforts like the *OECD AI Principles* (a set of intergovernmental AI guidelines endorsed by over 60 countries) (OECD, 2019) and the ongoing work on an EU–US alignment for AI standards.

**Data flows and sovereignty:** International collaboration on AI often runs into the issue of data sharing across borders. Data is the key input for AI, but laws like Europe's GDPR or various data localization requirements in countries can make sharing data challenging. Enhanced cooperation in trade and data agreements is essential to avoid unjustified restrictions that would reduce the potential benefits of AI diffusion (Brookings Institution, 2023). Protectionist measures (e.g., barring certain data from leaving a country) might stem from legitimate privacy or security concerns, but they can also hinder collaborative projects. Finding ways to enable cross-border data sharing for public-good AI projects—while respecting privacy (perhaps through anonymization or federated learning where data stays in country but models travel)—is a focus of international talks.

**Capacity building and inclusive collaboration:** A notable risk in global AI cooperation is that a few technologically advanced countries set the agenda, while developing countries are left out. International forums have been dominated by North America, Europe, and East Asia. However, many developing nations have a huge stake in AI for public goods (imagine AI to improve agriculture in Africa or education in South Asia) but lack resources or influence in rule-making. Inclusive multilateral approaches are needed to ensure AI benefits "the many, not just the few" (UNCTAD, 2024). This means involving developing countries in creating AI norms and providing support for their capacity building. For instance, UNESCO and other agencies are launching initiatives to train civil servants in AI across Africa (UNESCO, 2025), and to develop open-source AI tools that poorer countries can adapt. UNESCO's recent effort plans to train more than **15,000 civil servants and 5,000 judicial officials in Africa** on AI use and governance (UNESCO, 2025), building a knowledge base so they can both use AI effectively and shape its rollout ethically. Such efforts help level the playing field so international collaboration isn't one-sided.

**Joint R&D and knowledge sharing:** Nations are increasingly creating bilateral and multilateral partnerships for AI research related to public sector challenges. For example, the

Nordic Council and Baltic countries share best practices on AI in government as part of their digital cooperation. The Digital Nations (a group of leading digital governments including Canada, UK, Estonia, Israel, New Zealand, etc.) have an ongoing exchange on AI projects and open-source solutions—essentially a club for sharing what works and even co-developing tools. The European Union itself pools resources for AI research in areas like healthcare and green tech under programs like Horizon Europe. There have been proposals for a **Multilateral AI Research Institute** that would bring together top scientists and engineers from around the world to work on big societal challenges, with results shared openly (UNCTAD, 2024). Such an institute could focus on AI for climate, humanitarian response, or global health.

**Addressing global risks collaboratively:** There is also a defensive side to multilateral AI cooperation—managing risks that no single country can handle alone. For example, AI-enhanced cyberattacks may require collective security measures. The potential misuse of AI (like deepfakes destabilizing another country's politics, or autonomous weapons) has led to calls for global treaties or agreements. A Brookings study emphasizes reaffirming democratic principles and human rights in international AI forums, implicitly as a counter to authoritarian uses of AI (Brookings Institution, 2023). Likeminded democracies are indeed coordinating to put pressure against AI uses that violate human rights (such as mass surveillance or repression), concerned that a divergence could "split cyberspace into incompatible technology stacks" (Brookings Institution, 2023). Forums like the **Global Partnership on AI (GPAI)** bring together governments and experts from dozens of countries (including not just G7 but also India, Singapore, Senegal, Brazil, etc.) to discuss responsible AI innovation (OECD, 2025).

Concrete areas where enhanced collaboration is emerging include:

- **Regulatory sandboxes and standards:** Countries are comparing notes on AI regulatory sandboxes (safe testing environments) and even opening them up to foreign companies, to align approaches. There's early talk of recognizing each other's AI system certifications, much as pharmaceutical approvals sometimes are recognized across borders.
- **Ethical frameworks:** UNESCO's Recommendation on AI Ethics (2021) and the OECD principles serve as de facto global reference points (UNESCO, 2025). Over 40 countries in Africa worked together (with AU/UN help) on a *Continental AI Strategy* that aligns with these principles (UNESCO, 2025).

- **Multilateral funding for AI public goods:** The World Bank, regional development banks, and UN agencies are increasingly funding digital public goods including AI. For instance, the World Bank's GovTech initiatives encourage AI case studies and have facilitated South-South knowledge exchange in things like AI for tax collection (World Bank, 2024).
- **Trade agreements:** New trade deals (like the digital economy agreements Singapore has signed with partners) include provisions on AI, such as not forcing disclosure of source code but encouraging non-discrimination of AI-driven services. These are baby steps toward a consistent international trade regime for AI services.

**The bottom line:** International collaboration can create a level playing field that enables all countries to safely innovate with AI. By agreeing on baseline norms, sharing best practices, and pooling resources for big projects, countries can accelerate the benefits of AI for public goods globally. At the same time, collaboration provides a forum to manage disagreements—recognizing that some nations will prioritize privacy and human rights, for example, while others might emphasize innovation at all costs. Early cooperation allows space to negotiate those differences rather than let them spiral into conflict or a fragmented AI landscape. As one report put it, we need robust cooperation based on common principles and values as a foundation for each nation's successful AI development (Brookings Institution, 2023).

**Benefits (of multilateral collaboration):** Faster progress on AI solutions for global challenges (health, climate, etc.) through pooled data and talent; larger markets and interoperable systems that drive down costs and improve quality (e.g., a common standard for AI in medical devices means more competition and lower prices); avoidance of "AI nationalism" that could limit the flow of ideas and technology; collectively setting ethical guardrails to prevent harmful uses worldwide.

**Risks:** Geopolitical tensions spilling into the AI domain (e.g., differing values causing an "AI ethics schism"); domination of the agenda by a few big players, leaving others' needs behind; difficulty in reaching meaningful agreements (the pace of AI is fast, diplomacy is slow); potential lowest-common-denominator standards that don't actually solve problems (if agreements are too watered down).

**Actions:** Continue building international forums for AI in government (expand inclusion in GPAI and similar bodies); negotiate data-sharing arrangements for specific public-good objectives (e.g., a global database for pandemic detection); develop a common reference framework for assessing public AI systems (so an algorithm deemed trustworthy in one country can be recognized in another); invest in capacity building so all countries can participate (as UNESCO's training programs do); and pursue bilateral agreements as testbeds for larger coalitions (for example, a few countries aligning their procurement requirements for AI could set a model for others to join). In essence, treat AI governance as a new pillar of diplomacy - *"AI diplomacy"* - ensuring that as we connect our systems, we also connect our principles.

# Chapter 13: Impacts Across Governance, Workforce, Economy, and Citizens

The introduction of AI into government is not just a technical upgrade; it has far-reaching qualitative impacts on how society and the state interact. This section provides a cross-cutting analysis of AI's impact from multiple perspectives—governance quality, jobs and workforce, economic efficiency, and citizen outcomes—to assess the transformative potential and cautionary insights. Where possible, we include "back-of-the-envelope" estimations or statistics to ground these impacts in real terms.

**1. Governance Quality and Public Trust:** AI has the potential to enhance the quality of governance by enabling more evidence-based decisions, reducing corruption through automation, and improving transparency (for example, by making government data and processes more analyzable). When services improve and become more responsive, citizens often gain trust in institutions. Real-world anecdotes support this: when a city pre-emptively fixes potholes identified by sensors or AI analysis, citizens perceive a government that "cares and responds" (Global Government Forum, 2025). A study of civil servants worldwide found that 100% of those surveyed felt data and AI improved their ability to tackle irregularities in services, a shift that can increase internal accountability as well (Global Government Forum, 2025).

However, AI also poses a risk to governance quality if it is not well-managed. Opaque decision systems can undermine the **legitimacy** of public institutions. For example, if people believe an algorithm is wrongly denying them benefits or parole and they have no clarity or recourse, their trust erodes. There's also a risk of **"governing by algorithm"** where leaders deflect responsibility ("the computer says no"), weakening accountability. The net effect on governance quality will depend on the safeguards discussed throughout this paper: ensuring explainability, maintaining human oversight, and aligning AI use with rule of law principles. A noteworthy observation: in countries where digital government is very advanced, like *South Korea*, high service performance correlates with high citizen satisfaction (over **90%** of Koreans use digital public services and 93.8% report high satisfaction with those services) (Inter-American Development Bank [IDB] Gobernarte, 2024). This suggests that when AI is used to meet citizens' needs effectively, it can bolster the social contract. Conversely, high-profile failures or abuses of AI (like biased predictive policing systems revealed in some U.S. cities)

have prompted public outcry and moratoriums, showing that trust can be quickly lost. In summary, AI can **improve governance quality by making it more responsive and data-driven**, but it must be deployed with transparency and accountability to avoid a backlash that harms trust.

**2. Jobs and Workforce Transformation:** AI's impact on government jobs is complex—augmented productivity, changing roles, and the potential displacement of certain tasks. Current evidence suggests augmentation is the dominant theme in the short to medium term. According to an EY/Oxford survey of government executives in 2025, **58%** of organizations already using AI reported improved workforce productivity and even higher employee satisfaction (Tech Monitor, 2025). Automation of drudge work (like sorting forms, routing inquiries, compiling reports) frees public servants to focus on higher-value work. For example, UK government trials with coding assistants (AI that helps write software) found developers saved about **28 working days per year** each, roughly an hour per day of time saved (United Kingdom Government Digital Service, 2025). Extrapolating this, if even 25% of a government's IT staff gained an hour a day, thousands of work hours could be redirected to complex problem-solving annually.

Workforce transformation also means new kinds of jobs. Government will likely need more data analysts, AI system trainers, ethicists, and cybersecurity experts. Traditional roles may evolve: a benefits clerk becomes a case manager overseeing AI-flagged cases, or an urban planner becomes a curator of AI-generated scenarios. There will be a premium on digital and analytical skills. The public sector might face recruitment competition with the private sector for AI talent, emphasizing the need for upskilling existing staff and creating attractive career paths in civic technology.

There is concern about **job displacement**, especially in back-office or routine roles (data entry, basic information processing) where AI could automate tasks entirely. Estimates vary widely, but some studies have suggested that 10-15% of public sector tasks could be automated in the next decade with current AI—leading not necessarily to layoffs, but to attrition and redeployment. Governments often have the ability to retrain or shift employees rather than fire them, mitigating social impact. The key is proactive workforce planning: identifying roles likely to be affected and preparing those workers for new responsibilities. Unions and worker councils should be part of planning AI integration, to ensure it's seen as a tool that helps workers do their jobs better, not a stealth downsizing strategy.

Another perspective is the **spillover to the broader labor market**: by using AI in employment services (Chapter 8), governments can actually reduce unemployment duration and help workers adapt to AI-driven changes in the private sector. So government use of AI has a second-order effect on jobs economy-wide by improving labor market fluidity and training relevance.

In summary, the net impact on the government workforce can be positive—AI **can elevate the nature of public work**, making it more interesting and reducing burnout, if implemented with a "augmentation first" mindset. Rough estimations of efficiency gains (for instance, a BCG analysis posits up to 30–35% budget cost savings in high-volume administrative processes over ten years (Boston Consulting Group, 2025), much of which comes from labor hours saved) indicate significant capacity can be freed. The challenge is to reinvest those savings into improving services or tackling tasks that always get shelved due to lack of staff. If, on the other hand, AI is used crudely to cut headcount without maintaining service quality, both citizens and remaining staff will suffer, negating potential gains.

**3. Economic Efficiency and Fiscal Impact:** AI promises substantial improvements in governmental efficiency and economic returns on public spending. In concrete terms: - **Cost Savings:** Automation of processes can cut operational costs. A survey found **64%** of government executives see AI's potential for significant cost savings (Tech Monitor, 2025). We already have concrete examples: The city of Los Angeles, by using AI in routing waste collection trucks, reportedly saved millions in fuel and overtime. *Brazil* leveraged AI to optimize waste management routes in one major city and reduced costs by **45.4%** in that service (Virtasant, 2025). These are domain-specific, but add them up and the impact is large. If AI could save even 10% of the $10 trillion governments globally spend each year (a conservative scenario if applied broadly), that's $1 trillion freed for other purposes or deficit reduction. More realistically, certain high-volume domains (tax processing, claims handling, etc.) might see 20-30% cost reductions, while complex domains (healthcare, education) might reinvest AI gains into quality improvements rather than pure cuts.

- **Revenue Uptake:** On the revenue side, AI can increase tax collection by identifying evasion and improving compliance. Canada's tax agency already saved countless auditor hours and boosted yield. If AI tools globally enabled tax authorities to reduce the "tax gap" (lost revenue) by a few percentage points, that could mean tens

of billions in additional funds for public coffers. For instance, the US IRS estimates around $400 billion annual tax gap; a well-targeted AI fraud detection might recover even 5%, which is $20 billion per year.

- **Faster Service Delivery and Economic Productivity:** When permits are faster, businesses can start operations sooner, driving economic activity. When transport flows better due to AI (less congestion), you effectively give citizens and commerce time back (time is money). Some economists attempt to quantify the value of time saved: e.g., cutting average commute times by 10% in a big city through smart traffic systems could be equivalent to hundreds of millions of dollars in productivity gains annually for that metro area. AI's predictive maintenance in infrastructure can prevent costly disasters and downtime—for example, avoiding a major power outage by fixing things in advance saves not just repair costs but also economic losses from downtime.
- **Targeted Spending and Reduced Waste:** In welfare and social programs, AI can better target who truly needs benefits and reduce fraudulent or erroneous payments. If a nation's welfare budget is, say, $100 billion and fraud/leakage is 5%, that's $5 billion lost. Cutting that in half with AI saves $2.5 billion that can go either to more beneficiaries or back to the treasury. These "back-of-envelope" numbers illustrate why finance ministries are keen on AI.

It's important to note that economic efficiency should not come at the price of **equity**. Pure cost-benefit analysis might say, for instance, automating a service center saves money, but if it disenfranchises a segment of citizens who lack digital access, the broader social cost might outweigh the budget savings. So while AI can improve technical efficiency (doing more with less), policymakers must consider allocative efficiency (ensuring resources are distributed to maximize welfare). Done right, AI helps both: by saving money in one area, governments can afford to invest in under-served communities or emerging needs, improving overall welfare.

**4. Citizen Outcomes and Societal Impact:** Ultimately, the success of AI in the public sector will be judged by its impact on citizens' lives. Key outcome areas include:

- **Service Quality and Experience:** Citizens should experience faster, more convenient, and more personalized services. This can be measured in wait times, user satisfaction, uptake rates, and outcomes (like improved student learning or

health recovery rates). Already, digital government champions like Estonia or Denmark report high citizen satisfaction with e-services, and AI is poised to push that further by making services proactive (services find you, not vice versa). Imagine an "AI-augmented life journey": a child is identified as needing reading support by an AI tutor and gets it, a young adult finds a job faster through an AI job match, a senior gets health interventions early through AI health monitoring—all these improve individual life outcomes. Surveys in some countries show an increase in trust and satisfaction when governments introduced things like AI chatbots for 24/7 response, because people no longer felt ignored.

- **Equity and Inclusion:** A major promise of AI is to extend services to those who may be hard to reach. For example, rural populations can access specialist medical advice via AI telehealth systems where previously no specialists were available. AI translation can include non-native speakers in civic discussions. If governed properly, AI could help **close gaps** in service provision—though, as we've cautioned, it could also worsen them if not handled well. One qualitative impact is giving a voice to those who often don't speak up. AI summarization of public comments (like Taiwan's vTaiwan) shows even if thousands comment, everyone's input can count in the final summary, not just the loudest voices.
- **Safety and Well-being:** AI used in public safety (like disaster prediction, crime prevention targeted to hotspots, traffic accident reduction through smart signals) literally can save lives. If AI traffic management in a city cuts emergency response times by 20%, think of how many heart attack victims or accident victims that saves each year. Digital government analysts sometimes quantify these "quality of life years" gained. Similarly, AI in healthcare that catches diseases earlier extends lives and improves quality of life. These are harder to measure in dollars but are arguably the most important outcomes.
- **Democratic Engagement:** A subtle but important outcome is whether AI helps citizens feel more heard and connected to their government. If AI tools make participation easier and more impactful (government actually responds to the big issues raised by the public), one could see higher civic engagement rates. If AI handles millions of routine queries, officials have more time to engage on complex

cases personally, improving citizen relationships. Over time, this can strengthen democratic legitimacy in an era where many feel alienated from bureaucracies.

By way of estimation, let's consider a hypothetical mid-sized country with 50 million people after a decade of well-implemented AI in government: It might see GDP boosted by a few percentage points from efficiency (AI contributing to growth via productivity gains in public and private sectors), unemployment consistently perhaps 1 point lower than it would be (due to faster job matching and reskilling), life expectancy a few months higher (due to preventative health interventions and quicker emergency responses), and trust in government improving (say from 40% to 50% of the population expressing confidence, measured by surveys) thanks to more reliable services. These outcomes are speculative but align with the directions various studies have suggested.

Of course, there are dystopian possibilities if AI is misused—citizens alienated by automated, impersonal denial letters, or chilled by a feeling of constant surveillance, trust plummeting. It is a double-edged sword. But with proper governance, the qualitative analysis tilts optimistic: AI can make governance **more human-centered at scale**. Paradoxically, by automating the impersonal bureaucracy, it frees up humans to do the empathetic, complex work. For the economy, it can increase efficiency and free resources for new challenges. For the workforce, it can enrich public service roles. And for citizens, it can deliver tangible improvements in daily life and reinforce the sense that government is working for them.

In conclusion, the societal impact of AI in government will be as significant as the introduction of computers or the internet—perhaps more so, because it reaches into decision-making itself. The next sections provide a readiness assessment of different domains and case studies to illustrate where we stand on realizing these impacts.

# Chapter 14: Readiness for AI Adoption in Government

Not all domains of government are equally prepared for the adoption of AI. Some areas—like taxation, transport, and certain healthcare functions (e.g. diagnostics)—are already deploying mature, well-tested AI systems. Others—like judicial decision-making or high-stakes welfare eligibility—remain highly contested, with questions of fairness and legitimacy outweighing technical feasibility. To assess readiness, this study applies a qualitative scoring framework that evaluates each domain against four dimensions:

- **Technical Maturity** – How developed and proven are the AI tools for this domain? Are there commercially viable solutions already in use, or are most applications still at pilot stage?
- **Data Readiness** – Does the domain have access to large, reliable, and interoperable datasets? Are privacy and security issues manageable with available techniques?
- **Governance Fit** – How compatible is the use of AI with principles of fairness, transparency, and accountability in this domain? (For instance, AI in low-stakes, back-office tasks might be an easier fit than AI in roles that affect life, liberty, or fundamental rights).
- **Citizen Trust & Social Legitimacy** – To what extent are citizens likely to accept AI-driven decision-making in this area? Are appeal and accountability mechanisms clear to the public?

Each domain covered in this paper has been scored **High**, **Medium**, or **Low** readiness by examining real-world deployments, regulatory context, and societal acceptance as of 2025. "High" indicates the domain is ripe for scaling AI (with safeguards); "Medium" suggests promising opportunities but notable challenges to address; "Low" means significant caution and groundwork are needed before AI is broadly introduced.

**High Readiness domains** (e.g. taxation, transport, certain healthcare analytics) are those where technical solutions are already widespread, data is rich, and citizens have experienced visible benefits without major pushback. In these spaces, governments can scale AI relatively rapidly—provided that proper safeguards (transparency, opt-outs, audits) are in place to sustain trust.

**Medium Readiness domains** (e.g. education, social protection, labor/employment services, environmental monitoring) have clear opportunities, but either data is fragmented or the risks of bias and exclusion remain high. These domains benefit from continued pilots and a *"human-in-the-loop"* approach, where AI assists but human judgment remains central, until confidence in the technology and governance is higher.

**Low Readiness domains** (e.g. fully autonomous judicial decisions, real-time biometric surveillance in public) are often technically possible but face major governance and legitimacy barriers. Here, any experimentation must be tightly constrained, transparently managed, and accompanied by robust oversight and public dialogue. These are "not ready for prime time" and may require not only technical advances but new legal frameworks and trust-building before wider adoption.

It's important to note that readiness is not static. A domain can move from Medium to High as technical solutions mature, regulatory frameworks evolve, and citizen trust deepens. Conversely, early enthusiasm can be rolled back if scandals or misuse undermine legitimacy—as seen in some predictive policing pilots that were paused due to public controversy.

The readiness assessment for each domain covered in this paper is summarized below:

- **Education & Lifelong Learning – *Medium*:** AI tutors, automated grading, and recommendation systems are already in use (for example, Singapore's use of AI in SkillsFuture job matching, UNESCO pilot projects for AI in personalized learning), showing tangible benefits. However, data on student performance is sensitive and often siloed, and concerns over bias in adaptive learning and unequal access slow widespread scaling. Human teacher oversight remains essential both for pedagogy and for trust. Thus, the domain is in an expansion phase but with cautionary footings.
- **Healthcare & Public Health – *High*:** Diagnostic AI (imaging analysis), outbreak prediction models, and AI triage bots are relatively mature, with proven deployments in countries like China, the US, and Singapore. Data quality is high in many healthcare systems (though privacy is a paramount concern). Citizens generally show high acceptance where clear benefits are visible (e.g., faster cancer detection). Governance challenges (privacy, liability) exist but are being actively

addressed. This domain has momentum and a strong evidence base, indicating high readiness with continued governance innovation.

- **Social Protection & Welfare – *Medium*:** Fraud detection and eligibility automation are operational in several countries (Canada, various EU states use AI to flag welfare fraud; Latin American programs use predictive analytics for poverty prevention). The efficiency and inclusion gains could be significant. However, fairness concerns remain high due to past incidents of automated benefits being wrongfully cut (e.g., the "robodebt" scandal in Australia). Data is available but often fragmented across agencies. This domain requires careful piloting, transparency, and maintaining a human safety net (appeals, human caseworkers) — progress is happening, but society is cautious.
- **Public Safety, Justice & Borders – *Low to Medium***: This is a mixed bag. Border control AI (face recognition at airports, risk scoring for customs) is already quite mature and widely deployed — many countries use it and public acceptance is relatively higher for border security use-cases. But predictive policing and AI in core judicial decisions face legitimacy crises due to bias and due process issues. Within this broad domain, some subdomains are medium (e.g., AI at borders or AI for triaging legal documents), while others are low (sentencing algorithms, fully automated policing). On balance, because the high-risk areas dominate public debate, we rank the overall domain as Low-Medium readiness. It varies by subdomain and will advance only as fast as the most sensitive pieces allow.
- **Transport, Mobility & Infrastructure – *High*:** Smart traffic systems, predictive maintenance for infrastructure, and urban digital twins have been proven at scale (Dubai, Singapore, and China's City Brain were referenced as successes). Data is plentiful (sensors, GPS, etc.), and citizens generally welcome improvements in mobility and safety — making social acceptance high, since the benefits (shorter commutes, fewer accidents) are immediate and visible. This domain shows high technical and social readiness, with the caveat of ensuring equity (not letting AI prioritize service to wealthier areas) and cybersecurity.
- **Housing, Planning & Environment – *Medium*:** Automated permitting (e.g., Singapore's AI in construction permitting) and environmental monitoring AI are promising and have early deployments. These improve efficiency and enforcement.

However, risks of over-centralization and opaque denials must be managed. While technically many pieces exist (geospatial AI, satellite monitoring, etc.), adoption is moderate and contingent on proving that algorithms can be fair and explainable in regulatory contexts. Climate-related AI tools (like climate impact simulators) are technically advanced but not yet universally adopted in planning processes. So this domain is on the cusp but not all the way there—medium readiness.

- **Taxation & Revenue – *High*:** As discussed, pre-filled returns, AI-assisted customer service, and fraud detection AI are already live in numerous countries (Estonia, Spain, Canada, etc.). Citizens generally welcome the convenience and efficiency gains (few complain about an easier tax filing or catching big cheaters). Data here is abundant (financial records) and the technical solutions are well-tested. While transparency safeguards are needed (to prevent feeling of a tax "black box"), this is one of the most advanced domains. Thus, high readiness, evidenced by broad rollouts and positive feedback.
- **Labor & Employment Services – *Medium*:** AI job matching and reskilling guidance are operational in places like the Singapore and some Chinese provinces. They clearly add value, especially in dynamic economies, and technical feasibility is good (lots of data on jobs and skills exists). However, bias in algorithms and data privacy for worker profiles concern the public and experts. Some jobseekers might distrust algorithmic advice if it seems to pigeonhole them. Therefore, the domain is moving ahead but carefully; it's a medium readiness scenario where proof of fairness and building user confidence are the next steps.
- **Civic Participation & Democracy – *Low to Medium*:** AI for summarizing public consultations (as in Taiwan's vTaiwan or some EU pilots) is promising for scaling up democratic input. But the broader notion of AI in the democratic process is sensitive. Missteps (like AI-generated political messaging or deepfake use) can severely undermine trust. At this stage, while technical tools exist, public trust is quite low—people need assurances that AI will *support* and not manipulate democratic processes. Due to high risk and early stage experiments, readiness here is medium-low. Proceed with strong governance oversight and focus on transparency tools.

- **Cross-cutting enablers & governance – *Variable*:** This isn't a single domain but rather the foundation across domains. Some governments have strong digital backbones (e.g., Estonia's X-Road, as mentioned, which would be "High" readiness for supporting AI), while others have fragmented systems ("Low"). On procurement and capacity, many countries are still "Low" as they have not reformed rules or trained staff at scale. So this is variable: a country like South Korea or Denmark might score high on readiness due to decades of e-government work, whereas others have homework to do before AI can really take off.

In summary, **high-readiness domains** like healthcare, transport, and taxation offer "low-hanging fruit" where governments can push ahead with AI deployments now, delivering tangible wins (efficiency, better outcomes) to build momentum. **Medium-readiness domains** like education, welfare, housing, and labor services require measured approaches—experimentation with humans still heavily in the loop, and concurrent work on data integration and bias mitigation. **Low-readiness areas** such as AI in core justice decisions or pervasive surveillance should be considered mostly off-limits for now, except in controlled pilots, until trust and governance frameworks catch up.

By mapping readiness in this way, policymakers can **prioritize deployments** that are ripe for implementation while pacing those that require further safeguards or public dialogue. The aim is not to race ahead blindly in every domain, but to align AI adoption with where institutions and societies are ready-building on successes and learning from trials in more sensitive areas.

# Chapter 15: Global Case Studies of AI-Enabled Government

To illustrate how these themes play out in practice, this chapter presents case studies of governments that have been early movers in adopting AI for public goods and administration. We focus on a few leading examples, with additional glimpses of others, to draw comparative insights.

**Singapore: Precision and Integration at Scale** – Singapore has built one of the world's most integrated approaches to AI in government, leveraging its strong digital backbone (SingPass digital identity, a centralized GovTech agency, and the Smart Nation initiative). Nearly every sector of public service in Singapore has a flagship AI project: - *Citizen Services:* Virtual assistants like **Ask Jamie** (a multilingual chatbot deployed across agencies) and HealthBuddy (for healthcare inquiries) provide 24/7 help to citizens, handling millions of queries with consistent accuracy. - *Urban Management:* The Land Transport Authority uses predictive models to anticipate MRT train and bus ridership peaks, adjusting services dynamically. The **City Brain** pilot in parts of Singapore optimizes traffic flows in real time, as mentioned earlier. The **CORENET X** system streamlines construction permit approvals with AI rule-checking against building codes. - *Healthcare:* The Ministry of Health employs AI for patient triage (e.g., an AI symptom checker to advise whether to see a doctor) and for epidemiological tracking (as seen during COVID-19, where Singapore's analytic systems helped identify clusters and allocate resources swiftly). - *Education:* Singapore's National AI Strategy includes an "AI Education Personalized Learning" pilot, and the SkillsFuture job matching mentioned in Chapter 8 is a case where AI guides lifelong learning. - *Governance Enablers:* Singapore invests heavily in what it calls "digital trust" infrastructure – for example, **AI Verify**, a national framework/toolkit for testing AI systems for bias and explainability, which it shares with companies. The government publishes model AI governance frameworks and has created an Advisory Council on AI and Data to involve private sector and civil society input.

**Impact:** High citizen trust and strong cross-agency coordination allow Singapore to implement AI quickly. Singapore consistently ranks at the top of e-government surveys, and its residents generally expect tech-driven services. People can renew passports, apply for public housing, or check on their CPF pension through AI-assisted digital platforms. The tangible improvements (shorter wait times, proactive service offers) reinforce legitimacy. However, critics point out the risks of surveillance creep in such a top-down model; Singapore's extensive

use of sensors and cameras, combined with AI, has raised questions from privacy advocates about how far this could go. The government's response has been to emphasize cybersecurity and the voluntary nature of most data programs, but the debate illustrates the ongoing need to balance efficiency with civil liberties. Overall, Singapore demonstrates the *upper bound* of how far integration can go in a high-trust, well-resourced context.

**Estonia: AI in a Digital-First State** – Estonia's small size and long-standing e-governance ecosystem make it a pioneer in embedding AI across services. - *Judiciary:* As noted, Estonia piloted an AI "robot judge" for small claims up to a certain threshold, issuing proposed decisions that parties can appeal to a human judge. This garnered global attention as a bold move to deal with case backlogs, though it is limited in scope and carefully overseen by the Ministry of Justice. - *Taxation:* Estonia's pre-filled online tax returns (in place for years) are enhanced with AI chat assistance to help with any questions, making tax filing a 3-5 minute exercise for most citizens. AI-supported audit selection helps ensure relatively high compliance with minimal manpower. - *Cross-Government Data Backbone:* The **X-Road** platform allows all public (and some private) databases to interoperate securely. This means an AI system in, say, the welfare department can query data (with permission) from the tax office or education records to make holistic decisions. This interoperability greatly enhances AI potential and is a model internationally for digital integration. - *Civic Participation:* Estonia has e-participation platforms and has experimented with AI tools to cluster citizen proposals that come in via its "People's Assembly" online petition system. Given language constraints (small population, unique language), Estonia also uses AI translation to ensure services are available in both Estonian and Russian for its population.

**Impact:** Estonia shows how **data interoperability** is the single biggest enabler of AI readiness. Because citizens and politicians trust the digital system (99% of services are online, and citizens have digital ID cards enabling secure transactions), new AI solutions can plug into an existing ecosystem of data and services. Its small size and digital culture mean new systems scale quickly across the whole nation (a pilot in one department can be adopted government-wide in a year or two if successful). The challenges Estonia faces are also instructive: the **robot judge** pilot, for instance, illustrated the limits of automation in high-stakes decisions—public acceptance was cautious, and the government had to reassure that AI would not handle anything beyond minor disputes and that a human court review is always available. This reflects a theme:

Estonia is bold but still puts human accountability at the core (e.g., an AI decision isn't final until a person signs off in many cases).

**Dubai: AI as a Strategic National Showcase** – Dubai (and the UAE more broadly) positions AI both as a governance tool and a branding strategy to show itself as a "city of the future." - *Public Safety:* Dubai has invested heavily in AI surveillance and policing tools. Its "Oyoon" project integrates thousands of CCTV cameras with AI to monitor crowds for incidents, identify wanted suspects via facial recognition, and even detect traffic violations without human police present. Police robots and drones augment the human force for patrols in tourist areas. - *Transport:* Under Dubai's Autonomous Transportation Strategy, the city is testing AI-driven **autonomous vehicles**, including self-driving taxis and buses. Smart traffic management, similar to Singapore's, adjusts signal timings based on AI analysis of congestion. Dubai claims these efforts have significantly cut traffic light wait times and improved emergency vehicle response by dynamically clearing routes. - *Government Services:* The **DubaiNow** app is a one-stop portal integrating over 120 services (from paying utilities to renewing licenses), many of which are AI-powered for immediate approvals or tailored recommendations. The government even launched an AI-powered visa processing system that handles applications within minutes unless exceptions are flagged. - *Governance:* Notably, Dubai appointed a **Minister of State for AI** in 2017 – one of the first such ministerial positions in the world – and launched an AI strategy aiming to have AI across all government services by 2031. They created "Smart Dubai" office to coordinate AI initiatives and even an AI Lab in collaboration with IBM at one point.

**Impact:** Dubai excels at **top-down deployment** of AI, scaling tools rapidly across domains through strong political will and significant investment. In a relatively short time, it has implemented systems that other cities are still piloting. The visible effects – like flashy autonomous drones or generally efficient e-services – bolster Dubai's image as an innovative hub and likely improve user convenience. However, this comes with questions about transparency and civil liberties. Critics note that Dubai's political context (an unelected government with broad powers) means there is little public consultation or disclosure about how AI is used, especially in surveillance. There are concerns that in such a highly centralized model, it's hard to know if safeguards are in place. For example, how is facial recognition data stored and used? Who has access? The government assures it's for security and efficiency, but without independent oversight, these assurances rest on trust in leadership. The Dubai case thus

highlights that AI can be implemented **fast** when authority is concentrated, but it underscores the need for corresponding emphasis on ethics and oversight to prevent misuse. Nonetheless, Dubai is a case many developing cities look at for inspiration on leapfrogging infrastructure constraints using AI (e.g., managing traffic or utilities in rapidly growing urban centers).

**Nordic Countries: Human-Centric Digitalization** – Beyond Estonia (often grouped with Nordics in digital terms), countries like Finland, Denmark, and Sweden offer a slightly different approach focusing on ethical AI and co-creation. - Finland's **AuroraAI** program (mentioned earlier) is a noteworthy effort to create a nationwide AI-powered interoperability network. It aims to connect public and private services across life events (such as having a baby, career changes, elderly care) so that citizens can be proactively offered "the right service at the right time." The program explicitly foregrounds ethics and human-centric design, stating goals of well-being and avoiding hyper-personalization that might pigeonhole individuals (OECD, 2022) (European Parliament, 2021). While AuroraAI is still being gradually implemented, it exemplifies the Nordic tendency to marry technology with societal values (in this case, building a seamless welfare system that respects personal agency). - Denmark has used AI in government with a focus on trust and transparency. For example, the Danish Business Authority developed an AI system to help detect fraud in welfare and taxes but did so with an algorithm registry and close collaboration with Denmark's strong civil society to maintain a high trust environment (Denmark consistently ranks high in citizen trust in government). - Sweden employs AI for things like smart grids (energy) and has an innovation agency funding municipal AI pilots (for instance, AI to predict which students might need extra help, tested in some Swedish schools, but always with an opt-out and heavy data protection due to strict Swedish privacy laws).

**Impact:** These Nordic examples underscore the importance of **inclusion and ethics**. They haven't moved as lightning-fast as Singapore or Dubai in deployment, but they often achieve high outcomes in human well-being metrics. For instance, Nordic countries are exploring AI in healthcare but within robust public health systems that require explainability (a hospital in Sweden might use an AI decision support for radiology, but a radiologist will take time to explain the diagnosis to a patient—something that might be under-emphasized elsewhere). Their contribution to the global discourse includes efforts like Finland training a large portion of its population in AI basics via a free online course, "Elements of AI," to increase AI literacy.

This reflects a view that **societal readiness** (education, conversation about ethics) is part of technological readiness.

**United States: Pioneering Innovation amid Fragmentation** – The U.S. has a very mixed landscape given federal, state, and local levels: - *Use cases:* There have been many AI use cases: some cities did predictive policing (Los Angeles, Chicago, etc., though some rolled back after controversy), the Department of Veterans Affairs uses AI for patient health record analysis to predict risks for veterans, the IRS is exploring AI for fraud detection, and numerous states use AI chatbots for unemployment services (especially seen during the pandemic surge of claims). NASA uses AI for satellite imagery analysis relevant to climate and disasters, indirectly aiding FEMA. - *Federal vs Local:* Without a centralized mandate, progress is uneven. Some agencies are world leaders (the Defense Department, for better or worse, has advanced AI projects; the CDC used AI to model disease spread). Others lag or face bureaucratic inertia. - *Challenges:* The U.S.'s federal structure means one city's success doesn't automatically replicate. Political swings also affect support for government AI projects. One administration may fund AI for social services, the next may cut it citing different priorities or privacy worries. This fragmentation has limited scaling compared to smaller or more centralized nations.

**Impact:** The U.S. illustrates both innovation and **implementation challenges**. It boasts top AI talent and companies, so cutting-edge solutions exist (often pioneered in local pilots or academic labs), but scaling in government runs into legacy IT and trust issues, especially given higher skepticism of government in some communities. For example, some police departments halted AI use due to community pressure about bias. American discourse often highlights civil liberties—leading to for instance San Francisco briefly banning facial recognition use by city agencies. So in the U.S., the impact of AI in government has been spotty: transformative in some specific instances (like saving millions in fraud or improving some health diagnostics), but far from uniformly adopted. It serves as a caution that technology alone doesn't transform governance without addressing institutional and political factors.

**China: AI at Population Scale with Authoritarian Efficiency** – China presents a case of massive deployment of AI as a state tool, which yields certain public goods but at the cost of pervasive surveillance. - *City Management:* The **City Brain** initiative in several Chinese cities (Hangzhou was first) has shown remarkable results in traffic management – reportedly reducing congestion and improving emergency response markedly by coordinating across the

whole city's data feeds. - *Social Governance:* China infamously piloted a "social credit" system in various forms, using AI to aggregate citizen data to encourage or enforce certain behaviors. While often exaggerated in Western media, elements of this (like blacklisting debtors from luxury purchases) were implemented. It shows how AI can be used to integrate various public and private data sources for governance—something democracies would find hard due to privacy norms. — *Public Security:* China uses AI extensively for public security – facial recognition to catch fugitives (with claims of many arrests made by automated face matching), predictive policing akin to the American models but with even less transparency, and extensive censorship and public opinion monitoring online via AI. - *Public Services:* On the positive side, China's government hospitals are using AI diagnostics (some rural clinics use AI to help read medical scans where doctors are scarce), and government-affiliated research produced a widely-used AI weather forecasting system, etc. There's heavy investment in AI education tech as well (like Squirrel AI tutoring millions, mentioned in education chapter).

**Impact:** China demonstrates the **upper limits of AI-state integration** in terms of scale. The government's capacity to gather data on 1.4 billion people and apply AI has enabled things like very granular pandemic controls (as seen in their COVID response through health code apps). The efficiency and reach are undeniable – but it comes at the **cost of individual privacy and freedom**. Citizen outcomes in terms of convenience may be improved (people enjoy super-apps for government and commercial services combined, often with AI-driven personalization), and crime or fraud might be reduced through omnipresent monitoring. Yet, the model is not one that aligns with democratic or human rights values internationally. It highlights a fundamental trade-off: China's rapid AI deployment in governance is facilitated by public acceptance (or acquiescence) in exchange for order and services, whereas more liberal societies demand more consent and restraint, slowing deployment but protecting rights. The Chinese case pressures other countries to articulate what *responsible* AI in government means, lest the narrative "AI works best in authoritarian regimes" takes hold.

## Key Lessons from Case Studies

- **Integration Matters:** Countries that have invested in digital infrastructure and data integration (Singapore's whole-of-government approach, Estonia's X-Road, even China's national data systems) show that AI can scale and multiply benefits when systems talk to each other. Siloed data is a common stumbling block elsewhere.

Governments should prioritize "connective tissue" (ID systems, data exchanges) as much as the AI algorithms themselves.

- **Top-Down Ambition Accelerates Adoption:** Political will and clear strategies (Dubai's AI minister and strategy, China's national AI plan, Singapore's AI governance framework) can fast-track implementation and signal to agencies that AI is a priority. However, the flip side is that centralization without consultation can heighten risks of overreach. The fast adopters often have either strong central governments or cohesive policy agendas. Others can emulate the ambition while still incorporating checks—e.g., a national AI task force that sets goals but also includes ethicists.
- **Legitimacy is the Ultimate Constraint:** Even technically mature tools will face backlash if they conflict with public values or rights. The U.S. and European pause on predictive policing, and Estonia's cautious rollout of AI judges, show that governments must earn public trust for each new AI application. Legitimacy comes from transparency, engagement, and showing clear benefit. Fail that, and an AI program can be shelved due to public pressure regardless of its technical merits. Maintaining citizen trust (not just assuming it) is key to sustainable AI adoption.
- **Global Knowledge Transfer:** Governments are learning from each other. For instance, after seeing Estonia's success, countries like Uruguay and South Korea adopted similar digital ID and data exchange models. International bodies (like the OECD, World Bank) are codifying these lessons so latecomers can leapfrog. The case studies here suggest that while context differs, many principles (interoperability, pilot-testing, oversight mechanisms) are universally applicable.

With these insights from around the world, we now move to the conclusion, which charts practical pathways for different stakeholders—governments at various stages of AI adoption, private sector partners, and international collaborators—to responsibly scale AI for public good.

## Conclusion: Pathways for AI in Public Goods and Services

whether AI strengthens the socand services are delivered, from tax filing to traffic management, as we have seen. Across our analysis, one theme stands out: AI enables states to govern faster, with greater precision, and often at lower cost. But speed and efficiency are not ends in themselves. The true measure of success will be whether AI **strengthens** the social contract—making governance more transparent, inclusive, and responsive—or whether it risks eroding the legitimacy upon which government authority ultimately rests.

This conclusion outlines pathways for different stakeholders and summarizes the guardrails needed to ensure AI in public service is a force for good.

## Pathways for Private Organizations

Many of the innovations in AI for public goods will come from or in partnership with the private and non-profit sectors. For companies, investors, and civic tech organizations, the most ready domains (taxation, healthcare diagnostics, transport optimization, etc.) present immediate opportunities for collaboration. Private vendors are already supplying conversational AI for tax agencies, predictive analytics for hospitals, and city management platforms for urban authorities. These areas are ripe for **public-private partnerships**, innovative procurement models (like outcome-based contracts), and responsible scaling.

However, the lesson from case studies like Singapore is clear: **vendors must align with public sector values** of interoperability, transparency, and accountability. Governments are increasingly likely to exclude or downsize contracts with providers who offer only "black box" solutions. This means private firms should be proactive in building explainability features, agreeing to audits, and using open standards in their products. Investors should recognize that the long-term market in gov-tech AI will reward those who build ethical, inclusive solutions (e.g., AI that can explain decisions, or that is built on diverse datasets to avoid bias) because public clients face intense scrutiny. In short, the private sector has a significant role to play in innovation, but to succeed and scale in the public arena, **trustworthiness will be as important as technical prowess**.

## Pathways for Governments

For governments, strategic adoption of AI will look different depending on their starting point. We can distinguish between **brownfield** contexts (modernizing existing institutions) and **greenfield** contexts (designing new agencies or services from scratch with AI in mind).

- **Brownfield governments** – Most established governments should prioritize AI in **high-readiness domains first** (as identified in the readiness index). These include areas like healthcare diagnostics, tax administration, and transport management, where public trust is relatively strong and tools are technically mature. Early wins here can save money and improve visible services, building momentum and credibility. Gradual expansion into medium-readiness areas such as welfare, education, and labor services can follow, but **only with safeguards** in place—guaranteed rights to appeal automated decisions, independent algorithm audits to catch bias, and strong communication to the public about how and why AI is used. Brownfield adoption is essentially a modernization path: augment services one by one, retrofit governance processes to accommodate AI, and continuously involve citizens and employees in the transition to ensure buy-in.
- **AI-native governments** – In some cases (new agencies, smart city projects, or countries overhauling legacy structures entirely), there's an opportunity to build with AI at the center from day one. Dubai's creation of a Ministry of AI or Albania's experiment with an AI procurement advisor are examples of symbolic steps, but real value lies deeper: imagine designing a welfare agency whose business processes assume AI will handle initial triage, or a transportation authority that from inception uses digital twin simulations for planning. These **born-digital institutions** can leverage AI to skip over inefficiencies that incumbents struggle with. However, it is crucial that in doing so they bake in transparent design principles. An AI-native agency must avoid becoming an efficiency-driven but legitimacy-poor entity. Transparency, accountability mechanisms, and engagement of stakeholders should be part of the blueprint. For instance, if a country were to establish a new "Digital Benefits Agency" with AI determining eligibility in real-time, it should simultaneously establish an independent oversight board and a clear public portal explaining decisions, so that legitimacy is built in, not tacked on later.

No matter the approach, capacity building is key. Governments need to invest in **training public servants** about AI (both opportunities and limitations) so that there is internal capability to manage vendors, interpret AI outputs, and make informed policy choices. Likewise, engaging the public through consultations on AI ethics or pilots (as Canada and Finland have done) can preempt misunderstanding and resistance.

## Risks and Guardrails

Throughout this paper, we have repeatedly encountered certain risks. Here we distill them and the cross-cutting guardrails to address them:

- **Bias and Exclusion:** Perhaps the most pervasive risk is that AI systems, reflecting historical data or flawed design, could perpetuate or even exacerbate bias—be it in welfare, policing, or hiring. This could unjustly deny opportunities or benefits to certain groups, undermining fairness. **Guardrail:** Mandate bias testing and mitigation for all high-impact AI, involve affected communities in evaluating systems, and keep humans in the loop especially where subjective factors (like compassion or second chances) deserve weight beyond data patterns.
- **Surveillance and Privacy Erosion:** The flip side of AI's data hunger is the potential for a surveillance state. Public safety, transport, and even social services AI could intrude into personal lives in ways people find unacceptable (constant monitoring, profiling, etc.). **Guardrail:** Draw red lines (through laws or binding policies) that protect privacy—e.g., ban real-time facial recognition in public spaces without a court order, prohibit using personal data in one context (say welfare) for unrelated policing purposes, require Privacy Impact Assessments for new data initiatives, and deploy privacy-enhancing tech (like anonymization, federated learning) where possible.
- **Opaque Decision-Making:** If people do not understand how a decision was made, they cannot contest it or trust it. Opaque AI is a direct threat to accountability. **Guardrail:** Insist on explainability or at least interpretable results for any AI making consequential decisions. This might mean avoiding the most "black box" model if a slightly simpler model can be explained. It also means creating channels for citizens to ask for and receive explanations in plain language, as well as expert review panels that can audit algorithms under the hood.

- **Workforce Impacts and Pushback:** If AI is seen as a threat to jobs or professionalism, employees can resist or even quietly subvert deployments, and valuable organizational knowledge could be lost. **Guardrail:** Implement AI with a clear narrative and plan that it will augment, not replace. Offer retraining and even role guarantees where feasible to alleviate fear. Involve staff in design and deployment so they feel ownership. Measure and publicly report on how AI is affecting workload and job satisfaction, as Singapore has done internally.
- **Vendor Lock-in and Tech "Colonialism":** Many governments worry about dependency on a few big tech firms or foreign AI solutions that might not align with local needs. **Guardrail:** Strengthen procurement practices as discussed—demand open standards and data portability in contracts, consider open-source solutions (some nations are coming together to build open AI tools for exactly this reason), and foster domestic ecosystems by requiring knowledge transfer in projects.

Without these guardrails, the promise of AI making governance more inclusive could invert into a reality where decisions are faster, but fairness is eroded and citizens feel alienated. The stakes are high: get it wrong, and we could see tech-driven policy failures that set back not just AI adoption but trust in government overall. Get it right, and AI can genuinely help democracies deliver better for their people in a very tangible way.

## Actionable Steps for Policymakers

Drawing on everything in this paper, here are concrete steps governments should consider in the near term:

- **Prioritize High-Readiness Domains:** Focus initial efforts where AI is most likely to succeed and be welcomed. For example, deploy AI in **healthcare diagnostics**, **tax administration**, and **transport optimization** first. These offer visible wins (like shorter hospital waits, easier tax refunds, less traffic) and can build public confidence in AI's benefits. Meanwhile, hold off or carefully sandbox AI in controversial areas (like criminal sentencing or surveillance) to avoid early scandals.
- **Embed Safeguards in Medium Domains:** In areas like welfare, education, and labor where potential is matched by social sensitivity, implement AI gradually and with strict **"no harm"** safeguards. For instance, any AI in welfare should launch

alongside an improved appeal process and ombudsperson service, so people actually feel the system got fairer, not colder. Similarly, in education, keep teachers and parents informed and give them override power if an AI makes a recommendation about a student. Essentially, pair every AI rollout with a policy toolkit (appeals, audits, transparency reports, training) as a package.

- **Reform Procurement and Partnerships:** Update procurement rules to be agile and innovation-friendly. For example, use challenge-based procurements: instead of specifying a solution, specify a problem (like "reduce traffic accidents by 20%") and let companies propose AI solutions, judging them on outcomes and ethical compliance. Build Public-Private **sandboxes** where companies can pilot solutions with real government data under supervision—this speeds up learning for both sides. Also, collaborate with academia and civic tech groups, who can often provide more neutral assessments of algorithms or even help develop public sector open-source tools (as seen in some cities releasing open algorithms for public scrutiny).
- **Invest in Data Infrastructure:** Make foundational investments in things like a unified digital ID system, secure data exchange platforms, and cloud or high-performance computing accessible to agencies. Estonia's example and others show that without plumbing, shiny AI applications either won't function or won't scale. Governments might consider shared services approaches—a central AI unit that provides models or computing to all ministries, to avoid duplication and ensure consistency (the UK and France have set up central AI teams, for example). Also, address data quality issues: an AI is only as good as the data it learns from, so initiatives to clean and standardize data (boring as they sound) are vital precursors.
- **Codify Red Lines and Ethics:** Policymakers should proactively legislate or regulate in areas where they *know* they do not want AI used, or want it used only with strict controls. For example, ban real-time biometric identification in law enforcement except for extreme scenarios with court approval. Prohibit AI from making final decisions that affect fundamental rights (liberty, life) without human confirmation. Disallow use of AI for political surveillance or social scoring. These clear rules will both guide developers and reassure the public that some boundaries are firmly in place.

- **Maintain Hybrid Human-AI Models:** At least for the next decade, the goal should not be fully automated government (if that is ever desirable) but **augmented government**. Ensure that in high-stakes domains—welfare, justice, healthcare—AI systems are tools in the hands of skilled professionals, not autonomous agents. For example, an AI might draft a medical treatment plan, but a doctor must sign off and discuss it with the patient. Or an AI flags an unusual tax return, but an auditor reviews before action. This hybrid approach not only provides accountability and empathy, but also creates feedback loops where humans can catch and correct AI errors, making the systems better over time.

If governments follow these steps, they will increase the likelihood that AI deployment improves both **efficiency** and **equity**. Done carelessly, AI projects could waste money or cause harm and then face public rejection. But done thoughtfully, they can achieve what bureaucracy always aspired to but often struggled with: delivering public services that are universal yet personalized, expeditious yet fair.

## Closing Thoughts

AI can help governments evolve from reactive administrators to proactive stewards of public goods. It offers powerful tools to anticipate needs, tailor services, and allocate resources with a precision previously unimaginable. The most successful states in the AI era will not be those that simply adopt technology the fastest, but those that integrate it the most wisely — **aligning technological innovation with social values and public trust**.

Private actors will play a vital role in supplying tools and expertise, but governments must set the terms, insisting on systems that embody transparency, fairness, and accountability. Geopolitically, a nation's capacity to use AI for public good could become a new benchmark of development (just as electrification or internet access once were). Countries that get it right will likely enjoy not just more efficient services, but also higher citizen satisfaction and competitiveness, whereas those that stumble could see technology widen governance gaps.

The task ahead is not simply to automate government—it is to reimagine governance itself in an AI-driven world. This means rethinking processes, roles, and institutions with a technology-forward mindset while steadfastly protecting the human and democratic principles at the core of public service. If we succeed, AI will be remembered as the set of tools that helped fulfill the promise of government to its people: to deliver public goods faster, fairer, and more intelligently than ever before, and to do so in a way that reinforces the bond of trust between citizen and state.

# Bibliography


Boston Consulting Group. (2025). *How AI can cut through bureaucracy.* Boston Consulting Group. https://www.bcg.com/publications/2025/benefits-of-ai-in-government

Brookings Institution. (2023). *Strengthening international cooperation on AI.* Brookings Institution. https://www.brookings.edu/articles/strengthening-international-cooperation-on-ai

European Parliament. (2021). *AuroraAI Programme (Finland): Personalized and proactive public service delivery.* European Parliament. https://www.europarl.europa.eu/cmsdata/240174/AIDA%20workshop_28052021_Alanko.pdf

Global Government Forum. (2025). *Integrating data and AI to improve government services: Lessons from Jakarta.* Global Government Forum. https://www.globalgovernmentforum.com/the-golden-opportunity-integrating-data-and-ai-to-improve-government-services-and-renew-public-trust

Inter-American Development Bank (IDB) Gobernarte. (2024). *South Korea's digital transformation in the public sector.* IDB. https://blogs.iadb.org/administracion-publica/en/south-korea-digital-transformation-public-sector

Organisation for Economic Co-operation and Development (OECD). (2019). *OECD AI principles: Human-centered, fair, and transparent governance.* OECD Observatory of Public Sector Innovation. https://oecd-opsi.org/work-areas/ai

Organisation for Economic Co-operation and Development (OECD). (2021). *AI in regulatory enforcement: Applications in inspections and compliance monitoring.* OECD. https://www.oecd.org/en/publications/governing-with-artificial-intelligence

Organisation for Economic Co-operation and Development (OECD). (2024). *Governing with artificial intelligence: AI in regulatory design and delivery.* OECD. https://www.oecd.org/en/publications/governing-with-artificial-intelligence

Rest of World. (2025). *AI floods Brazil's courts: Over 140 AI projects tackling case backlog.* Rest of World. https://restofworld.org/2025/brazil-ai-courts-lawsuits

TechMonitor. (2025). *Public sector AI adoption lags despite widespread recognition of cost-saving potential.* TechMonitor. https://www.techmonitor.ai/digital-economy/ai-and-automation/ai-adoption-in-public-sector-lags

United Nations Educational, Scientific and Cultural Organization (UNESCO). (2025). *AI in Africa: UNESCO unveils new solutions for its development at G20 Summit.* UNESCO. https://www.unesco.org/en/articles/ai-africa-unesco-unveils-new-solutions-its-development-g20

Virtasant. (2025). *AI in government services: Efficiency and waste management optimization.* Virtasant. https://www.virtasant.com/ai-today/ai-in-government-services-efficiency