



Responsible Use of AI in Urban Service Delivery



Responsible Use of AI in Urban Service Delivery

Artificial Intelligence (AI) is rapidly reshaping how cities function. Across Asia, it is already helping governments respond faster to citizen queries through chatbots, ease congestion with smarter traffic systems, reduce waste through automated sorting, plan better with real-time data, and even predict floods and heatwaves. Yet despite this growing presence, many urban officials still see AI as complex, risky, or difficult to scale. The truth is that AI can be a powerful tool for better governance, but only when cities understand how to use it safely, responsibly, and effectively.

The paper on 'Responsible Use of AI in Urban Service Delivery' developed by GIZ and Artha Global offers cities that exact clarity.

AI brings real opportunities, but also real risks, particularly around privacy, discrimination, and long-term sustainability. Governments around the world, from the European Union to the United Nations, have already begun setting standards and guardrails. At the heart of these frameworks is a simple idea: i.e. AI must protect people's rights. For Asian cities looking to adopt AI, these principles provide a strong starting point, but they need to be adapted to local realities.

Conversations with officials, practitioners, and industry experts inferred that most AI projects fail not because the technology is weak, but because the surrounding systems are not ready. Cities struggle with scaling pilots, setting up the right digital infrastructure, managing vendors, and meaningfully involving citizens. These challenges can be solved, but they require clear, practical guidance.

The paper provides that guidance. It includes:

- A readiness checklist that helps cities assess whether they have the policies, infrastructure, skills, and governance structures needed to adopt AI responsibly.
- A vendor due diligence toolkit that helps governments pick the right partners (those who are transparent, ethical, and technically sound).
- A roadmap for scaling AI from small pilots to city-wide systems, with an emphasis on responsible procurement, citizen participation, and long-term sustainability.

Above all, the message is simple: AI will only strengthen cities if it is built on trust, human rights, and public value. Cities that invest today in strong institutions, transparent systems, and inclusive processes will be the ones that define the future of responsible urban innovation in Asia.

This paper shows you how to get there.

The two additional products (checklists) can be found in the end of the document.

Imprint

Published by: Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) GmbH
Dezember 2025

As a federally owned enterprise, GIZ supports the German Government in achieving its objectives in the field of international cooperation for sustainable development.

Registered offices
Bonn and Eschborn
E info@giz.de
I www.giz.de

Dag-Hammarskjöld-Weg 1-5
65726 Eschborn, Germany
T +49 61 96 79-0
F +49 61 96 79-11 15

Friedrich-Ebert-Allee 36 + 40
53113 Bonn
T +49 228 44 60-0
F +49 228 44 60-17 66

Cover Illustration: Generated by Gemini 3 Flash (Google AI) based on themes of responsible urban AI governance.

Authors

Artha Global

Mayura Gadkari, Principal
Sandeep Radhakrishnan, Senior Associate
Anushri Pundit, Analyst
Keith Fernandes, Research Intern

With thanks to Sridhar Ganapathy

Editors/Coordination

Luciana F. Maia
Samrat Banerjee
Dr. Fathim Rashna Kallingal

With thanks to Manjari Mishra

Responsible

turboclic – transformation – urban opportunities – climate change – an initiative by The Sector Networks Good Governance in Asia (SNGA) and Transport Environment Energy and Water in Asia (TUEWAS).

Speakers: Indira von Gierke (for SNGA) and Dr. Teresa Kerber (for TUEWAS).

These Sector Networks are financed through projects and programmes which GIZ is conducting on behalf of the German Government. This publication was financed through the innovation funds of TUEWAS.

Disclaimer

This paper is based on research conducted by Artha Global as part of a collaborative initiative with Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) GmbH. This publication is intended for non-commercial use. Excerpts may be reproduced or cited for research, policy, or educational purposes with appropriate attribution.

Report

Responsible Use of AI in Urban Service Delivery

AI Adoption Checklist

AI Vendor Due Diligence Checklist

December 2025

Written by:



Table of Contents

Acknowledgements	5
Executive Summary	7
What is AI – And Why Should Cities Care?	11
Natural Language Processing	12
Machine Learning	13
Computer Vision	17
Urban AI in Action:	20
Governance & Participation	20
Citizen Engagement	21
Land Use and Property Tax	22
Operations & Infrastructure	22
Traffic Management	22
Waste Management	23
Water Management	24
Urban Planning and Land Governance	25
Social Services	26
Wadhvani AI's Multi-Faceted Suite	26
AI's Social Impact in Asia	27
Environment & Climate	28
CivicDataLab's IDS-DRR	28
AI's Utility in Climate Resilience	29
Safety & Security	30
Traffic Safety	30
Women's Safety	30
GiveDirectly – Aid Targeting in Bangladesh	31
What to Watch Out For – Risks, Ethics & Bias	32
Privacy and Safety in the Context of AI	32
AI and the Challenge of Discrimination	33
Resource-Intensive Nature of AI	34
Ethical Principles	35
Data Readiness & Digital Infrastructure	37
Data Quality	37

Digital Infrastructure	38
Convergence	38
Scaling AI Responsibly	40
Primary Challenges to Scaling AI	40
Strategic Framework for Overcoming Scaling Challenges	42
Working with the Ecosystem – Procurement & Partnerships	44
Redefining Effective Urban AI Partnerships	46
Takeaways for Urban Stakeholders	46
AI Adoption Checklist for Cities	47
AI Vendor Due Diligence Checklist	50
Designing Trust – Participation, Transparency, and Inclusion	54
Designing Inclusive AI	54
Participatory AI Design	54
Recommendations	56
References	60

Further acknowledgements

This document draws on insights from interviews with key informants and has been shaped by their valuable participation. We extend our sincere gratitude to all those whose contributions and engagement have significantly enriched the project.

We extend our thanks to the representatives from the following organisations:

- Civic Data Lab
- Asian Development Bank
- eGov Foundation
- Wadhvani AI
- World Bank, and
- Mr Sameer Unhale, Joint Commissioner, Directorate of Municipal Administration, Government of Maharashtra

Executive Summary

Artificial Intelligence (AI) is rapidly reshaping how cities think, plan, and deliver services. Yet for most local governments, AI remains an unfamiliar terrain—filled with both promise and uncertainty. Recognising the inequal concentration of AI technologies in big tech companies, the white paper highlights the urgent need for Asian governments, cities, and local tech communities to take an active role in developing AI. By doing so, Asia can help create a fairer and stronger AI future—one that meets its unique urban challenges, encourages local innovation, protects control over important digital systems, and supports fair and balanced growth of cities. It aims to equip policymakers, city leaders, and their ecosystem partners with a practical roadmap for harnessing AI responsibly—to strengthen governance, improve service delivery, and make cities more liveable, inclusive, and resilient.

Why AI Matters for Cities

AI is not just about technology; it is fundamentally about decision-making. When deployed wisely, AI systems help municipalities shift from reactive, routine management to proactive, anticipatory governance. With AI, cities can see patterns previously invisible, respond before needs become crises, and shape more inclusive, equitable public services. But realising these outcomes is not automatic. Our research and interviews show that purposeful AI relies on something deeper than algorithms: physical and digital infrastructure, inclusive policies, institutional readiness, and above all, public trust. The question facing cities is no longer ‘Should we use AI?’ but ‘How can we do so in a way that genuinely serves and protects our communities?’

Where AI is Already Making a Difference

Drawing from real examples across Asian cities, this paper highlights how AI is already being used to transform urban systems in domains as varied as governance and participation, infrastructure and operations, social services, environment and climate, and safety and security across the south Asian context. AI is helping government address citizens’ grievances more effectively, deliver welfare to vulnerable citizens, and enhance disaster preparedness within ‘at-risk’ regions¹.

¹ Governance & Participation: AI-powered chatbots and grievance redressal tools in Mumbai and Bengaluru, and open-source frameworks such as ParakhAI, are helping governments respond faster and engage citizens more transparently.

Infrastructure & Operations: Smart traffic systems in Jaipur and Kanpur, predictive maintenance tools for water supply in Tashkent and Manila, and AI-led waste sorting in Tokyo are improving reliability and resource efficiency.

Social Services: Initiatives like Wadhwani AI’s maternal health models or Bihar’s disease prediction tools show how AI can extend public health and social protection to the most underserved.

Environment & Climate: From Jakarta’s flood prediction systems to Bhopal’s air quality dashboards, cities are using AI to better manage risk—but these depend on strong data foundations.

Safety & Security: Cities such as Pimpri-Chinchwad and Metro Manila are experimenting with AI for public safety and citizen reporting, testing new ways to make urban spaces safer.

The Challenges Holding Cities Back

Despite progress, three core barriers keep most municipalities from scaling or benefiting fully from AI:

1. **Data Foundations:** Many cities are held back by inconsistent, poor-quality, or inaccessible data—without which AI cannot operate reliably.
2. **Institutional Capacity:** Adoption is often limited by skill gaps, unclear roles, or fragmented leadership. The technology is only as effective as the teams, workflows, and politics behind it.
3. **Trust, Ethics & Accountability:** Fears of bias, discrimination, lack of transparency, and data breaches are real and can stall progress without firm safeguards. Successful use of AI depends as much on public values and buy-in as on technical solutions.

Building the Right Ecosystem

The paper underscores that AI adoption is a collective effort. It requires strong partnerships—between governments, start-ups, academia, and development partners. Models such as the AIRAWAT Centre of Excellence and the eGov Foundation show how India is beginning to build this ecosystem, while international efforts by the World Bank and ADB demonstrate the value of shared frameworks and responsible sandboxing².

Responsible AI and Public Trust

Responsible AI begins with transparency, privacy, and public trust. Cities must be clear about what data is being used, how decisions are made, and how citizens can question or correct errors. AI systems must be transparent, auditable, explainable, and inclusive reflecting the diversity of the communities they serve. Trust is not a by-product; it is a precondition for sustainable innovation.

Key Products of This Paper: Practical Tools for Leaders

Recognising the limitations of real-world municipal work, this paper delivers two signature outputs— the AI Adoption and AI vendor due diligence checklists- designed for direct use by city officials and policymakers, and a policy roadmap to scale AI pilots.

1. The AI Adoption Checklist for Cities

A hands-on tool to help cities assess readiness and guide responsible integration. The checklist asks, step by step:

- Is the proposed AI project aligned with city goals and backed by leadership?
- Is our data digital, reliable, and AI-ready?

² Sandboxes, in the context of AI, refer to controlled environments where AI systems can be tested, validated, and monitored without affecting real-world operations. These environments are isolated from the main systems to prevent any unintended consequences or security breaches. Sandboxes allow developers and researchers to experiment with AI models, run simulations, and fine-tune algorithms in a safe and controlled setting.

- Are the necessary digital infrastructure and IT teams in place?
- Are clear oversight and continuous feedback mechanisms available for staff and citizens?
- Have we built ethics and safety into the design—and are risks clearly communicated?
- Is there a plan for sustainability, scaling, and ongoing updates as city needs evolve?

This checklist makes it possible to catch red flags early, build cross-department coordination, and ensure city investments translate to real value for residents.

2. The AI Vendor Due Diligence Checklist

A straightforward framework to help public agencies screen external technology partners for reliability, safety, and fit:

- Is the vendor reputable and experienced in government work?
- Can they explain their system's function, outcomes, and data sources in clear terms?
- Does their approach to privacy, bias, and compliance meet city and national standards?
- Will the vendor provide upfront training, ongoing support, and flexibility for new needs?
- Are clear ethical guardrails in place—such as human review and the right to override—if something goes wrong?

Applying this checklist ensures transparency and minimizes risks before contracts are signed, creating a foundation for lasting, trustworthy partnerships.

3. Policy Roadmap: Moving from Pilots to Public Value

In addition to the hands-on checklists, the paper offers a concise menu of policy and operational recommendations—gathered from best practices, regional networks, and real-world casework:

- Make responsible AI adoption a top civic priority, grounded in public value, inclusion, and explicit human rights protection.
- Tear down data silos and treat data as a shared public asset, creating the infrastructure and accountability to support high-quality, trustworthy analytics.
- Build institutional and ecosystem readiness through dedicated AI task forces, partnerships with academia and civic tech, and a commitment to ongoing capacity-building.
- Centre co-creation by involving citizens, frontline staff, and especially marginalised groups early and throughout the AI lifecycle.
- Mandate responsible procurement and vendor management—requiring transparency, piloting ('safe trial environments'), and ongoing evaluation.
- Design for long-term inclusion, safety, and resilience by proactively managing bias, enabling adaptability, and embedding oversight and public engagement.
- Monitor, evaluate, and adapt continuously, allowing only solutions that prove their worth in local contexts to scale up and shape the city's digital future.

What is AI – And Why Should Cities Care?

The idea that machines can be capable of intelligent action is one that has occupied the minds of scholars and technologists alike since the mid-twentieth century. In 1956, a Dartmouth University workshop led by McCarthy, Minsky, Rochester, and Shannon (widely recognised as the genesis of artificial intelligence) set out to design machines whose behaviour would be judged ‘intelligent’ if exhibited by people. Since then, the number of definitions for artificial intelligence have proliferated, which has therefore been accompanied by a slight dissonance in society’s understanding of the concept. Some have emphasised rationality and problem-solving, others its capacity to learn from data, and still others its aspiration to replicate human skills. No single definition has achieved consensus, and perhaps that is precisely what has allowed AI to remain such a dynamic and generative field.

Artificial Intelligence (AI) refers to computer systems that can perform tasks usually requiring human intelligence, including recognising patterns, making decisions, learning from experience, and understanding natural language. This includes subfields like machine learning, where algorithms improve through data exposure, and deep learning, which relies on neural networks to analyse complex data such as images or speech. (Berente et. al., 2021)

The somewhat ambiguous nature of the concept has encouraged wide experimentation and application, which has in turn pushed computing into new domains with each technological advance. Given this experimentation, AI can be seen as less a static entity than a moving boundary, the point at which computational techniques begin to approximate or augment forms of human reasoning in ever more complex settings. For the purposes of this chapter, we treat AI as this evolving frontier: the set of computational tools and methods currently being mobilised to tackle decision-making challenges once thought to be the exclusive preserve of human judgment. This perspective is particularly useful in the municipal context, where urban governance increasingly relies on such tools to navigate complexity and deliver services - helping cities overcome resource constraints, streamline routine operations, anticipate citizen needs through predictive analytics, reduce manual errors, and enable faster, more equitable responses across diverse communities.

As cities adopt AI, they must also grapple with important ethical and governance questions to ensure technology use is transparent, accountable, and fair while guaranteeing that it respects citizen privacy and promotes inclusion. Responsible adoption is not optional but essential. Several tools, such as Natural Language Processing (NLP) and Machine Learning (ML), are utilised to analyse vast and diverse datasets at the city level, in addition to tools such as Computer Visioning, to tackle urban service delivery challenges.

Figure 1: The Map of Artificial Intelligence



Natural Language Processing

Over the past decade, Natural Language Processing (NLP) has emerged as one of the most powerful tools for working with the vast amounts of text data now available because of the internet. At its core, NLP sits at the intersection of artificial intelligence, machine learning, and deep learning, using computational methods to parse, interpret, and generate human language (Hirschberg & Manning, 2015).

Typical steps in an NLP workflow range from cleaning and structuring text to turning words into features that algorithms can analyse, to identifying patterns of meaning and association (Ghosh & Gunning, 2019). Because it applies systematic rules across large volumes of data, NLP can uncover insights with greater consistency and scale than traditional text analysis. In fact, experiments in fields such as medicine have shown that NLP techniques can reproduce and even extend the themes identified through manual coding (Guetterman et al., 2018).

Cities have begun to use these AI tools in different ways. A review of the literature shows five main areas of application: governance and management, public health, land use and zoning, mobility, and urban design. The first of these, i.e. governance and management, appears to account for the largest share of use cases, covering issues from citizen engagement and service

delivery to crime prevention and disaster response. Public health represents another significant strand of work, with solutions built for epidemic tracking, air quality, and the health impacts of the built environment. NLP has also been creatively incorporated to identify Industrial Land Uses (ILUs) and classify neighbourhoods by activity within a spatiotemporal matrix (Huang et. al., 2018; Fu et. al., 2018). Meanwhile, although smaller in number, mobility-focused solutions apply NLP to transportation patterns and user behaviour, and urban design is beginning to explore its potential for understanding how people experience and describe urban spaces (Cai, 2021).

These applications demonstrate NLP's ability to synthesise vast, diverse text datasets into insights that directly support urban policy and decision-making, enabling governments to more effectively understand and respond to evolving resident needs. For example, in cities like Mumbai, NLP is used to analyse citizen feedback from social media and complaint portals, accelerating grievance redressal and enabling proactive service delivery. This capability represents a significant shift from traditional data collection methods, which often rely on delayed or fragmented reporting. By integrating NLP into governance processes, city institutions can strengthen their responsiveness, build public trust, and better align interventions with residents' evolving needs.


Machine Learning

Machine learning (ML) is a branch of artificial intelligence in which algorithms improve their performance at a given task through experience with data, rather than being explicitly programmed for every outcome. By adapting their internal parameters during a process called training, these algorithms can learn from labelled examples (supervised learning), uncover hidden patterns in unlabelled data (unsupervised learning), or combine both approaches (semi-supervised learning). Another paradigm of machine learning is 'reinforcement learning', where an agent learns how to map situations to actions by interacting with an environment to maximise a numerical reward signal over time. In essence, the goal is to generalise beyond the training data and to recognise patterns, make predictions, or solve problems when confronted with new, unseen inputs (Sarker, 2021).

ML has become a cornerstone of data-informed urban governance. Its capacity to detect patterns and generate predictions from large, complex datasets enables city systems to shift from routine management to proactive decision-making. Applications now extend across a wide range of municipal functions, from optimising traffic flows and forecasting energy demand to refining waste collection schedules and assessing infrastructure vulnerabilities.

Figure 2: The Machine Learning Process

Machine learning process



When integrated into city operations, these models provide administrators with near-real-time intelligence that strengthens planning, resource allocation, and service delivery. The result is not only greater operational efficiency but also tangible improvements in citizen experience such as smoother commutes, more reliable utilities, and cleaner neighbourhoods. For policymakers, the challenge lies in embedding such tools within institutional processes and ensuring that insights derived from data translate into accountable, equitable action on the ground.

Machine Learning and Urban Resilience

Machine learning has emerged as a powerful tool in strengthening urban resilience, particularly in the face of natural and human-induced hazards. Scholars have explored resilience from multiple perspectives, ranging from the performance of buildings under seismic stress to the adaptability of healthcare facilities during crises. For example, Shah et al. (2019) developed computational frameworks to evaluate the seismic resilience of urban structures, while Bianchi et al. (2022) emphasised broader environmental, social, and economic dimensions. Others, such as Capolongo et al. (2020), examined resilience in critical infrastructure like medical centres. These studies illustrate that resilience is multidimensional, and machine learning is increasingly positioned as a core enabler of resilience-oriented urban planning.

The integration of machine learning into urban resilience strategies is driven by three major technological trends:

1. Proliferation of sensors and Internet of Things (IoT) devices has enabled cities to generate vast amounts of real-time data.
2. Availability of ML-specific hardware (such as GPUs) has made it easier and more affordable to process these data streams.

3. Advances in algorithms, especially in deep learning, have allowed for more complex analyses. Techniques like Convolutional Neural Networks (CNNs) have proven effective in flood simulation and early-warning systems, while Generative Adversarial Networks (GANs) are being used to create synthetic climate data, such as simulated typhoon paths, to inform preparedness planning (Chen et. al., 2025).

A typical machine learning workflow in the context of urban resilience begins with data collection, followed by model development and deployment. Depending on the problem, different approaches may be applied - supervised learning for tasks requiring labelled datasets, such as classification or regression; unsupervised learning for clustering and dimensionality reduction; and reinforcement learning for adaptive decision-making in uncertain environments. Each approach has its strengths and limitations, but together they provide a robust toolkit for addressing complex urban risks.

The need for these methods is particularly acute as cities face growing pressures from climate change, rapid urbanisation, and frequent disasters. International frameworks such as the Hyogo and Sendai strategies emphasise the importance of risk reduction, participatory governance, and sustainable urban planning. ML directly supports these objectives by providing cities with tools to anticipate hazards, enhance service delivery, and engage citizens in real-time decision-making. By enabling smarter infrastructure management, hazard forecasting, and resource allocation, machine learning helps align local governance with the Sustainable Development Goals, particularly SDG 11 on creating inclusive, safe, resilient, and sustainable cities.

Practical applications of machine learning in urban resilience are already visible across the globe. Early Warning Systems (EWS) have been deployed in flood-prone regions such as Zambia, where alerts generated through machine learning models help communities evacuate and enable governments to direct resources effectively. In parallel, open data innovation leverages publicly available datasets and citizen-generated content to identify vulnerabilities, monitor infrastructure, and improve mobility management. Big data analytics extends this potential further by processing massive streams of sensor, GPS, and mobile data to detect risks, predict failures, and guide adaptive responses. While these methods face challenges such as data quality, sensor reliability, and privacy concerns, they nonetheless represent a significant step toward data-driven resilience.

ML is emerging as a critical enabler of urban resilience. By integrating predictive analytics, pattern recognition, and adaptive learning into governance systems, cities can move from reactive crisis management to anticipatory risk mitigation. These tools allow administrators to identify vulnerabilities, model potential disruptions, and allocate resources more effectively — reducing the social and economic costs of disasters.

Beyond enhancing response capacity, machine learning supports continuous learning within city systems. It enables decision-makers to draw insights from past events, update planning models in real time, and adapt to evolving climate and infrastructure risks. When embedded in local governance frameworks, such data-driven foresight strengthens institutional preparedness and reinforces public confidence. The effective application of machine learning in resilience planning

thus demonstrates how technology, when guided by human judgment and inclusive design, can help build cities that are not only smarter and safer but also more equitable and sustainable in the long term.

This potential makes machine learning an indispensable tool in cities' efforts to fulfil the objectives of global agendas, such as the Sendai Framework and the UN Sustainable Development Goals, anchoring technology as a force multiplier for sustainable urban governance.

Examples of Data-Driven Urban Resilience Efforts

Region	Problem addressed	Solution implemented	Impact
Bihar	The <i>Visceral leishmaniasis</i> or Kala-azar vector strain	RBF-kernel-based-SVR models set up to monitor its spread	Assess non seasonal climatic changes that could contribute to its rapid spread. Provide timely medical intervention and services in the event of an outbreak.
Egypt	Effective agricultural water management in water-stressed developing countries	Random Forest Model (RFM)	Accurately estimate evapotranspiration Forecast vapor pressure deficit
Manila	Water leakage and theft	Asterra and Infracore, precision tools and AI based technology	Detected nearly 2,000 pipe leak cases Identify and correct problems pertaining to water loss and theft
Zambia	Flood-prone regions	Early Warning Systems (EWS)	Alerts generated help communities evacuate Enable governments to direct resources effectively.

NLP and ML's Harm Potential and Mitigation Strategies

NLP and ML systems often inherit (and in some cases amplify) the gender biases embedded in their training data. Because language reflects long-standing social norms and power structures, models learn stereotypical associations that link women with emotion, family, or morality, and

men with leadership, technical expertise, or agency (Sun et al., 2019; Leavy et al., 2020). These distortions appear in multiple forms: skewed representation of women in text corpora, gender-stereotyped word embeddings, misgendering in translation, and biased outputs in downstream tasks such as hiring algorithms, chat systems, or search (Sun et al., 2019). The ‘black box’ nature of many AI systems heightens the risk, making it difficult to understand how gendered assumptions shape decisions. In public domains such as planning, transport, and civic decision-making these systems can unintentionally reinforce structural inequalities and diminish the visibility or influence of marginalized groups (Sanchez et al., 2024).

Mitigating these risks requires interventions across the entire AI pipeline. At the data level, corpora should be systematically audited to identify gendered linguistic patterns, under-representation, and harmful associations, using frameworks informed by sociolinguistics and gender theory (Leavy et al., 2020). Data augmentation techniques, such as gender-swapping or name anonymisation, can help rebalance representations, while careful corpus curation reduces the likelihood of encoding historical stereotypes (Sun et al., 2019). At the model level, embedding debiasing, adversarial training, and constrained prediction techniques can reduce gender leakage and stereotyping, though these methods must be tested for unintended side effects (Sun et al., 2019). Transparent and explainable model designs, combined with ethical oversight, participatory processes, and continuous monitoring, help ensure that AI systems remain accountable and aligned with diverse gender realities (Sanchez et al., 2024). Ultimately, technical measures must be paired with interdisciplinary expertise and community engagement to meaningfully reduce gender bias and support equitable AI outcomes.

Computer Vision

Computer Vision (CV) is a field of artificial intelligence that enables computers to interpret and extract meaningful information from images and videos, automating the perception tasks that normally require human vision (Marasinghe et. al., 2024). Its role in urban planning can be explained as follows:

1. Automated Data Collection and Mapping

- a. Extracts information from street-view images, satellite imagery, and aerial photographs.
- b. Maps land use, monitors urban growth, and detects changes in the built environment at scale.

2. Objective, City-Wide Environmental Quantification

- a. Measures features such as green space, building density, sidewalk quality, and signs of urban decay.
- b. Overcomes the limitations of traditional surveys, which are time-consuming, costly, and prone to bias.

3. Data-Driven Identification and Prioritization of Urban Issues

- a. Statistically correlates physical environment metrics with public health indicators (e.g., linking lack of greenery to higher asthma prevalence).
- b. Analyses relationships between streetscape design and crime patterns.

- c. Connects urban form and aesthetic characteristics to economic vitality (e.g., property values, retail activity).
- d. Enables planners to pinpoint areas most in need of interventions and allocate resources more effectively.

4. Support for Planning and Design Processes

- a. Helps generate design alternatives and 3D visualizations.
- b. Enhances public participation through interactive and immersive VR-based planning tools.

5. Monitoring and Evaluation of Interventions

- a. Tracks compliance with planning regulations.
- b. Assesses long-term performance and spatial impacts of implemented policies or urban improvements.

Computer Vision's integration into urban planning must be approached strategically, as its effectiveness depends on addressing key challenges such as data quality issues, biases embedded in training datasets, significant privacy concerns arising from large-scale image analysis, and the high levels of technical expertise and computational resources required. Rather than replacing human planners, CV should be used to augment their work; when carefully and responsibly applied, it strengthens decision-making processes and supports the development of more sustainable, equitable, and liveable cities.

Cities today operate within complex data ecosystems, from administrative records and citizen feedback to sensor networks and service platforms. Technologies such as NLP and machine learning enable city leaders to interpret this information in real time, translating data into insights that guide planning, service delivery, and community engagement. When deployed responsibly, these tools extend far beyond efficiency gains: they can strengthen inclusivity by ensuring that decisions reflect diverse citizen experiences and needs. Beyond NLP machine learning, and Computer visioning AI encompasses a wide variety of tools and approaches, including computer vision, recommendation systems, robotic process automation, and more. Each of these offers unique capabilities of relevance to cities, depending on their priorities and data environments.

Artificial Intelligence is no longer a futuristic concept; it is actively redefining how urban systems function. Applications ranging from predictive maintenance of infrastructure to demand forecasting and environmental monitoring are already demonstrating tangible improvements in responsiveness and resource use.

For policymakers and administrators, the task ahead is to move beyond experimentation toward institutional adoption — building the governance capabilities, ethical frameworks, and technical fluency needed to harness AI effectively.

Urban AI in Action:

Current and Emerging Use Cases, Implementation, and Scaling Barriers

As discussed in the previous chapter, the use of Artificial Intelligence (AI) in cities is no longer just a futuristic idea; it is already reshaping how urban areas function, how services are delivered, and how governments connect with their citizens. This section takes a closer look at the role of AI in urban development by exploring both current applications and emerging possibilities. It also examines the barriers that make adoption and scaling difficult, while highlighting the strategies organisations are using to navigate them. Drawing on insights from key informant interviews with institutions such as the Asian Development Bank (ADB), World Bank, eGov Foundation, CivicDataLab, and Wadhvani AI, among others, the analysis offers a window into the diverse and evolving landscape of urban AI.

Cities worldwide are embracing AI to make everyday life better for their residents by working more efficiently, gaining real-time insights, and finding new ways to deliver services. For example, in Bengaluru, AI-driven adaptive signals and smart enforcement are reducing wait times at junctions and making travel along major corridors more predictable. This is implemented through the Bengaluru Adaptive Traffic Control System (BATCS). Singapore uses AI to predict water needs more accurately, ensuring a steady supply for homes and businesses (Brears, 2024). In Seoul, automated waste collection keeps neighbourhoods cleaner with less hassle (Aclima, 2017). These examples highlight how AI not only streamlines city operations but also helps governments connect more closely with citizens and bring faster, smarter improvements to urban services that people rely on every day.

Use Cases Across City Functions

In this chapter, we delve into a few of the urban sub-sectors to examine and document how AI is making a real difference across many parts of city life. This includes helping improve governance, infrastructure, social services, environmental management, public safety, and administration.


Governance & Participation

Artificial Intelligence is transforming urban governance by enhancing the ability of city systems to understand, anticipate, and respond to citizen needs that could sometimes be complex. Its applications now span a wide spectrum — from AI-powered chatbots that simplify public interactions and improve service access, to predictive analytics that guide infrastructure investments and resource allocation. These tools enable municipal leaders to identify patterns across departments, act on real-time information, and deploy resources where they are needed most.

Beyond operational gains, AI also holds potential to strengthen transparency and participation. By making citizen engagement more continuous and data-driven, it allows governments to respond to feedback with greater agility and accountability. For city administrations, the challenge lies not only in adopting these technologies, but in embedding them within governance frameworks that value inclusivity, trust, and equity — ensuring that smarter cities also become fairer and more responsive to those they serve.

Citizen Engagement

- In Nepal, work under the Digital Nepal agenda and various local governance programmes has helped municipalities roll out online grievance and issue-reporting systems. These platforms let residents quickly flag problems such as damaged roads, water supply issues, or broken streetlights via phone or internet, receive regular updates on status, and see when cases are resolved—strengthening accountability and trust between citizens and local government (2019 Digital Nepal Framework).
- Bengaluru has similar AI tools that speed up grievance redressal, so people spend less time waiting for responses (NITI Aayog, n.d.).
- CivicDataLab is co-creating ParakhAI — an open-source framework for Participatory Algorithmic Auditing of AI solutions, especially high-risk algorithmic decision-making systems (ADS)— with the Ministry of Electronics and Information Technology and Digital India Programme. This is developed to enforce the adoption of a continuous lifecycle with Public Consultation, Development, Deployment and Assessment as key steps built into the lifecycle. This ensures adequate avenues for impact groups to timely and effectively participate, collaborate and co-create the ADS system (CivicDataLab, n.d.).



- In Mumbai, city officials use AI to understand what people are saying in public feedback, helping them anticipate and respond to community needs sooner (Municipal Corporation of Greater Mumbai, Information Technology Department, n.d.).

Land Use and Property Tax

- In the Indian state of Tamil Nadu, AI supports land governance by tracking progressive land use changes and alerting officials to unauthorised constructions and encroachments (The Hans India, 2024).
- The DIGIT platform deployed in Andhra Pradesh, a southern state in India, is designed to manage various municipal services, including property tax payments, grievance resolution, and licensing. The platform helps cities establish clarity in accountability, hierarchies, and workflows, particularly in smaller towns where these structures might be less defined.

While DIGIT can function without AI, it serves as a foundation for generating real-time transactional data that AI models can consume to provide insights and assistance³.

- Ahmedabad's Municipal Corporation is working on developing an AI solution for fairer property tax assessments, helping improve city services funded by those taxes (Around Town Realty, 2025). Behind the scenes, AI is also helping government offices manage paperwork faster, so public services are delivered more quickly and transparently, making residents' interactions with the city easier and more reliable every day.

Operations & Infrastructure


Across many Asian cities, the growing integration of digital systems is transforming how essential services are managed and delivered. Artificial Intelligence now plays a critical role in this shift, automating routine functions and equipping administrators with real-time insights for more coordinated decision-making. In areas such as traffic management, water distribution, and solid waste collection, AI-driven tools are helping city teams optimise routes, balance supply and demand, and address service disruptions with greater speed and precision.

For residents, these advances translate into smoother commutes, more reliable utilities, and cleaner public spaces. More broadly, AI enhances coordination within city governance, enabling departments to respond more effectively and ensuring that essential services function more consistently for residents. In doing so, it reinforces a model of governance that is efficient, citizen-centred, and grounded in long-term sustainability.

Traffic Management

- In India, the Intelligent Traffic Management System (ITMS) in the city of Jaipur is transforming daily commutes by dynamically adjusting signal timings and using smart cameras to track violations (C-DAC, n.d.). This innovation helps reduce waiting times at intersections and improve traffic flow, addressing the challenges of a rapidly growing city with less reliance on manual control. The system is also expected to increase road safety by improving traffic rule enforcement and emergency response.
- In China, Hangzhou's 'City Brain' system is often highlighted worldwide for its remarkable use of real-time data—from sensors, traffic cameras, GPS devices, and public records—to optimise traffic signals, quickly identify incidents, and recommend the best routes for drivers. Since its introduction, the city has experienced a major drop in traffic congestion, making travel more reliable and efficient for everyone (The Atlas of Urban Tech, n.d.).
- Indian cities Surat and Kanpur are following this lead, adopting AI to tackle traffic issues. Surat synchronises signal timings and deploys thousands of AI-enabled CCTV cameras to monitor road conditions and enforce compliance, enhancing road safety and reducing bottlenecks (NMSC, 2025). In Kanpur, advanced AI systems work alongside academic institutions such as IIT Kanpur to manage traffic flow and interpret CCTV footage, enabling quicker responses to incidents and smoother movement across the city (Times of India, 2023).

³ eGov Foundation, interviewed by author, August 2025



Surat city using AI-enabled CCTVs for monitoring road conditions (Image Source: [Surat Municipal Corporation](#))

Waste Management

AI is revolutionising the world of waste management, bringing innovative solutions that enhance efficiency, effectiveness, and sustainability. In a time when the global waste crisis demands urgent action, AI emerges as a game-changing force. Far from just sorting robots, AI transforms waste management into a sophisticated, data-driven practice. Its impact is profound and multifaceted, with statistics revealing its extraordinary potential across various roles.

- Tokyo employs intelligent garbage trucks equipped with AI-assisted navigation and water recycling technology, allowing for smarter route planning and more consistent cleaning schedules (CSCTRUCK, 2024).
- Similarly, Ahmedabad, Delhi, and Bengaluru are making strides with AI-powered waste sorting, waste processing and real-time tracking systems (ICTPost, 2024). Innovations in sorting technology and applications that connect citizen reports with collection services lead to more efficient recycling and cleaner neighbourhoods.


Water consumption map of the Fergana region in Uzbekistan (Image Source: [ESRI](#))

Water Management

- Water supply management in Tashkent, Uzbekistan, has improved with the help of an AI-integrated system that monitors the infrastructure continuously. This setup detects leaks early and balances water distribution, significantly reducing losses and ensuring a steady supply to homes and businesses. Alongside this, training programmes are building local

capacity to operate cutting-systems for (flow DB,

- Maynilad Services, supplies Metro surrounding Philippines,



maintain and these edge the long term (2024).

Water which water to Manila and areas in the has taken a

leading role in using artificial intelligence to find leaks in underground pipes and keep the water network running smoothly. Using an AI system called Infrawise, developed by a Portuguese company, Maynilad analyses its vast network of pipes to pinpoint sections most likely to leak. Since launching a pilot in October 2023 covering over 1,700 kilometres of pipelines, Infrawise has detected vulnerabilities in nearly half of the length of the pipelines and found more than 1,500 actual leaks. The company plans to extend this coverage to another 1,500 kilometres. Beyond this, Maynilad also employs AI combined with satellite images to spot the unique signs of water leaking underground. This information is mapped out using Geographic Information Systems (GIS) to target repairs faster, reducing the need for slow, manual inspections. Thanks to these AI-powered tools, Maynilad has improved its ability to manage water loss, prioritise pipe repairs, and conserve precious water—especially important during times of scarcity (Maynilad Water Services Inc., 2024).



A Maynilad engineer is shown software on his tablet to aid an activity. Maynilad is now


intelligence to identify areas in its water distribution system with high likelihood of leaks, thus enabling the company to detect pipe leaks more quickly and efficiently. (Image source: [Maynilad Water](#))

using the Infrawise ongoing leak detection utilising artificial

Urban Planning and Land Governance

In the realm of urban planning and land governance, AI is helping cities become more proactive and organised.

- Authorities in Ho Chi Minh City use machine learning on satellite and geospatial data to monitor land use changes and identify illegal construction swiftly. An app empowers both officials and citizens to report zoning concerns, enabling timely interventions (Goldblatt et al., 2018).
- Singapore's Smart Nation initiative leverages AI, IoT sensors, and spatial data to support sustainable urban growth and monitor the health of structures (GOVTECH Singapore, 2024).
- Back in India, the Tamil Nadu government has proposed employing AI combined with drones and GIS to detect land and property encroachments and ensure compliance with planning regulations, contributing to safer, more orderly urban environments. With the pilot in Coimbatore, it will be developed from a database of existing plot boundaries and existing building footprints obtained from city master plans, local body building records and layout approvals. Using this data, the AI system will be able to identify changes in land use on both private and government properties, unauthorised constructions and encroachments, and alert officials to these land-use changes. The AI system will also be able to track deforestation and help preserve green cover (The Hans India, 2024).
- In Bengaluru, the Bruhat Bengaluru Mahanagara Palike (BBMP) has announced plans to use AI to detect unauthorised constructions across the city. The initiative reflects how urban administrations are increasingly turning to AI for governance efficiency and transparency. However, its success will depend on institutional readiness — including data quality, inter-departmental coordination, and safeguards to ensure fairness and accountability in enforcement (The New Indian Express, 2025).



Ho Chi Minh City incorporates geospatial data to help identify illegal construction (Image Source: [Goldblatt et al., 2018](#))

Social Services

As the use of Artificial Intelligence expands beyond city operations and infrastructure, its potential to strengthen social services is becoming increasingly evident. In areas such as public health, legal aid, and welfare delivery, AI offers tools that can make systems more responsive and equitable. Yet the true measure of its impact depends on how responsibly these technologies are designed, deployed, and governed.

Wadhwani AI's Multi-Faceted Suite

Across India and other parts of Asia, initiatives like those led by Wadhwani AI illustrate what this can look like in practice — where innovation is anchored in institutional partnerships, ethical safeguards, and a clear public purpose. Their approach underscores a broader principle: that the success of AI in social sectors will rest not only on technological advancement, but on the values and governance systems that guide its use⁴.

Wadhwani AI's work—spanning agriculture (cotton pest management), education (speech recognition for oral reading fluency), and health (maternal, newborn, and child health, tuberculosis)—consistently places the onus on real-world impact in the developing world. Their initiatives illustrate that AI's societal value emerges only when solutions target truly pressing problems, have accessible data, and are integrated with existing government programmes. The HealthVaani AI Assistant, for example, was conceptualised to help frontline workers (AWWs, ASHAs, ANMs) access actionable reports and beneficiary data across platforms, but was only built after a staged vetting of data quality, user needs, and institutional readiness—the same principles guiding their telemedicine collaborations, where AI-augmented symptom analysis always positions the doctor as the final arbiter.

⁴ Wadhwani AI, interviewed by author, August 2025

AI's Social Impact in Asia

- In Bihar, AI-driven kala-azar risk modelling supports but does not replace local health teams' judgement (Kumar et. al., 2024).
- ARMMAN's Kilkari predictive messaging is effective because it operates within established health system workflows and prioritises informed consent and privacy (ARMMAN, n.d.).
- Courtroom innovations like Adalat AI's legal chatbots and summaries thrive in environments where judiciary partners co-create and vet solutions that uphold transparency and procedural fairness.⁵



Singapore's smart estates must integrate modern technology to create sustainable ecosystems
(Image Source: [Alwi & Yun Ting, 2024](#))

Challenges persist—systemic data gaps, skill shortages, infrastructure deficits, and cultural reticence can all stymie scale. The approach of solution providers such as Wadhvani AI is to partner with anchor government agencies or leverage existing digital public goods, to structure in human oversight (“human-in-the-loop”), and to deploy only when there is a clear readiness and pathway to scale, which is echoed by successes elsewhere. Singapore’s SmartEstates (GovInsider, 2017) and Indonesia’s ClimateSmart Indonesia (Independent Observer, 2025) both show that achieving material and social benefits from AI is dependent on regulation, capacity building, and supportive policy environments. In Delhi, as municipal chatbots broaden digital access, their long-term value will rest on maintaining robust data privacy, user support, and clear accountability mechanisms (Times of India, 2025).

⁵ <https://www.adalat.ai/>

It is evident that building responsible pathways for AI in social services requires looking beyond technology adoption to the readiness of the institutions that use it. This means putting in place clear frameworks for testing and oversight, designing solutions in partnership with users, and investing continuously in the skills and systems that sustain them. When governance structures are designed to prioritise trust, safety, and accountability, technology becomes a genuine enabler of long-term impact.

For many Asian cities, this discipline will define the difference between AI that expands opportunity and AI that deepens divides. The goal is not simply to scale technology, but to shape its use around real human needs and institutional realities. When AI initiatives are grounded in empathy, context, and accountability, they move beyond experimentation to become instruments of genuine public value.


Environment & Climate

Artificial Intelligence is beginning to reshape how cities anticipate and respond to climate risks, offering a powerful complement to traditional approaches to urban resilience. By drawing on real-time data and predictive analytics, these tools allow city systems to act before hazards escalate, protecting lives and livelihoods in the process. In Jakarta, for example, an AI-powered flood management system uses sensor networks to monitor watersheds continuously, issue early warnings, and coordinate rapid mitigation efforts when water levels rise (IEEE Smart Cities, 2024).

CivicDataLab's IDS-DRR

AI solution provider CivicDataLab's Intelligent Data Solution for Disaster Risk Reduction (IDS-DRR) translates the pillars of the Sendai Framework—hazards, exposure, and governance—into a practical, decision-support platform⁶. Its dashboards help district administrators visualise flood risks, assess potential impacts on households and social services, and prioritise preventive action and capacity building. Together, these initiatives demonstrate how AI, when integrated thoughtfully into governance systems, can strengthen institutional preparedness and make urban resilience both data-informed and people-centred. CivicDataLab's co-creation approach, developed in partnership with stakeholders including officials in Assam, emphasises that AI tools must evolve through constant collaboration with government users and citizens to remain relevant and trusted. Assam's institutionalisation of the IDS-DRR platform into its flood risk planning process illustrates how the effective use of AI depends on embedding it within existing structures and workflows. However, their work also highlights challenges—such as the scarcity and non-machine-readable formats of key datasets, delays in real-time data availability, and the need for ground-truthing to ensure AI insights reflect local realities—all barriers that cities must navigate to realise AI's full potential.

⁶ CivicDataLab, interviewed by author, August 2025



Screengrab of the IDS-DRR platform developed by [CivicDataLab](#)

AI's Utility in Climate Resilience

- Elsewhere, Singapore's flood response prediction models, Kuala Lumpur's real-time air quality monitoring, and Shanghai's climate-responsive urban planning exemplify how AI-informed urban management extends across domains.
- In Bhopal, the city administration has deployed an integrated smart city solution powered by AI to strengthen its urban management systems. The network of "smart poles" combines AI, machine learning, and Internet of Things (IoT) technologies to give the city a real-time view of what is happening across multiple domains—traffic, public safety, waste collection, and environmental monitoring⁷.
- Andhra Pradesh's Transmission Corporation in 2020 had proposed the use of AI for accurate energy demand forecasting, enabling better resource allocation despite climate variability (IndiaAI, 2020). The model uses ML to meticulously plan power management, ensuring a more responsive and adaptive approach to the ever-changing dynamics of electricity consumption. It was developed in-house with technical assistance offered from Google and support from the World Bank.

Together, these examples highlight that the success of AI in addressing climate and environmental challenges rests as much on institutional readiness as on technological capability. Effective integration depends on systems that can manage data responsibly, operate transparently, and involve citizens meaningfully in both design and decision-making. Building truly adaptive and citizen-centric urban systems requires equal attention to these governance foundations, not just to data or model performance. When guided by these principles, AI can strengthen cities' ability to pursue climate resilience that is not only scalable but also sustainable and fair.

Safety & Security

As AI-enabled innovation becomes embedded across urban governance, infrastructure, and environmental management, its application to city safety and security offers cities an opportunity

⁷ Artificial Intelligence Use Case Compendium by MoHUA

to shift from reactive policing to anticipatory, citizen-centric protection—blending real-time awareness, predictive analytics, and community engagement to keep public spaces safer and more inclusive.

Traffic Safety

In Pimpri-Chinchwad, Maharashtra, more than 2,500 AI-enabled CCTV cameras feed into an Integrated Command and Control Centre that monitors key traffic corridors, high-incident zones, and critical urban assets. The system uses real-time facial and vehicle recognition, anomaly detection, and automated e-challan generation to support faster incident response and more consistent enforcement. Beyond improving traffic management and law enforcement efficiency, the initiative offers a scalable model for proactive urban safety — one that relies on strong data practices, coordination across agencies, and clear accountability for how information is used⁸.


Women's Safety

In Metro Manila, the SafeHer initiative illustrates how AI can strengthen inclusion and women's safety — areas that have long received limited attention in conventional urban policing. The platform integrates AI-based features such as scream and crash detection, automated SOS alerts, and live location sharing to support women commuters facing harassment or violence. Its community-based, invite-only design, combined with medical ID integration and direct links to emergency services, reflects a user-led approach that prioritises trust and safety. Crucially, SafeHer was developed in close partnership with local transport authorities, ensuring that technology complements — rather than replaces — existing social and legal support systems. The initiative offers an important lesson for cities: that responsible AI design must be rooted in collaboration, empathy, and institutional adoption if it is to deliver safety and confidence in public spaces⁹.

However, the effectiveness and legitimacy of such innovations depend less on the sophistication of the technology and more on the governance frameworks surrounding it. Transparent data use, auditable algorithms, and community feedback mechanisms are essential to ensure accountability and prevent bias or overreach. As these systems scale, embedding ethical safeguards and participatory oversight will be critical to maintaining public trust and ensuring that technology strengthens — rather than undermines — the social contract between citizens and the state.

⁸ Artificial Intelligence Use Case Compendium by MoHUA

⁹ <https://safeher.ph/>



Video grab of the AI-powered features (Image Source: [SafeHer App](#))

GiveDirectly – Aid Targeting in Bangladesh

GiveDirectly's AI-targeted aid pilot in Bangladesh¹⁰ was implemented to explore whether satellite imagery and mobile data could be used to identify and deliver cash assistance to the poorest households more efficiently than traditional survey-based methods. Partnering with the government and local telecom operators, the organisation used machine learning models trained on satellite images, phone metadata, and limited ground-truth data to predict household wealth and map poverty across rural districts. These predictive insights were then used to pre-select potential beneficiaries, who were subsequently contacted through mobile outreach and verification calls to confirm eligibility and consent before transfers were disbursed via mobile money. The implementation aimed to significantly reduce the cost and time required to target beneficiaries while improving coverage and accuracy compared to manual targeting methods.

On the ground, GiveDirectly worked closely with local enumerators and community representatives to validate AI predictions, ensure ethical standards, and address exclusion and inclusion errors. The team faced challenges related to incomplete or outdated telecom data, varying image resolution, and limited digital literacy among potential recipients, which sometimes complicated the verification process. However, the pilot demonstrated that AI-assisted targeting could identify poor households with reasonable accuracy at scale, even in data-scarce contexts. It also highlighted the importance of local adaptation, combining algorithmic efficiency with human oversight and community trust, as critical to both implementation success and ethical integrity.

¹⁰ GiveDirectly, interviewed by the author, October 2025

What to Watch Out For – Risks, Ethics & Bias

Privacy and Safety in the Context of AI


One of the central challenges in the use of artificial intelligence concerns privacy and security. AI systems depend on large volumes of personal data, and ensuring that this data is collected, processed, and stored in ways that respect individual consent and comply with relevant legal frameworks is essential. As Rössler (2005) notes, privacy violations can take different forms, such as unwarranted interference in personal actions, illicit surveillance, and unauthorised intrusion into private spaces, all of which may occur in connection with AI (Calo, 2010).

The risks are not abstract. For example, AI-driven robotic applications can be subject to cyberattacks, and a successful breach does not merely compromise data but may allow access to a person's physical living environment, which is the most intimate domain of privacy (Calo, 2012). Similar concerns have been raised around government use of AI in surveillance, which can pose serious risks to civil liberties (Gasser & Almeida, 2017). Unsurprisingly, public opinion reflects these anxieties: recent surveys show that a large majority of citizens fear that AI threatens their privacy (Soo, 2018).

The challenge of safeguarding privacy in AI thus spans in both technical and legal domains.

Figure 3:
safeguards

Privacy



Sources: Holdren & Smith, 2016; Krausová, 2017; Gasser, 2017

Recent national guidance is now giving shape to these concerns in the form of a clearer governance agenda for India. The India AI Governance Guidelines, released by the Government of India in 2025, acknowledge that AI can deepen risks such as deepfakes, misinformation, biased decision-making, loss of human oversight, and even national security vulnerabilities. At the same time, the guidelines aim to strike the right balance — encouraging innovation and experimentation, while ensuring that the necessary safeguards are in place to protect individuals, communities, and public institutions.

Technology think-tank iSPIRT argues that traditional consent-based data sharing is no longer enough in an AI ecosystem where data is repeatedly reused for analytics and model training.

They highlight India's Data Empowerment and Protection Architecture (DEPA) and the Account Aggregator (AA) framework as examples of how granular, auditable, and revocable data transactions can put individuals in control while reducing risks from broad, opaque consent practices. To support privacy-preserving AI, iSPIRT points to the growing use of privacy-enhancing technologies (PETs) such as differential privacy and Confidential Clean Rooms (CCRs), which allow organisations to collaborate and train models without exposing raw data. Building on this, they propose Data Collaboratives—neutral intermediaries that manage secure data-sharing contracts, onboarding, compliance, and value-sharing. A case from the pharmaceutical sector shows how sensitive clinical data can be pooled inside a Confidential Clean Room, protected with differential privacy, and used to train shared AI models without compromising confidentiality. Together, these approaches illustrate a path toward responsible, privacy-preserving data collaboration that can enable large-scale AI innovation across sectors¹¹.

AI and the Challenge of Discrimination

Another major concern in the deployment of artificial intelligence is the risk of discrimination. Research shows that AI applications, particularly those involving profiling, can produce unequal or unfair outcomes, undermining fundamental ethical principles such as equality and fairness (Thierer et al., 2017). Because AI systems are designed by humans and trained on human-generated data, they inevitably risk inheriting the biases and value judgments embedded in those inputs (Citron & Pasquale, 2014).

This risk becomes evident when AI models reproduce gender or racial prejudices present in language or behavioural data, sometimes even amplifying stereotypes rather than reducing them (Bass & Huet, 2017; Devlin, 2017). Such biases can manifest in hiring systems, predictive policing, financial services, or other sensitive domains, with significant implications for social justice and legality. Real-world inequities and injustices are ingrained in historical data, and if they are not considered when designing algorithms and solutions, they will be reinforced and exacerbated. For example, historical data may show higher crime statistics in low-income, marginalised neighbourhoods due to the historical oppression and injustices faced by these groups. However, an automated policing algorithm that is trained on this historical data will then over police these historically marginalised neighbourhoods, which can exacerbate the oppression they face moving forward. While the creators of an algorithm may have designed it not to consider race, for example, racial inequality is embedded into the data they are employing, so the algorithm actually must consider race and work to ensure that present inequalities in the data do not create results that reinforce them in the future (D'Ignazio & Klein, 2020).

In response, scholars and practitioners have begun developing methods within data mining and machine learning aimed at identifying and mitigating discriminatory patterns (Barocas, 2014; Bolukbasi et. al., 2016; Hajian & Domingo-Ferrer, 2013). These efforts include fairness-aware algorithms and auditing techniques designed to detect biased outputs.

¹¹ iSPIRT Foundation, 2025

Nonetheless, the task of fully preventing discrimination in AI remains unresolved. Despite early progress, ongoing research and policy innovation are needed to design systems that can not only recognise but also actively prevent unfair outcomes (Mittelstadt et al., 2016). Addressing this challenge is essential if AI is to serve society in ways that align with principles of justice and equality.

Resource-Intensive Nature of AI

Another ethical concern around AI is the extent of its resource-intensive nature. AI is often perceived as a cloud-based, bodiless technology, making it easy to ignore the massive resource extraction, energy and water consumption, waste production and exploitative labour that is required to power AI (Heilinger, 2022). A typical data centre that powers AI consumes as much electricity as 100,000 households; however, as AI becomes more widely deployed, large data centres are being constructed, which are estimated to consume 20 times that (IEA, 2025). Further, large data centres, such as those that support Google applications, consume approximately 2.1 million litres of water per day (for cooling purposes), and these data centres are expected to grow rapidly due to the exponential increase in computational demands for AI. These are just a few of the environmental impacts of AI, not including the energy and resources needed to produce the required hardware (Zhang, 2024).

It is, therefore, crucial that attention is given to the negative environmental consequences that come from running AI applications (Heilinger, 2022). To consider these environmental consequences, AI models should be evaluated based on their energy consumption and carbon footprint, alongside other evaluation factors.

Before AI is used to address a challenge, it must be considered whether it is a necessary and suitable means for solving the problem, considering all trade-offs. If the AI solution is deemed the best means for solving a problem, public agencies can work to adopt lightweight models and use techniques like model pruning, transfer learning, quantisation and use of energy-efficient hardware to minimise its environmental impact (Schwartz et al., 2020; Heilinger et al., 2023).

Ethical Principles

Addressing the issues of ethics, fairness, and bias in AI systems requires a set of overarching guiding principles that can help governments formulate policies and legislation. Numerous multilateral agencies and governmental bodies have published a set of principles, with little variance between perspectives. Among these, UNESCO's list of ethical AI principles based on core human rights principles serve as useful references as they are applicable universally:

UNESCO's Core Human-Rights-Centred Principles



Overview of regulatory frameworks

Characteristics	Asean AI Guide	EU AI	OECD	Model AI Governance Framework, Singapore
Goals	-Ethical designing of AI -Development and Deployment	-Create ethical and trustworthy AI -Position Europe as a global leader for AI standards	-Inclusive growth and development -Knowledge sharing	Responsible AI deployment and development
Principles	-Fairness and Equity -Safety and security -Transparency and accountability	-Classifies AI systems based on risk factors -Human-centric and focuses on ethicality	-Human-centric AI policies and values -Safety and security -Transparency and accountability	- Human-centric -Explainable and fair -Transparency
Strengths	Flexible approach	-Comprehensive framework set for global standards -Human oversight and transparency	-Evidence-based research and analysis -Forum for dialogue -Influential legal instruments and guidelines	-Agile governance -Compatible with global standards -Multi-stakeholder collaboration
Limitations	-Digital gaps -Non binding guidelines -Uneven policy focus	-Gaps and loopholes in prohibitions -High implementation costs -Stringent regulation leading to disadvantages	-Complex innovation policy -Data coverage gaps -Western-centric bias	-Copyright issues -Data coverage gaps -Voluntary in nature

Data Readiness & Digital Infrastructure

Artificial intelligence has the potential to transform urban governance, from optimising traffic and energy systems to improving healthcare and sanitation. Yet, for AI to deliver these outcomes, two interdependent pillars must be strong, i.e. data quality and digital infrastructure. Without precise, interoperable data and the infrastructure to process and store it efficiently, even the most sophisticated AI algorithms yield unreliable results.

Data Quality

High-quality data determines whether AI enhances or undermines public service delivery. In the case of India, where digital governance has already achieved scale (due to digital public infrastructure like Aadhaar, UPI, Ayushman Bharat, etc.), the problem of precision persists. A single erroneous data point can delay a pension, misroute a subsidy, or deny healthcare to vulnerable citizens. This shift from expansion to fidelity underscores a central truth: AI models are only as intelligent as the data they learn from.

In urban public services, where AI systems depend on vast streams of data from sensors, registries, and citizen interactions, poor data quality can create serious distortions, such as -

- Inaccurate targeting of beneficiaries in welfare schemes due to duplicate or missing records.
- Faulty AI predictions in areas like traffic management, waste collection, or air quality if sensor inputs are inconsistent or outdated.
- Erosion of public trust when citizens encounter errors or opaque decisions driven by flawed data.

To prevent this, NITI Aayog (the apex public policy think tank of the Government of India) proposes a 'Data Quality Scorecard' that is built on six measurable attributes of data quality, i.e. accuracy, completeness, consistency, timeliness, validity, and uniqueness. These attributes enable departments to treat data as a public asset, with stewardship roles, validation checks, and regular audits. NITI's Data Quality Scorecard and Maturity Framework offer practical instruments for agencies to measure, monitor, and improve data health over time (NITI Aayog, 2025).

For AI in urban governance, these mechanisms translate to trustworthy datasets and ensure, for example, that a smart water management system does not misallocate supply due to faulty meter readings, or that predictive policing tools operate on verifiable and bias-free records.

Some states, such as Tamil Nadu (2022), Assam (2022), and Maharashtra (2024) have drafted State Data Policies to address some of the challenges through a robust data governance framework and a dedicated State Data Authority responsible for monitoring and enforcement.¹²

¹² Maharashtra State Data Policy <https://tinyurl.com/2kknbk2y>

Tamil Nadu Data Policy https://tnega.tn.gov.in/assets/pdf/TN_DataPolicy_2022.pdf

Assam State Data Policy https://it.assam.gov.in/sites/default/files/public_utility/assam_state_data_policy_2022.pdf

These policies provide the framework and guidelines for publishing of open data and collation of data, secure access to non-open data to catalyse the research and policy making by line departments, and external, bonafide research and statistical agencies collaborating with the government to strengthen policy design and implementation.

Digital Infrastructure

While quality data fuels AI, digital infrastructure forms its circulatory system. The rise of AI is dramatically reshaping global data centre and infrastructure requirements. As AI workloads multiply, cities and governments must modernise their compute, storage, and network capacities to handle the demands of model training, real-time analytics, and secure data sharing. In the context of urban public service delivery, this means:

- **Edge and cloud integration:** AI-enabled traffic systems or emergency response tools rely on low-latency processing (often at the edge) combined with centralised cloud models for learning and optimisation.
- **Interoperability layers:** Following the principles of a Digital Public Infrastructure (DPI) approach, open APIs and shared schemas enable different departments (e.g., transport, health, waste management) to collaborate securely, creating a unified city data fabric.
- **Sustainable computing:** Given AI's energy intensity, cities must invest in green, efficient data centres and renewable-powered cloud services to ensure that AI growth aligns with environmental goals.

AI-ready infrastructure must balance computational power with scalability and sustainability. For governments, this translates to adopting AI-optimised cloud architectures, investing in secure data exchange platforms, and ensuring redundancy and resilience against system failures. Without these, even the best AI algorithms will falter under operational strain.

Convergence


In AI-enabled urban governance, data quality and digital infrastructure are not separate priorities but mutually reinforcing. High-quality, interoperable data enables AI models to deliver accurate insights, while robust infrastructure ensures that these models operate at scale, securely and efficiently. For example:

- A smart waste management system depends on reliable sensor data (quality) transmitted through resilient IoT networks (infrastructure).
- An AI-based health surveillance system must integrate hospital, pharmacy, and citizen health data using secure, standardised APIs (infrastructure) while maintaining accuracy and privacy (quality).

- A citizen service chatbot can only build trust if it accesses verified, updated records and operates on dependable digital platforms.

Governing bodies stress that ecosystems such as these must also institutionalise accountability, where each dataset should have a clear steward, and every system should embed validation at the point of capture. Additionally, sustainable AI infrastructure is not a one-time investment as it requires ongoing optimisation, capacity planning, and cross-sector coordination.

Figure 4: Importance of data quality and digital infrastructure



Scaling AI Responsibly

Building a Policy–Pilot–Programme Loop for Cities

The successful implementation of a single AI pilot project is a significant achievement, but the true transformative potential for cities is realised only when these solutions are scaled across departments and services. Moving from isolated proofs-of-concept to organisation-wide, platform-level AI is the central challenge for municipal governments.

However, scaling AI in cities is never a linear path; it's a process of learning, unlearning, and adapting—a loop where policies make way for experimentation, pilots generate insights which inform programmes, and feedback on/evaluation of the programmes reshape the policies. Responsible scaling is often shaped by this dialogue between innovation and governance, where cities continue to iterate while being anchored in public purpose.

Based on an analysis of cross-industry research, this section outlines the primary barriers to scaling AI in an urban context and provides a strategic framework for overcoming them.

Primary Challenges to Scaling AI

Urban governments face a unique convergence of technical, organisational, and societal hurdles when attempting to scale AI.

1. Socio-Technical Misalignment

Scaling AI is not merely a technological upgrade; it is a fundamental organisational transformation. Research highlights that failure often stems from treating AI as a purely technical project while neglecting the essential social components (Haefner et al., 2023). In a city context, this manifests as:

- **Technical Focus Only:** Investing in algorithms and data pipelines without a parallel investment in the workforce. This leads to a skills gap where city employees lack the training to use, interpret, or manage AI outputs effectively.
- **Resistance to Change:** AI can be perceived as a threat to jobs or a disruptive force to established workflows. Without a clear vision and change management, frontline staff and middle management may resist adoption, rendering even the most sophisticated tools useless.

2. Data Foundation Deficiencies

AI models are powered by data, and the public sector often struggles with the quality and accessibility of its data assets. This is consistently identified as the most significant barrier (WEF, 2023; Sjödin et al., 2021).

- **Data Silos:** Critical data is often locked within departmental silos (e.g., transportation, social services, utilities). An AI model for optimising city-wide resource allocation cannot function without integrated data.

- **Poor Data Quality and Interoperability:** Inconsistent data formats, legacy systems, and incomplete records make it difficult to build reliable and scalable AI models. Data collected for one purpose (e.g., permit applications) may not be "AI-ready" for another (e.g., predictive service demand modelling).
- **Privacy and Security Concerns:** As citizen data is involved, cities must navigate stringent data protection regulations. The fear of mishandling data can paralyse innovation and prevent the creation of the large, representative datasets needed to train effective and unbiased models.

3. The 'Pilot Purgatory' Trap

A common pitfall for cities is an abundance of successful pilots that never progress to full-scale deployment. This is characterised by a culture of experimentation that lacks a clear pathway to production (WEF, 2023). Cities may launch numerous small-scale AI projects for traffic management, waste collection, and building inspection, but fail to standardise the underlying infrastructure and processes needed to integrate these solutions into the core operations of the city, leading to fragmented impact and wasted resources.

4. Lack of Trust and Ethical Governance

For AI to be scaled in public services, it must earn the trust of both citizens and the government employees who use it.

- **Algorithmic Bias and Fairness:** If an AI system for allocating public housing or predicting building code violations is trained on biased historical data, it will perpetuate and potentially amplify existing inequalities. This erodes public trust and exposes the city to significant legal and reputational risk.
- **The 'Black Box' Problem:** The inability to explain how an AI model arrived at a decision (particularly in complex deep learning systems) creates accountability issues. A citizen has a right to know why their application for a service was denied by an algorithm.
- **Adoption Resistance:** As noted in the healthcare context, users will not adopt technology they do not trust (WEF, 2023). In a city, this includes both the civil servants expected to act on AI-driven insights and the citizens who are subject to its decisions.

Strategic Framework for Overcoming Scaling Challenges

To navigate these challenges, city leaders should adopt a structured, phased approach that balances technical and social investments.

1. Adopt a Socio-Technical, Phased Implementation Model

Cities should view AI scaling as a maturity journey with distinct phases, addressing both technical and social dimensions at each step (Haefner et al., 2023):

- Phase 1: Prove the Concept: Start with a high-value, feasible use case (e.g., AI-powered scheduling for public maintenance crews). Technically, focus on assembling the necessary data and using simple, interpretable models. Socially, build a central, cross-departmental AI team and ensure close collaboration between technical staff and domain experts (e.g., public works managers).
- Phase 2: Productise: Transition the successful pilot into a reliable, operational system. Technically, invest in a centralized data platform and standardized infrastructure to avoid redundant efforts. Socially, shift the AI growth vision to creating customer (citizen) value across services, and organize teams to support specific business lines (e.g., a dedicated AI team for the transportation department).
- Phase 3: Platformise: Achieve city-wide scale by creating reusable AI capabilities. Technically, democratize access to data and AI tools so non-experts in various departments can build solutions. Socially, establish horizontal teams focused on overarching capabilities (e.g., a 'City Forecasting Team') that serve all departments, and foster a culture of data-driven innovation.

2. Build a Robust and Ethical Data Foundation

A scalable AI strategy is built on a scalable data strategy.

- Break Down Data Silos: Establish a city-wide data governance council to mandate and manage data sharing between departments, with privacy and security as a core principle.
- Invest in Data Pipelines and Platforms: Prioritise investments in cloud infrastructure and data integration tools that clean, standardise, and make city data accessible for AI development.
- Ensure Representative and Unbiased Data: Proactively audit datasets for historical bias. Use techniques like synthetic data generation (WEF, 2023) to fill gaps and create more representative training data, ensuring AI services are fair for all demographic groups.

3. Design for Adoption and Trust from the Start

Technology must serve the people, not the other way around.

- Prioritise Human-Centric Design: AI tools should integrate seamlessly into existing employee workflows, augmenting their capabilities rather than complicating them. For citizen-facing AI, ensure interfaces are intuitive and accessible.
- Implement 'AI Model Cards': Inspired by healthcare proposals (WEF, 2023), cities should develop simple, public-facing documentation for their AI systems, explaining their purpose, data sources, performance, and how to appeal a decision.
- Foster AI Literacy and Co-Creation: Provide continuous training for city employees. Involve frontline staff and community representatives in the design and testing of AI systems to build trust and ensure the technology addresses real-world needs (Sjödín et al., 2021).

Challenges		Solutions
Socio-Technical Misalignment · Technical Focus Only · Resistance to Change		Adopt a Socio-Technical, Phased Implementation Model · Phase 1: Prove the Concept · Phase 2: Productionize · Phase 3: Platformize
Data Foundation Deficiencies · Data Silos · Poor Data Quality and Interoperability · Privacy and Security Concerns		Build a Robust and Ethical Data Foundation · Break Down Data Silos · Invest in Data Pipelines and Platforms · Ensure Representative and Unbiased Data
Lack of Trust and Ethical Governance · Algorithmic Bias and Fairness · The 'Black Box' Problem · Adoption Resistance		Design for Adoption and Trust from the Start · Prioritize Human-Centric Design · Implement 'AI Transparency Labels' · Foster AI Literacy and Co-Creation

Figure 5: Framework for overcoming scaling challenges

Working with the Ecosystem – Procurement & Partnerships

Enabling cities to work effectively with external AI providers and partners

As cities seek to responsibly adopt AI for urban transformation, the strength and structure of their external partnerships become pivotal—not only for technical success, but also for long-term sustainability and ethical integrity. This section highlights three key dynamics shaping ecosystem engagement in Indian and Asian contexts: institutional models exemplified by AIRAWAT, strategic guidance by multilaterals such as the ADB and World Bank, and lessons for city procurement and partnership strategies.

AIRAWAT: Catalysing AI Transformation at Scale

AIRAWAT, India's National Centre of Excellence for Artificial Intelligence in Sustainable Cities, led by IIT Kanpur, demonstrates how thoughtful partnership and operational design can take AI from pilot to scale. Unlike many innovation hubs, AIRAWAT works at the state level to provide a unified value proposition, collaborating with government agencies and data portals to securely steward urban datasets. Its approach is grounded in building trust, both as a technical centre and as a data custodian—an essential role in an era of growing data privacy and integrity concerns.¹³

Projects currently underway include collaborations with Surat Municipal Corporation (optimising city traffic, with scalable AI models) and with TCS and IIT Kanpur to Pioneer AI-Powered Urban Planning for Sustainable Cities (TCS, 2025). Plans also extend to regional analytics with Uttar Pradesh and property tax data triangulation with Punjab, where AIRAWAT's AI models help address leakage and improve state revenue efficiency. These use cases are prioritised based on potential for value creation and scalability, demonstrating that sustained partnerships—rather than one-off pilots—are the backbone of successful AI deployment in cities.

Recognising common constraints in city agencies—ranging from a lack of skilled human resources for GIS and data management to under-utilised smart city data platforms—AIRAWAT uses a build-operate-transfer (BOT) model. Initial implementation and demonstration occur under AIRAWAT's guidance, typically over 12–18 months, after which capacity is transferred to local government or partners. The model includes co-development with academic institutions to ensure the growth of local talent. Importantly, AIRAWAT strategically positions software-as-a-service (SaaS) solutions that lower the maintenance burden for states, further supporting sustainability.

The organisation also emphasises that AI readiness goes beyond algorithms—legacy data must be digitised and standardised, a process often best led by government, not external partners. AIRAWAT aligns new AI solutions with existing digital infrastructure, such as eGov's Digit and Upyog, supporting seamless integration with established platforms. By pursuing holistic system reform, AIRAWAT is moving from discrete pilots to helping states institutionalise the use of AI for routine decision-making, revenue management, and citizen services.

¹³ <https://airawat.org/>

Global Frameworks: ADB and World Bank

Comparatively, global development partners such as ADB and the World Bank bring both capital and policy guidance to the AI ecosystem in Asian cities—each with distinct approaches.

Asian Development Bank (ADB)¹⁴: ADB responds to government-driven demand by providing loan financing for AI-enabled service delivery, such as smart meters in Uzbekistan or pilot initiatives in South Asia. It operates within client countries' procurement frameworks, working with startups and civic tech as needed, but often relies on government-established standards. ADB is working to develop its own digital and AI frameworks, building on global experience and focusing on digital public infrastructure as a foundation for inclusive AI. Its approach is pragmatic: AI is seen as a system efficiency, not as a standalone policy or product.

World Bank¹⁵: The World Bank takes a hands-on and sustained approach to city AI initiatives—from project design through to post-loan monitoring and evaluation, often spanning 3–5 years. The Bank requires clear legal, social inclusion, and environmental objectives for all AI projects it supports, with compliance monitored by dedicated steering committees. Notably, the World Bank foregrounds four pillars for assessing responsible AI: trustworthiness, bias mitigation, environmental sustainability, and context-specific regulation. Sandboxing and auditability—where algorithms are scrutinised in controlled testbeds, and outcomes can be explained—are required to validate system integrity before scaling. The Bank also recognises that administrative and technical capacity, as much as data readiness, are essential for cities to effectively procure and manage AI projects.

Both ADB and the World Bank emphasise that effective procurement must look beyond cost or technology. They recommend that governments articulate AI-ready procurement specifications, demand interoperability, and build processes for co-design with vendors and civic stakeholders. Governments are encouraged to adopt vendor assessment checklists and include red flags related to opaque algorithms, insufficient data privacy guarantees, or a failure to commit to maintenance and skills transfer.

¹⁴ Asian Development Bank, interviewed by the author, August 2025

¹⁵ World Bank, interviewed by author, August 2025

Redefining Effective Urban AI Partnerships

Several shared challenges emerge from this ecosystem:

- **Data Fragmentation:** Both AIRAWAT and global agencies highlight that poor data quality, the absence of unified identifiers, and non-standardised legacy records are persistent barriers to city AI rollouts. Addressing these requires both top-down policy and bottom-up technical support—often offered by ecosystem partners.
- **Capacity Gaps:** Many cities, especially smaller or recently digitised ones, lack the administrative and technical resources to procure, maintain, and monitor complex AI systems. BOT and SaaS models, as deployed by AIRAWAT, reduce these burdens and enable gradual capacity-building.
- **Sustainability and Political Will:** AI pilots not aligned to the city's strategic goals, or lacking government buy-in, tend to stall at the prototype stage. Long-term relationships and government-led digitisation are key to breaking this cycle.

Takeaways for Urban Stakeholders

- Work with strong ecosystem partners who offer both technology and operational know-how. Leverage national Centres of Excellence like AIRAWAT, academic networks, and multilateral financing agencies to align technology with policy and on-the-ground needs.
- Write clear, AI-ready procurement specifications that set out ethical, privacy, governance, and interoperability requirements at the outset.
- Insist on transparency and co-creation in vendor selection, including use of sandboxes, due diligence checklists, and mock procurement exercises.
- Prioritise scalable solutions—pilot with long-term operational, financial, and capacity sustainability in mind.

In sum, ecosystem partnerships are the enablers of responsible, scalable AI in urban governance. By learning from AIRAWAT's integrative model and from multilateral partners' attention to policy, ethics, and capacity, Indian and Asian cities can chart a more confident path from pilot to impact—turning AI's promise into tangible value for citizens.

Below are two tools that the cities can use while getting on the AI journey, both while planning and procuring.

AI Adoption Checklist for Cities

Strategic Fit

- **Is the AI project aligned with your city goals and have leadership support?**
 - 1 = Low: No clear linkage to city strategy; leadership not engaged or uncertain.
 - 2 = Medium: Some alignment but not explicitly documented; leadership moderately supportive.
 - 3 = High: Strong fit with formal plans (e.g., SDGs, Smart City Mission); explicit senior leadership endorsement.
- **Are different departments and teams providing data aligned on the goals?**
 - 1 = Low: Departments work in silos; no shared understanding of objectives.
 - 2 = Medium: Informal alignment exists but roles/expectations unclear.
 - 3 = High: Cross-department coordination formally agreed; shared objectives documented.

Data Readiness

- **Is your data digital, reliable, and easy to use for AI?**
 - 1 = Low: Data mostly non-digital, outdated, inconsistent, or unavailable.
 - 2 = Medium: Some digital datasets exist but quality/format issues need cleaning.
 - 3 = High: Clean, digital, regularly updated datasets with clear metadata.
- **Is your IT setup (servers, internet) ready to support new tools?**
 - 1 = Low: Poor connectivity; outdated hardware; weak cybersecurity.
 - 2 = Medium: Basic infrastructure exists but needs upgrades or cloud support.
 - 3 = High: Reliable servers/cloud environment, stable internet, adequate cybersecurity.

Skills & Partners

- **Do your teams know how to use and manage the solution, or is training required and arranged?**
 - 1 = Low: No in-house skills; no training planned.
 - 2 = Medium: Some skills exist or training is tentatively planned.
 - 3 = High: Adequate skills already in place or formal training arrangements secured.
- **Are local partners/experts identified for support (academic institutions, civic tech)?**
 - 1 = Low: No partners identified.
 - 2 = Medium: Potential partners exist but engagement is informal.
 - 3 = High: Partners formally committed with defined roles.

Governance & Oversight

- **Is there a clear team/lead to oversee the project?**

- 1 = Low: No dedicated owner; responsibilities unclear.
- 2 = Medium: A responsible team is identified but not fully empowered.
- 3 = High: A fully empowered lead office/team with authority and accountability.

- **Is there a way for citizens and staff to give feedback and report issues?**

- 1 = Low: No feedback mechanism; no grievance redressal pathway.
- 2 = Medium: Basic channels exist (helpdesk, email) but not tailored to AI tools.
- 3 = High: Clear, accessible feedback loops; escalation process defined.

Ethics & Safety

- **How will you guarantee privacy, transparency, and fairness in the use of this tool?**

- 1 = Low: No data protection plan; no documentation of model decisions.
- 2 = Medium: Basic privacy practices; fairness considerations informal.
- 3 = High: Clear privacy protocols, risk assessments, auditability, public transparency (all mandated by legislation)

- **Are the intended uses and any limits of the system clear from the start?**

- 1 = Low: AI is proposed without a clear scope or boundaries.
- 2 = Medium: General use-case defined but limits not fully documented.
- 3 = High: Specific scope, boundaries, exclusions, escalation rules clearly defined.

Sustainability

- **Is there a plan post-pilot for scaling, funding, and supporting the system long-term?**

- 1 = Low: Short-term pilot only; no sustainability plan.
- 2 = Medium: Some budgeting or scale-up ideas exist.
- 3 = High: Defined multi-year plan with funding and responsibilities.

- **Can the system update as data needs change/updating?**

- 1 = Low: System is rigid; updates would be costly or technically difficult.
- 2 = Medium: Some ability to update but dependent on external contractors.
- 3 = High: Clearly planned update cycles; adaptable architecture; versioning in place.

Scoring and Interpretation

Total Score	Readiness Level	Recommendation

12–19	Low	Not ready. Strengthen data, governance, and skills first.
20–28	Moderate	Proceed only with a small pilot and capacity building.
29–36	High	Strong readiness. Proceed with adoption and scaling plan.

AI Vendor Due Diligence Checklist

Company Basics

- **Is the vendor reputable and experienced with city/government clients?**
 - 1 = Low: No prior government projects; limited track record.
 - 2 = Medium: Some relevant experience but limited to small pilots or private sector.
 - 3 = High: Strong record with cities/government; proven ability to work in public-sector contexts.
- **Does the vendor have credible references or use cases?**
 - 1 = Low: No verifiable references or case studies.
 - 2 = Medium: Some references exist but lack detail or measurable impact.
 - 3 = High: Strong, verifiable case studies with measurable outcomes and contacts available.

Transparency & Data

- **Can the vendor clearly explain how their AI works? (e.g. functions, outputs, etc.)**
 - 1 = Low: Explanations unclear or proprietary; vendor refuses to disclose model functioning.
 - 2 = Medium: General understanding provided but lacks detail on limitations.
 - 3 = High: Clear, accessible explanation of model logic, inputs, outputs, and limitations.
- **Will you have access to review how decisions are made (audit)?**
 - 1 = Low: No audit access; black-box system.
 - 2 = Medium: Partial auditability or conditional access.
 - 3 = High: Full audit logs, documentation, dashboards, and model interpretability tools.
- **Are data sources for training the system reliable and legally compliant?**
 - 1 = Low: Data provenance unclear; legal compliance not assured.
 - 2 = Medium: Data partially validated but documentation incomplete.
 - 3 = High: Fully documented, legally compliant, high-quality training datasets.

Privacy & Security

- **Are personal and city data protected, with clear rules for data use and storage?**
 - 1 = Low: No clear data governance; potential misuse of data.
 - 2 = Medium: Basic protections exist but retention/sharing rules are unclear.
 - 3 = High: Comprehensive data governance; strict limits on use/retention; clear contractual terms.
- **Is data encrypted and stored securely (ideally locally or as per regulations)?**

- 1 = Low: Weak security; unclear storage location; no encryption guarantees.
- 2 = Medium: Some encryption and access controls but storage does not fully meet local rules.
- 3 = High: Strong encryption, strict access controls, compliant storage (local or regulated cloud).

- **When and how does the vendor get rid of the collected data?**

- 1 = Low: No deletion policy or only deletes on request.
- 2 = Medium: General deletion policy but unclear timelines or verification.
- 3 = High: Clear, time-bound deletion processes with verifiable proof (logs/certificates).

Compliance

- **Does the solution comply with local and national rules (e.g., data privacy, digital governance)?**

- 1 = Low: No evidence of compliance; ambiguous legal footing.
- 2 = Medium: Some compliance commitment but gaps remain.
- 3 = High: Fully compliant; documentation and legal validation provided.

- **Are there terms for adapting to future policy changes?**

- 1 = Low: No flexibility; contract locks the city into current model.
- 2 = Medium: Some adaptability but may incur additional cost or delays.
- 3 = High: Clear contractual clauses for updates, compliance changes, and policy alignment.

Support & Training

- **Will the vendor provide training for your staff?**

- 1 = Low: No training or only minimal onboarding.
- 2 = Medium: Basic training offered but not continuous.
- 3 = High: Structured training program; documentation, materials, and recurring sessions included.

- **Is ongoing support included in the contract?**

- 1 = Low: Support only during deployment; no long-term commitment.
- 2 = Medium: Limited support (email/helpdesk) with unclear service legal agreements SLAs.
- 3 = High: Full support with SLAs, response times, dedicated staff, and escalation processes.

Ethics

- **Does the vendor check for bias and avoid unfair or unsafe uses?**

- 1 = Low: No bias testing; no documentation of mitigations.

- 2 = Medium: Some testing performed but not systematic or transparent.
- 3 = High: Regular bias audits, documented risk mitigation, published responsible-AI practices.

- **Can city staff review or override AI decisions if necessary?**

- 1 = Low: No override option; AI acts autonomously.
- 2 = Medium: Some override capability but not integrated into workflows.
- 3 = High: Clear oversight mechanisms; manual review and override fully supported.

Long-Term Fit

- **Will the system be able to grow/adapt with your city's needs?**

- 1 = Low: Rigid system; scaling requires major redevelopment.
- 2 = Medium: Some modularity or upgrade options but limited.
- 3 = High: Scalable architecture; regular updates; extensible to new datasets/modules.

- **Are contract terms flexible and clear about responsibilities?**

- 1 = Low: Ambiguous terms; unclear ownership of data/models; lock-in risk.
- 2 = Medium: Terms are partially clear but negotiation required.
- 3 = High: Transparent, flexible contract with clear roles, IP clauses, data ownership, exit plans.

Scoring and Interpretation

Total Score	Vendor Suitability	Recommendation
16–26	Low	High risk. Do not contract; seek alternatives.
27–38	Medium	Proceed only with strict controls and pilot testing.
39–48	High	Vendor is reliable and suitable for long-term partnership.

Designing Trust – Participation, Transparency, and Inclusion

Designing Inclusive AI

If AI is not intentionally designed for inclusivity, it has the potential to reinforce existing inequalities and injustices. Often, it is positioned as a tool to remove humans' individual biases from decision-making processes; however, algorithms can reinforce systemic biases and exacerbate existing inequalities. Algorithms are trained on historical data, and while data is often thought to be unbiased, resulting in unbiased decision-making by algorithms, this is not necessarily true.

To begin designing inclusive AI solutions, the biases and historical and cultural context of the training data must be acknowledged and considered in the algorithm design. To mitigate biased outputs, developers should compile counter data in the face of missing data or institutional neglect of certain groups within a database and should analyse data for possible inequitable outcomes that could stem from it before using it as training data. Furthermore, algorithms should be audited by journalists and academics to show how the benefits they produce are differentially distributed (D'Ignazio & Klein, 2020).

Beyond examining the data that AI algorithms are based on, inclusive design must also involve participatory methods of design.

Participatory AI Design

Participatory AI design means involving a range of stakeholders in the phases of planning and designing an AI solution and provides an opportunity to involve the wider public in the development and deployment of AI systems. Participation can empower communities that have been historically marginalised and tend to face disproportionately negative impacts from AI systems. Engaging these communities to participate in design can allow for their representation, which might be missing from the datasets being used (Birhane et al., 2022). It is therefore a crucial aspect of designing inclusive AI systems.

Bratteteig & Verne (2018) identify three key challenges in participatory AI design:

1. Designers often struggle to fully understand or explain how machine learning systems make decisions due to their technical complexity.
2. Evaluating AI designs meaningfully requires long-term engagement and discussion of possible futures rather than quick assessments.
3. Since AI systems learn continuously from user interaction, biases in training data or limited participation (especially from marginalised groups) can reinforce inequality unless inclusive, participatory approaches are adopted.

This means that stakeholders must be involved in the planning and design process early enough that the use of AI has not yet been decided upon. (Birhane et al., 2022). This ensures that the solution is being developed directly out of the need expressed by participants rather than out of any preconceived notions of what they need.

Citizen Science is an aspect of participatory AI design that provides people with an opportunity to partner in the development of Urban AI. It incorporates crowdsourced data produced by ordinary citizens to further fuel the training processes of ML systems (Abdul-Rahman et al., 2025; Edinburg ReproducibiliTea, 2023). Citizens should be given the opportunity to annotate, evaluate, and interpret the data, helping to identify gaps or possible misrepresentations (Birhane et al., 2022). This empowers them to contribute more often to effective decision-making, creating more transparent AI systems. It also allows for the framework of an inclusive data system to be produced, building renewed trust towards the use of AI systems in development.

Participatory design aims to develop solutions based on community values and interests, and to minimise instances of commercial actors co-opting research and datasets to produce solutions that do not benefit historically marginalised or underrepresented groups (Birhane et al., 2022). An example of Participatory AI incorporating Citizen Science into its framework is the GeoAI Digital Platform, monitored by the UNDP Accelerator Lab in Bihar, India. The platform helps detect hotspots of air pollution using satellite imagery and object-detection algorithms (BPSCB et al., 2023). By incorporating both human and artificial intelligence, citizen scientists across the globe volunteered to help detect brick kilns responsible for air pollution. This data was incorporated into the GeoAI system, helping detect 47,000 kilns and combat air pollution hotspots in Bihar, Haryana, Punjab, and Rajasthan. (BPSCB et al., 2023).

Recommendations

AI is no longer a futuristic ambition for cities—it is a practical, transformative tool for urban governance, infrastructure, social services, and environmental sustainability. Yet, as demonstrated by evidence from Indian and Asian cities, AI's promise will only be fully realised when accompanied by strong policy, ethical safeguards, institutional readiness, and robust ecosystem partnerships. As cities across Asia begin adopting Artificial Intelligence (AI) to improve service delivery, policymakers must recognise that success depends less on technological sophistication and more on institutional readiness, ethical governance, and ecosystem collaboration.

1. Anchor AI Adoption in Public Value, Human Rights, and Ethical Governance

- Define AI initiatives with clear public-interest goals centred on equity, inclusivity, and sustainability—not only for the sake of adopting a technology.
- Uphold core human rights in AI systems, embedding transparency, fairness, privacy protections, and human oversight throughout design, deployment, and operations.
- Use internationally recognised ethical frameworks (UNESCO, OECD, Singapore, EU) as standards, adapting to local law and societal context.
- Mandate ethical impact assessments for high-risk AI deployments and promote transparent communication to build public trust.

2. Invest Strategically in Data Foundations and Digital Infrastructure

- Prioritise the creation and maintenance of robust, secure digital infrastructure—cloud, IoT, data platforms—to support AI scalability.
- Implement city-wide standards for data quality, accuracy, interoperability, and stewardship, utilising frameworks like India's NITI Aayog Data Quality Scorecard.
- Treat urban data as a public asset, appointing data stewards, mandating regular data audits, and ensuring open, machine-readable formats with mechanisms for error correction.
- Integrate AI initiatives with existing digital public goods platforms (e.g. DIGIT, Upyog and the like) to avoid fragmented pilots and encourage synergy across departments.

3. Build Institutional and Ecosystem Readiness

- Adopt phased AI adoption strategies that begin with feasible, high-impact pilots and scale through organisational and technical maturity.
- Establish AI governance structures such as task forces within municipalities for strategy, oversight, and cross-departmental coordination.

- Foster long-term partnerships with academia, research centres (e.g. AIRAWAT), civic technology actors, and multilateral organisations (ADB, World Bank) to leverage expertise, financing, and regulatory guidance.
- Prioritise capacity-building programs to equip government staff, especially frontline workers, with AI literacy and data management skills to complement technological tools.

4. Integrate Responsible Procurement and Vendor Partnerships

- Develop AI-ready procurement specifications covering model transparency, expected benefits, ethical considerations, bias mitigation, compliance requirements, and contractual flexibility for updates.
- Implement vendor due diligence processes that assess company's track record, data security protocols, audits, human oversight capabilities, and support services.
- Require vendor co-creation and iterative development models, including sandbox environments to test AI efficacy and fairness before full deployment.
- Avoid procurement solely based on cost; emphasise innovation with accountability and sustainability in long-term engagement models.

5. Centre Co-Creation and Public Participation

- Involve citizens, frontline staff, and vulnerable populations throughout the AI lifecycle—from problem definition to design, pilot, deployment, and evaluation.
- Institutionalise participatory auditing to enable community oversight and grievance redress mechanisms.

6. Design for Inclusion, Safety, and Sustainability

- Proactively identify and mitigate bias through audits, the use of representative or synthetic data, and a human review in high-impact systems.
- Embed sustainability by ensuring AI systems are maintainable, upgradable, and financially feasible beyond initial pilots, with sensitivity to local resource constraints.
- Balance technological advances with ethical vigilance to prevent discrimination, privacy violations, or unintended harms, particularly for marginalised groups.

7. Monitor, Evaluate, and Adapt Continuously

- Define clear, measurable success indicators for AI projects, such as service efficiency, user satisfaction, inclusion, error rates, and environmental impact.

- Establish adaptive governance processes for periodic review, public reporting, and iterative improvements to AI models and deployment practices.
- Learn actively from both successes and failures—scale projects only when evidence supports impact and readiness.

8. Embed AI Into Broader Urban Transformation and Regional Collaboration

- Align AI adoption with comprehensive city development goals—including resilience, sustainability, and digital transformation agendas.
- Incorporate AI governance into city budgeting, planning, and sectoral policies to ensure coordination and resource flow.
- Encourage cross-city and regional collaboration for shared learning, standard-setting, joint capacity-building, and innovation diffusion.

AI offers cities a unique opportunity to leapfrog towards more inclusive, efficient, and resilient governance—but only if deployed with intentionality, accountability, and empathy. The path forward lies in strengthening institutions, embedding ethics, and centring people in every stage of the AI journey. Cities that invest in trust, transparency, and partnerships today will set the benchmark for responsible urban innovation tomorrow.

References

- Abdul-Rahman, G., Zwitter, A. & Haleem, N. A systematic literature review on the role of artificial intelligence in citizen science. *Discov Artif Intell* 5, 165 (2025). <https://doi.org/10.1007/s44163-025-00437-z>
- Aclima. (2017, October 11). How Seoul used the IoT to improve waste management and collection. aclima.eus. <https://www.google.com/url?q=https://aclima.eus/en/noticia/how-seoul-used-the-iot-to-improve-waste-management-and-collection/&sa=D&source=docs&ust=1762347734266325&usg=AOvVaw0aogiuiVpU3k140DGhBEsi>
- ARMMAN. (n.d.). Kilkari. <https://armman.org/kilkari/>
- Around Town Realty. (2025, February 17). A step-by-step guide to property tax in Ahmedabad: Everything you need to know. AroundTown Realty Pvt Ltd. <https://aroundtownrealty.in/a-step-by-step-guide-to-property-tax-in-ahmedabad-everything-you-need-to-know/>
- Barocas, S., & Selbst, A. D. (2016). Big Data's Disparate Impact. *California Law Review*, 104(3), 671. <https://doi.org/10.15779/z38bg31>
- Bass, D., & Huet, E. (2017). Researchers combat gender and racial bias in artificial intelligence. Bloomberg. com, Dec, 4.
- Bengaluru Traffic Police. (n.d.). Adaptive Traffic Control Systems (ATCS). Retrieved November 5, 2025, from <https://btp.karnataka.gov.in/214/adaptive-traffic-control-system-%28atcs%29/en>
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial Intelligence. *MIS Quarterly*, 45(3), 1433–1450. <https://doi.org/10.25300/misq/2021/16274>
- Bianchi, S., Ciurlanti, J., Overend, M., Pampanin, S. (2022). A probabilistic-based framework for the integrated assessment of seismic and energy economic losses of buildings, *Eng. Struct.* 269, 114852, <https://doi.org/10.1016/j.engstruct.2022.114852>.
- Birhane, A., Isaac, W., Prabhakaran, V., Diaz, M., Elish, M. C., Gabriel, I., & Mohamed, S. (2022, October). Power to the people? Opportunities and challenges for participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization* (pp. 1-8).
- Bolukbasi, T., Chang, K. W., Zou, J. Y., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, 29.
- Bratteteig, T. & Verne, G. (2018). Does AI Make PD Obsolete? Exploring Challenges from Artificial Intelligence to Participatory Design. In *Proceedings of the 15th Participatory Design Conference: Short Papers, Situated Actions, Workshops and Tutorial 2*, 8 (2018), 1–5. <https://doi.org/10.1145/3210604.3210646>

- Brears, R. C. (2024). Smart water management in Singapore. Medium.
<https://medium.com/mark-and-focus/smart-water-management-in-singapore-38e51bf1f14>
- BSPCB, UNDP, University of Nottingham, & Development Alternatives. (2023). GeoAI for Brick Kilns in Bihar: Learnings & Recommendations. Undp.org.
https://www.undp.org/sites/g/files/zskgke326/files/2023-12/geoai_for_brick_kilns_v8_web_pages_002_0.pdf
- Cai, M. (2021). Natural language processing for urban research: A systematic review. *Heliyon*, 7(3), e06322. <https://doi.org/10.1016/j.heliyon.2021.e06322>
- Calo, M. R. (2012) "Robots and Privacy," in *Robot Ethics: The Ethical and Social Implications of Robotics* (Patrick Lin, George Bekey, and Keith Abney, eds.). Cambridge: MIT Press.
- Calo, R. (2010). Peeping HALs: Making Sense of Artificial Intelligence and Privacy, 2 *European J. Leg. Studies* 168, <https://digitalcommons.law.uw.edu/faculty-articles/1148>
- C-DAC. (n.d.). Intelligent transportation systems.
https://www.cdac.in/index.aspx?id=pe_its_Intelligent
- Capolongo, S., Gola, M., Brambilla, A., Morganti, A., Mosca, E.I., & Barach, P. (2020). COVID-19 and healthcare facilities: a decalogue of design strategies for resilient hospitals, *Acta. Bio. Medica.: Atenei. Parmensis.* 91 (9–S) (2020) 50.
<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8023092/>
- Chen, Y., You, W., Ou, L., & Tang, H. (2025). A Review of Machine Learning Techniques for Urban Resilience Research: The Application and Progress of Different Machine Learning Techniques in Assessing and Enhancing Urban Resilience. *Systems and Soft Computing*, 7, 200269. <https://doi.org/10.1016/j.sasc.2025.200269>
- CSCTRUCK. (2024, December 8). Industrial cleaning revolution: Industrial sweeper trucks deployed in Tokyo! Brand Tank Truck, Sewer Truck, Tow Truck Manufacturer, Supplier, Factory - CSCTRUCK. <https://www.csctrucks.com/news/industrial-cleaning-revolution-industrial-swe-82785327.html>
- Citron, D. & Pasquale, F. (2014). The Scored Society: Due Process for Automated Predictions, 89 *Washington Law Review*, 1–34.
- D'Ignazio, C., & Klein, L. F. (2020). Data feminism. In *The MIT Press eBooks*.
<https://doi.org/10.7551/mitpress/11805.001.0001>
- Devlin, H. (2018, February 14). AI programs exhibit racial and gender biases, research reveals. *The Guardian*. <https://www.theguardian.com/technology/2017/apr/13/ai-programs-exhibit-racist-and-sexist-biases-research-reveals>
- Fu, C., McKenzie, G., Frias-Martinez, V., & Stewart, K. (2018). Identifying spatiotemporal urban activities through linguistic signatures. *Computers, Environment and Urban Systems*, 72, 25-37.

flow DB. (2024, April 29). Revitalising Tashkent's water sector – Deutsche Bank. flow – Deutsche Bank. <https://flow.db.com/trade-finance/revitalising-tashkent-s-water-sector#!>

Gasser, U. and Almeida, V. A. F. (2017, December) "A Layered Model for AI Governance," in IEEE Internet Computing, vol. 21, no. 6, pp. 58-62, doi:10.1109/MIC.2017.4180835.

Ghosh, S., & Gunning, D. (2019). Natural Language Processing Fundamentals: Build intelligent applications that can interpret the human language to deliver impactful results. In Google Books. Packt Publishing.
https://books.google.co.in/books/about/Natural_Language_Processing_Fundamentals.html?id=i8-PDwAAQBAJ&redir_esc=y

Goldblatt, R., Kaiser, K., Tran, H. T., & Vu, K. (2018, March 13). Artificial intelligence for smart cities: Insights from Ho Chi Minh city's spatial development. Medium. <https://medium.com/world-of-opportunity/artificial-intelligence-for-smart-cities-insights-from-ho-chi-minh-citys-spatial-development-e95f3064b617>

GovInsider. (2017, March 8). Exclusive: AI takes over Singapore's property estate. <https://govinsider.asia/intl-en/article/ai-takes-over-singapores-property-estate>

GOVTECH Singapore. (2024, December 30). Smart nation 2.0: Initiatives in Singapore. Government Technology Agency of Singapore (GovTech Singapore).
<https://www.tech.gov.sg/technews/our-enhanced-smart-nation-vision-paving-the-way-for-a-new-digital-era>

Guetterman, T. C., Chang, T., DeJonckheere, M., Basu, T., Scruggs, E., & Vydiswaran, V. V. (2018). Augmenting Qualitative Text Analysis with Natural Language Processing: Methodological Study. Journal of Medical Internet Research, 20(6), e231.
<https://doi.org/10.2196/jmir.9702>

Haefner, N., Parida, V., Gassmann, O., & Wincent, J. (2023). Implementing and scaling artificial intelligence: A review, framework, and research agenda. Technological Forecasting and Social Change, 197, 122878. <https://doi.org/10.1016/j.techfore.2023.122878>

Hajian, S., & Domingo-Ferrer, J. (2013). Direct and indirect discrimination prevention methods. In Discrimination and privacy in the information society: Data mining and profiling in large databases (pp. 241-254). Berlin, Heidelberg: Springer Berlin Heidelberg.

Heiling, J., Kempt, H., & Nagel, S. (2023). Beware of sustainable AI! Uses and abuses of a worthy goal. AI And Ethics, 4(2), 201–212. <https://doi.org/10.1007/s43681-023-00259-8>

Huang, L., Wu, Y., Zheng, Q., Zheng, Q., Zheng, X., Gan, M., ... & Zhang, J. (2018). Quantifying the spatiotemporal dynamics of industrial land uses through mining free access social datasets in the mega Hangzhou Bay region, China. Sustainability, 10(10), 3463.

ICTpost. (2024, November 1). Cities across India are harnessing AI to revolutionize waste management! ICTpost. <https://ictpost.com/ai-the-game-changer-for-urban-waste-management-in-india/>

IEEE Smart Cities. (2024, July). Smart cities with AI for flooding management a case study of Jakarta. IEEE Smart Cities - Announcements, Events, Webinars & More. https://smartcities.ieee.org/newsletter/june-july-2024/smart-cities-with-ai-for-flooding-management-a-case-study-of-jakarta?utm_source=chatgpt.com

IndiaAI. (2020, December 28). Andhra Pradesh forecasts energy needs using AI. indiaai.gov.in. <https://indiaai.gov.in/news/andhra-pradesh-forecasts-energy-needs-using-ai>

Independent Observer. (2025, April 27). Indonesia Launches Novel AI-Powered Public Health Platform to Tackle Climate-Driven Diseases. observerid.com. <https://observerid.com/indonesia-launches-novel-ai-powered-public-health-platform-to-tackle-climate-driven-diseases/>

International Energy Agency (IEA). (2025). Energy and AI: How much energy will AI use? IEA. Retrieved from <https://www.iea.org/topics/artificial-intelligence>

iSPIRT Foundation. (2025, November 6). *Privacy in the Age of AI: New Frameworks for Data Collaboration-Part-1*. ProductNation. <https://pn.ispirt.in/privacy-in-the-age-of-ai-new-frameworks-for-data-collaboration-part-1/>

iSPIRT Foundation. (2025b, November 6). *Privacy in the Age of AI: New Frameworks for Data Collaboration-Part-2*. ProductNation. <https://pn.ispirt.in/privacy-in-the-age-of-ai-new-frameworks-for-data-collaboration-part-2/>

Kumar, S., Srivastava, A., & Maity, R. (2024). Modeling climate change impacts on vector-borne disease using machine learning models: Case study of Visceral leishmaniasis (Kala-azar) from Indian state of Bihar. *Expert Systems with Applications*, 237, 121490.

Marasinghe, R., Yigitcanlar, T., Mayere, S., Washington, T., & Limb, M. (2023). Computer vision applications for urban planning: A systematic review of opportunities and constraints. *Sustainable Cities and Society*, 100, 105047. <https://doi.org/10.1016/j.scs.2023.105047>

Maynilad Water Services Inc. (2024, July 3). Maynilad pilots new AI technology for pipe condition monitoring. Maynilad Water Services, Inc. <https://www.mayniladwater.com.ph/maynilad-pilots-new-ai-technology-for-pipe-condition-monitoring/>

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2). <https://doi.org/10.1177/2053951716679679>

Municipal Corporation of Greater Mumbai, Information Technology Department. (n.d.). MyBMC Assist Whatsapp Chatbot. MyBMC - Welcome to BMC's Website.

<https://www.mcgm.gov.in/irj/go/km/docs/documents/Director%20IT/Case%20Study/MyBMC%20Assist%20Case%20Study.pdf>

NITI Aayog. (2025, June). India's Data Imperative: The Pivot Towards Quality, GHOSH, D., & GARG, S. (2025). India's data imperative: the pivot towards quality. In *NITI Frontier Tech Hub's Quarterly Insights* [Report]. <https://niti.gov.in/sites/default/files/2025-06/FTH-Quarterly-Insight-june.pdf>

NITI Aayog. (n.d.). Bengaluru: Adopting AI to enhance Law enforcements services. ABP. <https://abp.championsofchange.gov.in/content/1102bengaluru-adopting-ai-to-enhance-law-enforcements-services>

NMSC. (2025, September 10). AI CCTV cameras launch with smart storage and detection. nextmsc.com. <https://www.nextmsc.com/blogs/ai-cctv-camera-expands-with-smarter-surveillance-features>

Rössler, B. (2005). *The value of privacy*. Polity, Cambridge

Sarker, I. H. (2021). Machine Learning: Algorithms, Real-World Applications and Research Directions. *SN Comput Sci*. doi: 10.1007/s42979-021-00592-x. Epub 2021 Mar 22. PMID: 33778771; PMCID: PMC7983091

Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63. <https://doi.org/10.1145/3381831>

Shah, S.A., Seker, D.Z. , Rathore, M.M., Hameed, S., Yahia, S.B., Draheim, D. (2019). Towards disaster resilient smart cities: can internet of things and big data analytics be the game changers? *IEEE Access* 7. 91885–91903. <https://ieeexplore.ieee.org/abstract/document/8759905>

Sjödin, D., Parida, V., Palmié, M., & Wincent, J. (2021). How AI capabilities enable business model innovation: Scaling AI through co-evolutionary processes and feedback loops. *Journal of Business Research*, 134, 574–587. <https://doi.org/10.1016/j.jbusres.2021.05.009>

Soo, Z. (2018, March 5). "The increasing use of artificial intelligence is stoking privacy concerns in China." *South China Morning Post*. <https://www.scmp.com/business/companies/article/2135713/increasing-use-artificial-intelligence-stoking-privacy-concerns>

Subedi, B., & Joshi, B. (2024). Classification of grievances in hello Sarkar using supervised machine learning. *Journal of Science and Technology*, 4(1), 35-40. <https://doi.org/10.3126/jost.v4i1.74564>

TCS. (2025, September 3). TCS, IIT Kanpur partner to pioneer AI-powered urban planning for sustainable cities. *Tata Consultancy Services: Building Perpetually Adaptive Enterprises*.

<https://www.tcs.com/who-we-are/newsroom/press-release/tcs-iit-kanpur-partner-ai-powered-urban-planning-sustainable-cities>

The Atlas of Urban Tech. (n.d.). Hangzhou city brain. <https://atlasofurbantech.org/cases/chn-hangzhou/>

The Hans India. (2024, November 11). Tamil Nadu to detect land encroachments using AI, pilot study in Coimbatore. Andhra Pradesh Breaking News, Telangana News, Hyderabad News Updates, National News, Breaking News. <https://www.thehansindia.com/technology/tech-news/tamil-nadu-to-detect-land-encroachments-using-ai-pilot-study-in-coimbatore-920810>

The New Indian Express. (2025, March 25). AI to deliver property documents at doorsteps in Bengaluru. <https://www.newindianexpress.com/cities/bengaluru/2025/Mar/25/ai-to-deliver-property-documents-at-doorsteps-in-bengaluru>

Times of India. (2023, December 10). AI-driven traffic control system launched at city's five intersections | Kanpur news - Times of India. The Times of India. <https://timesofindia.indiatimes.com/city/kanpur/ai-driven-traffic-control-system-launched-at-citys-five-intersections/articleshow/105872078.cms>

Times of India. (2025, October 10). Delhi plans AI governance: Soon, get birth certificate; pay property tax via WhatsApp | Delhi news. The Times of India. <https://timesofindia.indiatimes.com/city/delhi/soon-get-birth-cert-pay-property-tax-via-whatsapp/articleshow/124432936.cms>

Thierer, A., O'Sullivan, A.C., and Russell, R. (2017). "Artificial Intelligence and Public Policy." Mercatus Research, Mercatus Center at George Mason University, Arlington, VA.

World Economic Forum. (2023, June). Scaling Smart Solutions with AI in Health: Unlocking Impact on High-Potential Use Cases. <https://www.weforum.org/publications/scaling-smart-solutions-with-ai-in-health-unlocking-impact-on-high-potential-use-cases/>

Zhang, M. (2024, June 21). Data center water Usage: A Comprehensive guide. Dgtl Infra. <https://dgtlinfra.com/data-center-water-usage/>

Responsible Use of AI in Urban Service Delivery

AI Adoption Checklist

AI Vendor Due Diligence Checklist

December 2025

Executive Summary

Artificial Intelligence (AI) is rapidly reshaping how cities think, plan, and deliver services. Yet for most local governments, AI remains an unfamiliar terrain—filled with both promise and uncertainty. Recognising the unequal concentration of AI technologies in big tech companies, the white paper highlights the urgent need for Asian governments, cities, and local tech communities to take an active role in developing AI. By doing so, Asia can help create a fairer and stronger AI future—one that meets its unique urban challenges, encourages local innovation, protects control over important digital systems, and supports fair and balanced growth of cities. It aims to equip policymakers, city leaders, and their ecosystem partners with a practical roadmap for harnessing AI responsibly—to strengthen governance, improve service delivery, and make cities more liveable, inclusive, and resilient.

Why AI Matters for Cities

AI is not just about technology; it is fundamentally about decision-making. When deployed wisely, AI systems help municipalities shift from reactive, routine management to proactive, anticipatory governance. With AI, cities can see patterns previously invisible, respond before needs become crises, and shape more inclusive, equitable public services. But realising these outcomes is not automatic. Our research and interviews show that purposeful AI relies on something deeper than algorithms: physical and digital infrastructure, inclusive policies, institutional readiness, and above all, public trust. The question facing cities is no longer ‘Should we use AI?’ but ‘How can we do so in a way that genuinely serves and protects our communities?’

Where AI is Already Making a Difference

Drawing from real examples across Asian cities, this paper highlights how AI is already being used to transform urban systems in domains as varied as governance and participation, infrastructure and operations, social services, environment and climate, and safety and security across the south Asian context. AI is helping government address citizens’ grievances more effectively, deliver welfare to vulnerable citizens, and enhance disaster preparedness within ‘at-risk’ regions¹.

The Challenges Holding Cities Back

Despite progress, three core barriers keep most municipalities from scaling or benefiting fully from AI:

1. **Data Foundations:** Many cities are held back by inconsistent, poor-quality, or inaccessible data—without which AI cannot operate reliably.
2. **Institutional Capacity:** Adoption is often limited by skill gaps, unclear roles, or fragmented leadership. The technology is only as effective as the teams, workflows, and politics behind it.

¹ **Governance & Participation:** AI-powered chatbots and grievance redressal tools in Mumbai and Bengaluru, and open-source frameworks such as ParakhAI, are helping governments respond faster and engage citizens more transparently.

Infrastructure & Operations: Smart traffic systems in Jaipur and Kanpur, predictive maintenance tools for water supply in Tashkent and Manila, and AI-led waste sorting in Tokyo are improving reliability and resource efficiency.

Social Services: Initiatives like Wadhwani AI’s maternal health models or Bihar’s disease prediction tools show how AI can extend public health and social protection to the most underserved.

Environment & Climate: From Jakarta’s flood prediction systems to Bhopal’s air quality dashboards, cities are using AI to better manage risk—but these depend on strong data foundations.

Safety & Security: Cities such as Pimpri-Chinchwad and Metro Manila are experimenting with AI for public safety and citizen reporting, testing new ways to make urban spaces safer.

3. **Trust, Ethics & Accountability:** Fears of bias, discrimination, lack of transparency, and data breaches are real and can stall progress without firm safeguards. Successful use of AI depends as much on public values and buy-in as on technical solutions.

Building the Right Ecosystem

The paper underscores that AI adoption is a collective effort. It requires strong partnerships—between governments, start-ups, academia, and development partners. Models such as the AIRAWAT Centre of Excellence and the eGov Foundation show how India is beginning to build this ecosystem, while international efforts by the World Bank and ADB demonstrate the value of shared frameworks and responsible sandboxing².

Responsible AI and Public Trust

Responsible AI begins with transparency, privacy, and public trust. Cities must be clear about what data is being used, how decisions are made, and how citizens can question or correct errors. AI systems must be transparent, auditable, explainable, and inclusive reflecting the diversity of the communities they serve. Trust is not a by-product; it is a precondition for sustainable innovation.

Key Products of This Paper: Practical Tools for Leaders

Recognising the limitations of real-world municipal work, this paper delivers two signature outputs— the AI Adoption and AI vendor due diligence checklists- designed for direct use by city officials and policymakers, and a policy roadmap to scale AI pilots.

1. The AI Adoption Checklist for Cities

A hands-on tool to help cities assess readiness and guide responsible integration. The checklist asks, step by step:

- Is the proposed AI project aligned with city goals and backed by leadership?
- Is our data digital, reliable, and AI-ready?
- Are the necessary digital infrastructure and IT teams in place?
- Are clear oversight and continuous feedback mechanisms available for staff and citizens?
- Have we built ethics and safety into the design—and are risks clearly communicated?
- Is there a plan for sustainability, scaling, and ongoing updates as city needs evolve?

This checklist makes it possible to catch red flags early, build cross-department coordination, and ensure city investments translate to real value for residents.

2. The AI Vendor Due Diligence Checklist

A straightforward framework to help public agencies screen external technology partners for reliability, safety, and fit:

- Is the vendor reputable and experienced in government work?
- Can they explain their system's function, outcomes, and data sources in clear terms?
- Does their approach to privacy, bias, and compliance meet city and national standards?
- Will the vendor provide upfront training, ongoing support, and flexibility for new needs?

² Sandboxes, in the context of AI, refer to controlled environments where AI systems can be tested, validated, and monitored without affecting real-world operations. These environments are isolated from the main systems to prevent any unintended consequences or security breaches. Sandboxes allow developers and researchers to experiment with AI models, run simulations, and fine-tune algorithms in a safe and controlled setting.

- Are clear ethical guardrails in place—such as human review and the right to override—if something goes wrong?

Applying this checklist ensures transparency and minimizes risks before contracts are signed, creating a foundation for lasting, trustworthy partnerships.

3. Policy Roadmap: Moving from Pilots to Public Value

In addition to the hands-on checklists, the paper offers a concise menu of policy and operational recommendations—gathered from best practices, regional networks, and real-world casework:

- Make responsible AI adoption a top civic priority, grounded in public value, inclusion, and explicit human rights protection.
- Tear down data silos and treat data as a shared public asset, creating the infrastructure and accountability to support high-quality, trustworthy analytics.
- Build institutional and ecosystem readiness through dedicated AI task forces, partnerships with academia and civic tech, and a commitment to ongoing capacity-building.
- Centre co-creation by involving citizens, frontline staff, and especially marginalised groups early and throughout the AI lifecycle.
- Mandate responsible procurement and vendor management—requiring transparency, piloting ('safe trial environments'), and ongoing evaluation.
- Design for long-term inclusion, safety, and resilience by proactively managing bias, enabling adaptability, and embedding oversight and public engagement.
- Monitor, evaluate, and adapt continuously, allowing only solutions that prove their worth in local contexts to scale up and shape the city's digital future.

AI Adoption Checklist for Cities

Strategic Fit

- **Is the AI project aligned with your city goals and have leadership support?**
 - 1 = Low: No clear linkage to city strategy; leadership not engaged or uncertain.
 - 2 = Medium: Some alignment but not explicitly documented; leadership moderately supportive.
 - 3 = High: Strong fit with formal plans (e.g., SDGs, Smart City Mission); explicit senior leadership endorsement.
- **Are different departments and teams providing data aligned on the goals?**
 - 1 = Low: Departments work in silos; no shared understanding of objectives.
 - 2 = Medium: Informal alignment exists but roles/expectations unclear.
 - 3 = High: Cross-department coordination formally agreed; shared objectives documented.

Data Readiness

- **Is your data digital, reliable, and easy to use for AI?**
 - 1 = Low: Data mostly non-digital, outdated, inconsistent, or unavailable.
 - 2 = Medium: Some digital datasets exist but quality/format issues need cleaning.
 - 3 = High: Clean, digital, regularly updated datasets with clear metadata.
- **Is your IT setup (servers, internet) ready to support new tools?**
 - 1 = Low: Poor connectivity; outdated hardware; weak cybersecurity.
 - 2 = Medium: Basic infrastructure exists but needs upgrades or cloud support.
 - 3 = High: Reliable servers/cloud environment, stable internet, adequate cybersecurity.

Skills & Partners

- **Do your teams know how to use and manage the solution, or is training required and arranged?**
 - 1 = Low: No in-house skills; no training planned.
 - 2 = Medium: Some skills exist or training is tentatively planned.
 - 3 = High: Adequate skills already in place or formal training arrangements secured.
- **Are local partners/experts identified for support (academic institutions, civic tech)?**
 - 1 = Low: No partners identified.
 - 2 = Medium: Potential partners exist but engagement is informal.
 - 3 = High: Partners formally committed with defined roles.

Governance & Oversight

- **Is there a clear team/lead to oversee the project?**
 - 1 = Low: No dedicated owner; responsibilities unclear.
 - 2 = Medium: A responsible team is identified but not fully empowered.
 - 3 = High: A fully empowered lead office/team with authority and accountability.

- **Is there a way for citizens and staff to give feedback and report issues?**
 - 1 = Low: No feedback mechanism; no grievance redressal pathway.
 - 2 = Medium: Basic channels exist (helpdesk, email) but not tailored to AI tools.
 - 3 = High: Clear, accessible feedback loops; escalation process defined.

Ethics & Safety

- **How will you guarantee privacy, transparency, and fairness in the use of this tool?**
 - 1 = Low: No data protection plan; no documentation of model decisions.
 - 2 = Medium: Basic privacy practices; fairness considerations informal.
 - 3 = High: Clear privacy protocols, risk assessments, auditability, public transparency (all mandated by legislation)
- **Are the intended uses and any limits of the system clear from the start?**
 - 1 = Low: AI is proposed without a clear scope or boundaries.
 - 2 = Medium: General use-case defined but limits not fully documented.
 - 3 = High: Specific scope, boundaries, exclusions, escalation rules clearly defined.

Sustainability

- **Is there a plan post-pilot for scaling, funding, and supporting the system long-term?**
 - 1 = Low: Short-term pilot only; no sustainability plan.
 - 2 = Medium: Some budgeting or scale-up ideas exist.
 - 3 = High: Defined multi-year plan with funding and responsibilities.
- **Can the system update as data needs change/updating?**
 - 1 = Low: System is rigid; updates would be costly or technically difficult.
 - 2 = Medium: Some ability to update but dependent on external contractors.
 - 3 = High: Clearly planned update cycles; adaptable architecture; versioning in place.

Scoring and Interpretation

Total Score	Readiness Level	Recommendation
12–19	Low	Not ready. Strengthen data, governance, and skills first.
20–28	Moderate	Proceed only with a small pilot and capacity building.
29–36	High	Strong readiness. Proceed with adoption and scaling plan.

AI Vendor Due Diligence Checklist

Company Basics

- **Is the vendor reputable and experienced with city/government clients?**
 - 1 = Low: No prior government projects; limited track record.
 - 2 = Medium: Some relevant experience but limited to small pilots or private sector.
 - 3 = High: Strong record with cities/government; proven ability to work in public-sector contexts.
- **Does the vendor have credible references or use cases?**
 - 1 = Low: No verifiable references or case studies.
 - 2 = Medium: Some references exist but lack detail or measurable impact.
 - 3 = High: Strong, verifiable case studies with measurable outcomes and contacts available.

Transparency & Data

- **Can the vendor clearly explain how their AI works? (e.g. functions, outputs, etc.)**
 - 1 = Low: Explanations unclear or proprietary; vendor refuses to disclose model functioning.
 - 2 = Medium: General understanding provided but lacks detail on limitations.
 - 3 = High: Clear, accessible explanation of model logic, inputs, outputs, and limitations.
- **Will you have access to review how decisions are made (audit)?**
 - 1 = Low: No audit access; black-box system.
 - 2 = Medium: Partial auditability or conditional access.
 - 3 = High: Full audit logs, documentation, dashboards, and model interpretability tools.
- **Are data sources for training the system reliable and legally compliant?**
 - 1 = Low: Data provenance unclear; legal compliance not assured.
 - 2 = Medium: Data partially validated but documentation incomplete.
 - 3 = High: Fully documented, legally compliant, high-quality training datasets.

Privacy & Security

- **Are personal and city data protected, with clear rules for data use and storage?**
 - 1 = Low: No clear data governance; potential misuse of data.
 - 2 = Medium: Basic protections exist but retention/sharing rules are unclear.
 - 3 = High: Comprehensive data governance; strict limits on use/retention; clear contractual terms.
- **Is data encrypted and stored securely (ideally locally or as per regulations)?**
 - 1 = Low: Weak security; unclear storage location; no encryption guarantees.
 - 2 = Medium: Some encryption and access controls but storage does not fully meet local rules.
 - 3 = High: Strong encryption, strict access controls, compliant storage (local or regulated cloud).
- **When and how does the vendor get rid of the collected data?**

- 1 = Low: No deletion policy or only deletes on request.
- 2 = Medium: General deletion policy but unclear timelines or verification.
- 3 = High: Clear, time-bound deletion processes with verifiable proof (logs/certificates).

Compliance

- **Does the solution comply with local and national rules (e.g., data privacy, digital governance)?**
 - 1 = Low: No evidence of compliance; ambiguous legal footing.
 - 2 = Medium: Some compliance commitment but gaps remain.
 - 3 = High: Fully compliant; documentation and legal validation provided.
- **Are there terms for adapting to future policy changes?**
 - 1 = Low: No flexibility; contract locks the city into current model.
 - 2 = Medium: Some adaptability but may incur additional cost or delays.
 - 3 = High: Clear contractual clauses for updates, compliance changes, and policy alignment.

Support & Training

- **Will the vendor provide training for your staff?**
 - 1 = Low: No training or only minimal onboarding.
 - 2 = Medium: Basic training offered but not continuous.
 - 3 = High: Structured training program; documentation, materials, and recurring sessions included.
- **Is ongoing support included in the contract?**
 - 1 = Low: Support only during deployment; no long-term commitment.
 - 2 = Medium: Limited support (email/helpdesk) with unclear service level agreements SLAs.
 - 3 = High: Full support with SLAs, response times, dedicated staff, and escalation processes.

Ethics

- **Does the vendor check for bias and avoid unfair or unsafe uses?**
 - 1 = Low: No bias testing; no documentation of mitigations.
 - 2 = Medium: Some testing performed but not systematic or transparent.
 - 3 = High: Regular bias audits, documented risk mitigation, published responsible-AI practices.
- **Can city staff review or override AI decisions if necessary?**
 - 1 = Low: No override option; AI acts autonomously.
 - 2 = Medium: Some override capability but not integrated into workflows.
 - 3 = High: Clear oversight mechanisms; manual review and override fully supported.

Long-Term Fit

- **Will the system be able to grow/adapt with your city's needs?**
 - 1 = Low: Rigid system; scaling requires major redevelopment.
 - 2 = Medium: Some modularity or upgrade options but limited.

- 3 = High: Scalable architecture; regular updates; extensible to new datasets/modules.

- **Are contract terms flexible and clear about responsibilities?**

- 1 = Low: Ambiguous terms; unclear ownership of data/models; lock-in risk.
- 2 = Medium: Terms are partially clear but negotiation required.
- 3 = High: Transparent, flexible contract with clear roles, IP clauses, data ownership, exit plans.

Scoring and Interpretation

Total Score	Vendor Suitability	Recommendation
16–26	Low	High risk. Do not contract; seek alternatives.
27–38	Medium	Proceed only with strict controls and pilot testing.
39–48	High	Vendor is reliable and suitable for long-term partnership.