From Automation to Autonomy:

Agentic AI in Public Governance & Mobility Systems

Discussion Paper

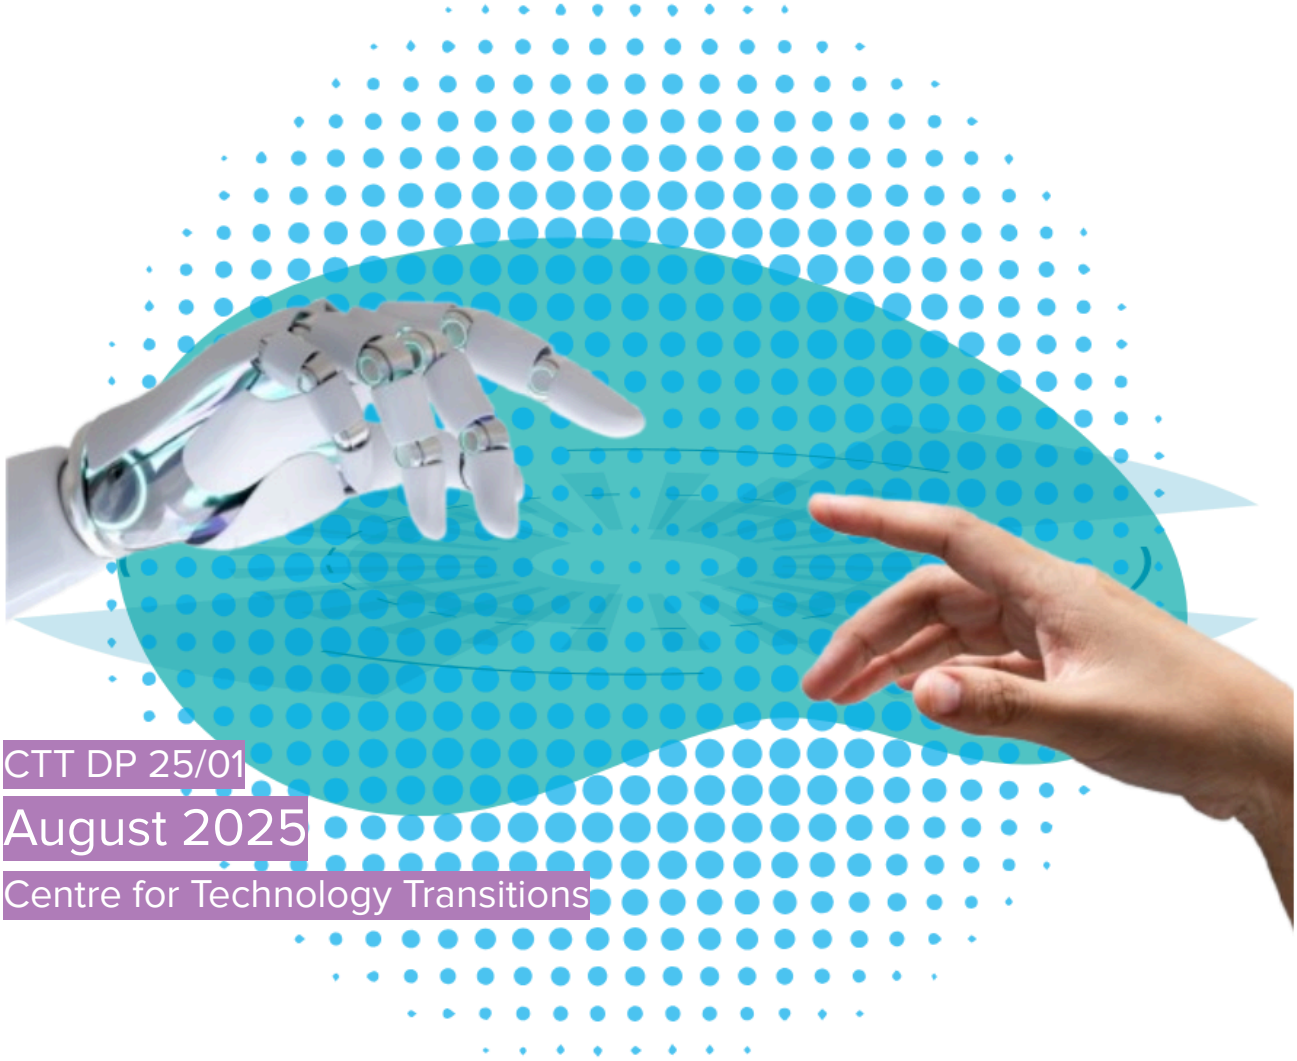

CTT DP 25/01

August 2025

Centre for Technology Transitions

Table of Contents

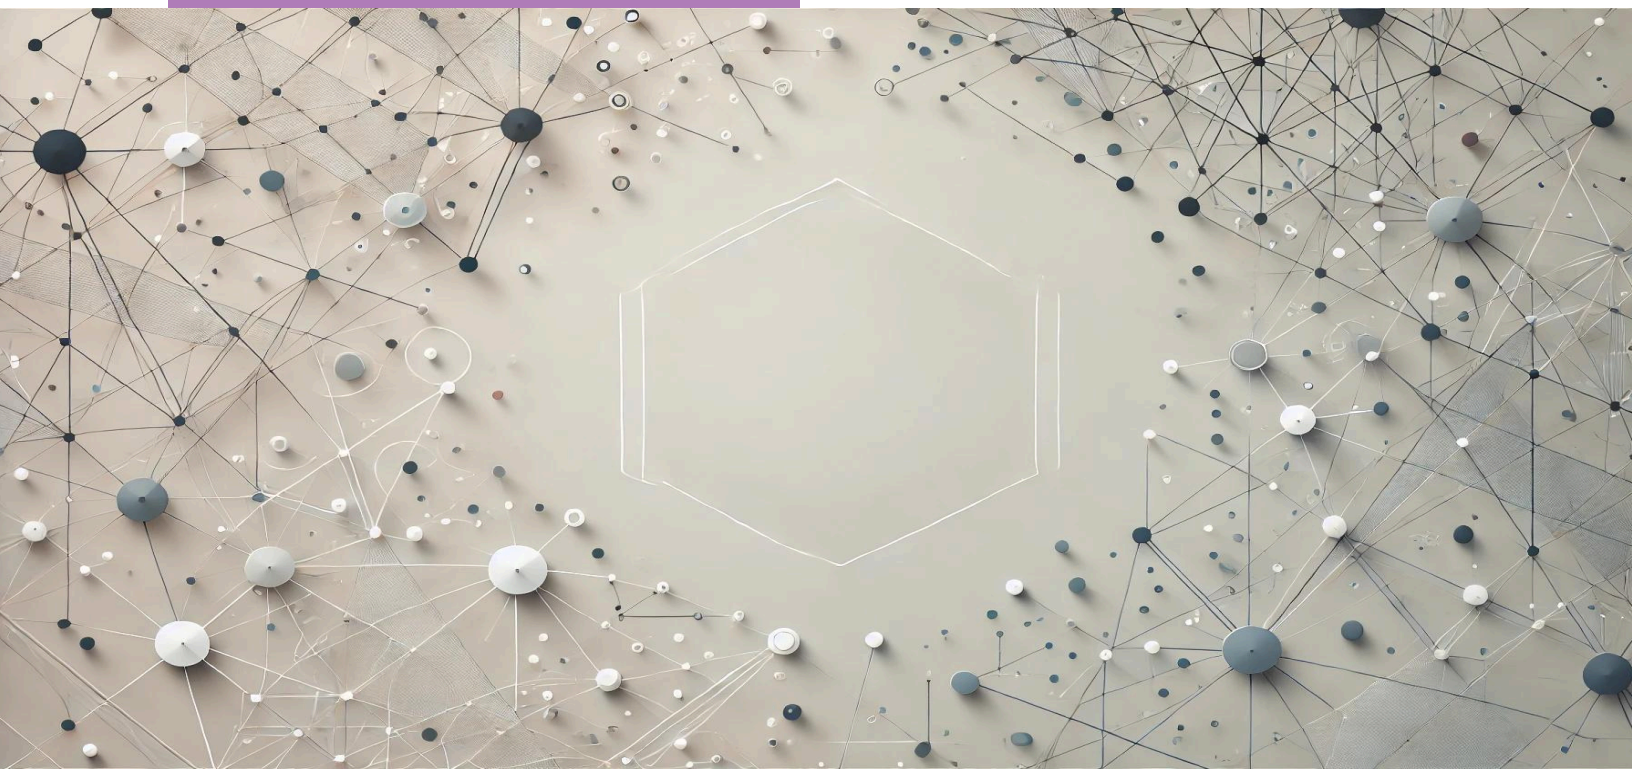

From Automation to Autonomy: *Agentic AI in Public Governance & Mobility Systems*

Discussion Paper

Abstract

The rise of Agentic AI systems, capable of autonomous goal-setting and real-time decision-making, promises to transform public governance, from personalized citizen services (e.g., Estonia's Bürokratt) to crisis response optimization. Yet its adoption surfaces existential questions about ethics, accountability, and regulatory adaptability. This discussion paper analyzes two high-stakes case studies: Tesla's Full Self-Driving (FSD) system and Bürokratt's proactive welfare assistance, to expose tensions between AI autonomy and human oversight.

We highlight three critical dilemmas:

1. **The Liability Gap:** When Agentic AI errs (e.g., a Tesla collision or biased benefit denial), existing legal frameworks lack mechanisms to assign blame across developers, users, or the AI itself.
2. **Transparency vs. Efficiency:** Agentic systems' "black-box" nature (e.g., Bürokratt's predictive algorithms) clashes with GDPR's "right to explanation," risking public trust.
3. **Ethical Trade-offs:** Autonomous decisions in edge cases (e.g., FSD's "trolley problem" scenarios) reveal unresolved conflicts between utilitarian logic and human moral intuition.

The paper concludes with questions and invitations for further discussions, including the need for dynamic auditing protocols and adaptive liability models, and calls for piloting these approaches in sandbox environments.

1. Why Agentic AI Matters Now

As Generative AI continues to evolve, a new frontier is emerging with profound implications: Agentic AI. This paradigm marks a shift from content generation to autonomous goal-directed action. While **Generative AI focuses on the "what"** (e.g., producing text, images, or code), **Agentic AI is concerned with the "how"**. It enables systems capable of perceiving, deciding, and executing tasks with minimal human intervention. Unlike traditional automation, Agentic AI operates with contextual awareness, adaptive learning, and varying degrees of goal-setting autonomy - hallmarks of agency in machine systems.

To illustrate the distinction, consider the task of designing a 3-day travel itinerary. Such a user might receive different outputs from Gen AI, Agentic AI, and a hybrid model.

The GenAI Approach			
Gen AI sees this as a text generation problem; it predicts the most likely sequence of words to produce a coherent itinerary based on its training data (e.g., travel blogs, booking sites, user reviews). It would take the following steps			
Input Understanding: Accepts a prompt like "Create a 3-day itinerary for Tokyo focusing on food and history."	**Pattern Recognition:** Uses its training data to recall common Tokyo attractions (e.g., Sensō-ji, Tsukiji Market), meal recommendations, and typical time allocations.	**Output Generation:** Produces a structured itinerary with: Day 1: Breakfast at Tsukiji, visit Sensō-ji, lunch at a ramen shop, etc. Day 2: ... Day 3: ...	**Limitations:** May suggest unrealistic transitions (e.g., cramming too much in one day). Lacks real-time data (e.g., doesn't check if a restaurant is closed on Tuesdays). No user feedback loop; it won't ask clarifying questions.

The Agentic AI Approach				
Agentic AI treats the task as a goal-oriented process, dynamically interacting with tools (e.g., APIs, databases) and the user to optimize the plan. It would take the following steps.				
Clarify Goals: Asks follow-ups like *"Do you*	**Tool Use:** Queries real-time APIs (Google	**Optimization:** Adjusts the itinerary to avoid	**Execution:** Could book hotels/ flights via	**Iteration:** Revises based on user feedback

prefer luxury or budget dining?" or *"Will you use public transport?"*	Maps for travel times, OpenTable for reservations). Checks weather forecasts or event calendars.	overbooking or conflicts (e.g., accounting for jet lag).	integrations (if permitted).	(e.g., *"Swap the museum for a bike tour"*).

The Hybrid Approach			
The Hybrid Model combines Gen AI's creativity with Agentic AI's dynamic problem-solving. It uses LLMs for content generation but delegates precision tasks to agents. It would take the following steps.			
Draft Generation: Creates a rough itinerary using Gen AI (e.g., suggests neighborhoods/ activities).	**Validation & Enhancement:** Verifies distances via Google Maps API. Checks attraction hours/ scrape recent reviews. Optimizes route order to minimize travel time.	**Personalization:** Adjusts for user preferences (e.g., adds vegan options if detected in chat history).	**Multimodal Output:** Generates a mix of text, maps, and booking links.

2. Agentic AI in Action: Transformative Potential in Governance and Mobility

While the theoretical promise of Agentic AI is compelling, its true impact is best understood through practical application. **This paper focuses on two domains undergoing profound transformation - public governance and mobility systems**. These sectors are not only data-rich and decision-intensive, but also deeply intertwined with societal well-being. From autonomous vehicles navigating dynamic urban environments to AI agents coordinating government services across ministries, Agentic AI is redefining what machines can perceive, decide, and do.

Agentic AI can revolutionize public governance by enhancing decision-making, service delivery, and citizen engagement (Wirtz et al., 2019; Eggers et al., 2020). By analyzing real-time data (e.g., crime stats, health records), it predicts trends like disease outbreaks and optimizes emergency responses (Taddeo & Floridi, 2018). It automates bureaucratic processes, such as verifying welfare claims or personalizing tax assistance, while reducing delays (Sun & Medaglia, 2019). Additionally, it fosters

transparency by logging decisions and auditing resource allocation, exposing biases or inefficiencies (Veale et al., 2018). For citizens, it acts as a 24/7 interface, crowdsourcing feedback for urban projects or explaining policy decisions in plain language.

However, its success depends on ethical safeguards (e.g., bias audits), human oversight, and secure infrastructure (Cath, 2018). Without these, risks like over-automation or hacked systems could undermine trust. If implemented responsibly, Agentic AI could make governance more adaptive, evidence-based, and inclusive - transforming how policies are designed and services delivered (Zhang et al., 2021).

The following sections explore this shift through concrete use cases. We examine how two very different systems - Tesla's Full Self-Driving technology and Estonia's Bürokratt - embody agentic principles in practice. One reimagines mobility through autonomous navigation; the other reinvents public service delivery through personalized AI-driven engagement. Together, they offer critical insights - and urgent questions - on what agency means in machines, and what it demands of us in return.

3. Case Study: Agentic AI in Tesla's Full Self-Driving System

3.1. Overview

Tesla's Full Self-Driving (FSD) system uses agentic AI in autonomous vehicles. It has the ability to make content-sensitive decisions, set goals and act proactively with a degree of autonomy, therefore making decisions like a human. This system assesses its environment continuously and therefore takes real-time driving decisions.

3.2. Problem Context

For autonomous driving, we need AI that can operate in a dynamic, open-world environment; handle ambiguous and unpredictable human behaviors and situations; make independent decisions about the driving goals (i.e. when to stop, slow down and overtake); and acquire and comply with human-centered rules of operation (i.e. courtesy yielding to pedestrians, actual pedestrian behavior). Traditional rule-based AIs (like those found in traffic control systems) never had to operate in this dynamic type of environment and cannot exhibit agentic behaviour (i.e., autonomy, adaptability, and proactivity).

3.3. Tesla's Agentic AI Architecture

Tesla uses vision-based neural networks that have learned from real-world data. The perception system relies on all camera inputs instead of lidar/radar and computes the environmental information provided by the camera inputs.

Tesla's planning system makes use of multiple-agent trajectory prediction, which means predicting the actions of other agents (vehicles, pedestrians, etc.). The AI chooses safe and efficient paths through the environment using dynamic agentic planning models.

The AI will adjust to current conditions on the road, unexpected objects, or aberrant traffic behaviors. The AI will adjust the plan to modify its behavior upon encountering construction zones, erratic drivers, or emergency vehicles.

3.4. Key Agentic Behaviors

Tesla's Full Self-Driving (FSD) system demonstrates qualities of agency in its ability to respond to real-world driving situations that aren't just based on a set of predetermined rules. For instance, while the system may encounter an intersection that lacks traffic control signs or markings, it will not simply wait for something defined in the rulebook; it will determine whether the nearby vehicles will yield right-of-way based on an assessment of their behaviour.

In a similar vein, the system will use a pedestrian's body language (such as looking at the driver or walking speed) to determine whether the vehicle should stop and yield. Even in parking lots and construction sites - situations where GPS and/or lanes may not be trusted or available - the system will rely on its understanding of the context to make decisions that enable it to operate safely. Each of these instances exemplifies the system's behavior and abilities as agentic, stateful AI, as it has the potential to make decisions and operate contextually.

3.5. Key Questions

1. **Responsibility and accountability:** Who is liable in the event that agentic AI gets it wrong?
2. **Ethical decisions:** If faced with a potential collision, what does the AI do to determine lives saved or what is acceptable priority?
3. **Explainability:** FSD is a black box in many respects, making auditing or prediction of its behavior complicated.
4. **Regulation:** Agentic AI disrupts accepted regulatory paradigms that were fundamentally rooted in human operation.

3.6. Discussion: Tesla's FSD and the "Trolley Problem" - Why It Matters

The "trolley problem" (a classic ethical dilemma about choosing between two harmful outcomes) isn't just theoretical for Tesla's FSD - it exposes unresolved tensions in agentic AI (Lin, 2016; Nyholm, 2018):

1. **Real-World Stakes:** In edge cases (e.g., a sudden pedestrian vs. swerving into oncoming traffic), the AI must make split-second decisions with moral implications. Unlike humans, its choices are pre-programmed or learned from data, raising questions:
    a. Should it prioritize passenger safety over pedestrians?
    b. How are these rules encoded, and who decides them? (Bonnefon et al., 2016)
2. **Legal Gray Zones:** Current liability laws don't account for AI agency. If a Tesla injures someone, is the fault with the programmer's training data, the car owner, or the AI itself? This parallels debates around "AI personhood" in legal systems (Gurney, 2018; Pagallo, 2013).
3. **Public Trust:** Transparency in how such decisions are made is critical. Tesla's opaque "black box" algorithms fuel skepticism, highlighting the need for explainable AI in life-or-death contexts (Burrell, 2016; Wachter et al., 2017)..

Implication: Tesla's struggles mirror broader societal challenges: autonomous tech outpaces our ethical and regulatory frameworks, demanding interdisciplinary solutions (e.g., ethicists + engineers co-designing decision trees) (Rahwan et al., 2019; Borenstein et al., 2017).

4. Case Study: Estonia's Bürokratt - Agentic AI for Digital Governance

4.1. Overview

Estonia has developed Bürokratt, an agentic virtual assistant, utilizing AI to enhance the relationship with citizens for many government services. Bürokratt incorporates several databases and services and provides tailored and anticipatory governance assistance that builds on traditional e-government with autonomous goal-setting and decision-making capabilities.

4.2. Problem Context

Governance today faces significant challenges, including

1. Services provided fragmented across departments.
2. Bureaucratic delay to receive benefits (e.g. pensions, permits).
3. Low levels of citizen engagement given complex processes and interfaces.

To reduce these challenges, Estonia sought an AI system that would:

1. Engage with citizens in a proactive way.
2. Aggregate siloed information across ministries.
3. Improve decision-making accuracy while reducing reliance on human involvement.

4.3. Agentic AI Features of Bürokratt

1. **Integration Among Agencies:** Bürokratt connects services between ministries, municipalities, and registries, providing citizens with one point of access.
2. **Proactive Communication:** Bürokratt has the advantage of not having to wait for citizens to contact it first. It can predict when an individual will have a need for services based on events that happen in their life (birth, retiring) and initiate offers for them (e.g., enroll a child, claim benefits).
3. **Context-Sensitive Decision-Making:** The AI will review individual citizen data (income, age, job, residence) to independently make relevant decisions on the best options for the citizen (e.g., best subsidy, offer of service).
4. **Communication and Interface Competence:** Bürokratt can communicate in multiple languages and adapts its responses to a citizen based on their digital literacy so that the experience is participatory.

4.4. Key Agentic Behaviours

Bürokratt demonstrates agentic behaviour through its capacity for context-sensitive decision-making. By accessing and synthesizing information from consolidated government databases, it can interpret an individual's socio-economic context - such as income level, employment status, and major life events - and autonomously initiate appropriate services. For instance, if a citizen loses their job, Bürokratt can detect this change, identify relevant welfare benefits or reskilling programmes, and proactively recommend them - without requiring any manual input. This ability to perceive context and act independently distinguishes Bürokratt from conventional rule-based systems, elevating it from a reactive chatbot to a truly agentic AI system.

4.5. Key Questions

While Bürokratt offers efficiency and transparency, by reducing administrative burdens and automating logs and audit trails, it comes with some unanswered questions.

1. **Data protection:** Personalized vs. GDPR & privacy legislation.
2. **Automated over-reliance:** Key human decisions remain ambiguous, as human decision-making is needed for complexity or ethical processes.
3. **AI bias and explainability:** Auditing AI has to continue so that different demographics are treated fairly.

4.6. Discussion: Estonia's Bürokratt and GDPR - The Privacy Tightrope

Bürokratt's hyper-personalized governance clashes with GDPR's strict privacy protections, illustrating a core tension in public-sector AI:

1. **Data Hunger vs. Privacy:** To predict citizen needs (e.g., automatic pension enrollment), Bürokratt requires real-time access to sensitive data (income, health records). GDPR mandates minimal data collection and explicit consent - but proactive AI thrives on maximal, anticipatory data use.
    a. **Example:** If Bürokratt suggests a disability benefit, does analyzing the citizen's medical history without direct request violate GDPR's purpose limitation principle?
2. **Anonymization Trade-offs**: Aggregating data to avoid profiling weakens personalization. Estonia's solution (e.g., federated learning) must balance utility with anonymity.
3. **Right to Explanation:** GDPR grants citizens the right to know how AI decisions affect them. Bürokratt's complex algorithms might struggle to provide simple, actionable explanations (e.g., *"Why was my housing subsidy denied?"*).

Implication: Estonia's model tests GDPR's flexibility. Its success could redefine how democracies reconcile AI efficiency with fundamental rights, setting a global precedent.

5. Conclusion: Designing the Agentic Future

Agentic AI signals a shift in how we define intelligence, autonomy, and public value in the digital age. From navigating complex roadways to delivering personalized governance, its promise lies in systems that don't just compute but act - independently, contextually, and sometimes, unpredictably.

But with this autonomy comes a cascade of unresolved questions:

1. What kind of oversight is meaningful when systems make their own decisions?
2. Can explainability keep up with real-time adaptation?
3. Should agentic systems mirror human judgment or resist it?

As governments, companies, and citizens consider the role of Agentic AI, we are no longer just asking what machines can do, we must ask what they should do, and who decides. This paper offers a starting point, not a conclusion. We invite interdisciplinary dialogue, from policymakers to engineers, ethicists to end-users, to shape an agentic future that is not only intelligent, but also just, accountable, and human-centered.

Credits

About OMI Foundation

OMI Foundation Trust is a new-age policy research and social innovation think tank operating at the intersection of mobility innovation, governance, and public good. Mobility is a cornerstone of inclusive growth providing the necessary medium and opportunity for every citizen to unlock their true potential. OMI Foundation endeavours to play a small but impactful role in ushering meaningful change as cities move towards sustainable, resilient, and equitable mobility systems which meet the needs of not just today or tomorrow, but the day after.

OMI Foundation houses four interconnected centres that conduct cutting-edge evidence-based policy research on all things sustainability:

1) The Centre for Technology Transitions is dedicated to transforming India's innovation ecosystem through a systems approach. It aims to position India as a global leader in ethical, inclusive, and sustainable technological innovation.

2) The Centre for Future Mobility supports the leapfrog of cities to a sustainable future anchored in the paradigms of active, shared, connected, clean, and AI-powered mobility.

3) The Centre for Clean Mobility catalyses the adoption of electric vehicles, future fuels, and renewable energy within the mobility ecosystem as a key climate strategy of cities.

4) The Centre for Inclusive Mobility promotes safe, accessible, reliable, and affordable mobility for all. It paves the road for the future of work and platform economy to fulfil the modern promise of labour.

Authors

Jagriti Arora,
Lead, Centre for Technology Transitions

Jagriti approaches tech policy with curiosity and a readiness to question the status quo. With a degree in urban planning and additional training in technology policy through short-term courses, she brings a multidisciplinary perspective to complex challenges. Her experience spans policy research in mobility, the platform economy, Human-Computer Interaction (HCI), political consulting, and pedagogy, guided by research and first principles thinking.

Himani Upadhyay,
Research Volunteer

Himanshi is a Research Voluteer at the Center for Technology Transitions, OMI Foundation. She is a second year Public Policy student at the National Law School of India University, Bengaluru. Her interest lies at the intersection of technology, policy and social impact.

Editor

Aishwarya Raman, Executive Director, OMI Foundation Trust

Suggested Citation

Arora, J. & Upadhyay, H. (2025, August). *From Automation to Autonomy: Agentic AI in Public Governance and Mobility Systems.* Centre for Technology Transitions Discussion Paper (CTT DP) 25/01. OMI Foundation.

Disclaimer

Neither OMI Foundation Trust nor its donor(s) nor any other party associated with this report will be liable for any loss or damage incurred by the use of this report.

Image Credits

Image Credits: Cover image created by the lead author on Canva and freely available images on Raw Pixel. 'Table of Contents' image generated by DALL·E (via ChatGPT).

Contact

For feedback and collaborations, email: comms@omifoundation.org

Copyright

Copyright © 2025 OMI Foundation

OMI FOUNDATION

https://omifoundation.org

https://x.com/OMIFoundation

https://linkedin.com/company/omi-foundation

https://www.instagram.com/omi.foundation

https://www.facebook.com/omifoundation.org

https://www.youtube.com/@OMIFoundation2023

Scan the QR code to view our website