

Building Trust by Design: What Espoo (Finland), Copenhagen (Denmark) and Amsterdam (Netherlands) Teach Regions about Civic AI



DOI reference: 10.1080/13673882.2025.12533896

By Mahshid Cheraghi (e-mail; LinkedIn), PhD candidate in sustainability and urban regeneration at the Departamento de Urbanística y Ordenación del Territorio, Escuela Técnica Superior de Arquitectura - Universidad Politécnica de Madrid, Universidad Politécnica de Madrid, Spain

Across Europe, artificial intelligence is moving from pilot to practice in urban planning and public services. Yet technical capability does not automatically lead to public trust. For many regions, the decisive factor is whether citizens trust the systems that influence decisions affecting their lives. This trust depends less on

performance metrics alone and more on whether processes are transparent, inclusive, and accountable. Examining Espoo, Copenhagen, and Amsterdam offers insight into three distinct pathways of civic AI adoption, each highlighting how different governance choices shape public confidence. The experiences of these cities illustrate both the opportunities and the risks of adopting AI in public governance, and that building legitimacy is as essential as achieving efficiency. This article draws from my doctoral project on citizen-centred smart urban regeneration.

- **Research question:** *How do local governance choices (explainability, lifecycle transparency, and redress) shape public trust in civic AI?*
- **Method:** *a comparative case study (Espoo, Copenhagen, Amsterdam) using document analysis and policy materials, complemented by practitioner insights from public sources.*

This article focuses on three practical steps regions can adopt now: (1) plain explanations for any AI system; (2) transparency over time through a public algorithm register (a public webpage listing each system, its purpose, data used, responsible contact, and risks); and (3) feedback and redress (a clear route for questions or appeals with promised response times) so people can ask questions or challenge outcomes.

Espoo: Explainability and feedback by design

Espoo has pioneered participatory mapping through its “**My Espoo on the Map**” initiative. Thousands of residents have contributed input on issues ranging from mobility to green space, with the city demonstrating how these entries inform planning outcomes. What distinguishes Espoo is not only the scale of engagement but the feedback loop: participants see how their contributions are analysed and reflected in subsequent plans. This closes the gap between citizen input and policy response, making participation tangible. Espoo exemplifies a model where explainability and responsiveness are built into the planning cycle, ensuring that civic AI tools are experienced as legitimate rather than opaque. Importantly, the city frames citizen input not as an optional add-on but as a central element in the design of policies, signalling that residents’ voices matter in shaping the future of their communities.

What this means in practice. People can follow their idea from input to

outcome, rather than wondering if it gets lost in a black box.

How it maps to the three steps.

- Plain explanations: short, non-technical updates alongside the map.
- Transparency over time: published updates that track changes from draft to decision.
- Feedback and redress: contact points for follow-up questions and clarifications.

Copenhagen: High capability, limited formal contestation

Copenhagen's "**Solutions Lab**" has positioned the city as a testbed for innovative AI and data-driven applications. Pilot projects in energy optimisation, mobility, and environmental monitoring demonstrate a strong technical capacity and the ability to deploy new systems rapidly. However, while outputs are visible, the city has invested less in institutionalised mechanisms for public scrutiny of these systems. Unlike Espoo, Copenhagen lacks a city-wide algorithm register (a single page listing each system with its purpose, data, risks, and contact person) or standardised disclosure tool. As a result, citizen involvement remains tied to specific projects rather than embedded at the system level. This approach demonstrates capability but risks leaving legitimacy contingent solely on performance. It highlights a broader challenge for technologically advanced cities: the risk of prioritising innovation and efficiency without investing equally in transparency and public accountability.

What this means in practice. Residents may hear about an innovative tool, but they often lack a single source to look it up, learn how it works in plain terms, and know whom to contact.

How it maps to the three steps.

- Plain explanations: visible for some pilots, but not consistent.
- Transparency over time: a gap without a register that is kept up to date.
- Feedback and redress: a gap without a routine route to ask questions or appeal outcomes.

Amsterdam: Institutionalising transparency through an Algorithm Register

Amsterdam offers a contrasting path by embedding transparency through its "Algorithm/AI Register", launched in 2020 and later aligned with the Dutch

national register. The register documents high-risk and experimental systems, with entries that describe purpose, datasets, oversight, and potential risks. Citizens can consult these records to understand how algorithms are used in decision-making and offer feedback. Importantly, the register is maintained as a living instrument, updated as systems evolve. By institutionalising disclosure and linking it to policy frameworks such as the Digital City Agenda, Amsterdam ensures that transparency is not optional but routine. This positions the city as a leader in demonstrating how system-level governance instruments can anchor legitimacy in civic AI. It also demonstrates how regions can transition from pilot projects to sustainable governance frameworks, offering a replicable model for other municipalities seeking to strike a balance between innovation and accountability.

What this means in practice. A resident can search the register, read a clear summary, and send feedback to the listed contact. Staff can point to a single source of truth and update it when the system evolves.

How it maps to the three steps.

- Plain explanations: built into each register entry.
- Transparency over time: updates are part of everyday work.
- Feedback and redress: contact details and links to feedback channels are listed with each entry.

Overall findings as this research continues

- Trust follows design, not performance: participants respond to visible feedback loops more than to technical metrics.
- System-level disclosure beats pilot-by-pilot PR: Amsterdam's register institutionalises transparency; Copenhagen's project-based disclosures don't scale.
- Low-cost explainability matters plain-language updates and named contacts reduce "black box" (AI-supported public service) perceptions even for simple tools.
- Compliance ≠ legitimacy: aligning with EU AI Act obligations helps, but citizen-facing explainability and redress drive acceptance in practice.

Comparison at a glance: Espoo operationalises explainability and feedback;

Copenhagen demonstrates capability without consolidated disclosure; Amsterdam embeds transparency by default via a living algorithm/AI register, offering three distinct routes to trust.

Making trust operational

Taken together, these cases demonstrate that trust is not a byproduct of AI performance, but rather an outcome of governance choices. Espoo builds legitimacy through visible responsiveness, Copenhagen demonstrates the risks of focusing on technical excellence without sufficient disclosure, and Amsterdam leverages transparency to create infrastructure. For other regions, the lesson is clear: capability must be paired with institutions that provide citizens with clarity, transparency, and involvement. Without these features, even the most sophisticated technologies risk being viewed with suspicion or indifference, undermining their potential impact.

Practical tools for regions

Lessons for other urban regions: Start with model cards, a public algorithm register, and a redress route with response-time commitments; keep all three updated as systems evolve.

- **Model cards:** one-page, non-technical summaries of a system and its limits that explain systems in plain language, covering purpose, inputs, limits, and a contact point.
- **Decision pathway maps:** simple notes showing how inputs become decisions that show how inputs become outputs; three or four boxes are enough for most services.
- **Redress routes:** Clearly defined routes for questions or appeals with promised response times, allowing citizens to ask questions or appeal outcomes. Post the route next to any tool that uses AI.
- **Feedback response-time commitments,** for example, “we will reply within 10 working days,” enable consultation to become an ongoing dialogue.
- **Algorithm registers:** public webpages listing each system, its purpose, the data it uses, the responsible party, and known risks, which document systems across their lifecycle, ensuring transparency is systematic rather than ad hoc.

For any new system, publish a model card (Clarity), add an entry to the register, and include a short decision pathway (Traceability). Additionally, post a redress route with response-time commitments (Involvement).

The CTI framework

We track adoption with a CTI (Clarity-Traceability-Involvement) check. Clarity means plain explanations that a non-specialist can read in five minutes. Traceability means a public trail, register entry, decision pathway, and update notes that others can follow. Involvement means people can take part and seek redress, with published response times. By scoring systems against these three pillars, regions can identify weaknesses, benchmark progress, and ensure that trust is maintained as systems scale. Is there a plain-language model card? A register entry with update history? A working redress route with a promised response time? If yes to all three, trust can scale.

Regional implications

These cases resonate beyond the municipal scale. The European Union is currently developing the AI Act (in force since August 1, 2024; prohibitions & AI literacy apply from February 2, 2025; GPAI model obligations from August 2, 2025; most high-risk rules from August 2, 2026). The EU is also strengthening digital governance frameworks that require transparency, explainability, and human oversight. The practices of Espoo, Copenhagen, and Amsterdam illustrate how local administrations can anticipate and align with these requirements, creating compliance pathways that also enhance legitimacy. By embedding algorithm registers, redress mechanisms, and participatory mapping, cities are effectively localising EU-level principles into everyday governance. For regional planners, this demonstrates that building civic trust is not separate from meeting regulatory expectations but integral to it. Moreover, aligning local innovation with European strategies enables regions to share best practices, benchmark their performance, and strengthen collective resilience in the face of technological change.

Conclusion

AI in regional governance is here to stay, but its acceptance depends on more than technical outputs. Espoo demonstrates how participatory explainability fosters confidence, Copenhagen illustrates the risks of relying on capability without disclosure, and Amsterdam shows how registers can institutionalise transparency. For regions aiming to harness AI responsibly, the message is clear:

treat trust as civic infrastructure: pair every AI system with a model card, a public register entry, and a redress route with response-time commitments, and track them with CTI. Cities differ, but these three commitments travel well. Only by treating legitimacy as a core design principle will smart systems achieve the durable consent needed for regional transformation.

References

- City of Espoo. (2024). My Espoo on the Map [online]. Available at: <https://espoo.fi>
- OECD. (2023). Smart City Data Governance: Copenhagen Case Study [online]. Available at: <https://oecd.org>
- City of Amsterdam. (2023). Digital City Agenda and Algorithm Register [online]. Available at: <https://algoritmeregister.amsterdam.nl>
- Government of the Netherlands. (2022-). National Algorithm Register [online]. Available at: <https://algoritmes.overheid.nl>
- Haataja, M., van de Fliert, L., & Rautio, P. (2020). Public AI Registers: Realising AI transparency and civic participation in government use of AI. City of Helsinki & City of Amsterdam. Available at: <https://ai.hel.fi>
- Københavns Kommune. (2020). Kodeks for anvendelse af kunstig intelligens.
- Datatilsynet. (2023, November 22). Kommunes hjemmel til AI-løsning til identifikation af borgere med behov for vedligeholdende træning og rehabiliterende indsats. <https://www.datatilsynet.dk/english>