IEAI Technical University of Munich
School of Social Sciences and Technology
Institute for Ethics in Artificial Intelligence

Research Brief – Q4 2025

Discussing AI and Ethics with Local Communities: Insights from Amsterdam

By Catarina Fontes and Milou Jansen

Urban systems incorporating AI provide new ways to monitor and manage cities, enhancing local governments' ability to intervene, optimize and reshape urban processes. However, these systems are also socio-technical constructs, drawing inherently from, and reproducing, value-laden representations of urban life. Public participation in these systems fosters dialog, enabling deliberation on values and trade-offs and improving AI literacy. It also serves as a channel for communicating citizens' priorities, needs and preferences to policymakers, and when objectives are aligned with anticipated impacts and expectations it can yield pivotal and rewarding outcomes for all stakeholders.

This Research Brief depicts how a participatory process focused on deliberation and value-mapping can help identify critical issues related to data collection and use, factors influencing technology acceptance, and key implementation safeguards. The Brief also highlights persistent challenges, including balancing representativeness and inclusiveness, ensuring informed participation and linking input to policy design.

Introduction: Why and How to Engage the Public in the Making and Governing of Complex Systems

AI is increasingly used by cities to support governance through techniques like data mining, clustering, classification, pattern recognition and prediction, helping to automate or inform public sector processes (Wirtz et al., 2019). This reflects a governance model built on the belief that legal, organizational and professional rules can be decoded and encoded, with technology viewed as a neutral and objective provider of evidence-based insights for policymaking. However, cities are more than static artifacts, being complex, adaptive systems shaped by human interaction and planning (Portugali, 2016). Similarly, AI systems designed to model and manage them are also socio-technical constructs, drawing inherently from, and reproducing, value-laden representations of urban life (Batty, 2016; Caldarelli et al., 2023).

As citizen participation is known to be essential to effective governance, this also is a very relevant task for effective AI governance. Participation strengthens the legitimacy of decision-making, improves public policies and services, and fosters inclusion and diversity. Engaging citizens helps authorities address complex issues by bringing in diverse perspectives, filling idea gaps and shaping more equitable priorities. It supports better decisions through relevant feedback, co-created solutions and policies that reflect real needs. Meaningful participation supports policy implementation, reveals blind spots, helps prevent conflict, builds trust and increases legitimacy. For citizens, participation is more than a benefit. It is a fundamental right and lever for transparency and empowerment, and promotes shared responsibility (OHCHR, 2019; UN-Habitat, 2025).

The challenge remains how to best involve citizens to accurately determine their values and priorities for AI and its use in public governance, while ensuring that the process is meaningful and rewarding for stakeholders involved, generating real impact. In this brief, we present a methodology to address this challenge, drawing on a participatory process conducted in the city of Amsterdam. The results show that, despite certain shortcomings and challenges, defining clear and transparent objectives within the participation process is paramount to ensure that stakeholders' expectations align with anticipated outcomes. Furthermore, the meaningfulness of the process should rest on both community-oriented benefits and the added value it brings to context and value-sensitive policy design (in this case) concerning the adoption of AI systems to support urban management and governance.

The participatory process was designed to strengthen AI literacy, awareness of ethical implications of AI, enhance participants' capacities to engage meaningfully in deliberations on AI and contribute to community empowerment.

Bridging Representativeness and Inclusiveness in Participatory Policymaking: a Methodological Approach

The participatory methods presented in this Research Brief were developed in collaboration with the city of Amsterdam. They were intended to demonstrate how representativeness and inclusiveness can be effectively bridged in participatory policymaking with a focus on digital rights. By integrating large-scale, representative consultation with targeted, inclusive engagement, the goal was to use this approach to enable an effective value-mapping process that reveals diverse local perspectives, priorities and concerns surrounding the adoption of AI systems in urban governance. The participatory process was designed with a dual focus: policy-oriented and community-oriented, and to bridge the principles of representativeness and inclusiveness in policymaking. In addition, the process was designed to strengthen AI literacy, awareness of ethical implications of AI, enhance participants' capacities to engage meaningfully in deliberations on AI, and contribute to community empowerment. Today, few online experiences are untouched by these systems. From social media feeds to search results to shopping platforms, personalization engines influence what users encounter and how they interact with digital spaces. Yet many remain unsure about the

---

Source Title Page Image: Catarina Fontes - Workshop in Amsterdam *"Digital Rights in the eyes of Amsterdam citizens"*

Project page: www.ieai.sot.tum.de/research/ethics-for-the-smart-city-applied-socio-technical-frameworks-to-assess-the-implementation-of-ai-related-solutions/

mechanisms behind the content shown to them or the data that informs these decisions (McDonald & Cranor, 2010). The involvement of artificial intelligence in these processes adds further complexity and makes their influence even harder to detect or interpret.

From a policy perspective, the process aimed to:

- Reflect on deliberative democracy, representativeness and inclusion in public participation.
- Advance methods and recommendations for improved governance through participatory practices.
- Support cities in the ethical and value-aligned implementation of AI-enabled systems.

From a community perspective, the objectives were to:

- Gather public opinions on the use of personal data in AI systems.
- Explore how people interact with AI in urban and public settings.
- Promote AI literacy and awareness.
- Foster public deliberation on the social and ethical implications of AI.

To achieve both representativeness and inclusiveness, the process combines quantitative and qualitative methods: a large-scale online survey to reach a broad cross-section of residents and a series of targeted workshops designed to engage underrepresented and vulnerable groups.

Large-scale Online Survey

The purpose of the online survey was to capture local values, attitudes toward privacy and the meaninging attached to personal data. It explored how residents perceive privacy boundaries and their willingness to make trade-offs, such as sharing personal data in exchange for individual or societal benefits.

The survey consisted of four main sections:

1. Local perceptions of privacy and intrusion.
2. Attitudes toward personal data, including risk perception and trade-offs.
3. Responses to hypothetical, value-based scenarios involving AI technologies.
4. Socio-demographic profiling to ensure diversity in responses.

It examined multiple dimensions, including spatial contexts (e.g.: public spaces, digital environments), types of data collection (e.g.: sensors, public records, smart devices), perceived risks and benefits, legal and rights-based frameworks (e.g.: GDPR, AI Act), and levels of trust in institutions and technologies. The design enabled a nuanced understanding of public attitudes toward data use in AI-enabled urban governance. A total of 331 complete responses were collected in this process.

Targeted Workshops

The goal of the workshops was to ensure inclusiveness by engaging citizens from vulnerable groups, particularly elderly residents and youth organization representatives. A total of 30 participants took part in structured sessions designed to promote reflection, dialog and collective learning about AI and its implications for urban governance.

Figure 1. *Cards used to prompt group deliberation focusing on ethics-based scenarios for the implementation of video surveillance systems enhanced with AI features in cities.*

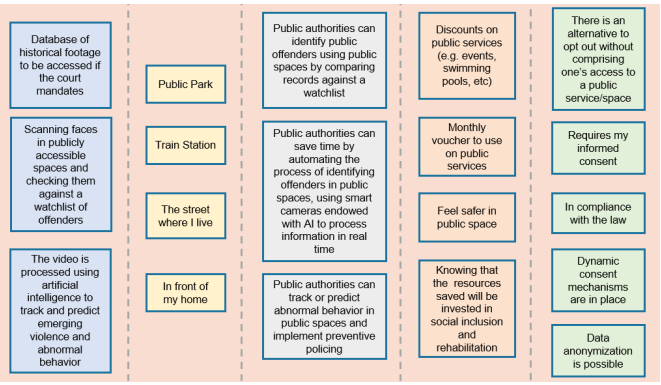

Source: Data collected during workshops with residents in Amsterdam (September 2023).

Each workshop was divided into four parts, combining individual reflection and group-based activities:

1. Survey completion and revision: Participants began by completing the online survey, with opportunities to revisit and update their responses during the session.
2. Discussion of AI systems: Participants examined three AI applications (recommendation tools, surveillance

technologies and AI-generated content), using visual prompts and thematic cards to identify risks, opportunities and ethical considerations.

3. Ethics-based scenario building: Using a five-level analytical framework (data generation, data collection process, potential benefits, compensations and incentives, and ethical and legal safeguards), participants collaboratively constructed ethical scenarios for three urban AI systems: air quality monitoring, digital identity in public transport and video surveillance in public spaces. These exercises helped identify ethical trade-offs and red lines in AI deployment.
4. Reflection on sustainability goals: Participants selected relevant sustainable development goal (SDG) themes and discussed how AI technologies could both advance and hinder sustainability objectives.

Figure 2. *Activity related to AI and sustainability.*

Source: Data collected during workshops with residents in Amsterdam (September 2023).

Adapting to an Empirical Ground: Amsterdam as a Case Study

Key to the methodological approach presented above was to adapt the process to the unique context of the city and population in question. Amsterdam is an early mover in this space and has its own history in terms of embracing the use of AI in public administration.

The City of Amsterdam developed a comprehensive strategy, supported by policy frameworks and governance tools, to ensure the ethical and transparent deployment of AI and maintain public trust in municipal decision-making. Their prioritization of developing such a strategy was based in their recognition that while AI-based systems present opportunities for public administration to enhance responsiveness, efficiency and innovation, they also raise critical concerns related to fairness, accountability and democratic legitimacy (OECD, 2025). Trust in government institutions is essential for strengthening participation and legitimacy through public acceptance and policy adherence, two key dimensions of government–citizen interaction. This trust dimension is equally central to the adoption and use of AI in cities, where the primary challenge lies in upholding citizens' rights and values (Schiff et al., 2022).

The city's *Vision on AI* (City of Amsterdam, 2024) along with complementary policy documents such as the *Grip on Algorithms Playbook* (El Yassini et al., 2022) and the *Fairness Handbook* (City of Amsterdam, 2022), together constitute a comprehensive local framework for algorithm governance grounded in European context. This framework introduces a set of measures to ensure the responsible use of AI, including the public disclosure of algorithms through an algorithm register, bias analysis, mitigation procedures, mechanisms for contesting automated decisions and auditing protocols to safeguard accountability. Collectively, these initiatives demonstrate the city's commitment to institutionalize digital rights, such as transparency, non-discrimination and procedural fairness, within its administrative processes.

Although participatory steps informed the drafting of these policy documents, recent evaluations reveal shortcomings in the implementation of participation. Engagement practices have often been fragmented, taking place at later stages of policy cycles and serving primarily consultative or legitimizing purposes rather than enabling co-creation and allowing for real impact. This limits opportunities for citizens to meaningfully influence how AI systems are designed, deployed and monitored[1]. In response, initiatives such

[1] See: https://publicaties.rekenkamer.amsterdam.nl/participatieparagraafbestuurlijk-rapport/

as the Digital Rights House seek to embed digital rights awareness and participation more deeply within local institutions, fostering continuous dialog between citizens, policymakers and civil society organizations. Against this backdrop, Amsterdam provides a highly relevant empirical setting to test participatory methodologies that move beyond academic inquiry toward practical impact.

In collaboration with the Digital Rights House, the City of Amsterdam, and in dialog with UN-Habitat, a research team from the Technical University of Munich employed the methodological approach presented above with Amsterdam residents. The findings highlight how local communities perceive privacy risks associated with specific AI-enabled systems, how they connect these concerns to data collection processes supporting AI in public administration and what ethical trade-offs and red lines emerged from such deliberations.

Involving Local Communities in Value-mapping: Findings from a Participatory Approach and Implications for Context-sensitive and Value-aligned AI

The findings from survey and workshop results can be structured around three evolving themes: (1) privacy and personal data, (2) AI acceptance and (3) sustainable development.

On the Value of Privacy and its Relation to Personal Data

Privacy was deeply valued among Amsterdam residents, shaping attitudes toward data collection and data sharing, also in the context of AI-enabled systems. Survey respondents expressed concern about privacy violations, particularly those involving social media tracking, audio sensors and AI-driven surveillance. Lawfulness alone did not appear to alleviate these concerns. Rather, transparency, consent and trust in the entities managing data were more decisive factors for public acceptance. These findings echo broader scholarly debates on the ethical foundations of smart city governance, which emphasize that privacy protection and meaningful consent are critical for legitimacy (Kitchin, 2016).

Respondents' limited willingness to trade personal data for monetary or community rewards reinforces that privacy should not be understood as a tradable commodity, but as a social and moral value embedded in power relations, and essential to human autonomy and dignity (Fuchs, 2011). Data sharing was perceived as more acceptable when linked to public benefits, such as health, safety or transport efficiency, yet income and health data, for instance, were considered highly sensitive and largely off-limits. Regarding consent, participants highlighted the importance of informed consent and dynamic consent management, with possibility to withdraw an active renewal. This points to the preference of ongoing agency and data ownership control, rather than one-off authorization. Results also revealed that one-third of participants were unfamiliar with the GDPR, highlighting the need to improve literacy on data protection rights and to raise public awareness about data governance and the everyday implications of algorithmic decision-making.

Complementary, workshop discussions around accountability and governance underscored that questions of "who decides" acceptable trade-offs are as important as the technical safeguards themselves. Different stakeholders, citizens, policymakers, researchers and technology providers, bring diverse expectations and have different explanatory requirements, reinforcing the importance of transparency on underlying purposes and objectives as well as anticipated impacts.

The use of video surveillance and facial recognition systems in public spaces underscores the tension between policies that rely on technology to support public safety objectives and the overexposure of communities to (automated) surveillance.

With regards to surveillance systems, respondents expressed strong resistance to unconsented data collection, even in public spaces. Findings emphasize that transparency should extend beyond simple signage and that individuals should retain the option to opt out without compromising their right to access public spaces. The "right to know" presumes that available information enables an informed voice in decisions and the assessment of decisions made by those in power (Florini, 2007). The use of video surveillance and facial recognition systems in public

spaces underscores the tension between policies that rely on technology to support public safety objectives and the overexposure of communities to (automated) surveillance. This tension has led to community backlash and prompted many cities to ban the use of real-time, remote biometric identification, a position also supported by the European Data Protection Board (EDPB, 2021) in its response to the EU AI Act proposal.

On AI Technology Acceptance: Awareness, Tradeoffs and Redlines

Overall, collective goals and goods, such as public health, sustainability, equitable resource allocation and urban safety, were seen as stronger justifications for AI systems than efficiency or user convenience. Awareness of personal harms, especially privacy violations or denied accessibility, was considered more immediate than broader societal risks (in line with Zimmermann et al., 2019). Discussions about automated video surveillance in public spaces provide a clear example: participants focused on risks to privacy and the lack of any possibility to opt out, rather than on the broader implications of mass surveillance and its impacts on power relationships. Nonetheless, a few groups justified drawing a red line at the automation of public surveillance, arguing that it could lead to an excessive concentration of power in government hands. Similarly, when discussing AI-generated content, some groups contrasted potential benefits with the high risks of misuse as a tool for political manipulation.

When respondents assessed hypothetical scenarios focusing on three AI use cases in cities, results revealed that air quality sensors in public spaces were highly acceptable (94–96%), while AI-enabled facial recognition and surveillance in residential areas were widely rejected. Acceptance declined as technologies became more intrusive, highlighting the tradeoff between potential benefits and risks, e.g. involving biometric identification and public surveillance enforced in residential areas or semi and private spaces. The collection and use of any biometric data were deemed unjustifiable to enable enhancement features for any system discussed in the workshops. Technologies relying on remote biometric identification to support public surveillance raise concerns about proportionality, necessity and due process, given the risks of discrimination, privacy violations and erosion of democratic values (Waelen, 2023; Fontes et al., 2022).

Moreover, across all AI use cases discussed, centralized monitoring using sensors or cameras in private or semi-private spaces was firmly rejected. The installation of sensors in private or semi-private spaces was considered a red line, regardless of potential benefits. As a close reference, literature on smart home technologies indicates that users expect contextual integrity, meaning that data governance must align with domestic norms such as invisibility to third parties, limited use for user-approved purposes and transparent control mechanisms (Zimmermann et al., 2019).

When discussing AI-generated content, some groups contrasted potential benefits with the high risks of misuse as a tool for political manipulation.

Another clear red line identified, concerned predictive systems for law enforcement. Indeed, the deployment of predictive algorithms in the public sector entails significant risks related to bias, transparency and legal accountability. The opacity of automated decisions can lead to unfair outcomes that may go unnoticed and proceed without effective recourse (Waller & Waller, 2020), potentially causing harm to individuals or communities when predictions are incorrect or discriminatory (see Meijer et al., 2021).

The risk of AI misuse with high impacts on individual rights, social cohesion and societal well-being was echoed in workshop discussions related to generative AI and synthetic content. Participants highlighted that fake or manipulated media constitute one of today's most pressing information security threats (Botha & Pieterse, 2020), with risks including the erosion of trust, social and political manipulation, misinformation, and psychological and reputational harm (Chesney & Citron, 2019).

Finally, incentive-based policies were generally viewed as undesirable. Incentives, like nudges, can influence human behavior and steer decision-making. The first ethical concern is voluntariness i.e. whether incentives entail coercion, manipulation or unfair pressure that undermines autonomy. Second, there is the risk of fostering an "incentive culture", where

voluntariness as well as civic and moral motivation are jeopardized or underestimated (Grant, 2006). Furthermore, financial incentives were seen as a red line, as participants perceived them as potentially discriminatory. Policies based on incentivizing citizens to accept technologies were considered undesirable.

On AI and Sustainable Development

Within the workshop, the connection between AI and sustainability was part of the focus. However, participants struggled to link AI with sustainability goals or the common good. Although the SDGs were introduced as a framework, participants often framed sustainability through lived experience (Fontes et al., 2024). This gap mirrors broader debates in literature, where the alignment between AI-driven innovation and sustainability remains largely aspirational (Vinuesa et al., 2020; Dwivedi et al., 2021). While AI shows potential to advance SDGs in areas such as climate modelling, energy efficiency and health diagnostics, it can also create tensions among different SDGs (Cowls et al., 2021).

Workshop discussions focused on envisioning how AI could help while still considering implications and trade-offs. For SDG 2 (Zero Hunger), it was suggested that AI could reduce food waste and improve distribution but feared benefits would be unevenly distributed, excluding those most in need. In healthcare (SDG 3), AI was viewed as potentially cost-saving yet opaque and privacy-intrusive, echoing debates on algorithmic accountability and the loss of human oversight (Jobin et al., 2019).

SDGs 1 (No Poverty) and 5 (Gender Equality) were discussed through a lens of discrimination and social inclusion; AI could foster equity through transparent pay and public services yet concerns about data sharing and privacy pointed to the commodification of personal information. Education (SDG 4) was especially contested, while AI was seen as potentially expanding access to education and personalized learning, it was emphasized that unequal access is a foundational issue stemming from digital divides and highlighted the importance of human interaction in learning (as shown in Holmes et al., 2022).

Discussions on clean energy (SDG 7) and sustainable cities (SDG 11) led to contesting current governance structures, which are often at the root of the problem, while technological transitions could entail significant costs. These concerns reflect ongoing debates in which gains in efficiency may be offset by increased energy and material demands associated with automation and digitalization, positioning AI as both an enabler and a disruptor of sustainability transitions (Vinuesa et al., 2020; Dwivedi et al., 2021).

When discussing the benefits of AI, participants struggled to link AI with sustainability goals or the common good.

Discussion and Remaining Challenges

The participatory methods developed in collaboration with the city of Amsterdam illustrate both the potential and the challenges of bridging representativeness and inclusiveness in participatory policymaking. By combining large-scale, representative consultation with targeted engagement of vulnerable or underrepresented groups, the approach aimed to enable value-mapping and capture diverse local perspectives on AI adoption.

The results indicate that establishing clear and transparent objectives for the participation process is key to ensuring that stakeholders' expectations correspond to the anticipated outcomes. Specifically, our findings show that a participatory process focused on deliberation and value-mapping can help identify critical issues related to data collection and use, factors influencing technology acceptance, and key safeguards that should be implemented.

Community empowerment emerges as a direct impact, highlighting the importance of AI literacy and addressing information asymmetries. Our experience working with vulnerable and often underrepresented groups demonstrates that their interest in contributing to policy and city-making can overcome informational and literacy barriers. The meaningfulness of participants' engagement grew as group deliberations progressed, demonstrating the importance of complementing surveys with initiatives that promote group discussions and deliberation. The added value to policy design lies in providing evidence-based, empirically grounded recommendations to fine-tune AI systems and tailor policies to local priorities, needs and preferences, thereby fostering acceptance, legitimacy, inclusion and trust.

At the same time, the project highlighted several persistent challenges that must be addressed to make

such processes more effective. Although the findings presented can be valuable for prioritizing needs and legitimizing the introduction of AI in specific city systems, it is important to examine how limitations encountered during implementation of the participatory process may have constrained the process' ability to fully bridge representativeness and inclusiveness, potentially undermining its meaningfulness. Addressing these limitations is essential for improving future participatory efforts. These include the following limitations:

Bridging representativeness and inclusiveness: Current practices require revaluation to truly balance representativeness and inclusiveness. Defining what constitutes a representative sample and determining which community groups to include often receives insufficient attention and encounters practical limitations on the ground. In this project, we suggested using population demographics and census data to ensure representativeness. Despite this, the 331 survey responses collected fell short of fully representing the city's population. At the local level, more attention should be paid to structural approaches that define and achieve representativeness for large-scale consultations.

Ensuring inclusiveness: Partnerships with local networks, such as public schools, libraries and civil society organizations, proved crucial for engaging underrepresented communities. Some populations, however, remain difficult to reach, and participation can skew toward those who are already interested or digitally literate. Mobilizing the public is further complicated by citizens' trust in governance, perceived relevance of AI issues and time constraints. While incentives were considered undesirable in this context, maintaining the voluntariness of participation, along with integrating it into broader cultural and social engagement strategies, appears key to meaningful involvement.

Supporting informed participation: Effective deliberation depends on participants' understanding of AI systems, ethical considerations and potential trade-offs. Digital, AI and rights literacy remain limited in many urban populations. Facilitators play a critical role in tailoring explanations to participants' needs, but achieving a baseline of understanding across diverse groups requires oriented design choices, prioritizing education and enabling iterative feedback prior and during participatory events (Fontes et al., 2024). Participatory processes should explicitly aim to assess and improve literacy, ensuring that informational asymmetries are minimized, namely during engagements.

Linking participation to policy design: The relationship between participatory processes and policy design is not inherently structural. This makes it challenging to extract value beyond the immediate communicative interface. Participation has intrinsic value in fostering communication between policymakers and affected communities, raising awareness, and promoting legitimacy and trust. However, its impact is significantly enhanced when participatory input is explicitly linked to specific policy agendas and integrated across the policy cycle. Moreover, different methodological approaches and scopes produce different types of inputs, which must be carefully calibrated to the objectives and timing of a particular policy.

Although this research focused on creating a space for deliberation on AI and its ethical implications for cities and communities, AI itself may offer opportunities to enhance participatory processes. Digital platforms are widely used by cities to scale and mainstream certain forms of public engagement. However, there remains significant potential to explore how AI could make participatory processes more effective, meaningful, accessible, inclusive and mainstream (Fontes & Mahajan, Forthcoming).

References

Batty, M. (2016). Complexity in city systems: Understanding, evolution, and design. In A planner's encounter with complexity (pp. 99-122). Routledge.

Botha, J., & Pieterse, H. (2020, March). Fake news and deepfakes: A dangerous threat for 21st century information security. In ICCWS 2020 15th International Conference on Cyber Warfare and Security. Academic Conferences and publishing limited (p. 57).

Caldarelli, G., Arcaute, E., Barthelemy, M., Batty, M., Gershenson, C., Helbing, D., ... & Fernández-Villacañas, J. L. (2023). The role of complexity for digital twins of cities. Nature Computational Science, 3(5), 374-381.

Chesney, R., & Citron, D. (2019). Deepfakes and the new disinformation war: The coming age of post-truth geopolitics. Foreign Aff., 98, 147.

City of Amsterdam. (2024). Amsterdam's Vision on AI. Available at: https://www.amsterdam.nl/innovatie/amsterdamse-visie-ai/.

City of Amsterdam. (2022). The Fairness Handbook. Available at: https://openresearch.amsterdam/en/page/87589/the-fairness-handbook.

Cowls, J., Tsamados, A., Taddeo, M., & Floridi, L. (2023). The AI gambit: leveraging artificial intelligence to combat climate change—opportunities, challenges, and recommendations. Ai & Society, 38(1), 283-307.

Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., ... & Williams, M. D. (2021). Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. International journal of information management, 57, 101994.

El Yassini, S., Van Eemeren, A., & Prinsen, L. (2022). Grip on algorithms Approach and tools for a responsible use of algorithms in Amsterdam. City of Amsterdam.

European Data Protection Board (EDPB). (2021). Guidelines 05/2021 on the Use of Facial Recognition Technology in the Area of Law Enforcement.

Florini, A. (2007). The right to know: transparency for an open world. Columbia University Press.

Fontes, C., Al Haddad, C., Antoniou, C., & Lütge, C. (2024). Assessing the collective challenges of AI: An applied framework focused on people centeredness and public participation. Urban Governance, 4(4), 245-254.

Fontes, C., Hohma, E., Corrigan, C. C., & Lütge, C. (2022). AI-powered public surveillance systems: why we (might) need them and how we want them. Technology in Society, 71, 102137.

Fontes, C., & Mahajan, S. (Forthcoming). AI Strengthening Democratic Engagement: Exploring Existing Digital Tools for Public Participation and Future Directions.

Fontes, C., & Perrone, C. (2021). Ethics of surveillance: harnessing the use of live facial recognition technologies in public spaces for law enforcement. Technical University of Munich, 1-11.

Fuchs, C. (2011). Towards an alternative concept of privacy. Journal of Information, Communication and Ethics in Society, 9(4), 220-237.

Grant, R. W. (2006). Ethics and incentives: a political approach. American Political Science Review, 100(1), 29-39.

Holmes, W., Porayska-Pomsta, K., Holstein, K., Sutherland, E., Baker, T., Shum, S. B., ... & Koedinger, K. R. (2022). Ethics of AI in education: Towards a community-wide framework. International Journal of Artificial Intelligence in Education, 32(3), 504-526.

Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. Nature Machine Intelligence, 1(9), 389–399.

Kitchin, R. (2016). The ethics of smart cities and urban science. Philosophical transactions of the royal society A: Mathematical, physical and engineering sciences, 374(2083), 20160115.

Meijer, A., Lorenz, L., & Wessels, M. (2021). Algorithmization of bureaucratic organizations: Using a practice lens to study how context shapes predictive policing systems. Public administration review, 81(5), 837-846.

OHCHR (2019). Guidelines for states on the effective implementation of the right to participate in public affairs. OHCHR/UN Publications.

Portugali, J. (2016, May). What makes cities complex?. In Complexity, cognition, urban planning and design: post-proceedings of the 2nd delft international conference (pp. 3-19). Cham: Springer International Publishing.

Schiff, D. S., Schiff, K. J., & Pierson, P. (2022). Assessing public value failure in government adoption of artificial intelligence. Public Administration, 100(3), 653-673.

UN Human Settlements Programme (UN-Habitat). (2025). International Guidelines on People-Centred Smart Cities. https://unhabitat.org/sites/default/files/2025/03/international_guidelines_on_people_centred_smart_cities_clean_4.3.2025.pdf.

Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., ... & Fuso Nerini, F. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. Nature communications, 11(1), 233.

Waelen, R. A. (2023). The struggle for recognition in the age of facial recognition technology. AI and Ethics, 3(1), 215-222.

Waller, M., & Waller, P. (2020). Why predictive algorithms are so risky for public sector bodies. Available at SSRN 3716166.

Wirtz, B. W., Weyerer, J. C., & Geyer, C. (2019). Artificial intelligence and the public sector—applications and challenges. International Journal of Public Administration, 42(7), 596-615.

Zimmermann, V., Gerber, P., Marky, K., Böck, L., & Kirchbuchner, F. (2019). Assessing users' privacy and security concerns of smart home technologies. i-com, 18(3), 197-216.