AI Governance in Smart Cities: Religious Freedom and Fundamental Rights

Daniela Tarantino
Department of Political and International Sciences, University of Genoa, Italy

daniela.tarantino@unige.it

Abstract: Artificial intelligence increasingly structures urban mobility, security, welfare, planning and participation. This article asks when AI-enabled smart-city governance may interfere with, or indirectly burden, freedom of religion or belief and which safeguards should govern that risk. It is principally a doctrinal and functional comparative legal study, supported by a bounded documentary examination of Hassan, Bridges, SyRI, the CJEU automated-decision cases and Amsterdam's crowd-monitoring infrastructure. The materials are selected to compare recurring mechanisms of identification, classification, allocation and procedural opacity across different jurisdictions and evidentiary settings. Their role is to provide factual anchoring by documenting systems, allegations as characterised by courts and legal findings; they do not support prevalence estimates or constitute empirical research in the social-scientific sense. The analysis identifies three forms of technologically mediated interference. Direct interference uses religion, belief or a belief-related site as a criterion; indirect interference arises through neutral proxies or optimisation rules that create particular disadvantage; ambient interference results when traceability, inaccessible participation and repeated administrative friction cumulatively chill or obstruct protected conduct. The article further identifies four mechanisms through which those burdens arise—identification, classification, allocation and procedural opacity—and tests them against legality, necessity, proportionality, equality, data-protection and effective-remedy standards. Its original contribution is a Religious-Pluralism Stress Test that translates those standards into lifecycle questions for design, procurement, deployment, monitoring and review. The framework requires authorities to define purpose and responsibility, examine sensitive inference and proxy effects, include less visible communities without compelling disclosure of belief, retain meaningful human override and provide timely challenge and suspension. The study concludes that the decisive legal inquiry is functional: whether an AI system, alone or cumulatively, materially changes the practical conditions of conviction, worship, collective organisation, non-belief or spatial access. The framework is exploratory and non-exhaustive. It offers a legally grounded starting point for further comparative and empirical research without treating documented risk as proof of widespread harm.

Keywords: Artificial intelligence governance, Smart cities, Freedom of religion or belief, Religious pluralism, Algorithmic accountability, Indirect discrimination

1. Introduction: Research Problem, Question and Contribution

Artificial intelligence is reshaping public decision-making. Here, a smart city means an urban environment in which connected sensors, data infrastructures and algorithmic systems support or structure public functions such as mobility, security, welfare, planning and participation (Kitchin, 2014; UN-Habitat, 2021). These systems are legally relevant not only when they issue a final decision, but also when data collection, classification, ranking or optimisation alters access to public space and services. Unlike a single administrative database, this infrastructure can operate continuously across streets, transport networks and service portals. Its distinctive risk is therefore cumulative: separate systems may jointly determine which persons and communities are seen, classified and heard.

Freedom of religion or belief is a demanding test because it protects inner conviction, public manifestation, collective organisation and access to the places and procedures through which religious or non-religious life is exercised. Urban AI may affect these dimensions by tracing attendance at places of worship, classifying gatherings, allocating routes or premises, inferring sensitive beliefs, or excluding less visible communities from digital participation. The problem is distinctive because many practices are spatial, collective and time-sensitive, while affiliation may be inferred from conduct rather than declared. A technically neutral intervention may consequently affect both privacy of belief and the practical ability to manifest it.

The research question is: under what conditions can AI-enabled urban governance constitute an interference with, or discriminatory burden on, freedom of religion or belief, and which safeguards should govern that risk? The article uses doctrinal and problem-oriented comparative analysis to connect Article 9 ECHR, the EU Charter, the GDPR, the AI Act and selected constitutional standards with concrete surveillance, scoring and sensor-based deployments.

The argument is that the central problem is not only error but regulation by design. Code, proxies, thresholds and optimisation objectives can change the probability, cost, visibility or contestability of exercising a right without an express prohibition (Hildebrandt, 2015; Rouvroy and Berns, 2013; Yeung, 2018). The article's

contribution is a three-tier account of direct, indirect and ambient interference and an operational lifecycle test. It is deliberately exploratory rather than exhaustive and seeks to open new lines of legal and empirical inquiry.

2. State of the Art and the Unresolved Legal gap

Critical smart-city scholarship shows that data-driven urbanism redistributes visibility, authority and resources, often privileging what is measurable and already represented in administrative data (Greenfield, 2013; Hollands, 2015; Kitchin, 2017; Vanolo, 2014). People-centred approaches accordingly require public purpose, inclusion and accountable data governance (OECD, 2023; UN-Habitat, 2021). Studies of municipal procurement and participation also warn that transparency may be lost behind proprietary systems and that citizen-centred language may conceal one-way data extraction (Brauneis and Goodman, 2018; Cardullo and Kitchin, 2019; Tarantino, 2026). The contested issue is thus not efficiency alone, but who defines urban objectives and whose practices are legible to the system.

Algorithmic-regulation scholarship adds that systems can rank, recommend, delay or exclude before any formal administrative act, while responsibility is fragmented among authorities, vendors and data intermediaries (Burrell, 2016; Calo and Citron, 2021; Cohen, 2019; Yeung and Lodge, 2019). Data-justice research further explains how neutral proxies, missing groups and historical datasets reproduce disadvantage (Barocas, Hardt and Narayanan, 2023; Selbst et al., 2019; Taylor, 2017). Opacity may be technical, organisational or legal, and it can undermine technological due process even where no single actor intends discrimination (Citron, 2008; Hildebrandt, 2020; Mittelstadt et al., 2016). Crawford (2021) further situates apparently local systems within wider relations of data, labour and institutional power.

Research specifically linking AI and freedom of religion or belief remains limited. Ashraf (2022) addresses online moderation and privacy, while Temperman (2023) maps broader AI risks to religious freedom. Work on governing urban religious diversity rarely integrates algorithmic public administration (Martínez-Ariño, 2020). Recent legal work on religion and digital transition confirms the need to connect general AI governance with the doctrinal structure of religious freedom rather than treating religion as merely another data category (Annicchino, 2025; Tarantino, 2026). What remains unsettled is the threshold at which a data-mediated inconvenience becomes a rights-relevant interference and how that threshold should shape procurement and review.

The unresolved gap is therefore a comparative legal framework for deciding when urban data practices become a cognisable burden on religious freedom and for translating that finding into operational duties. This article addresses that gap without claiming an exhaustive survey of technologies or cities.

3. Methodology: Doctrinal and Documentary Case Analysis

The study is principally a doctrinal and functional comparative legal analysis, supported by a bounded documentary examination of selected cases and deployments. It reconstructs standards from Article 9 ECHR, Articles 7, 8, 10, 21 and 47 of the EU Charter, the GDPR, the AI Act and selected constitutional jurisprudence. Comparison is functional: materials from different jurisdictions isolate recurring mechanisms of identification, classification, allocation and procedural opacity rather than rank national systems. The documentary materials comprise judgments, allegations as characterised by courts and official deployment records. Their role is evidentiary and instrumental: they establish the existence and operation of selected systems and test rights-relevant mechanisms, but do not constitute empirical research in the social-scientific sense or estimate the frequency of harm.

Primary materials were included if they met at least one of three criteria: they governed public-sector AI; addressed religious manifestation or equality; or officially documented a deployment capable of affecting urban rights. Materials supported only by promotional description or speculative commentary were excluded. Secondary literature was selected for direct relevance to smart-city governance, algorithmic regulation, discrimination, data protection or religion. The analysis covers material available to 31 July 2026 and uses neither interviews nor municipal microdata. Selection sought variation in jurisdictions, technologies, legal issues and evidentiary status. The sources were cross-checked across legislation, judgments, official datasets and specialised scholarship so that no conclusion rested on a single source type.

Four principal comparators serve distinct functions. Hassan supplies a surveillance baseline based on judicially recognised allegations concerning mosques and Muslim communities; Bridges concerns live facial recognition and ex ante legality and equality safeguards; SyRI, opaque public-sector scoring; and Amsterdam, ordinary crowd-monitoring sensors. SCHUFA and Dun & Bradstreet Austria supplement explanation and contestation.

Evidence is distinguished as adjudicated findings, allegations judicially recognised as plausible, and official infrastructure records. Amsterdam is retained as a limiting comparator: it establishes deployment, not an Article 9 violation, thus testing risk without presuming unlawfulness. Non-religion cases are used only where the transferable mechanism is specified. This bounded set supports comparison of legal mechanisms, not prevalence claims or statistical generalisation.

Table 1: Comparative materials and their analytical function

Comparator	Documented feature	Legal issue	Analytical function
Hassan (2015)	Alleged monitoring of mosques and Muslim communities	Religion-directed surveillance	Baseline for alleged stigma and chilling effects
Bridges (2020)	Live facial recognition in public space	Legality, equality and data protection	AI-specific ex ante safeguards
SyRI (2020)	Linked-data fraud-risk indication	Opacity and proportionality	Proxy classification comparator
Amsterdam sensors	Crowd-monitoring sensors and cameras	Purpose and contextual inference	Ex ante comparator; no established Article 9 violation

The analytical sequence is: identify the public function and responsible actors; identify the protected practice or belief-related interest; map data, proxies and system influence; assess legality, necessity, proportionality, equality and procedural fairness; and specify an effective remedy. A burden may be direct, indirect or cumulative, but it must be evidenced and attributable rather than merely imagined.

4. Freedom of Religion or Belief as a Multidimensional and Spatial Right

Article 9 ECHR protects thought, conscience and religion, including manifestation alone or in community, in public or private, through worship, teaching, practice and observance. The forum internum is absolute; restrictions on manifestation require lawfulness, a legitimate aim and necessity in a democratic society (Kokkinakis v. Greece, 1993). Its negative aspect also protects persons against coercion to disclose their beliefs (European Court of Human Rights, 2026). The right is individual, collective and institutional (Doe, 2011; Ferrari, 2015). These dimensions make it especially sensitive to systems that infer conviction or alter public manifestation without a conventional prohibition.

Religious communities must be able to organise and operate autonomously. Strasbourg jurisprudence treats unjustified interference with recognition or organisation as capable of violating Article 9 (Hasan and Chaush v. Bulgaria, 2000; Metropolitan Church of Bessarabia v. Moldova, 2001). Italian constitutional law likewise protects denominations, worship and religious institutions, while secularism operates as a guarantee of pluralism (Constitutional Court, judgment no. 203/1989).

The spatial dimension is crucial: worship, processions, festivals and charitable activities depend on streets, transport, premises and planning procedures. The Italian Constitutional Court has required fair access to worship facilities, including for communities without agreements under Article 8(3) (judgments nos. 195/1993, 63/2016 and 254/2019). Urban technologies may therefore burden a right without formally banning it.

A functional concept of interference follows. Direct interference uses religion or belief as a criterion. Indirect interference arises when a neutral proxy or optimisation rule creates a particular disadvantage. Ambient interference results when traceability, sensor density, inaccessible participation and repeated administrative friction cumulatively chill or obstruct protected conduct. Each category still requires evidence, attribution and proportionality analysis.

5. Deployment-Based Analysis: From Urban Data to Rights Effects

5.1 Surveillance, Identification and the Chilling of Religious Practice

Hassan v. City of New York makes the surveillance pathway concrete. At the pleading stage, the Third Circuit held that allegations of mapping, photographing and monitoring mosques, organisations and Muslim-owned businesses plausibly described stigma, altered behaviour and institutional harm. The judgment did not make final findings after trial, but it provides a judicially recognised account of how belief-directed surveillance may burden worship and collective life.

Bridges concerned live automated facial recognition in public space. The Court of Appeal found inadequate legal constraints on watchlists and locations, deficiencies in the data-protection impact assessment and failure properly to discharge the public-sector equality duty. Although the case did not concern religion, it shows why automated identification around gatherings or places of worship requires foreseeable limits, equality analysis and meaningful oversight. Its relevance is strongest at the governance stage: watchlist criteria, deployment locations and equality assessment must be defined before scanning begins, because later review cannot fully reverse the deterrent effect of identification in public space.

Amsterdam's official open-data service documents crowd-monitoring infrastructure comprising 3D and Wi-Fi sensors, cameras and beacons. It is an ex ante infrastructure comparator, not evidence of an established Article 9 violation. Its relevance lies in the possibility that ordinary movement data, when linked with time, location or event information, may reveal attendance at religiously significant places and enable secondary uses beyond the original purpose. Purpose limitation is therefore central. Data collected to manage flows should not be repurposed to profile communities, predict religious attendance or support enforcement without a new legal basis, necessity assessment and safeguards against contextual inference (Edwards, 2016).

5.2 Risk Scoring, Service Allocation and Proxy Discrimination

The second pathway is automated classification. In SyRI, linked administrative data produced fraud-risk indications; the District Court of The Hague held that the legal framework lacked sufficient transparency and struck an unfair balance under Article 8 ECHR. The case illustrates why opaque proxies and data linkages can impede proportionality review even without an expressly religious variable. For religious pluralism, the concern is not that SyRI used religion, but that linked data and undisclosed indicators illustrate how residence, household composition, language or service use may operate as proxies. A claimant cannot demonstrate indirect disadvantage if the variables and decisional weight remain inaccessible.

A religious-freedom impact may arise if scores influence permits, premises, welfare or services supporting collective life. SCHUFA and Dun & Bradstreet Austria reinforce the procedural point: where automated scoring significantly shapes a decision, affected persons need sufficiently specific information to understand and challenge it. Formal human involvement is inadequate if officials routinely defer to an unexplained output.

5.3 Mobility, Crowd Management, Planning and Participation

The third pathway concerns allocation of urban space. Traffic optimisation, crowd management, event-risk tools and digital planning platforms may determine routes, opening times, premises and participation opportunities. A neutral system can disadvantage practices occurring at fixed times, requiring collective movement or lacking a large digital footprint. The appropriate counterfactual is whether substantially the same access and opportunity would have existed under a transparent, human-administered process. Authorities should examine effects by site, practice, time and community rather than relying only on aggregate accuracy or overall traffic efficiency.

Across these deployments, four mechanisms recur: identification links persons to belief-related places or networks; classification assigns risk or priority; allocation distributes services and space; and procedural opacity impedes notice, reasons and challenge. These mechanisms transform plausible scenarios into legally testable questions. The doctrinal chain is consequently traceable: a documented system feature affects a protected interest through one of these mechanisms; the resulting burden is then tested under legality, proportionality, equality, data protection and effective-remedy standards. This chain distinguishes evidence-based risk analysis from conjecture.

6. Legal Assessment: Legality, Equality, Data Protection and AI Regulation

Public-sector AI is rights-relevant infrastructure. Legality requires an accessible and foreseeable basis defining purpose, scope, data, responsible actors and limits, especially for surveillance or significant automated decisions. Procurement documents cannot replace legislation where serious interference is involved, and outsourcing does not remove public responsibility.

Proportionality requires a concrete objective, evidence of suitability and necessity, consideration of less intrusive alternatives and balancing of benefits against the severity and distribution of burdens. The assessment must cover belief-related inference, worship, collective autonomy, spatial access and cumulative effects across connected systems. Where benefits are uncertain or the burden concerns the forum internum, a particularly demanding justification is required. Authorities should document why less intrusive data, lower retention, narrower geography or non-automated alternatives are insufficient.

Equality review must address both explicit religious classifications and neutral proxies producing particular disadvantage. The applicable legal basis must be identified precisely: Article 21 of the Charter applies within EU-law scope; Article 14 ECHR operates with another Convention right; and Protocol No. 12 applies where ratified. Religion-specific and intersectional effects should be tested before and after deployment. Testing should include false positives, false negatives, underrepresentation and the distribution of administrative delay. Formal parity is insufficient where the system repeatedly imposes heavier practical costs on a minority practice.

Personal data revealing religious or philosophical beliefs fall within Article 9(1) GDPR. Location, event attendance, language or networks may reveal belief even without a field labelled religion. Controllers must assess purpose limitation, minimisation, legal basis, retention, security and linkage; Article 35 requires an impact assessment for likely high-risk processing, while Article 22 governs qualifying solely automated decisions. The assessment should cover not only collected fields but also reasonable inferences after linkage, especially where movement data are persistent, interoperable or shared across public functions (Edwards, 2016).

The AI Act adds lifecycle obligations. Article 5 prohibits specified biometric categorisation that uses biometric data to deduce or infer religious or philosophical beliefs, subject to its express exceptions. For high-risk systems, Articles 9–15 impose risk-management, data-governance, documentation, logging, transparency, human-oversight, accuracy and robustness duties, while Article 27 requires a fundamental-rights impact assessment for specified deployers (Regulation (EU) 2024/1689). The application timetable was amended in July 2026: Regulation (EU) 2026/1744 deferred the Annex III high-risk rules to 2 December 2027 and the Annex I rules to 2 August 2028. These dates affect the AI Act duties, not the independently applicable ECHR, Charter, GDPR or constitutional standards. The Act's legislative impact assessment treated public-sector uses and fundamental-rights risks as central to the regulatory design (European Commission, 2021).

Council of Europe instruments likewise require legality, transparency, accountability and remedies throughout design, procurement, deployment and review (Council of Europe, 2020; 2024). A final human signature cannot cure opacity where officials lack the information or authority to question the system.

7. An Operational Governance Model: The Religious-Pluralism Stress Test

The operational contribution is a Religious-Pluralism Stress Test embedded in fundamental-rights and data-protection assessments before procurement and deployment. It does not privilege religion; it uses a multidimensional right to reveal weaknesses also affecting groups whose practices are collective, spatial, episodic or poorly represented in data. The authority should publish a concise decision record stating the affected functions, evidence considered, accommodations adopted and residual risks accepted. A system should not proceed where its purpose is vague, sensitive inference is unnecessary, meaningful explanation is impossible or timely remedies cannot be guaranteed.

Table 2: Lifecycle governance matrix for freedom of religion or belief

Stage	Required safeguard	Religion-or-belief question
Purpose and data	Define authority, necessity, sources, proxies, retention and linkage	Could the system reveal belief or affect worship, autonomy or spatial access?
Design and validation	Test thresholds, missing groups, errors and less intrusive alternatives	Are minority, informal, migrant and non-religious communities represented fairly?
Procurement and deployment	Allocate responsibility; secure documentation, notice, reasons and human override	Can the authority explain, correct, suspend and terminate the system?
Participation	Use accessible, non-tokenistic consultation without compelled disclosure	Can less visible communities influence criteria and accommodations?
Monitoring and remedy	Audit differential and cumulative effects; provide rapid challenge and exit	Can harm be corrected before a time-sensitive practice loses practical value?

Participation must include minority, migrant, informal, non-institutionalised and non-religious actors, not only established denominations. It should occur while objectives remain open, avoid compelled disclosure of belief, and link feedback to documented design changes rather than symbolic consultation.

Procurement contracts should specify purpose, prohibited secondary uses, data sources, audit and explanation rights, retention, deletion, override, remedies and exit. Authorities must retain the capacity to suspend; trade secrecy cannot defeat accountability or judicial review.

Assessment must be iterative as data, thresholds and urban behaviour change. Monitoring should examine complaints, differential errors, proxy drift and cumulative outcomes. Material updates, new datasets or changed purposes should trigger reassessment; the response may be accommodation, redesign, suspension or non-deployment.

8. Discussion: When Does Urban AI Transform Religious Freedom?

The materials support three forms of interference. Direct interference deliberately uses religion, a religious site or a belief-related category to target, rank or exclude. Indirect interference arises when neutral data or optimisation rules disproportionately burden a practice or community. Ambient interference occurs when traceability, inaccessible participation and repeated low-level friction cumulatively chill or obstruct protected conduct. Hassan informs the first category, SyRI and data-justice scholarship the second, and sensor infrastructures the third. These are evidentiary categories, not automatic findings of unlawfulness: direct use of belief demands the clearest authority and justification, while indirect and ambient burdens require contextual and cumulative proof.

The typology develops the claim of regulation by design: a system regulates when its objectives, data, categories and institutional use systematically alter the probability, cost, visibility or contestability of exercising a right, even without a binding command. Public-law analysis should address that functional effect while preserving ordinary requirements of proof and attribution. The article's novelty lies in reorganising established principles around mechanisms and lifecycle decisions, so that municipalities, auditors, courts and affected communities can apply them.

Procedural safeguards are substantive. Bridges links legality and equality assessment to lawfulness; SyRI shows that opacity obstructs proportionality review; and the CJEU scoring cases connect explanation with effective challenge. Because fixed-time worship, festivals, processions or consultations may pass before ordinary review concludes, an effective remedy must identify who can override the system, what information is disclosed and whether interim relief is available.

This legal-comparative study is exploratory and non-exhaustive. Its documentary materials establish selected deployments and legally plausible mechanisms, not the incidence of harm across cities. Vendor documentation is uneven, and several comparators concern privacy or equality rather than Article 9 directly. The typology therefore generates propositions for later testing against procurement files, impact assessments, complaints, audit logs, interviews and comparative municipal data. Its purpose is to structure further research, not infer widespread harm.

9. Conclusions: Answering the Research Question

AI-enabled urban governance can burden freedom of religion or belief when it identifies belief-related conduct or locations, classifies through explicit or proxy variables, allocates opportunities unequally, or makes public power difficult to understand and contest. The decisive inquiry is functional: whether the system, alone or cumulatively, materially changes the practical conditions of conviction, worship, collective organisation, public manifestation, non-belief or spatial access.

The comparators anchor each mechanism without conflating evidentiary status: Hassan supplies judicially recognised allegations of religion-directed surveillance; Bridges, ex ante safeguards for automated identification; SyRI, opacity in scoring; the CJEU cases, explanation and challenge; and Amsterdam, the need to assess sensor purpose and linkage before reuse. They document mechanisms, not prevalence.

Lifecycle accountability is therefore required. Legality, necessity, proportionality, non-discrimination, minimisation, explanation, participation, human oversight, audit and effective remedies must govern design and procurement as well as final decisions. The Stress Test turns those duties into questions about protected practices, proxies, missing groups, cumulative effects and time-sensitive remedies.

Freedom of religion or belief tests whether a smart city can govern pluralism without converting difference into risk, invisibility or administrative friction. The framework is a legally grounded and non-exhaustive starting point for further comparative and empirical research.

Ethics Declaration: No ethics approval was required because this doctrinal and comparative legal study involved no human participants, personal datasets or experiments.

AI Declaration: AI tools supported English translation, language editing, structural revision, source checking and formal compliance review. The author determined the research question, legal method and argument and remains responsible for verifying the sources, analysis and conclusions.

References

Annicchino, P. (2025) L'interazione tra diritto e religione nella transizione digitale. Turin: Giappichelli.

Ashraf, C. (2022) "Exploring the impacts of artificial intelligence on freedom of religion or belief online", The International Journal of Human Rights, 26(5), pp. 757–791. doi: 10.1080/13642987.2021.1968376.

Barocas, S., Hardt, M. and Narayanan, A. (2023) Fairness and Machine Learning: Limitations and Opportunities. Cambridge, MA: MIT Press.

Brauneis, R. and Goodman, E.P. (2018) "Algorithmic transparency for the smart city", Yale Journal of Law & Technology, 20, pp. 103-176.

Burrell, J. (2016) "How the machine 'thinks': Understanding opacity in machine learning algorithms", Big Data & Society, 3(1), pp. 1-12.

Calo, R. and Citron, D.K. (2021) "The automated administrative state: A crisis of legitimacy", Emory Law Journal, 70, pp. 797-846.

Cardullo, P. and Kitchin, R. (2019) "Being a 'citizen' in the smart city: up and down the scaffold of smart citizen participation in Dublin, Ireland", GeoJournal, 84(1), pp. 1-13.

Citron, D.K. (2008) "Technological due process", Washington University Law Review, 85(6), pp. 1249-1313.

City of Amsterdam (n.d.) Crowd Monitoring Systeem Amsterdam (CMSA): sensor locations and context. Amsterdam Datapunt API Documentation, available at: https://api.data.amsterdam.nl/v1/docs/datasets/cmsa.html (accessed 15 July 2026).

CJEU, Dun & Bradstreet Austria, Case C-203/22, judgment of 27 February 2025.

CJEU, SCHUFA Holding, Case C-634/21, judgment of 7 December 2023.

Cohen, J.E. (2019) Between Truth and Power: The Legal Constructions of Informational Capitalism. Oxford: Oxford University Press.

Council of Europe (2020) Recommendation CM/Rec(2020)1 of the Committee of Ministers to member States on the human rights impacts of algorithmic systems. Strasbourg: Council of Europe.

Council of Europe (2024) Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law. Strasbourg: Council of Europe.

Court of Appeal of England and Wales, R (Bridges) v. Chief Constable of South Wales Police, [2020] EWCA Civ 1058, judgment of 11 August 2020.

Crawford, K. (2021) Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence. New Haven: Yale University Press.

District Court of The Hague, NJCM c.s. v. State of the Netherlands (SyRI), ECLI:NL:RBDHA:2020:865, judgment of 5 February 2020.

Doe, N. (2011) Law and Religion in Europe: A Comparative Introduction. Oxford: Oxford University Press.

ECtHR, Hasan and Chaush v. Bulgaria, App no 30985/96, judgment of 26 October 2000.

ECtHR, Kokkinakis v. Greece, App no 14307/88, judgment of 25 May 1993.

ECtHR, Metropolitan Church of Bessarabia and Others v. Moldova, App no 45701/99, judgment of 13 December 2001.

Edwards, L. (2016) "Privacy, security and data protection in smart cities", European Data Protection Law Review, 2(1), pp. 28-58.

European Commission (2021) Impact Assessment Accompanying the Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence, SWD(2021) 84 final.

European Convention on Human Rights, Article 9.

European Court of Human Rights (2026) Article 9 – Freedom of thought, conscience and religion: Case-Law Guide. Strasbourg: Council of Europe (accessed 15 July 2026).

European Union (2012) Charter of Fundamental Rights of the European Union, Articles 7, 8, 10, 21 and 47.

European Union (2016) Regulation (EU) 2016/679 (General Data Protection Regulation), especially Articles 9, 22 and 35.

European Union (2024) Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), especially Articles 5, 9–15 and 27.

European Union (2026) Regulation (EU) 2026/1744 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence.

Ferrari, S. (ed.) (2015) Routledge Handbook of Law and Religion. London: Routledge.

Greenfield, A. (2013) Against the Smart City. New York: Do Projects.

Hildebrandt, M. (2015) Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology. Cheltenham: Edward Elgar.

Hildebrandt, M. (2020) Law for Computer Scientists and Other Folk. Oxford: Oxford University Press.

Hollands, R.G. (2015) "Critical interventions into the corporate smart city", Cambridge Journal of Regions, Economy and Society, 8(1), pp. 61-77.
Italian Constitutional Court, judgment no. 195/1993.
Italian Constitutional Court, judgment no. 203/1989.
Italian Constitutional Court, judgment no. 254/2019.
Italian Constitutional Court, judgment no. 63/2016.
Kitchin, R. (2014) "The real-time city? Big data and smart urbanism", GeoJournal, 79(1), pp. 1-14.
Kitchin, R. (2017) "Data-driven urbanism", in Kitchin, R., Lauriault, T.P. and McArdle, G. (eds) Data and the City. London: Routledge.
Martínez-Ariño, J. (ed.) (2020) Governing Religious Diversity in Cities: Critical Perspectives. London and New York: Routledge.
Mittelstadt, B.D., Allo, P., Taddeo, M., Wachter, S. and Floridi, L. (2016) "The ethics of algorithms: Mapping the debate", Big Data & Society, 3(2), pp. 1-21.
OECD (2023) Smart City Data Governance: Challenges and the Way Forward. Paris: OECD Publishing.
Rouvroy, A. and Berns, T. (2013) "Algorithmic governmentality and prospects of emancipation: Disparateness as a precondition for individuation through relationships?", Réseaux, 177(1), pp. 163-196.
Selbst, A.D., Boyd, D., Friedler, S.A., Venkatasubramanian, S. and Vertesi, J. (2019) "Fairness and abstraction in sociotechnical systems", Proceedings of the Conference on Fairness, Accountability, and Transparency, pp. 59-68.
Tarantino, D. (2026) "Smart Cities, Smart Rights? Some Reflections on Religious Pluralism, Urban Technologies and New Frontiers of Legal Protection in Europe", Consulta Online, 2026/II, pp. 616-632.
Taylor, L. (2017) "What is data justice? The case for connecting digital rights and freedoms globally", Big Data & Society, 4(2), pp. 1-14.
Temperman, J. (2023) "Artificial Intelligence and Religious Freedom", in Quintavalla, A. and Temperman, J. (eds), Artificial Intelligence and Human Rights. Oxford: Oxford University Press, pp. 61–75. doi: 10.1093/law/9780192882486.003.0004.
UN-Habitat (2021) Centering People in Smart Cities: A Playbook for Local and Regional Governments. Nairobi: UN-Habitat.
United States Court of Appeals for the Third Circuit, Hassan v. City of New York, 804 F.3d 277 (3d Cir. 2015), judgment of 13 October 2015.
Vanolo, A. (2014) "Smartmentality: The smart city as disciplinary strategy", Urban Studies, 51(5), pp. 883-898.
Yeung, K. (2018) "Algorithmic regulation: A critical interrogation", Regulation & Governance, 12(4), pp. 505-523.
Yeung, K. and Lodge, M. (eds) (2019) Algorithmic Regulation. Oxford: Oxford University Press.