

POLITECNICO
MILANO 1863

AI Governance in Urban Planning: A Transatlantic Perspective

School of Architecture, Urban Planning, Construction Engineering
Masters of Science in Urban Planning and Policy Design

Author: Michael Meyers Moudry

Student ID: 10970016
Advisor: Davide Ponzini
Academic Years: 2025-2026

Abstract

Urban planning is at a turning point due to advanced artificial intelligence (AI) tools that promise to enhance data analysis and automate routine design and documentation tasks (Sanchez, 2025). At the same time, this rapid AI adoption has outpaced existing governance processes, raising ethical and regulatory concerns (Wasserman & Flaxman, 2023). Planning organizations have begun issuing principles, guidelines, and training programs to help practitioners adapt. This thesis examines how leading planning bodies on in North America and Europe are framing AI governance in urban planning. Specifically, it compares the American Planning Association and the Association of Collegiate Schools of Planning (ACSP) in the U.S. and Canada with the Association of European Schools of Planning (AESOP) and the European Council of Spatial Planners (ECTP-CEU) in Europe. Using a qualitative comparative framework, the study analyzes published policies, statements, and educational initiatives to map out institutional strategies. The paper follows this outline: after a conceptual framework and literature review, it describes methods, analyzes U.S. and European cases separately, and concludes with discussion.

Key-words: artificial intelligence, urban planning, AI governance, transatlantic comparison, planning institutions, smart cities.

Abstract in italiano

La pianificazione urbanistica si trova a un punto di svolta a causa degli strumenti avanzati di intelligenza artificiale (IA) che promettono di migliorare l'analisi dei dati e automatizzare le attività di routine legate alla progettazione e alla documentazione (Sanchez, 2025). Allo stesso tempo, la rapida adozione dell'IA ha superato i processi di governance esistenti, sollevando preoccupazioni di natura etica e regolatoria (Wasserman & Flaxman, 2023). Le organizzazioni di pianificazione hanno iniziato a elaborare principi, linee guida e programmi di formazione per aiutare i professionisti ad adattarsi. Questa tesi esamina come i principali enti di pianificazione negli Sati Uniti e in Europa stiano inquadrando la governance dell'IA nella pianificazione urbanistica. In particolare, confronta l'American Planning Association e l'Association of Collegiate Schools of Planning (ACSP) negli Stati Uniti e Canada con l'Association of European Schools of Planning (AESOP) e l'European Council of Spatial Planners (ECTP-CEU) in Europa. Attraverso un framework comparativo qualitativo, lo studio analizza politiche pubblicate, dichiarazioni e iniziative educative per mappare le strategie istituzionali. Il lavoro segue questa struttura originale: dopo un framework concettuale e una revisione della letteratura, descrive i metodi, analizza separatamente i casi statunitensi ed europei, e si conclude con una discussione.

Parole chiave: intelligenza artificiale, pianificazione urbanistica, governance dell’IA, confronto transatlantico, istituzioni di pianificazione, smart cities.

Contents

Abstract ..... **i**

Abstract in italiano ..... **ii**

Contents ..... **iii**

Introduction **1**

1 Conceptual Framework and Comparative Approach **3**

1.1. United States: APA (American Planning Association) and ACSP (Association of Collegiate Schools of Planning) 4

1.2. Europe: EU Institutions and Planning Networks 6

1.3. International Planning Institutions 8

2 AI Policy Application in the Urban Planning Field **10**

2.1. Policy and Regulatory Frameworks 10

2.1.1. United States Policy Landscape 10

2.1.2. European Union Landscape 12

2.1.3. International Frameworks and Global Standards 12

2.2. AI Applications in Urban Planning 13

2.2.1. Spatial Data Analysis and Modelling 13

2.2.2. Generative Design and Optimization 14

2.2.3. Community Engagement and Decision Support 14

2.2.4. Administrative Efficiency 15

2.3. Ethics, Equity, and the Public Interest in AI 16

2.3.1. Bias and Fairness 16

2.3.2. Transparency and Accountability 18

2.3.3. Privacy and Data Governance 19

2.3.4. Equity and Public Inclusion 20

2.4. Planning Association Publications and Initiatives 20

3 Research Design and Methodology **23**

3.1. Research Scope and Case Selection 23

3.2. Data Collection Methods 24

3.3. Uses of Digital Tools and Artificial Intelligence Assistance 24

4 Responses to AI in the United States Planning Context **27**

4.1. American Planning Association (APA) Initiatives 27

4.1.1. Policy Guidance and Foresight Efforts 27

4.1.2. Professional Development and Ethical Stances 28
4.1.3. Advocacy and External Collaboration 30
4.2. Association of Collegiate Schools of Planning (ACSP) Initiatives 31
4.2.1. Research Promotion 31
4.2.2. Curriculum and Pedagogy 32
4.2.3. Collaboration and Discourse 34
5 Responses to AI in the European Planning Context 36
5.1. European Commission and EU-Level Initiatives 36
5.1.1. Regulatory Framework: The EU AI Act 37
5.1.2. Strategic Initiatives: AI for Urban Innovation 38
5.2. European Council of Spatial Planners (ECTP-CEU) Initiatives 39
5.2.1. Policy Guidance and Professional Standards 40
5.2.2. Advocacy and Cross-Border Coordination 40
5.3. Association of European Schools of Planning (AESOP) Initiatives 41
5.3.1. Research Promotion and Scholarly Discourse 41
5.3.2. Curriculum and Pedagogy 42
5.3.3. Collaboration and Dialogue with Practice 43
6 UN-Habitat and Global Perspectives 45
6.1. People-Centred Smart Cities Framework 45
6.1.1. Principles and the Rights-Based Approach 45
6.1.2. International Guidelines and Co-Creation 46
6.1.3. Influence on City and National Policy 47
6.2. Global Assessments, Capacity-Building and Pilot Projects 48
6.2.1. Global Assessment of Responsible AI in Cities 48
6.2.2. Toolkits and Capacity-Building for Local Officials 49
6.2.3. Pilot Projects and Knowledge Exchange 50
7 Comparative Analysis and Discussion 52
7.1. Institutional Role and Approach: Professional Associations vs. Public Bodies 52
7.2. Ethical and Policy Alignment: Similarities and Differences 54
7.3. Influence on Practice and Identification of Gaps 56
7.4. Transatlantic Lessons and Opportunities for Mutual Learning 58
Conclusion 61
Bibliography 64
List of Figures 71
Acknowledgements 72

Introduction

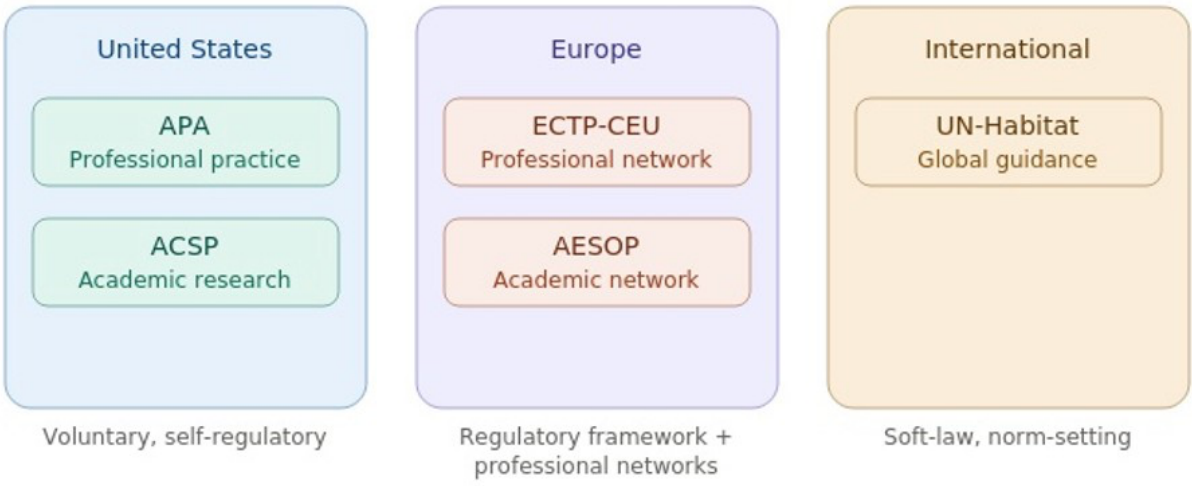
Urban planning as a practice is undergoing rapid technological transformation due to the rise of advanced artificial intelligence (AI). Recent advances in machine learning and data analytics offer planners powerful new tools for data-driven modeling, generative design, and predictive forecasting (Sanchez et al., 2023). Meanwhile, generative AI systems such as ChatGPT can automate tasks that were once manually intensive – for example, drafting zoning memos or summarizing technical reports in seconds (Sanchez, 2025). These capabilities suggest both significant efficiency gains and fundamental shifts in professional roles. At the same time, experts note that the unprecedented speed and scope of AI development have outstripped traditional oversight, with AI's impact in planning already exceeding existing educational and regulatory frameworks (Wasserman & Flaxman, 2023). This context raises an urgent governance question about how planning institutions can guide the use of AI so that it serves equity, transparency, and the public interest rather than undermining them.

Professional and academic planning organizations are central actors in answering this question. In the United States, for example, the American Planning Association (APA) and the Association of Collegiate Schools of Planning (ACSP) play key roles— APA disseminates best practices, issues policy statements, and provides training on emerging tools, while ACSP influences how future planners are educated and what research is prioritized. Across Europe, analogous bodies include the Association of European Schools of Planning (AESOP) and the European Council of Spatial Planners (ECTP-CEU). Unlike the U.S., however, European planning institutions operate in a milieu of supranational regulation. This can be seen by observing the European Union, which has negotiated the world's first comprehensive AI law (the AI Act), which imposes risk-based requirements on "high-risk" AI applications and reflects a strongly human-centric approach (European Commission, 2019, 2024). In short, the U.S. context tends to favor voluntary, principle-based guidance from professional bodies, whereas the European context is increasingly shaped by binding regulations and transnational initiatives. A transatlantic comparison of planning associations can thus reveal how these different institutional and cultural environments influence the governance of AI in cities.

This thesis will develop a comparative framework to analyze these issues at the institutional level. It asks: How are planning associations in the U.S. and EU governing, guiding, or adapting to the rise of AI in urban planning? The scope of inquiry includes the four association bodies (APA/ACSP vs. AESOP/ECTP-CEU). UN-

Habitat is also examined, though as the global normative backdrop against which the transatlantic comparison takes place rather than as a third parallel case. The study will systematically review policy documents, professional guidelines, educational programs, and public statements to identify the governance mechanisms, policy principles, and capacity-building efforts each institution promotes. The focus is on general strategies – such as codes of ethics, upskilling programs, or research agendas – rather than on any specific AI tool or city. By examining both common themes (for example, an emphasis on fairness and accountability in AI) and points of divergence (for example, legally mandated rules versus self-imposed guidelines), the research will remain exploratory and refrain from drawing conclusions prematurely.

1 Conceptual Framework and Comparative Approach



"Figure 1: Comparative Framework of Planning Institutions by Governance Context"

This thesis is structured around a comparative framework. The primary axis of comparison is transatlantic, between planning institutions in the United States and Europe. UN-Habitat is examined alongside them as the global normative context that shapes both:

1. United States (Professional and Academic bodies): The American Planning Association (APA) and the Association of Collegiate Schools of Planning (ACSP).

2. Europe (Policy and Guiding bodies): The European Union's governing framework (exemplified by the European Commission's initiatives and relevant EU legislation), the European Council of Spatial Planners (ECTP-CEU) and AESOP as professional/academic networks,

3. The United Nations Human Settlements Programme (UN-Habitat) as a global agency influencing planning norms.

The framework captures a range of governance mechanisms, from voluntary professional guidance and educational initiatives (APA and ACSP) to formal regulatory action at the EU level, with UN-Habitat's soft-law guidelines providing the global normative backdrop. This mix allows the research to identify how AI governance in planning is being approached from multiple angles and scales.

1.1. United States: APA (American Planning Association) and ACSP (Association of Collegiate Schools of Planning)



"Figure 2: APA Logo"

The APA is the foremost professional organization for urban planners in the United States, with a mission to support planning practice and uphold ethical standards in the profession. As such, it wields significant influence through its resources, publications, and advocacy. In recent years, the APA has begun to actively engage with the topic of AI in planning. Through its Foresight Initiative, for example, APA has identified AI as an emerging trend that could reshape planning practice. APA convened experts and task forces to study the implications of AI, signaling to its membership that this is a critical issue for the future of planning; this led to APA's Research KnowledgeBase now including a curated collection on "Artificial Intelligence and Planning," which compiles case studies, policy guidance, and background resources on how local and regional agencies can use AI (American Planning Association, 2024). By cataloging these resources, APA is effectively providing a knowledge hub for planners curious or cautious about AI's role. Additionally, APA offers training programs – "Upskilling Planners" modules for example – that help practitioners develop new competencies, including digital and data skills relevant to AI.

On the policy and ethics front, the APA has not issued a formal "AI policy guide" to date, but it has integrated AI considerations into its commentary on technology and planning. Key questions being addressed within APA circles include the opportunities and risks that AI tools present for planners, as well as how planners should address issues of algorithmic bias and data ethics. APA's existing ethical principles (as outlined in the AICP Code of Ethics) mandate that planners "seek social justice" and serve the public interest (American Institute of Certified Planners, 2021), and the APA is beginning to interpret these principles in the context of AI. While the APA cannot regulate technology use, its influence is indirect yet powerful, exercised through the publication of best practice reports, the hosting of conference sessions on AI, encouragement for members to approach AI thoughtfully, and advocacy at the

national level for responsible technology use. At this conceptual stage, APA represents a self-regulatory approach – it provides guidance and sets a tone, but compliance is voluntary, shaped by professional norms and peer learning rather than law.



"Figure 3: ACSP Logo"

If the APA embodies the voice of practitioners, the ACSP represents the academic side of the planning field. As a consortium of university planning programs and educators in the U.S. (and Canada), ACSP shapes what future planners learn and what research gets prioritized. Because academia is often where new ideas and critical perspectives originate, its response to the rise of AI matters. A notable response has been their establishment of a Task Force on Urban Planning and AI in 2024, which recognizes the need to systematically examine AI's burgeoning role in planning scholarship and education (University of Florida Department of URP, 2024). Creating this task force – initiated by planning scholars and approved by ACSP's governing board – indicates a concerted effort to integrate AI into the core agenda of planning academia. The Task Force's mandate includes assessing how AI can be incorporated into planning curricula, what research gaps exist, and how to guide faculty and students in using AI tools responsibly. In parallel, academic journals and conferences under ACSP's umbrella have begun highlighting AI. Publications such as the Journal of Planning Education and Research (JPER) (an ACSP-affiliated journal) announced a special issue call aiming to "foster a comprehensive understanding of the opportunities, challenges, and implications of AI in the field of urban planning" (JPER, 2024) – an indication of how quickly scholarly interest in the topic has grown.

Early surveys and commentaries reveal a mix of excitement and caution in academia, as shown by Ye et al. (2024), which found that among U.S. planning faculty there is high awareness of AI tools like ChatGPT, coupled with reservations about over-reliance and concerns about maintaining critical thinking skills in an AI-influenced pedagogy. What these findings show is that educators are trying to figure out how to use AI's advantages — for data analysis, as a teaching aid — without weakening the analytical rigor and ethical awareness that planning programs are meant to build. ACSP's role is to steer the scholarly conversation on AI and prepare the next generation of planners, which means asking hard questions about algorithmic bias, equity, and whether AI tools actually work as promised. It also means providing spaces — conference sessions, workshops, the task force — where planners and AI researchers can talk to each other. ACSP is, at its core, a knowledge and capacity-

building institution: its work is in research and education, pushing the planning field to engage critically with AI rather than absorb it passively.

ACSP and APA feed off each other in useful ways — ideas debated in academia filter into practice through APA's channels, and practitioners' needs can spur academic research. With both organizations prioritizing AI in their respective domains, the U.S. planning community is mobilizing on multiple fronts to address AI governance, albeit in a non-coercive, exploratory fashion.

1.2. Europe: EU Institutions and Planning Networks

In the European context, the governance of AI in urban planning must be understood as a multi-level effort. Unlike the U.S., which has a single dominant professional body (APA) and a mostly voluntary approach, Europe's institutional environment is fragmented across EU regulatory bodies, transnational planning associations, and global organizations that each influence norms in different ways. This paper will focus on the European Commission and EU Regulatory Initiatives, planning and professional networks including ECTP-CEU and AESOP.



"Figure 4: European Commission Logo"

As the executive arm of the EU, the European Commission has taken a proactive yet precautionary stance on AI governance. Most prominently, it drafted and championed the EU Artificial Intelligence Act (European Commission, 2024) – the world's first comprehensive AI regulation, which as of 2024 has been moving through the legislative process. It categorizes AI systems by risk level (unacceptable risk, high-risk, limited risk, minimal risk) and imposes requirements on high-risk systems related to transparency, fairness, and human oversight (European Commission, 2024). For urban planning, this matters because several planning-related AI applications could fall under "high-risk" categories – such as algorithms used in managing critical infrastructure, public service delivery, or systems that significantly impact citizens' rights. If a city deploys an AI tool for optimizing transportation investment or predicting areas of decline, and that tool influences resource allocation, it may be subject to the Act's provisions (European Commission, 2024). Beyond the AI Act, the EU has a broader digital strategy for cities, seen in projects like "CitiVerse," which equips planners with digital twins and AI-driven simulations (European Commission, 2025). Beyond hard law, the EU also promulgates soft law in the form of guidelines –

the *Ethics Guidelines for Trustworthy AI* (European Commission, 2019) – which outline principles like accountability, transparency, privacy, and societal benefit. The European Commission, then, operates through a top-down governance mode that can directly influence planning practice through law (e.g., mandating risk assessments for AI systems used by city authorities) and through continent-wide programs that showcase how AI can be used for public good.



"Figure 5: ECTP-CEU Logo"

As mentioned, Europe does not have a single professional institute equivalent to APA, due to the nation-by-nation organization of the planning profession; however, the European Council of Spatial Planners (ECTP-CEU) serves as a federation of national planning associations across Europe, providing a platform for coordination and sharing best practices. While ECTP-CEU itself issues charters and visionary documents (such as, on sustainable planning or participatory democracy), it is more of an umbrella network than a rule-setting body. Currently, ECTP-CEU has acknowledged digital transformation as a theme in European planning, including the use of AI, but mostly in broad strokes (e.g., highlighting the need for planners to adapt to new IT and AI tools). The organization's true value lies in bringing together professionals from different countries to discuss common challenges, and AI has become one such challenge as planning practice across Europe increasingly engages with it.



"Figure 6: AESOP Logo"

Similarly, AESOP (Association of European Schools of Planning) is the academic counterpart, connecting planning scholars and educators across Europe. AESOP has begun to actively engage with AI as well; in 2025 for instance, AESOP organized a webinar series on "AI in Planning and Planning Education," indicating a concerted

effort to understand how AI will affect planning curricula and research (AESOP, 2025). Through AESOP, European academics share findings and pedagogical strategies, likely mirroring many of the concerns seen in ACSP, such as how to teach planning students about AI, how to encourage research in urban AI that is relevant to European contexts, and how to critically assess AI's societal implications. The role of these networks (ECTP-CEU and AESOP) in AI governance is more informal– they are not creating laws, but they guide the conversation and could potentially develop guidelines or position papers that influence national associations. Conceptually, including ECTP-CEU and AESOP in this comparative framework acknowledges that professional norms and educational standards in Europe are set through collective processes that cross national borders, even as the EU provides the regulatory umbrella. Any analysis of Europe's response to AI in planning would be incomplete without considering these professional and academic communities, because they provide the expertise and on-the-ground perspective that inform policy implementation.

1.3. International Planning Institutions

"Figure 7: UN Habitat Logo"

Expanding the lens beyond Europe per se, the thesis also considers UN-Habitat, the United Nations agency focused on cities and sustainable urban development, as an institution that shapes AI governance norms at the global level. UN-Habitat operates globally (including in both developing and developed countries), and while it does not have regulatory power, it issues guidance, frameworks, and fosters international cooperation on urban issues. In recent years, UN-Habitat has championed a "people-centered smart cities" approach, which is highly relevant to AI. This approach emphasizes that digital technologies (AI included) should be used to advance inclusion, equity, and human rights in urban contexts – effectively a values-based counterweight to more techno-centric or market-driven smart city narratives (UN-Habitat, 2022). In 2022, UN-Habitat and partners launched the People-Centered Smart Cities Initiative (UN-Habitat, 2022), outlining principles to ensure that technology deployment in cities – such as AI for public services – remains focused on improving

residents' quality of life, engaging communities, and avoiding harm. UN-Habitat then released a Global Assessment of Responsible AI in Cities in 2024 (UN-Habitat, 2024). The report found that city leaders worldwide are aware of both the opportunities and risks posed by AI, and it urged municipalities to embed governance into their broader digital strategies. Its core argument was that local officials need the capacity to make informed choices about AI, which means investing in training, oversight mechanisms, and partnerships that ensure new tools are evaluated against the public interest. UN-Habitat is also directly involved in capacity-building projects by partnering with tech companies or universities on pilot projects like using AI to predict urban growth patterns or to enhance climate resilience planning (UN-Habitat, 2024). These projects serve as case studies that can be shared through UN-Habitat's networks (like the World Urban Forum) to demonstrate responsible AI use. While UN-Habitat doesn't impose rules, it creates soft governance — model policies, toolkits, and best practice guidelines — that local and national governments, and indeed planning associations, can adopt or adapt (UN-Habitat, 2022 and 2024).

With the key players and the conceptual framing established, the next step is to ground the inquiry in existing scholarship. What follows is a literature review organized around four areas that inform the thesis: current applications of AI in urban planning, the ethical and equity issues AI raises in the public sector and planning, the evolving policy frameworks in the U.S., Europe, and internationally, and the publications and initiatives of the planning associations themselves.

2 AI Policy and Application in the Urban Planning Field

Building on the institutional categories established in Chapter 1, this literature review surveys recent research and commentary — primarily from 2018 to 2025 — on topics that bear directly on AI governance in urban planning. The review is organized into four thematic subsections: Policy and Regulatory Frameworks, AI Applications in Urban Planning, Ethics, Equity, and the Public Interest in AI, and Planning Association Publications and Initiatives. These four areas were chosen because they map onto distinct but overlapping dimensions of the problem. The first addresses the concrete rules or guidelines being developed to steer AI's use. The second deals with the technical potential of AI in planning and what practitioners are actually doing with it. The third takes up the value-based questions AI provokes — questions about who benefits, who is harmed, and what standards should apply. The fourth looks at how the profession itself, through its associations and publications, is grappling with all of the above. Across these areas, most authors arrive at a broadly similar position: AI holds real potential to improve planning outcomes, but that potential comes with significant risks that demand critical evaluation and human oversight. This thesis shares that orientation — open to innovation, but attentive to what can go wrong and skeptical of narratives that treat technological adoption as inherently beneficial.

2.1. Policy and Regulatory Frameworks

The first part of the literature review turns to how AI is being governed and guided across multiple levels of authority. On one side sit public policy frameworks — federal guidance in the United States, binding regulation in the European Union, and global ethical standards advanced by bodies like UNESCO and the OECD — all of which create legal and normative expectations for how AI should be developed and deployed. On the other side are the professional planning institutions themselves, including APA, ACSP, and AESOP, which take these broader frameworks and translate them into guidance, ethics codes, and practitioner-oriented publications. The relationship between these two sides is not always seamless; policy frameworks tend to be written in general terms that do not always map neatly onto the specifics of planning practice. But taken together, they define the environment in which planners encounter and use AI.

2.1.1. United States Policy Landscape

As of mid-2025, there is no single omnibus federal law governing AI in the United States. The approach has instead been piecemeal — a mix of guiding principles, voluntary frameworks, and sector-specific guidelines that leave considerable

discretion to individual agencies and practitioners. One of the more influential high-level policy documents is the White House Office of Science and Technology Policy's Blueprint for an AI Bill of Rights, released in 2022 (OSTP, 2022). The Blueprint is not legally binding, but it has served as an important statement of values and intended protections. It lays out five broad rights: first, Safe and Effective Systems, meaning AI should be tested for safety and potential harms before deployment; second, Protection from Algorithmic Discrimination; third, Data Privacy; fourth, Notice and Explanation, so that people know when an AI is being used and can understand its outcomes; and fifth, Human Alternatives, Consideration, and Fallback, meaning the right to opt out of automated decisions and retain access to human review (OSTP, 2022).

None of this is specific to planning, but the principles resonate strongly with issues central to the field. The "right to explanation," for instance, reinforces an expectation that communities should be able to demand clarity when AI informs zoning or infrastructure decisions. Similarly, protections against algorithmic discrimination align with longstanding concerns over equity in housing, transportation, and policing — concerns that predate AI but are amplified by it (OSTP, 2022).

Another significant federal initiative is the AI Risk Management Framework, released in 2023 by the National Institute of Standards and Technology (NIST, 2023). The RMF provides voluntary guidelines for managing risks across the AI lifecycle, from initial development through deployment and ongoing operation. Its emphasis falls on mapping potential harms, measuring risks like bias or lack of explainability, and establishing governance processes to ensure ongoing accountability (NIST, 2023). Although designed for use across many sectors, the framework carries clear implications for public agencies. A city government could, for example, adopt the RMF to guide how it procures AI tools or to structure internal oversight of vendors selling algorithmic products to planning departments.

Beyond the federal level, there is growing activity at state and local scales. Several states have created AI task forces or passed laws touching on AI use in government. Illinois's AI Video Interview Act (2019), while focused narrowly on hiring tools, illustrates a broader trend toward state-level oversight of automated systems. At the municipal scale, New York City has enacted ordinances requiring bias audits for algorithms used by city agencies — a pioneering, if still limited, move in local AI governance (NYC Council, 2021).

What emerges from this landscape is a preference for principles and experimentation over uniform rules (Jasanoff, 2005). This is consistent with the broader American regulatory tendency to encourage innovation through light-touch oversight until clearer evidence of harm appears — an approach with obvious strengths in terms of flexibility but equally obvious weaknesses when it comes to protecting communities from poorly understood technologies. For urban planning, the practical result is a fragmented governance environment. Planners must keep track of evolving federal

guidance, adopt best practices like those in the NIST framework where they seem applicable, and comply with whatever state or local mandates happen to exist in their jurisdiction. There is no single rulebook, and the profession is largely left to figure things out on its own.

2.1.2. European Union Landscape

Europe has taken a markedly different path. Rather than relying primarily on voluntary frameworks, the EU has pursued binding regulation alongside broader ethical guidance. The centerpiece of this effort is the EU Artificial Intelligence Act, passed by the European Parliament in 2024, which introduces mandatory requirements for providers and users of "high-risk" AI systems (European Parliament, 2024). Urban planning intersects with several of the Act's risk categories — critical infrastructure, access to essential services, and law enforcement among them — meaning that city governments and planning agencies in EU member states are directly subject to its provisions. A municipality deploying AI for water distribution or housing allocation, for instance, would need to ensure that the system satisfies requirements for accuracy, transparency, and human oversight (Veale & Borgesius, 2021).

There are legitimate concerns about whether these compliance checks will introduce administrative burdens that smaller municipalities cannot easily absorb. But scholars also emphasize the upside: by 2026, AI systems used in European cities should meet a baseline standard of trustworthiness that is simply absent in contexts where adoption is governed only by voluntary norms (European Commission, 2024). Whether this promise is fulfilled in practice remains to be seen, but the intent is clear.

Well before the Act's provisions take full effect, European institutions had already set expectations through softer means. The Ethics Guidelines for Trustworthy AI, issued by the EU's High-Level Expert Group in 2019, emphasize human agency, technical robustness, privacy, transparency, diversity and non-discrimination, societal well-being, and accountability. These guidelines do not carry the force of law, but they have served as a normative bridge — orienting professional practice and governance debates around fairness and accountability while policymakers worked out the binding provisions of the AI Act. In this sense, Europe's regulatory approach has been layered: ethical principles first, then legislation, with each reinforcing the other.

2.1.3. International Frameworks and Global Standards

International bodies beyond the EU have reinforced this trajectory in important ways. The OECD Principles on AI (2019) and UNESCO's Recommendation on the Ethics of AI (2021) both provide soft-law frameworks that stress inclusivity, sustainability, and human-centric values. UN-Habitat has gone further than most by explicitly linking AI to the Sustainable Development Goals. Its Global Assessment of Responsible AI in Cities (2024) urged municipalities worldwide to integrate AI into their broader digital

strategies and, crucially, to build institutional capacity for governance — something many cities currently lack (UN-Habitat, 2024). None of these frameworks carry legal force, but they shape expectations and serve as reference points that national governments frequently cite when aligning local strategies with global best practices.

Sector-specific policies illustrate how these general principles get translated into practice. European transport agencies, for instance, have issued safety frameworks for autonomous vehicles that directly shape how cities plan streets and mobility systems (European Transport Safety Council, 2020; USDOT, 2020). At the municipal scale, Barcelona's Digital Data Charter (2019) shows how a single city can apply data governance principles to balance innovation with safeguards on privacy and accountability. These are not comprehensive solutions, but they represent concrete attempts to operationalize what would otherwise remain abstract commitments.

Europe's model, then, combines binding regulation, ethical guidance, and local adaptation into something reasonably coherent. Compared with U.S. planners, who operate in a more fragmented and permissive environment, European planners work within a clearer and more prescriptive framework — one where AI is encouraged but always under explicit boundaries of transparency, accountability, and citizen protection (Kitchin, 2021). The EU also provides support through funding and capacity-building programs, helping municipalities comply with new requirements while still innovating responsibly (European Commission, 2025). The result is a governance environment that is dense and sometimes bureaucratic, but internally consistent — AI adoption is steered by legal mandates and normative commitments that, at least in theory, work together to safeguard rights and maintain public trust (Jasanoff, 2005).

2.2. AI Applications in Urban Planning

Research on AI applications in urban planning has expanded rapidly over the past several years. Much of this expansion reflects genuine optimism about "smart city" innovations, but it also reflects the pragmatic challenges of implementing these technologies in messy, real-world settings. The literature covers a wide range of use cases, which can be grouped into a few broad categories.

2.2.1. Spatial Data Analysis and Modelling

Urban planning has always been data-intensive, and AI — especially machine learning — offers new ways to work with large spatial datasets. Machine learning algorithms can detect patterns and make predictions using data on land use, transportation, demographics, and other variables. Planners have been developing urban growth models and traffic forecasting tools that use AI to sift through historical data and predict future trends (Batty, 2018; Yeh & Shi, 2021). AI-enhanced planning support systems also allow simulation of various scenarios — say, how a city might evolve

under different zoning policies or development pressures. Computer vision has been used to evaluate urban design qualities by analyzing street-level imagery for walkability or the presence of green space (Naik et al., 2014; Zhang et al., 2018). Clustering algorithms, an unsupervised AI technique, have been applied to land suitability analysis, helping planners group areas with similar characteristics or identify outliers that might warrant special attention (Malczewski & Rinner, 2015; Xie & Batty, 2003). What these approaches share is the ability to process more data and surface patterns that might otherwise go unnoticed — though the question of whether those patterns are meaningful, or merely artifacts of the data, is not always addressed as carefully as it should be.

2.2.2. Generative Design and Optimization

Moving from analysis to creation, AI techniques are also being used to generate and optimize design alternatives for urban layouts. Planners have experimented with genetic algorithms, neural networks, and other optimization approaches to produce street network designs, site layouts, or transit route configurations that meet specified criteria — for example, minimizing average travel time, maximizing green space, or balancing development pressures against conservation goals (Ligmann-Zielinska & Jankowski, 2014). Out of this work has emerged the idea of an "algorithmic planner" — AI systems that can produce dozens or even hundreds of plan options for human planners to evaluate and refine (Moreno, Allam, Chabaud, & Jones, 2020). A concrete example might be an AI system generating multiple land-use zoning maps, each fulfilling a different goal such as reducing energy consumption or improving accessibility, allowing planners to explore a wider solution space than they could manage manually. Scholars working in this area tend to stress that AI should augment the human planner's vision and expertise rather than replace it, but they also raise open and difficult questions about whether reliance on AI dulls creative thinking or instead frees planners to focus on broader strategic concerns (Allam & Jones, 2020). The honest answer is probably that it depends on how the tools are used and how much critical distance planners maintain.

2.2.3. Community Engagement and Decision Support

Another domain where AI is making inroads is communication and participation. Natural language processing (NLP) and related technologies have enabled new ways to interface with the public. AI chatbots can be deployed on city websites to answer common questions about zoning rules or development projects, functioning as a kind of virtual planning assistant available around the clock (Chen & Gascó-Hernández, 2024). Some planning departments have experimented with AI to summarize citizen comments from public consultations
— using NLP to identify common themes or sentiments across thousands of written submissions, helping staff digest public feedback more efficiently than they could by reading every comment individually (Sanchez et al., 2023 and 2024).

Predictive models are also being used in decision support. One application focuses on housing markets, where AI can analyze patterns of change to identify neighborhoods at risk of gentrification and potential displacement of low-income residents — insights that could allow cities to intervene earlier with affordable housing policies (Reades et al., 2019). Another involves infrastructure, where predictive maintenance algorithms draw on sensor data to anticipate which water mains or bridges are most likely to fail, giving planners a way to prioritize repairs and capital investments. These applications position AI as a decision support tool — augmenting planners' ability to foresee problems and respond proactively rather than reactively.

The literature reports generally positive outcomes from pilot projects in this space, but also flags real challenges. AI models might not capture the nuance of community sentiment — there is a difference between what people say in a survey and what they mean, and algorithms are not always good at that distinction. Chatbots, however useful for routine inquiries, cannot replicate the trust and empathy that come from face-to-face interaction with a human planner during a contentious rezoning hearing. Used wisely, AI can handle informational tasks and sift through complex data, potentially freeing planners to focus on the direct community interaction and deliberation where human judgment is irreplaceable. But "used wisely" is doing a lot of work in that sentence, and the literature does not always grapple with what happens when the conditions for wise use are absent.

2.2.4. Administrative Efficiency

Urban planning involves a substantial amount of administrative process — reviewing development applications, issuing permits, ensuring compliance with regulations, and so on. AI is being used to speed up and improve consistency in some of these routine tasks. A number of municipalities have adopted, or are piloting, AI software to screen building permit applications or development proposals for completeness and basic compliance with zoning codes. Early systems can scan digital submissions to flag missing information or highlight potential code violations. When these tools work as intended, the benefit is straightforward: staff spend less time on routine checks and more on tasks requiring judgment, negotiation, and professional expertise.

Environmental impact assessments, which often involve checking proposals against numerous regulatory criteria, are another area where AI might assist — quickly cross-referencing databases of environmental regulations or past cases to identify potential issues. But these administrative AIs require oversight. If an AI erroneously approves or rejects an application because of a flaw in its algorithm or training data, the consequences could have legal or equity implications. A wrongly denied permit is not just a technical error; it affects someone's livelihood or housing. Planners need to remain "in the loop," double-checking AI outputs particularly in borderline or high-stakes cases. Agencies are developing protocols to ensure that accountability stays with human officials and that algorithmic decisions are transparent and explainable to

applicants, though how well these protocols work in practice is not yet well documented.

Taken as a whole, the academic and professional literature on AI applications in planning tells a story of cautious progress. Studies document gains in efficiency, analytical depth, and even creativity — AI-driven traffic optimization has produced measurable reductions in congestion in some cities, and generative design experiments have yielded ideas for public space layouts that human planners later refined and improved. But nearly every study or pilot comes with caveats. The point made repeatedly is that AI should augment rather than replace human planners, who bring contextual knowledge, ethical judgment, and a breadth of perspective that current AI systems simply cannot replicate.

Recent scholarship has also started to document how planners and educators themselves feel about these tools. Ye et al. (2024) found high awareness of tools like ChatGPT among U.S. planning faculty, but also real caution about over-reliance. Faculty worried that students might use AI to shortcut critical thinking — a concern that echoes what practitioners express too. Many planners are excited about using AI for data analysis but wary of trusting algorithms blindly, and this tension between enthusiasm and wariness runs through much of the literature. Across these discussions, the recurring call is to keep human judgment, community values, and professional ethics at the center of any AI-assisted process (Sanchez et al., 2024). The field treats AI as a powerful addition to the planner's toolkit — but one that requires training to use well and judgment about when not to use it at all. That outlook goes a long way toward explaining why governance has become so pressing: without supportive frameworks, there is little to ensure these tools are used responsibly.

2.3. Ethics, Equity, and the Public Interest in AI

The ethical dimension of AI in urban planning is probably the most intensely discussed topic in the literature, and for good reason. Planning, at its core, is a value-driven profession. It deals with distributing resources, shaping communities, and balancing competing interests while trying to uphold principles like equity and inclusion. Introducing AI into that process does not resolve these tensions — it complicates them. The literature is working through these complications in real time, and what follows covers the main lines of concern, connecting them to planning theory and practice where possible.

2.3.1. Bias and Fairness

AI systems are only as good as the data and algorithms behind them, and a substantial body of research has shown that bias can enter at multiple stages — from data collection to model design to interpretation of outputs. In urban planning, this could manifest in predictive models or decision-support tools that inadvertently perpetuate

historical inequities. Consider an AI model trained to identify "high-risk" neighborhoods for crime or blight based on historical data. Such a model might label predominantly minority or low-income neighborhoods as problematic simply because those areas were underserved or over-policed in the past, not because they are inherently more dangerous or more blighted. Scholars have termed this pattern "algorithmic redlining," drawing an analogy to the discriminatory lending maps of the 20th century (O'Neil, 2016). Cathy O'Neil's work, though not written specifically about planning, is frequently cited in planning literature to illustrate how algorithms that appear objective can produce deeply unjust outcomes.

More recent planning scholarship has picked up these themes with increasing specificity. Sanchez et al. (2024), for example, provide cases of machine learning models for housing market predictions that undervalued minority neighborhoods, and routing algorithms for public transit that systematically favored affluent districts. Planning associations have begun to recognize and respond to these concerns. APA's ethical guidelines already emphasize social justice, and they are increasingly being interpreted to cover AI usage — meaning that a planner who deploys an AI tool is expected, under the APA's Code of Ethics, to ensure that the tool does not violate principles of fairness or cause harm to disadvantaged groups (American Planning Association, 2025).

But ensuring fairness requires more than technical fixes applied after the fact. It demands governance practices that are baked into how AI tools are selected, designed, and monitored. Bias audits and algorithmic impact assessments are designed to reveal inequities before they become entrenched in policy. Inclusive design processes ensure that affected communities help define the criteria by which systems are judged — not just the outputs. What this signals is a shift in how bias is understood: not purely as a technical flaw to be debugged, but as a social and institutional challenge that requires political and ethical engagement.

In practical terms, this might look like a city adopting an AI tool for infrastructure investment decisions while also involving community representatives in reviewing whether the tool's criteria align with local priorities and values. It might also look like a city choosing to abandon a tool entirely. Kitchin (2021) describes cases where planning departments dropped predictive policing software after equity analyses revealed biased results — a decision that required courage and institutional willingness to absorb short-term costs. These cases reinforce that fairness is not a lofty abstraction but a concrete standard, and that AI tools need to be tested and monitored for equitable outcomes with the understanding that systems should be revised or discontinued when they fall short.

2.3.2. Transparency and Accountability

Traditional planning processes usually involve public hearings, written reports, and official decisions that are documented and available for scrutiny — a paper trail that provides at least a baseline of transparency. When AI enters the picture, especially as a decision-support tool, it can complicate this transparency in ways that are not always immediately obvious. Many AI models, particularly advanced ones like deep learning neural networks, are often described as "black boxes" because their internal logic can be opaque even to the people who built them (Burrell, 2016; O'Neil, 2016). If a planner cannot explain why an algorithm recommended a particular course of action, it becomes very difficult for the public, elected officials, or even the planner's own supervisors to evaluate that recommendation. The literature on algorithmic transparency argues that in domains affecting the public, AI systems need to be explainable (European Commission, 2019; NIST, 2023). Stakeholders should be able to understand at least the rationale behind an AI's recommendation — which factors were most influential, how reliable the prediction is, and what uncertainties remain.

Several approaches to strengthening accountability appear in the planning literature. One recommendation is to favor interpretable models — such as decision trees or rule-based systems — in contexts where explanation is crucial, rather than defaulting to opaque neural networks just because they might offer marginal gains in predictive accuracy (NIST, 2023). Another is to formalize documentation practices so that whenever an AI tool informs a planning decision, the record makes this explicit. A staff report might note, for instance, "This recommendation was informed by X AI model, which projected Y outcome" (American Planning Association, 2025). This kind of documentation does not solve the black box problem, but it at least makes the use of AI visible and traceable.

Transparency, though, is only part of the challenge. Liability questions also arise. If an AI-assisted decision leads to harm — a wrongful permit denial, a misallocation of services, a planning recommendation that disproportionately affects a vulnerable community — who bears responsibility? The planner who relied on the output? The software vendor who sold the system? The agency that approved its use? The law remains largely unsettled on these questions (Veale & Borgesius, 2021), but the prevailing ethical view within the profession is that accountability must remain with the human official who made the final decision (American Institute of Certified Planners, 2021). Planning associations reinforce this stance by urging members to use AI as a support tool rather than a substitute for professional judgment (American Planning Association, 2025). The difficulty, of course, is that as AI tools become more sophisticated and their outputs more convenient, the line between "support" and "substitute" may blur in ways that are hard to police through ethics codes alone.

2.3.3. Privacy and Data Governance

AI in planning often relies on large and sometimes sensitive datasets — mobility information derived from cell phone signals, social media analysis to capture neighborhood sentiment, property records used to forecast development trends. These applications illustrate both the promise of data-driven insight and the challenges of safeguarding privacy. The two sides of the Atlantic handle this quite differently. In Europe, the General Data Protection Regulation provides a comprehensive framework that requires agencies to establish a legal basis for processing personal data, implement safeguards, and uphold individual rights. Planning departments in EU member states frequently conduct Privacy Impact Assessments before introducing new technologies — for instance, before deploying AI-enabled traffic management cameras, they must verify that data is anonymized and that the scope of collection is proportional to the intended purpose.

The U.S., by contrast, operates within a fragmented regulatory environment where no single federal law covers data privacy comprehensively. Local initiatives have stepped in to fill some of the gaps. Several cities have banned or restricted AI-based surveillance technologies, most notably facial recognition in public spaces, citing concerns about civil liberties and the absence of adequate oversight. Professional associations have also started addressing these issues directly. APA's guidance on managing big data urges planners to de-identify personal information and obtain informed consent wherever feasible (American Planning Association, 2020). International frameworks filter down into practice as well. UNESCO's 2021 Recommendation on AI ethics, with its emphasis on privacy and human oversight, has become a reference point for planners engaged in global initiatives, often through UN-Habitat programs.

Across these different contexts, the literature consistently returns to a central tension. Data-driven planning offers powerful new capacities — the ability to model complex systems, identify patterns across vast datasets, and anticipate problems before they become crises. But these capacities rest on access to data that may be personal, sensitive, or collected without full awareness on the part of the people it describes. Planners are acutely aware that a single high-profile privacy breach or misuse of data can erode public trust and jeopardize support for innovative tools more broadly. The literature points to data minimization — using only the information strictly necessary for the task at hand (European Commission, 2019) — and proactive transparency about how data is collected and applied (OECD, 2019) as key safeguards. Privacy concerns also intersect with equity in important ways: limiting data collection can reduce opportunities for bias, since fewer sensitive attributes in a dataset means less potential for an algorithm to reinforce discriminatory patterns (O'Neil, 2016; Kitchin, 2021). Strong data governance — anchored in legal frameworks like the GDPR and international guidance such as the OECD's principles on AI (2019) — emerges as a necessary foundation for responsible AI use in planning.

2.3.4. Equity and Public Inclusion

Planning's core mission is to serve the public interest and foster equitable, inclusive communities. AI could either advance or undermine these goals, depending on how it is designed, deployed, and governed. Critics warn of a drift toward what might be called "technocratic planning" — a mode in which algorithms and technical experts drive decisions without sufficient public input, sidelining the community voices that planning is supposed to elevate. Imagine a city ranking neighborhood improvement projects through an AI model that relies only on quantitative data. Even if the model accounts for socio-economic need in some fashion, without meaningful resident engagement the process risks becoming opaque and exclusionary — a sophisticated version of top-down decision-making dressed up in the language of objectivity.

At the same time, AI does offer genuine possibilities for broadening participation. Translation tools and chatbots can help reach non-native speakers or residents who cannot attend evening meetings. Participatory simulations — AI-driven platforms that let citizens adjust scenarios, such as adding a park or changing density, and instantly see the projected impacts — can make planning more accessible and interactive than traditional hearings allow. Pilot projects suggest that AI-powered platforms can extend consultations beyond the usual suspects, drawing in participants who are typically excluded from planning processes. But professional and international bodies stress that these tools must support democratic dialogue, not replace it. UN-Habitat, for example, emphasizes "responsible and inclusive urban AI," framing technology as a means to enhance public engagement rather than bypass it (UN-Habitat, 2022).

The ethical and equity concerns discussed throughout this section are not peripheral to the AI-planning debate — they are shaping its agenda. AICP's longstanding emphasis on transparency and accountability is now being interpreted to cover AI use, while scholars warn that without careful design these tools may deepen existing disparities rather than close them (Allam & Jones, 2020; Kitchin, 2021). Preventing this requires practices like bias audits, transparency requirements, participatory design, and privacy safeguards — mechanisms that align AI with planning's public-interest mission rather than allowing it to drift toward efficiency for its own sake. These principles are filtering into institutional practice, visible in APA's resources on ethical AI and in the EU's insistence that fundamental rights anchor its regulatory framework (American Planning Association, 2025; European Commission, 2024). Across these perspectives, a shared principle emerges: AI in planning should remain human-centric — augmenting professional judgment, widening participation, and advancing equity rather than automating decisions that belong in the public sphere.

2.4. Planning Association Publications and Initiatives

Publications from APA, ACSP, and AESOP illustrate a clear shift: AI has moved from a marginal interest to a mainstream concern within the planning community. APA's

Planning magazine has featured both success stories and cautionary tales, reflecting a growing awareness across the profession of what AI might offer and what it might cost (Walker, 2024). The Planning Advisory Service published Planning With Artificial Intelligence (American Planning Association, 2020), which provided an early knowledge base for practitioners trying to make sense of a fast-moving field. More recent foresight reports from APA have placed greater emphasis on ethics and implementation challenges — a sign that the conversation has matured beyond initial excitement (American Planning Association, 2024, 2025). On the academic side, ACSP outlets, including newsletters and conference sessions, have explored AI's implications for planning education and research. Roundtables on how tools like ChatGPT are reshaping skill requirements have become common features at ACSP conferences (ACSP, 2024). In Europe, AESOP has convened webinars and thematic groups focused on case studies and teaching practices that reflect the region's particular regulatory and academic context (AESOP, 2025). What these publications share is a recognition that AI is now part of routine professional and academic discourse. It is no longer treated as a niche concern for the technically inclined; it is a topic that touches every part of the planning enterprise.

This practitioner and academic literature serves an important bridging function between abstract policy frameworks and on-the-ground planning practice. It translates the broad principles articulated by OSTP, the European Commission, or UN-Habitat into advice that is concrete enough to inform decisions about zoning codes, comprehensive plans, and participatory processes (American Planning Association, 2025). It also surfaces the everyday challenges that high-level policies tend to overlook — the risks of working with poor-quality data, for instance, or the lessons from failed pilot projects that prompt course corrections at the local level (American Planning Association, 2024). Across these discussions, the prevailing attitude is one of interest tempered by caution. Planners see AI as a way to improve efficiency and engagement, but they keep returning to the need for accountability, equity, and public trust — concerns that no amount of technical sophistication can make disappear (Walker, 2024).

These association-driven outputs extend and reinforce the themes identified in the broader policy frameworks. They make clear that AI in planning is not simply a technical innovation to be absorbed passively — it is fundamentally a governance challenge, raising questions about ethics, accountability, and institutional capacity that the profession cannot afford to sidestep (Kitchin, 2021; Allam & Jones, 2020). The policy frameworks aim to balance efficiency and innovation against risks such as bias, inequity, and the erosion of public trust (European Commission, 2019; OSTP, 2022; NIST, 2023). The U.S. context leans toward flexibility and voluntary adoption (OSTP, 2022; American Planning Association, 2025), while Europe emphasizes precaution and binding regulation (European Commission, 2024; Veale & Borgesius, 2021). But in both cases, planners are expected to develop competence in technology and governance

simultaneously, and neither side has fully figured out how to make that happen (American Planning Association, 2020; AESOP, 2025).

This review has mapped the major contributions and debates in a field that is still taking shape. Policy frameworks and professional publications converge in portraying AI as both an opportunity and a challenge — a characterization that, while accurate, also risks becoming a cliché if it is not backed by specifics (Walker, 2024; ACSP, 2024). The chapters that follow attempt to provide those specifics. Chapter 3 describes the research design used to analyze the institutional documents introduced in this chapter, and the subsequent empirical chapters assess how well practice aligns with the ideals described here — and where transatlantic exchange might strengthen the profession's readiness for what comes next (UN-Habitat, 2024; OECD, 2019).

3 Research Design and Methodology

3.1. Research Scope and Case Selection

This research adopts a qualitative comparative research design to examine how major planning institutions in the United States, Europe, and international governance contexts are responding to the emergence of artificial intelligence in urban planning. The methodological approach reflects the exploratory and interpretive nature of the research question, which seeks to understand institutional positioning, governance framing, and normative responses rather than to measure causal outcomes or quantitative impacts.

The units of analysis consist of institutional actors that influence planning practice, education, or policy discourse. These include:

- The American Planning Association (APA)
- The Association of Collegiate Schools of Planning (ACSP)
- The European Council of Spatial Planners (ECTP-CEU)
- The Association of European Schools of Planning (AESOP)
- United Nations–Habitat (UN-Habitat)

Cases were selected using purposive sampling based on three criteria:

1. Institutional influence over planning practice or governance
2. Engagement with AI-related discourse or initiatives
3. Representation of distinct governance environments

This selection allows for a structured transatlantic comparison capturing variation across:

- Professional self-regulation
- Academic agenda-setting
- Formal regulatory governance
- Soft-law global guidance

3.2. Data Collection Methods

Data collection relied on qualitative document-based research. Materials analyzed included:

- Policy frameworks and legislation
- Institutional strategy reports
- Ethical guidelines and codes of practice
- Educational initiatives and curricular statements
- Public communications and position papers
- Conference outputs and published commentaries

Sources were obtained through institutional repositories, organizational publications, academic databases, and publicly accessible archives. Selection prioritized documents that articulated institutional positions, governance approaches, or programmatic responses to AI in planning contexts.

The research process emphasized interpretive engagement rather than mechanical aggregation of content. Documents were reviewed iteratively, allowing emerging themes to inform subsequent source selection and refinement of analytical categories.

3.3. Uses of Digital Tools and Artificial Intelligence Assistance

Digital tools, including generative artificial intelligence systems, were used in a limited supportive capacity throughout the research workflow. Their use is disclosed here in the interest of methodological transparency.

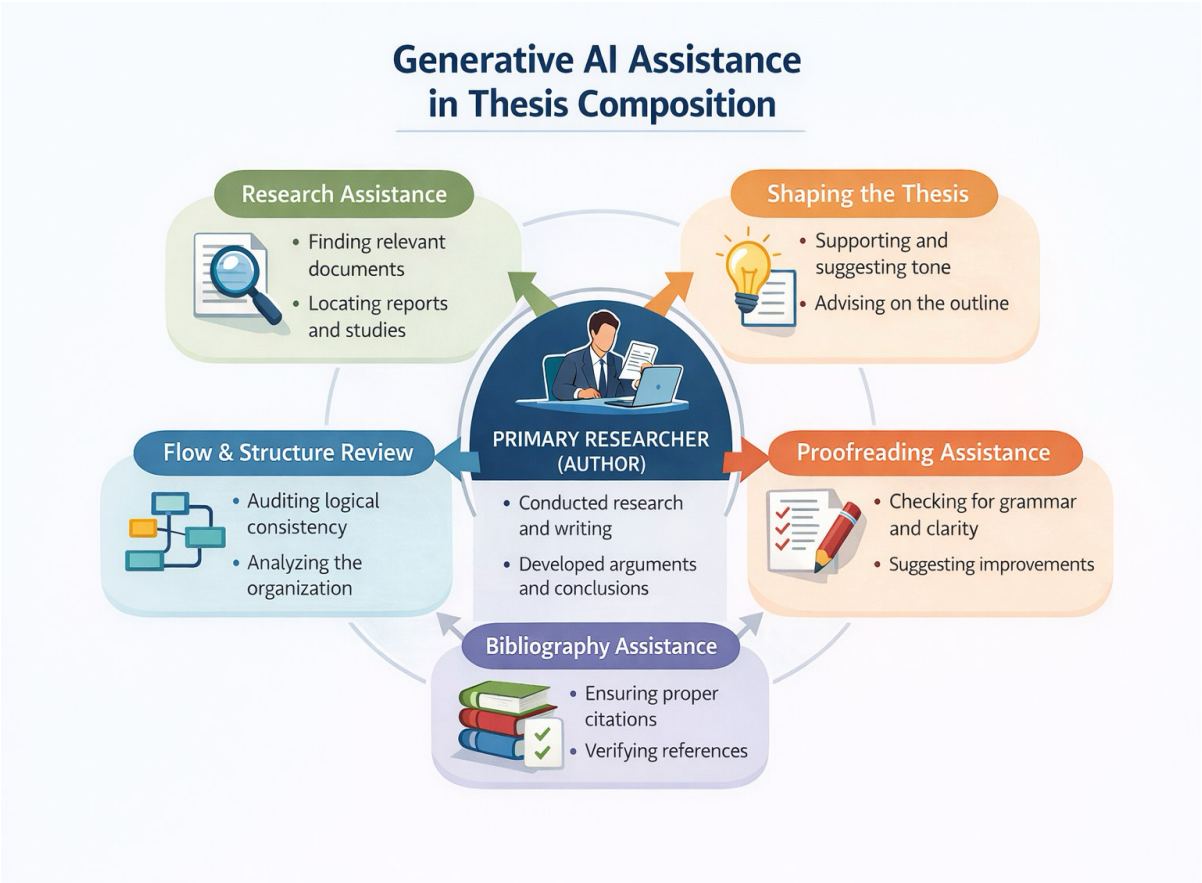
Artificial intelligence tools were employed to assist with:

- Identifying potential literature leads during early review stages

- Suggesting alternative academic phrasing or tone adjustments
- Reviewing grammar and stylistic clarity
- Supporting citation formatting consistency
- Providing structural feedback on document organization

These tools were not used to generate research arguments, conduct analysis, select cases, interpret findings, or determine conclusions. All conceptual framing, data collection decisions, comparative analysis, and written content reflect the author's independent scholarly judgment.

AI assistance functioned as an auxiliary editorial and exploratory aid comparable to grammar software or reference management tools, while intellectual responsibility for the research process and outputs remained fully human-directed. Figure 8 illustrates the division of responsibilities between the primary researcher and generative AI tools throughout the thesis composition process.



"Figure 8: Generative AI Assistance in Thesis Composition"

This chapter outlines the research approach used in this thesis. It describes the selection of institutional cases, the use of document-based data collection, and the

qualitative comparative analysis applied to interpret governance responses to AI in planning. The following chapters use this framework to examine institutional approaches in the United States and Europe and compare their implications for planning practice.

4 Responses to AI in the United States Planning Context

Drawing on the comparative framework outlined in Chapter 1, this chapter examines how the major U.S. planning institutions – the American Planning Association (APA) and the Association of Collegiate Schools of Planning (ACSP) – have been responding to the rise of AI in urban planning. These two bodies cover both professional practice (APA) and academia (ACSP), offering a comprehensive view of the "planning community" in the United States. The chapter is organized into two main sections, one for APA and one for ACSP, each further broken down into key areas of activity. Throughout, the emphasis is on identifying the initiatives, guidelines, and narratives these organizations have developed around AI, and how they frame the opportunities and challenges. As we will see, the U.S. approach is characterized by proactive yet non-regulatory strategies – focusing on education, ethical guidance, and thought leadership, rather than formal rules.

4.1. American Planning Association (APA) Initiatives

The APA (founded in 1978), with its tens of thousands of members across the country, serves as the voice of the planning profession in the U.S. Its responses to AI are significant because they can shape practitioner awareness and norms nationwide. In recent years, the APA's recognition of AI as a transformative force in planning has grown considerably. Rather than issuing strict directives, the APA's approach has been to inform and equip planners so they can thoughtfully integrate AI into their work. Key areas of APA's engagement with AI include foresight reporting and knowledge dissemination, professional development and ethics, and advocacy and partnerships.

4.1.1. Policy Guidance and Foresight Efforts

One of the APA's early steps in addressing AI was to include it in its Foresight practice. The APA Research Department produces trend reports and scenario studies to help planners anticipate future changes. In the *APA 2025 Trend Report for Planners*, artificial intelligence was highlighted as an emergent trend poised to reshape how planners analyze data and make decisions (American Planning Association, 2025). This report does not create policy per se, but it signals to practitioners that AI is something to pay attention to. By identifying AI as a trend, APA essentially put the topic on the professional agenda. Planners reading the report encounter scenarios where, for example, permit reviews might be partially automated or predictive models might guide infrastructure investments.

Complementing these reports, APA has developed a Research KnowledgeBase collection on "Artificial Intelligence and Planning" (American Planning Association, 2024). This online resource hub compiles case studies, best practices, academic articles, and external resources about AI applications in planning. By curating such a collection, APA acknowledges both the interest among planners and the knowledge gap – planners are looking for guidance on what AI means for them. The KnowledgeBase includes examples of cities using machine learning for zoning code enforcement, tools for analyzing big datasets in comprehensive plans, and even explorations of generative AI tools that can help write planning documents. The point is access: busy practitioners can find what they need in one place instead of combing through scattered sources.

APA has also started weaving AI-related guidance into its *Planning Advisory Service (PAS)* publications. For example, a PAS Memo by Wasserman & Flaxman (2023) entitled "Artificial Intelligence and Planning Practice" provides an accessible overview of AI technologies and discusses their implications for daily planning tasks. It outlines use cases like using AI for predictive analytics in transportation planning, but also cautions about data biases and the need for human oversight. This kind of advisory content stops short of formal policy, but it frames AI in a way that aligns with planning values – as a tool that can enhance analysis and efficiency if used responsibly (Wasserman & Flaxman, 2023). The widespread distribution of such content through APA's channels means thousands of planners are reading about AI in a context tailored to their profession.

Overall, the APA's policy guidance on AI has been informational and future-oriented. Through trend reports and curated knowledge, APA is preparing its members for changes rather than imposing rules. The association's message is essentially, "AI is coming (in fact, it's here in some places); here's what it might mean for you, and here's where you can learn more." This fits what APA is. It cannot regulate; it can only inform, convene, and nudge.

4.1.2. Professional Development and Ethical Stances

Beyond written guidance, the APA's engagement with AI extends into training and ethical discourse, reflecting its mission to uphold professional competency and integrity. One notable effort is the inclusion of AI topics in APA's professional development offerings. Over the past couple of years, APA conferences and webinars have featured sessions on AI, where practitioners share experiences with applications like AI-based mapping and community engagement tools. These sessions are well-attended, indicating planners' curiosity and concern about AI's role. APA has also offered webinar series (often in partnership with its Technology Division or external experts) on subjects like "Introduction to Machine Learning for Planners." Such

training doesn't make one an AI specialist, but it raises the baseline understanding so planners can better communicate with data scientists or evaluate AI products being pitched to their agencies.

Hand-in-hand with skill-building is the conversation about ethics. The American Institute of Certified Planners (AICP), affiliated with APA, maintains a Code of Ethics that all certified planners must follow. While the code (last updated comprehensively in 2021) does not mention AI explicitly, it enshrines values like serving the public interest, fighting bias, and using a factual basis for decisions (American Institute of Certified Planners, 2021). APA has started to interpret these existing ethical principles in light of AI. A blog post on APA's website titled "Planning Ethics and Generative AI" discusses scenarios like using ChatGPT to draft a public report or employing algorithms for community profiling, and asks how a planner should navigate such situations (American Planning Association, 2024). The guidance from APA's Ethics Committee emphasizes that planners remain accountable for work done with AI assistance – for example, if AI summarizes community input, the planner must ensure it's accurate and unbiased (American Planning Association, 2024). The blog also warns against over-reliance on tools that planners don't fully understand, aligning with the principle of not "misrepresenting facts" or analysis.

The framing matters because it refuses to treat AI as magic. A model is a planning method like any other, and it carries the same duty of care. The APA has not (as of now) issued a formal addendum to the Code of Ethics specifically on AI, but by publishing commentaries and hosting discussions, it is socializing the profession to new dilemmas. For example, if an algorithm recommends a policy that could unintentionally disadvantage a minority group, APA's ethical guidance would remind the planner of the duty to check that outcome, much as they would scrutinize a traditional analysis, in order to "seek social justice" per the code (American Institute of Certified Planners, 2021).

Additionally, APA is encouraging the sharing of practitioner experiences. Through its magazine *Planning*, which often features case stories, APA has highlighted early adopters and lessons learned, including planners managing community concerns around transparency when introducing AI tools into routine workflows. This kind of peer knowledge helps others foresee issues like the need to explain AI decisions to the public.

APA's stance is cautiously optimistic. AI is welcome for the efficiency and insight it can add, but not at the expense of the profession's core values. By folding AI into its training and ethics work, APA is quietly redefining what counts as competent, ethical practice.

4.1.3. Advocacy and External Collaboration

While APA mainly guides its own members, it also has a role in advocating for planning interests in broader policy discussions. In the realm of AI, this has meant engaging with governmental initiatives and collaborating with academia and industry when appropriate. For example, when the White House Office of Science and Technology Policy released the *AI Bill of Rights* blueprint in 2022, APA engaged with the resulting policy conversation, echoing principles of fairness, transparency, and privacy and bringing in the context of local government planning systems (OSTP, 2022). Engagement of this kind brings the perspective of urban planners – who deal with public data and community trust every day – into national conversations on AI governance. This kind of advocacy is subtle but important, as it raises awareness among federal policymakers. Tools such as AI-driven zoning analytics or predictive policing have direct planning implications, and national guidelines should consider the on-the-ground realities planners face.

APA has also positioned itself as a partner in cross-sector work on AI, supporting collaborations with academic institutions and technology groups on applied projects. These kinds of collaborations allow planners to be *co-creators* or at least test drivers of new technology, rather than passive recipients. APA has also been involved in inter-disciplinary workshops alongside engineering associations and AI developer communities, where the planning perspective sits alongside technical and policy expertise on the ethical use of AI in public infrastructure. By stepping outside the immediate planning world, APA represents the values and concerns of planners in these wider tech-policy dialogues.

On the international stage, APA has informal ties with organizations like the Global Planners Network and even UN-Habitat. While not directly setting global policy, APA's participation in forums like the World Urban Forum (often via its representatives or members) means it can share the U.S. planning community's approach to AI and also bring back global insights to its members.

In terms of formal advocacy, APA has not lobbied for specific AI legislation – consistent with the U.S. context where AI regulation is still nascent and often sector-specific. However, APA does advocate for funding and support, as it has backed federal smart city grants and research initiatives that include AI components, arguing that planners should be involved in the development of smart city technologies to ensure that human and community factors are considered. This aligns with APA's long-standing advocacy for planning in federal programs – ensuring that technology isn't pursued for its own sake but serves larger urban policy goals.

APA's external engagement on AI is about making the planner's voice heard in the technology domain and ensuring planners have a seat at the table in shaping how AI tools are developed and governed. By collaborating across sectors, APA also helps demystify AI for its members – as they see their organization involved in these

conversations, it reinforces that AI is part of mainstream planning practice, not an isolated tech niche.

4.2. Association of Collegiate Schools of Planning (ACSP) Initiatives

ACSP is the network of university planning programs and faculty, and thus represents the educational and research arm of the planning field in the U.S. Its engagement with AI revolves around integrating AI into planning curricula, promoting research on planning and AI, and fostering dialogue between scholars and practitioners. Academia often serves as an incubator for new ideas, and ACSP's actions can shape what new planners know about technology as they enter the workforce, as well as what research gets prioritized. Over the past few years, ACSP has moved from sporadic discussion of AI to more structured initiatives, recognizing that AI is altering the landscape of planning scholarship and pedagogy.

4.2.1. Research Promotion

One of ACSP's key moves was to explicitly encourage and organize research on AI in planning. In 2024, ACSP established a *Task Force on Urban Planning and AI* – a working group of planning professors dedicated to examining AI's implications for both planning practice and education (University of Florida Department of URP, 2024). The formation of this task force (initiated by faculty interest and sanctioned by ACSP's Governing Board) signaled that AI is no longer a fringe topic in planning academia. The Task Force's mandate covers several objectives including assessing how AI can be responsibly incorporated into planning curricula, identifying research gaps where planning scholars could contribute (for example, studying the impacts of AI on community engagement or equity in cities), and recommending ways for ACSP to support its members in this domain.

Following its creation, the Task Force quickly became a catalyst for scholarly discourse. At ACSP's annual conference, special sessions were convened on AI – including workshops and roundtables. Topics ranged from "AI in Planning Theory: Do We Need New Paradigms?" to "Using Big Data and Machine Learning in Planning Research." ACSP's journal, the *Journal of Planning Education and Research (JPER)*, in fact announced a call for a special issue on AI in late 2024. The call for papers explicitly sought studies that could "foster a comprehensive understanding of the opportunities, challenges, and implications of AI in the field of urban planning." This special issue (forthcoming) is a clear sign of high scholarly interest and the desire to build a coherent body of knowledge (indeed a number of submissions were expected from both U.S. and international researchers).

By promoting and legitimizing research on AI, ACSP performs a vital function by creating a feedback loop with practice. Academic research can take a broader and more critical view – for example, exploring whether AI tools truly improve planning outcomes or simply add new layers of complexity. The findings from such research then inform how educators teach future planners and how organizations like APA draft their guidelines. One early piece of scholarship in this vein is the article by Sanchez, Brenman, & Ye (2024) which examines ethical concerns of AI in urban planning — bias, transparency, the planner's accountability for algorithm-driven decisions. The fact that this article appeared in JAPA (the Journal of the American Planning Association) indicates collaboration between academic inquiry and professional dialogue. ACSP encourages its members to contribute to these debates, bridging the gap between theory and practice.

ACSP also acts as a convener beyond traditional journals. It has supported the creation of research clusters and interest groups on technology and smart cities within its community, providing the platform and network so that an economist studying AI in transportation and a design scholar looking at AI in public participation can find each other and possibly collaborate.

ACSP's research-focused initiatives around AI indicate an embrace of innovation in the intellectual realm of planning. They demonstrate an understanding that planning knowledge must evolve alongside technological change – not by unquestioningly accepting tech, but by critically examining it. ACSP is ensuring that there is a growing body of planning scholarship that grapples with AI – which in turn enriches the content taught to students and offered to practitioners for guidance.

4.2.2. Curriculum and Pedagogy

The impact of AI on planning education has become a pertinent issue for ACSP member schools. After all, today's planning students will be practicing in an era where AI is commonplace, so their training should prepare them for that reality. ACSP has therefore encouraged programs to update their curricula and share pedagogical strategies related to AI.

Several planning schools have begun introducing new content. For example, some universities now offer a course on "Planning Analytics and AI" or incorporate modules on machine learning in existing quantitative methods classes. ACSP facilitates the exchange of these ideas by hosting workshops for faculty. In 2025, AESOP (the European association) and ACSP held a joint webinar series titled "AI in Planning and Planning Education" (AESOP, 2025). Through such international collaboration, educators presented how they are teaching about AI – whether it's hands-on training in GIS-based AI tools, or seminar discussions on ethical scenarios. ACSP's role here is not to dictate a standard curriculum, but to provide a forum where educators who are early adopters or experimenters can mentor others. For instance,

the Department of Urban Planning at a major university might share its syllabus and assignments for a "Smart Cities" class that covers AI case studies, and ACSP disseminates this through its network.

Another concrete development involves student engagement with AI tools. Some planning programs have started incorporating tools like data mining of social media or using AI to generate planning scenarios as part of studio projects. ACSP recognizes these experiments by featuring them in its communications (newsletters or annual conference education sessions). This signals that using AI in studios or research is not just a novelty, but an emerging best practice to expose students to. ACSP also keeps an eye on outcomes, highlighting student work that engages substantively with AI-related questions. Recognition of this kind incentivizes both students and faculty to push further on the integration of technology into planning education.

A key challenge that ACSP and planning educators face is ensuring that while teaching new tech skills, they don't lose the core planning competencies. The literature has noted a concern that over-reliance on AI tools could potentially erode critical thinking or qualitative sensitivities among new planners (Ye et al., 2024). ACSP addresses this by promoting a balanced pedagogy. An example assignment might have students run an AI model to identify "high-risk" neighborhoods for disaster impacts, and then critically analyze whether the model's definition of risk aligns with social equity goals or whether it might be reflecting historical biases. This encourages a mindset where future planners see AI as an aid, not a replacement for human judgment.

ACSP's discussions (through its Committee on Academic Affairs, for instance) have also touched on faculty development. Not all planning faculty are versed in AI or data science, so ACSP considers ways to upskill educators too – maybe through targeted short courses or collaborations with computer science departments. Some universities have begun hiring faculty with expertise in urban data analytics, and ACSP tracks these trends to inform its member programs about evolving faculty profiles.

ACSP's work on curriculum is, at bottom, about modernizing planning education without hollowing it out. The goal is to produce graduates who are literate in AI and comfortable working with data and algorithms, yet deeply grounded in planning's humanistic and ethical foundations. By acting as a clearinghouse for educational innovation, ACSP helps the whole community elevate its teaching to meet the demands of the coming decades.

4.2.3. Collaboration and Discourse

Finally, ACSP recognizes that bridging academia and practice is crucial, especially with fast-moving subjects like AI. Knowledge shouldn't stay siloed in universities. Thus, ACSP has taken steps to enhance dialogue between scholars and practitioners on AI in planning.

One avenue has been through joint initiatives with APA. While APA and ACSP are separate, they often overlap (many APA members are alumni of ACSP schools or even current educators). In recent years, ACSP and APA have coordinated on conference content – for example, a technology-themed track at the APA National Conference might be co-curated by academics from ACSP. This ensures that the latest research findings (say, on how an AI tool performed in a case study city) are presented directly to practitioners, and conversely, practitioners' questions and experiences inform academic inquiry. A planner might stand up at a conference and say, "We tried using an AI-based engagement platform and got low participation from certain neighborhoods – why might that be?" Such questions can spark research by an academic in the audience or lead to a collaborative study.

Additionally, ACSP's Task Force on AI has included a few practitioner representatives and even technologists from outside the traditional planning sphere. By doing so, ACSP created a small community of practice that spans beyond academia. For instance, an AI entrepreneur working on civic tech could share what tools are being developed, while a city planning director shares the pain points and needs. ACSP acts as the convener for these discussions, in settings like pre-conference workshops. The outcome is mutual learning, in which scholars get a reality check on what's actually needed in the field, and practitioners gain exposure to cutting-edge research (and might volunteer their city as a "living lab" for studies).

Moreover, ACSP's member schools often have outreach programs to local governments (planning clinics, applied research centers). Through ACSP, there's an informal network where these programs exchange how they're assisting municipalities on AI issues. One planning school's lab might be helping their city analyze equity impacts of a new algorithm; another might be designing curriculum for municipal staff training. ACSP circulates these stories in its newsletters or at meetings, implicitly encouraging more schools to form such partnerships.

The synergy between ACSP and practice is reinforced by transatlantic collaboration with AESOP, where planning academics from both regions exchange views with practitioners on shared challenges like smart city development and AI use. ACSP, representing U.S. planning academia, can then align some of its priorities — including

what the Task Force should focus on next — with the needs practitioners express, for example guidance on using AI tools for public participation without excluding those with limited tech access.

ACSP's collaborative efforts ensure that the academic response to AI is not happening in a vacuum. By engaging with APA and practitioners directly, ACSP helps ensure that research and educational efforts translate into practical benefit, and likewise that practice-informed questions drive scholarly exploration. The U.S. planning community, through APA and ACSP together, is thus mobilizing on multiple fronts – knowledge-sharing, education, ethics, and advocacy – to address AI. This multi-faceted but largely voluntary and network-driven approach sets the stage for comparison with Europe's more policy-driven approach, which we turn to in the next chapter.

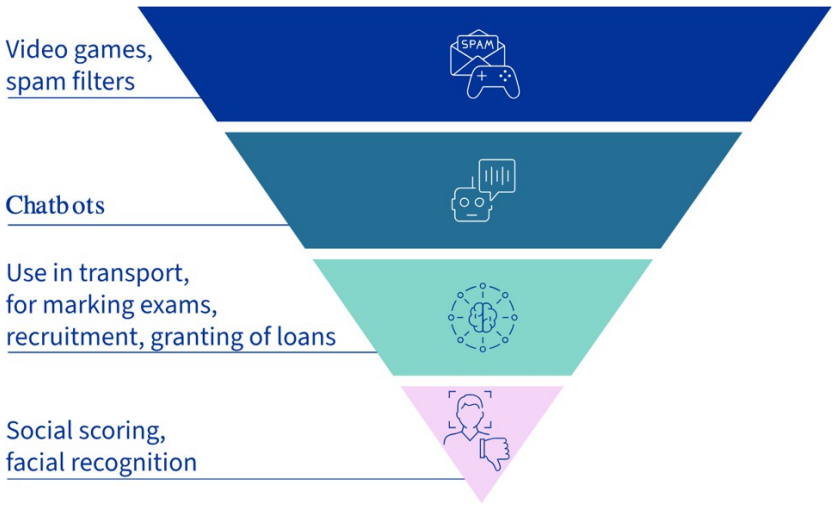
5 Responses to AI in the European Planning Context

Following the analysis of U.S. institutions in Chapter 4, this chapter turns to the European context, which presents a different governance landscape. Europe's approach to AI in urban planning is deeply intertwined with government policy and international frameworks. This chapter examines how key European institutions are addressing AI's integration into planning: the European Union's regulatory stance and supportive programs, the role of European planning networks (ECTP-CEU and AESOP), and their combined influence on professional practice. These actors form a multi-level response, from hard law and formal guidelines to professional guidance and academic agenda-setting. Throughout, there is a pronounced emphasis on ensuring AI aligns with public values like fundamental rights, sustainability, and inclusivity — a perspective shaped by Europe's policy environment and the profession's commitment to democratic planning.

5.1. European Commission and EU-Level Initiatives

The European Union has emerged as a global leader in attempting to regulate AI. For urban planning, this means that some uses of AI by cities and planning agencies may soon be governed by EU law, and planners must be aware of new compliance requirements. The centerpiece of Europe's approach is the EU's Artificial Intelligence Act, complemented by various strategic programs to harness AI for urban innovation while safeguarding ethical standards.

5.1.1. Regulatory Framework: The EU AI Act



"Figure 9: EU AI Act Risk Classification Pyramid. Source: Council of the European Union (2024)."

As discussed in the literature review (Section 2.1), the EU finalized negotiations on the Artificial Intelligence Act in 2024, the world's first comprehensive legislation focused on AI (European Parliament, 2024). This Act takes a risk-based approach, categorizing AI systems into tiers (unacceptable risk, high-risk, limited, minimal) and imposing obligations accordingly (see Figure 9). As illustrated, the inverted pyramid arranges these categories by the volume of AI systems affected: the broad top tier represents the vast majority of AI systems that carry minimal risk and face no mandatory requirements, while banned practices occupy the narrow base, covering only a small number of applications. The regulatory burden increases as one moves downward through the tiers. For urban planning applications, the critical category is "high-risk AI systems." These include algorithms used in areas that can significantly affect people's lives or rights — which can encompass planning-related tools. For example, an AI system that helps decide who gets access to public housing, or one predicting areas of a city for code enforcement crackdowns, could fall under high-risk if those decisions carry implications of fairness and rights. Under the AI Act, any high-risk AI system must meet requirements for transparency, accountability, accuracy, and human oversight. In practice, if a city planning department in, say, Amsterdam or Barcelona deploys an AI tool to guide infrastructure investments or police deployment, it would need to ensure the tool is registered, documented, and subject to human review in decision-making (European Parliament, 2024). The Act even mandates that users of AI be notified when they are interacting with an AI system, to prevent (for instance) chatbots masquerading as human public officials.

For planners, this legislation introduces a new dynamic: legal compliance becomes a component of AI governance. Where a U.S. planner might adopt an AI tool based on internal guidelines, a European planner may also have to verify that the tool adheres to the AI Act. Certain AI practices could be illegal or require modification; "unacceptable risk" systems (like social credit scoring, which is banned) would be off-limits entirely. If a planning-related AI inadvertently did something resembling social scoring — imagine an algorithm ranking neighborhoods by desirability in a way that stigmatizes certain populations — that would be a serious red flag.

The AI Act's influence reaches beyond the letter of the law. It reinforces the ethos of "human-centric AI" that the EU had already articulated in its earlier ethics guidelines (European Commission, 2019). The EU has explicitly framed the Act as a means of ensuring that AI in Europe respects fundamental rights and democratic values. So, when a city uses AI for traffic flow optimization, the expectation is that this happens transparently, without violating privacy, and with recourse if harm occurs. Planning agencies in EU countries are already consulting with legal teams about how to audit algorithms for bias or implement human-in-the-loop mechanisms for compliance. This cascading of broad regulation into local public sector practice is somewhat new territory for planning; the profession has plenty of regulations (land use laws, environmental assessments), but now the tools planners use are themselves becoming objects of regulation.

The AI Act had not yet fully come into force at the time of writing (enforcement is expected a couple of years after passage). But European planning institutions — national associations, ECTP-CEU — have been preparing by informing their members. We see the European Council of Spatial Planners discussing the AI Act in their meetings, translating what it means for a planning office. They may issue guidance along the lines of: "if you use an AI-based urban mobility model, here's what you need to document and how to involve a human decision step." The EU's regulatory approach, in short, is pushing the principle that planners cannot just plug-and-play AI — they must govern it as a matter of legal duty.

5.1.2. Strategic Initiatives: AI for Urban Innovation

Regulation is one side of the coin; the other is the EU's active promotion of AI as a tool for innovation and sustainability in cities. The European Commission has launched and funded several initiatives to encourage beneficial AI use in urban planning and management. A prime example is the "CitiVerse" project under the EU's digital strategy (European Commission, 2025). CitiVerse is envisioned as an ecosystem using AI, digital twins (virtual 3D models of cities), and extended reality technologies to help cities plan and engage with stakeholders in immersive ways. For planners, digital twins enhanced with AI could enable simulation of different development scenarios or infrastructure changes with realistic detail and real-time data. The Commission has been funding pilot projects where cities create these digital urban twins. This

illustrates the EU's dual approach: enabling the technology while shaping it through funded pilots that align with EU values. Cities receiving EU support for digital twins are typically required to incorporate features like open data accessibility or citizen participation tools, tying tech deployment to transparency and inclusion goals.

Investment also flows through programs like Horizon Europe, the EU's research and innovation funding framework. Calls have been issued for projects on AI in mobility, energy-efficient smart city planning, and climate resilience. Horizon Europe has funded a cluster of initiatives that use AI for climate adaptation planning, each analyzing data on heat islands and flooding to guide urban design changes. The GreenInCities project pairs AI and data-driven models with nature-based solutions to mitigate urban heat, with Barcelona, Prato, Nova Gorica, Athens, and Helsinki piloting interventions that smaller follower cities then adapt and scale. UP2030 applies AI to generate high-resolution temperature maps that let planners target heat-island hotspots with greater precision. In Italy, the USAGE project has used AI to process aerial and hyperspectral imagery of Ferrara, mapping urban materials and heat distribution to inform local climate policy. Taken together, these projects serve dual purposes: solving local problems while producing demonstrations that can be replicated across Europe if they succeed (European Commission, 2025).

The EU facilitates knowledge exchange platforms as well. The EU Smart Cities Marketplace (formerly the European Innovation Partnership on Smart Cities) now includes discussions and resources on AI in urban services. European Commission workshops bring together city leaders, tech providers, and researchers to discuss case studies — using AI to optimize public transport routes or waste collection, for example. Through these initiatives, the Commission shapes a narrative in which AI serves the objectives of the European Green Deal (reducing emissions through smarter transport systems, say) while emphasizing knowledge sharing and avoidance of known risks.

Europe also ties AI to its urban policy frameworks indirectly. The New Leipzig Charter (an EU urban policy document) and the Urban Agenda for the EU both stress digital transition as central to future cities. AI is increasingly featured as a means to an end — the end being sustainable and inclusive urban development. In practical terms, a city might mention "using AI for better energy demand management in buildings" as one action within a broader climate strategy.

The EU's strategic initiatives position AI as an enabler of urban innovation that must be cultivated responsibly. Planners in Europe therefore find themselves not only having to comply with new rules but also presented with opportunities (and funding) to experiment with AI in pursuit of policy objectives. It creates a dynamic environment where a planner might pilot an AI tool through an EU program and simultaneously be figuring out how to comply with the AI Act's requirements for that same tool.

5.2. European Council of Spatial Planners (ECTP-CEU) Initiatives

If the European Commission provides the regulatory and funding architecture for AI governance, ECTP-CEU serves as the professional voice of European planners — analogous, in some respects, to the APA's role in the United States. As a federation of

national planning associations across Europe, ECTP-CEU coordinates professional standards and facilitates the exchange of best practices across borders. Its engagement with AI has been growing, though it operates through guidance and consensus-building rather than binding directives.

5.2.1. Policy Guidance and Professional Standards

ECTP-CEU periodically issues charters and declarations on major trends affecting the planning profession. In recent years, digital transformation — including AI — has become a recurring theme. In 2022, ECTP's General Assembly included a statement encouraging member associations to help planners "adapt to new digital tools such as AI, while upholding the human-centric values of our profession" (ECTP-CEU, 2023). While broad, this sets a tone and empowers national associations (in France, Italy, Germany, and elsewhere) to develop more detailed guidelines tailored to their own regulatory contexts.

ECTP-CEU's approach mirrors some of APA's strategies: it does not create binding rules, but it shapes the professional conversation. By issuing position papers and hosting discussions at events like the Biennial of European Towns and Town Planners, ECTP-CEU signals to its member associations that AI is a priority issue requiring professional attention. Some national institutes have responded by producing their own guidance. The Royal Town Planning Institute (RTPI) in the UK, which often collaborates with ECTP despite the UK's departure from the EU, published introductory guides on AI for planners (Mills, 2025). ECTP amplifies these resources by sharing them across its network — a guide from the UK or the Netherlands on AI ethics in planning might be translated or presented at an ECTP event for others to learn from.

Where ECTP-CEU differs from APA is in its relationship to regulation. APA operates in a context where AI governance is largely voluntary; ECTP-CEU operates within the shadow of the AI Act. This means that part of ECTP-CEU's role is translating EU-level requirements into practical guidance for planning offices. A national planning association in Spain or Poland, for example, may look to ECTP-CEU for help understanding what the AI Act means for their members' day-to-day practice — what documentation is needed, how to structure human oversight, and how to evaluate vendor claims about AI tools.

5.2.2. Advocacy and Cross-Border Coordination

ECTP-CEU's value as a transnational federation lies in its ability to connect planners across very different national contexts. AI governance challenges in a Nordic country with strong digital infrastructure look quite different from those in a Southern or Eastern European context where digital capacity may be more limited. ECTP-CEU provides a platform where these varied experiences can be shared.

At the European level, ECTP-CEU has engaged with EU consultations on digital strategy and smart cities. Its representatives participate in stakeholder events organized by the European Commission, ensuring that the perspective of practicing planners — who deal with the practical consequences of AI deployment in communities — is represented in policy discussions. This advocacy role parallels APA's engagement with the White House OSTP blueprint and other federal initiatives in the U.S., though it plays out within a more formalized consultation structure in Europe.

ECTP-CEU has also begun coordinating with city networks like Eurocities on questions of responsible technology use. These collaborations help connect the profession-level conversation about AI ethics with the city-level conversation about smart city implementation, ensuring that the principles planners espouse are reflected in the tools municipalities actually procure and deploy.

5.3. Association of European Schools of Planning (AESOP) Initiatives

AESOP is the academic counterpart to ECTP-CEU, connecting planning scholars and educators across Europe — much as ACSP does in the United States. Its engagement with AI revolves around integrating AI into planning curricula, promoting research on planning and AI, and fostering dialogue between European academics and practitioners. As with ACSP, AESOP's actions shape what new planners learn about technology and what research directions the field prioritizes. Over the past few years, AESOP has moved from occasional discussion of AI to more structured initiatives, recognizing that AI is reshaping planning scholarship and pedagogy across the continent.

5.3.1. Research Promotion and Scholarly Discourse

AESOP has taken deliberate steps to position AI as a priority research theme within European planning academia. Its thematic groups — particularly the group on "Technologies in Planning" — have made AI a regular discussion point at AESOP Congresses and annual conferences. These sessions bring together scholars from across Europe working on topics ranging from AI-driven spatial analysis to the ethics of algorithmic decision-making in public services.

A particularly visible initiative was AESOP’s engagement with European-level discussions on responsible AI procurement in cities, where scholars from its network have contributed to policy-relevant outputs alongside European urban research bodies. AESOP members played a key role in shaping this work, bringing an academic perspective that reinforced calls for impact assessments and public engagement whenever AI is introduced into planning processes. This kind of advisory work —

where academics contribute directly to policy-relevant outputs — illustrates the cross-pollination between AESOP's scholarly community and European governance institutions.

European planning journals affiliated with AESOP's network have also begun featuring AI-focused scholarship. Research on topics like algorithmic bias in housing allocation, AI-assisted climate adaptation, and the use of machine learning in transport planning has appeared with increasing frequency. While AESOP does not publish its own journal in the way ACSP is linked to JPER, it influences the research agenda through conference themes, calls for papers, and the professional networks its members maintain. The growing volume of European planning research on AI parallels the trend seen in U.S. journals like JAPA and JPER, suggesting a field-wide shift in priorities.

5.3.2. Curriculum and Pedagogy

The question of how to teach planning students about AI has become a pressing concern for AESOP member schools, just as it has for ACSP programs in the U.S. European planning programs face a particular challenge: their students will practice in a regulatory environment shaped by the AI Act, meaning they need to understand not just how AI tools work but also how to navigate compliance requirements and evaluate ethical implications within a legally binding framework.

In 2025, AESOP organized a webinar series on "AI in Planning and Planning Education" — a direct effort to support faculty across Europe in adapting their teaching (AESOP, 2025). This series covered practical topics: how to introduce AI tools in studio courses, how to design assignments that develop critical thinking about algorithmic outputs, and how to teach data ethics within a European regulatory context. The webinar format made these resources accessible to educators across the continent, from well-resourced universities in Western Europe to smaller programs in Eastern and Southern Europe that may have less in-house expertise on AI.

A notable dimension of AESOP's educational efforts is its collaboration with ACSP. In 2025, the two organizations held a joint webinar series where educators from both sides of the Atlantic shared pedagogical strategies for teaching AI in planning. This transatlantic exchange allowed European educators to learn from ACSP's experience with the Task Force on Urban Planning and AI, while U.S. educators gained insight into how European programs are preparing students for the AI Act's compliance requirements. The joint sessions reflected a shared recognition that AI in planning education is not a national issue but a global one.

Some European planning schools have already started introducing new course content. Programs in the Netherlands, Germany, and the UK have added modules on planning analytics and AI or have embedded machine learning components into existing quantitative methods courses. Others are experimenting with AI in design

studios — having students use generative tools to produce planning scenarios and then critically evaluate the outputs for bias, feasibility, and alignment with community values. AESOP tracks and disseminates these innovations through its network, functioning as a clearinghouse in much the same way ACSP does for U.S. programs.

As in the U.S., European educators face the tension between embracing AI's potential and preserving the critical thinking skills that planning education is meant to develop. AESOP's approach has been to promote a balanced pedagogy — encouraging experimentation with AI tools while insisting that students learn to interrogate the assumptions, data, and outputs of those tools. Faculty development is also on AESOP's agenda; not all planning educators across Europe are trained in data science, and AESOP has begun exploring ways to support faculty upskilling through workshops and cross-departmental collaborations.

5.3.3. Collaboration and Dialogue with Practice

AESOP recognizes that academic research on AI must connect to the realities of planning practice if it is to have meaningful impact. To this end, the organization has pursued several avenues for bridging the gap between scholarship and professional application.

One important channel is AESOP's relationship with ECTP-CEU. Where ACSP and APA in the U.S. coordinate through overlapping memberships and joint conference programming, AESOP and ECTP-CEU share a similar dynamic in Europe. Academics from AESOP's network contribute to ECTP-CEU's discussions on professional standards and digital transformation, ensuring that research findings inform the guidance that reaches practicing planners. At the same time, practitioners' experiences — their frustrations with AI vendors, their struggles with data quality, their need for practical tools — feed back into the research questions AESOP scholars pursue.

AESOP has also engaged with broader European policy processes. Its members sit on advisory panels for EU-funded research programs and contribute to consultations on the EU's digital strategy for cities. This gives European planning academics a more direct channel to policy influence than their U.S. counterparts typically have — a reflection of the EU's more structured approach to stakeholder engagement. When the European Commission develops guidelines for AI use in public services, AESOP members are often among the experts consulted, bringing a perspective grounded in both planning theory and empirical research.

At the national level, AESOP's member schools frequently run outreach programs to local governments — applied research centers, planning clinics, and advisory services. Through AESOP's network, these programs share how they are helping municipalities navigate AI-related questions. One university's urban analytics lab might be assisting a city in evaluating the equity impacts of a predictive algorithm; another might be designing training workshops for municipal planning staff on responsible data use.

AESOP circulates these experiences through its conferences and communications, building a collective evidence base that strengthens the European planning community's capacity to engage with AI thoughtfully.

The synergy between AESOP and ECTP-CEU, combined with the regulatory backdrop of the EU and the global guidance provided by organizations like UN-Habitat (discussed in Chapter 6), means that European planning institutions are approaching AI governance from multiple directions simultaneously. The result is a denser, more connected response than the fragmented U.S. landscape. The connections are concrete rather than abstract: AESOP academics contribute to ECTP-CEU's discussions on professional standards, ECTP-CEU in turn works to translate the AI Act's requirements into practical guidance for national associations (ECTP-CEU, 2023), and AESOP members participate in the EU consultations that shape the digital strategy itself. Research, professional guidance, and legal requirements are not simply operating in parallel; each feeds into the others through these channels, even if the resulting coordination remains uneven.

6 UN-Habitat and Global Perspectives

The previous two chapters traced how American and European institutions have responded to AI. This chapter widens the lens to the global level, focusing on UN-Habitat, which was introduced briefly in Section 1.3. Unlike APA or the European Commission, UN-Habitat does not represent a single profession or regulatory bloc. It operates across very different urban contexts, from well-resourced capitals to cities with little internal digital capacity, and its authority is advisory rather than legal. That is a real limit. But its reach is also real. When UN-Habitat frames AI as a question of rights and sustainable development rather than efficiency, that framing shows up in national strategies, in municipal policy documents, and in the training materials of planning associations around the world. This chapter looks at two parts of UN-Habitat's work on AI in cities: the People-Centered Smart Cities framework, launched in 2022, and the more recent turn toward assessment, capacity-building, and pilot projects that put the framework's principles into practice.

6.1. People-Centred Smart Cities Framework

UN-Habitat launched the People-Centered Smart Cities initiative in 2022 with a coalition of partner agencies and urban think tanks (UN-Habitat, 2022). The starting point was a critique. The dominant smart city discourse had become corporate, tech-led, and detached from the residents it claimed to serve. The initiative tried to shift that center of gravity. Digital infrastructure and AI are treated not as ends in themselves but as means to outcomes planners have always cared about: equity, sustainability, inclusion, and the protection of basic rights. AI falls inside this framing as one category of technology among several, though one with outsized stakes given how fast municipal governments are procuring it. The framework is deliberately broad. That breadth lets it adapt to cities of very different sizes, income levels, and institutional capacities.

6.1.1. Principles and the Rights-Based Approach

The framework organizes itself around a compact set of principles: human rights, inclusion, participation, privacy, transparency, and accountability (UN-Habitat, 2022). Each of these has a direct line of application to AI systems in urban contexts. The rights principle asks whether a given AI deployment respects the protections enshrined in international instruments, including privacy, non-discrimination, and freedom from arbitrary decision-making. The inclusion principle requires that systems do not leave out the people who most need urban services, such as informal residents, migrants, or the elderly. Participation means communities have a role in deciding whether and how

AI is used in their neighborhoods, not just in being notified after the fact. Transparency and accountability ask that decision logic be explainable and that someone, ultimately a human, can be held responsible when things go wrong.

The principles themselves are not unique. Several appear in one form or another in the OECD and EU guidance discussed in Section 2.1. UN-Habitat's framing is distinct because it grounds AI governance in human rights explicitly. AI governance is treated as a rights question rather than a matter of good practice, and that has practical consequences. When a city procures an AI tool for allocating public housing or identifying code violations, a rights-based frame requires it to ask whether the tool can discriminate, whether residents have any recourse, and whether the procurement process itself was open. These are familiar questions in planning. The framework just makes them non-optional. It also echoes the values-driven counterweight to techno-centric narratives noted in Section 1.3, giving planning associations a vocabulary for pushing back against vendor-driven deployments. The question UN-Habitat returns to is whether an AI application actually helps the people in a city, especially those at the margins, and whether it upholds their rights.

The rights orientation also addresses something the broader AI governance discussion often sidesteps: the global distribution of AI harms. Cities in the Global South are often testing grounds for technologies developed and marketed from elsewhere, and the populations least able to contest a poorly designed system may be the ones most exposed to it. By insisting on rights as the organizing concept, UN-Habitat is trying to prevent a two-tier AI governance landscape in which wealthy cities get compliant tools and everyone else gets whatever is being sold.

6.1.2. International Guidelines and Co-Creation

Flowing from the principles is the more concrete instrument: the International Guidelines on People-Centered Smart Cities, which UN-Habitat has been developing with member states and partner networks (UN-Habitat, 2022). These guidelines are non-binding, following the pattern familiar from other UN instruments, but they are specific enough to be operational. They cover what good practice looks like across the lifecycle of a digital or AI-driven city project, from initial scoping and procurement through deployment, monitoring, and eventual replacement.

A recurring theme across the guidelines is co-creation. This is not co-creation in the vague sense the term sometimes carries in corporate materials, but something more structured: residents and affected communities should be involved in the design, monitoring, and evaluation of smart city tools that directly shape their lives. For AI specifically, that might mean neighborhood working groups that help define what a predictive model should actually optimize for, advisory boards with seats for populations most likely to be affected by an algorithmic decision, or community-led audits of a deployed system's outcomes. The guidelines also push for open data

practices that allow external scrutiny, for algorithms to be auditable rather than black-boxed, and for sustained investment in digital literacy so that residents can engage meaningfully with decisions they are asked to weigh in on.

This emphasis on co-creation ties directly to the participatory themes developed in the literature review, particularly Section 2.3.4 on equity and public inclusion, where the question of whether AI broadens or narrows the circle of participation was left open. UN-Habitat's answer is that broadening is possible but not automatic. It requires intentional design, consistent resourcing, and a willingness to treat residents as partners rather than endpoints of consultation. Planning associations that adopt portions of the guidelines, particularly in the Global South, typically adapt the co-creation provisions to their local context, sometimes modifying them substantially to reflect existing participatory traditions and constraints.

6.1.3. Influence on City and National Policy

A framework without uptake is just a document. The more interesting question is where the People-Centered Smart Cities framework has actually traveled, and the evidence suggests its influence is uneven but real. In countries with limited internal capacity to develop their own AI governance from scratch, the guidelines have been adopted more or less wholesale, offering a ready-made vocabulary and set of commitments that national planning bodies and municipal governments can point to when they confront a new AI procurement decision. Several national planning associations across the Global South have incorporated elements of the framework into their own policy positions, sometimes explicitly and sometimes by quiet borrowing.

In Europe and North America, where institutions already have their own frameworks, the influence is less about substitution and more about reinforcement. UN-Habitat's stance gives European planners additional backing when they argue that the AI Act's compliance requirements should be read through a rights lens rather than a purely technical one. It also provides U.S. planners, who work without equivalent legislative scaffolding, a global reference point when advocating for equity considerations in AI use. That dynamic helps explain the informal ties APA maintains with UN-Habitat and its participation in forums like the World Urban Forum, noted in Section 4.1.3. The framework functions here as moral weight rather than legal force, but moral weight is not nothing: it changes what it is respectable to argue, and in professional contexts that can be decisive.

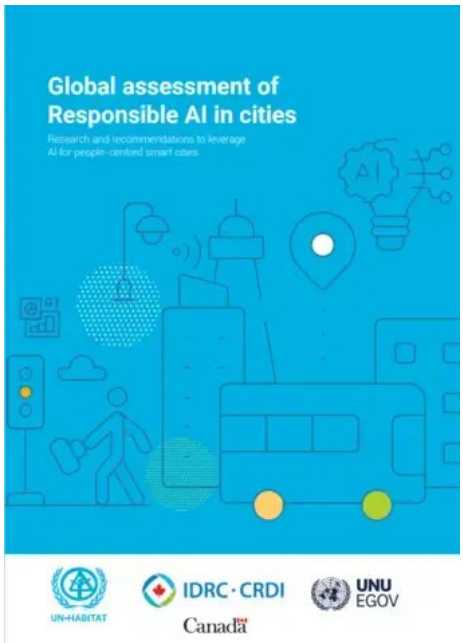
There are limits worth naming. UN-Habitat is an intergovernmental agency, and its guidance reflects the compromises that come with that. Some critics argue the principles are broad enough to be endorsed by almost anyone, which can make them hard to operationalize in a politically contested situation. The framework also depends heavily on member states and cities choosing to engage with it, and not all do. Where political will is absent, the guidelines can sit on a shelf even in cities that have formally

endorsed them. Still, the fact that such a framework exists and has a recognized institutional home matters. It provides a reference point that did not exist a decade ago, and over time the weight of reference points tends to accumulate.

6.2. Global Assessments, Capacity-Building and Pilot Projects

If the People-Centered Smart Cities framework represents UN-Habitat's normative contribution, the work examined in this section represents its empirical and operational contribution. Since 2023, UN-Habitat has moved increasingly from articulating principles toward diagnosing what is happening on the ground, equipping city officials with tools to act on the principles, and running pilot projects that test whether the theory actually holds up in practice. Taken together, this second track is where the framework meets the messy realities of municipal AI adoption, and where its gaps and contradictions become most visible.

6.2.1. Global Assessment of Responsible AI in Cities



"Figure 10: Global Assessment of Responsible AI in Cities. Source: UN Habitat"

The 2024 Global Assessment of Responsible AI in Cities, produced by UN-Habitat with partner research institutions, is the most comprehensive stock-take to date of how municipal governments around the world are using AI and what governance measures they have in place (UN-Habitat, 2024). The report surveyed cities across every region, drawing on submissions from municipal officials, civil society organizations, and academic researchers. Its findings are not uniformly optimistic.

One of the clearest patterns is a mismatch between awareness and capacity. Most city leaders reported that AI was on their agenda, and many had already deployed or were considering deploying AI systems for services such as transport management, utility forecasting, or administrative permitting. Yet far fewer reported having internal staff capable of meaningfully evaluating those systems. A mid-sized city might be pitched a crime prediction tool by a vendor and recognize the political appeal, but the planning and IT teams responsible for vetting it may have neither the data science expertise nor the ethical frameworks to scrutinize the model behind the interface. This mismatch between awareness and capacity is one of the report's central concerns, and it is particularly acute outside the world's wealthiest municipalities.

The assessment balances these concerns with concrete successes. Some cities used AI to plan more efficient waste-collection routes, cutting fuel and labor. Others used predictive maintenance to catch failing water pipes early, before they turned into service outages in lower-income neighborhoods. But the report doesn't shy away from failures either. In some cities, officials had to pull AI tools after their effects on particular neighborhoods turned out to be lopsided. In others, an algorithm meant to allocate public services just relearned the old pattern of underinvestment — sending resources where they had always gone and skipping the places that had always been skipped, because that was what the training data showed it.

6.2.2. Toolkits and Capacity-Building for Local Officials

The Global Assessment's findings have shaped UN-Habitat's subsequent work, which has moved strongly toward practical capacity-building. The logic is straightforward. Principles and reports are useful, but what city officials most urgently need are concrete resources they can use when a vendor walks into their office with a pitch deck or when a council member asks whether an AI tool would solve a policy problem.

In response, UN-Habitat and its partners have developed a growing library of toolkits and training programs aimed at local officials. These cover terrain that had previously been scattered across academic literature and consultancy reports. There are materials on responsible procurement, including how to write tender requirements that demand transparency and auditability from vendors rather than accepting proprietary opacity. There are guides on community involvement in tech projects, drawing on participatory planning traditions but adapted to AI's particular challenges around explainability. There are templates for oversight mechanisms, such as model documentation for deployed AI systems, municipal ethics committees, and impact assessment frameworks similar to the environmental reviews planners already conduct.

The audience is intentionally broad, but UN-Habitat has paid particular attention to cities with limited resources. Some of the toolkits are explicitly designed to be usable by small municipal teams without dedicated data scientists, and several national planning associations in the Global South have partnered with UN-Habitat to adapt and localize the materials, translating them and adjusting examples to reflect regional contexts. This pattern echoes the capacity-building theme that runs through several chapters of this thesis, including the emphasis ACSP and AESOP place on upskilling in Chapters 4 and 5. UN-Habitat's contribution here is to extend that logic to the parts of the world where professional training infrastructure is thinnest, which is also where the governance stakes are in some ways highest.

6.2.3. Pilot Projects and Knowledge Exchange

The third strand of UN-Habitat's operational work involves direct pilot projects carried out with local governments, universities, and occasionally private technology partners. These pilots test whether the principles articulated in the framework can actually survive contact with real urban problems, and they function simultaneously as case studies that can be shared back through UN-Habitat's networks.

The portfolio spans a range of urban challenges. One strand uses AI to map informal settlements from aerial imagery. In Durban, South Africa, a UN-Habitat team developed a tool called BEAM that automatically identifies building footprints in settlements where the built environment changes faster than manual surveys can track, reducing the time needed to map the entire city from an estimated 1,320 working days to roughly 72 hours. The resulting maps allow planners to direct infrastructure and services toward these areas rather than responding only after informal growth has already taken hold (Tonnarelli & Mora, 2025). Another applies machine learning to climate vulnerability data, helping cities prioritize adaptation investments in neighborhoods most exposed to flooding or heat stress. One documented UN-Habitat pilot focused on urban heat islands, where an AI model proposed greening interventions, but crucially the pilot did not stop at the model output. A community workshop followed, where residents reviewed the model's suggestions and adjusted them based on local knowledge about shade patterns, informal uses of public space, and priorities the data had not captured (UN-Habitat, 2024). That combination of algorithmic analysis and resident interpretation is what the framework means, in practice, by people-centered AI.

These pilots serve two purposes at once. They produce local outcomes, which matter to the communities involved, and they produce transferable knowledge, which matters to the wider network of cities and planning bodies that UN-Habitat convenes. Dissemination happens through several channels, including the Urban Agenda Platform, regional UN-Habitat offices, and most visibly through the World Urban Forum, which brings together city leaders and planners from around the world. Failures are shared alongside successes, which is unusual in institutional reporting and a deliberate choice aimed at shortening learning curves. A municipal planner from one city can hear directly from a counterpart in another about what went wrong in a predictive policing pilot, and adjust their own thinking before committing to a similar

procurement. That kind of learning is not available through vendor channels, and it is

one of the more distinctive contributions UN-Habitat makes to the governance landscape.

Taken together, the assessment, the toolkits, and the pilots give UN-Habitat something the framework alone could not: a working loop between diagnosis, response, and evidence. UN-Habitat's global perspective puts a justice-oriented narrative at the center of AI in urban planning, and it insists that every city, regardless of income, deserves guidance for navigating this terrain. Its authority remains soft. That is a limit, but it has an upside too. UN-Habitat cannot compel anyone to do anything, but it also can't easily be captured by vendors or governments, and it operates on the long horizons that planning depends on. Chapter 7 returns to these themes when it considers what American and European planning institutions can learn from the global perspective UN-Habitat offers, and where the limits of that perspective lie.

7 Comparative Analysis and Discussion

Having detailed the U.S. response (Chapter 4), the European response (Chapter 5), and the global context provided by UN-Habitat (Chapter 6), this chapter turns to a direct comparison and discussion of what these approaches mean. The goal of this chapter is to tease out how different institutional roles and cultural-political contexts shape the governance of AI in urban planning. Key dimensions are examined including: the nature of institutions (professional associations vs. governmental bodies vs. international agencies), the ethical and policy principles espoused, the actual influence on planning practice, and the gaps or challenges that remain. Finally, lessons are theorized that might be shared across the Atlantic and revisit the core research question: in what ways are planning institutions guiding AI's integration into planning, and what strategies are emerging from their efforts?

7.1. Institutional Role and Approach: Professional Associations vs. Public Bodies

One striking difference in these two contexts is who is taking the lead and what powers they hold. In the United States, the heavy lifting is done by non-governmental professional and academic associations (APA and ACSP). These organizations rely on persuasion, guidance, and knowledge dissemination – they cannot mandate behavior, but they can strongly influence it by setting professional norms. The tone is collegial and service-oriented where APA and ACSP offer resources and encouragement, trusting that planners and educators will voluntarily rise to the challenge of AI. The assumption underlying this is that the planning profession values ethical practice and will self-regulate if given the right tools and understanding (which historically has been true to an extent, as seen in how planners incorporate sustainability or social equity over time largely through normative guidance).

In contrast, the European approach is anchored by governmental and regulatory bodies – the European Commission and Parliament through the AI Act, and by extension national governments that will enforce it. These institutions operate with legal authority as they can require or prohibit certain actions. This doesn't mean there isn't voluntary initiative in Europe (there is, through ECTP-CEU, AESOP, etc.), but the presence of legislation creates a baseline of expected behavior. If APA's mode is self-regulation through ethics, the EU's mode is formal regulation through law. Each type of institution naturally leverages its strengths. APA publishes a trend report (because that's what a professional body can do), whereas the EU enacts a law (because only a governmental entity can do that).

This difference has deep historical and cultural roots. The U.S., with its strong tradition of professional autonomy and skepticism of top-down mandates (especially in planning, which is mostly locally controlled), finds it more palatable to handle new issues like AI through professional guidance. Planning associations in the U.S. often pride themselves on their Codes of Ethics and certification as a way to maintain standards without needing government rules. Europe, with a tradition of public welfare safeguards and precautionary regulation (especially visible in areas like data privacy with GDPR), is more comfortable setting binding rules to steer technology in alignment with public values.

An interesting manifestation of these differing approaches is how they conceive accountability. In the U.S. context, if an AI planning tool causes a controversy (say, a biased outcome in a city plan), accountability would likely be dealt with through public outcry, possibly litigation, or internal professional review (an APA ethics complaint perhaps, or just reputational damage to the planners involved). In Europe, under the AI Act, accountability is more codified. There could be investigations by regulatory bodies, fines, or orders to withdraw a system, in addition to any public or professional fallout. The *system* of accountability is thus more institutionalized in Europe.

However, both approaches aim for the same outcome, AI that is used responsibly to improve planning outcomes. Common goals do exist – transparency, avoidance of bias, public interest – but they are pursued via different mechanisms. APA might issue a reminder that bias is unethical (American Planning Association, 2024), while the EU might require a bias risk assessment in the deployment of an AI system (European Parliament, 2024). One is a moral suasion, while the other is a legal requirement.

Another difference in institutional approach is scale and scope. APA and ACSP focus purely on planning and planners. The EU's approach to AI is general rather than sector-specific, since the AI Act is a horizontal framework that regulates AI by risk category across all sectors, from health and finance to transport, rather than a rulebook written for any single profession. It was not designed with planning in mind, even though it clearly reaches planning. Its framework character matters in a second sense as well, because the Act sets out a structure of risk categories and obligations that still depends on harmonized technical standards and national implementation to become fully operational, which means planners are left not with a finished set of planning-specific rules but with a broad regulatory scaffold they must interpret and apply to their own work. UN-Habitat's approach is broad as well – covering "cities" which includes many aspects beyond the narrow planning profession (e.g., governance, infrastructure, etc.). So, the planning-specific bodies (APA, ACSP, ECTP-CEU, AESOP) often serve as translators or advocates within those broader frameworks to ensure the nuances of planning are addressed.

In the interplay between these institutions, one can see that collaboration could be beneficial. The U.S. could benefit from some formal standards (it's been noted by observers that voluntary guidelines sometimes aren't enough to change behavior). Conversely, European professional networks might take a page from APA's book in

terms of engaging practitioners more directly with training and discussion – not relying only on compliance. In fact, that is happening. ECTP-CEU and national associations are stepping up education precisely because they know a law alone

doesn't implement itself; planners need to understand and internalize it. And in the U.S., there are starting discussions in policy circles about whether certain AI applications (like facial recognition or predictive policing) should be regulated or banned by law – which shows a possible convergence in the future.

7.2. Ethical and Policy Alignment: Similarities and Differences

When comparing the ethical principles and policy themes emphasized on both sides, we find a strong overlap in values, albeit with some differences in emphasis and language. Both contexts recognize the critical issues of bias and fairness, transparency, accountability, privacy, and public engagement — the same themes identified in the literature review (Section 2.3). These have become near-universal concerns in AI governance (Kitchin, 2021; O'Neil, 2016).

In the U.S. discourse (APA/ACSP), one sees a lot of mention of equity and justice. For instance, APA's code and writings stress serving the public interest and not harming disadvantaged groups (American Institute of Certified Planners, 2021; Sanchez et al., 2024). This aligns with the broader American planning ethics focus on social justice that dates back to advocacy planning movements. So APA frames AI ethics often around avoiding "digital redlining" or ensuring AI doesn't perpetuate segregation or inequality – basically fitting AI into historic equity commitments of planners.

The European policy language also cares about fairness but often uses the term "fundamental rights compliance." For example, the AI Act and EU guidelines speak of avoiding discrimination as a legal imperative, echoing human rights language (European Commission, 2019, 2024). Privacy is extremely prominent in Europe (perhaps more so than in U.S. planning discussions) because of GDPR and a cultural weight on data protection – a planner in Europe must be very careful with personal data in a way a U.S. planner might not be mandated to (though ethically they're encouraged). So, while both talk about not reinforcing bias, a European document might explicitly reference legal non-discrimination mandates and rights to explanation, whereas a U.S. one might talk about "unintentional bias" and the moral duty to avoid it, referencing cases where things went wrong (like the predictive policing example in New York mentioned in the Introduction).

Transparency and explainability are championed by both, but again the framing differs slightly. APA encourages transparency as a practice – advising planners to be open with the public about how decisions are made, to maintain trust (American Planning Association, 2024). The EU, via the AI Act, could *require* that AI systems reveal certain information (like giving a user the right to an explanation of an algorithmic decision). The difference is voluntary best practice vs. legal right. But in

an ideal scenario, they converge. Planners should explain anyway, law or not, and likely will do so more as this becomes the norm globally.

Human oversight is another principle both share. The phrase "human-in-the-loop" shows up in both U.S. and European discussions. APA's guidance implies planners should not cede their decision authority entirely to a machine (Wasserman & Flaxman, 2023). The EU's Act literally requires human oversight for high-risk AI. UN-Habitat's guidance similarly says tech should augment, not replace, human decision-making (UN-Habitat, 2022). So we see a very aligned stance that AI should be a tool, not an autonomous decider in planning. This consensus likely comes from a shared understanding of planning as inherently a human, democratic process – you can't automate public deliberation or value judgments fully.

One difference in ethical discourse is the scale of consequences considered. European/global discussions often connect AI ethics to broader existential or societal considerations – for instance, aligning with sustainable development goals, climate action, and long-term resilience. UN-Habitat explicitly ties AI use to sustainability and inclusion outcomes, and as Chapter 6 showed, frames the question as a matter of rights rather than good practice. APA's discussions, while acknowledging broader impacts, often read as more immediate and operational (e.g., don't violate ethics code tenets, be fair to communities). It's subtle, but the global view tends to incorporate, say, climate justice or global inequality into the narrative of AI, whereas the U.S. view is typically more focused on national/local community impacts. This might be because UN-Habitat's remit is global development, and EU's is union-wide rights, whereas APA's is narrower. However, ACSP's academic angle does bring in big-picture thinking, referencing societal implications akin to those global concerns (Allam & Jones, 2020 emphasize capacity and structural issues, for example).

Another aligned area is capacity-building as an ethical imperative. Both contexts realize that without educating planners and stakeholders, principles remain paper tigers. ACSP and AESOP, along with UN-Habitat, all emphasize training and literacy as key. This is seen as an ethical need too – if people don't understand AI, power imbalances worsen. So there's a shared drive to democratize knowledge of AI.

In terms of formal policy alignment, the U.S. *AI Bill of Rights* blueprint (OSTP, 2022) closely resembles European approaches. It articulates principles related to safe and effective systems, protections against algorithmic discrimination, data privacy, transparency and explanation, and the preservation of human alternatives. While neither APA nor ACSP has issued an explicit "AI policy for planners," their existing guidance, taken together, addresses many of the same concerns. This pattern suggests a convergence at the level of underlying principles, even if implementation mechanisms differ substantially.

Both the United States and Europe are grappling with a shared fundamental challenge: how to harness the benefits of AI, such as efficiency, predictive capacity, and creativity,

while minimizing risks related to bias, opacity, and the erosion of human control. The U.S. approach relies primarily on professional ethics and innovation, whereas the European approach emphasizes legal frameworks and structured public programs. Despite these institutional differences, both converge on a common ethical position, namely that AI in planning should be fair, transparent, accountable, and oriented toward the public good. This convergence is significant, as it suggests that planning institutions across different governance systems are aligning around a shared set of core principles for AI governance. Such alignment provides a foundation for meaningful international dialogue, including exchanges between APA and European counterparts and coordination through UN forums.

7.3. Influence on Practice and Identification of Gaps

Ultimately, the value of these institutional responses will be judged by their impact on actual planning practice. Are planners in city departments or consulting firms doing things differently because of these initiatives? And where are the gaps — things not yet being addressed?

In this section, I present six gaps derived from the analysis in both Europe and the U.S. In the U.S., since APA and ACSP's efforts are relatively recent, we can observe some early influence. Many practicing planners are now *aware* of AI as a topic due to APA's communications. Some larger planning agencies have begun issuing internal guidelines or cautionary notes on algorithm use, reflecting the broader environment APA contributes to. Some planning consulting firms have begun showcasing their AI capabilities, but also note ethical commitments, reflecting APA's influence. Planning educators are indeed adding content on AI, meaning new graduates are coming out with at least some familiarity – a direct impact of ACSP community discussions. However, a gap in the U.S. appears to be the *lack of formal oversight or evaluation*. If a small city buys an AI tool, there's no standard process to evaluate it other than whatever the vendor provides or the city's own due diligence. APA's materials suggest questions to ask, but there's no certification or independent audit mechanism specifically endorsed by planning institutions.

This hints at a larger gap, specifically, the capacity to vet AI tools. Neither APA nor ACSP can certify an AI tool as "planner-approved" (nor have they tried to yet). Planners often lack deep technical expertise to interrogate AI models. So there's a risk that despite all the guidance, some flawed or biased AI system could slip into practice and do damage before being caught — a scenario O'Neil (2016) warned about broadly. The European side might be better positioned here, since the AI Act will require conformity assessments for high-risk AI. But even that may rely on technical standards not yet fully developed for all contexts. Planning-specific assessment protocols are not yet defined. So both contexts share a gap in practical assessment frameworks tailored to planning. There isn't, for example, a checklist specifically for a "land use suitability

AI" to ensure it's fair. Planning institutions could collaborate with data scientists to create such tailored evaluation guides.

Another gap is resources for smaller communities. Large cities or well-funded agencies might pilot AI with expert help; smaller towns might be left out or, worse, adopt tools without proper understanding. APA, through its state chapters and perhaps the Planning Advisory Service, could aim to reach those smaller jurisdictions. UN-Habitat's 2024 Global Assessment, discussed in Chapter 6, makes the same point at a global scale: city leaders are aware of AI, but most lack the in-house capacity to evaluate the tools being sold to them. The risk of a new digital divide is real, where advanced cities benefit and others lag. That's a gap in equity on a global scale. Efforts like UN-Habitat's toolkit try to close it, but implementation is tough.

We should consider interdisciplinary integration as well. Are planning institutions bringing in AI experts sufficiently? ACSP's task force includes planners primarily; maybe they consult computer scientists, but planning orgs are still dominated by planners (naturally). There might be a gap in cross-disciplinary collaboration. Conversely, tech industry often creates urban-focused AI without planners at the table. Bridging that gap is something institutions could do more. For instance, APA could formally partner with an AI professional society to co-develop ethical standards for urban AI, ensuring both technical and planning perspectives. Some early partnership events have happened, but more structural collaboration could be beneficial.

On the ground, another measure of influence is whether planners feel "better prepared" for AI. Surveys (like Ye et al., 2024) indicate many educators and planners are cautious – which is appropriate, but it might also mean they aren't sure how to proceed. The cautious optimism in literature suggests a lot of unanswered questions remain in practice. For example, if a planning director is given a budget and asked, "shall we invest in an AI traffic model or hire another human analyst?" do they have enough guidance to make that call? Possibly not yet; cost-benefit analyses and case evidence are still emerging.

A very practical gap both contexts face is the lack of case evaluations: we have anecdotal stories of AI use, but few comprehensive evaluations of outcomes. Did using AI in city X's master plan genuinely produce a better plan or more equitable outcomes? There's sparse empirical evidence so far, as this is all new. Planning institutions could spearhead or fund studies to fill this gap – this might fall to the academic side (ACSP/AESOP) to do, with APA/ECTP amplifying results to practice.

Finally, there's the issue of public trust and engagement. Are the public being involved in shaping how AI is used in planning? Institutions have voiced the need (e.g., UN-Habitat's emphasis on participation), but in practice, examples are few. One might consider it a gap that we do not yet have good models for citizen oversight or citizen representation in AI governance in planning. Perhaps in the future, planning commissions or citizen advisory boards will include algorithm auditors or community representatives specifically to review AI tools. That concept hasn't materialized yet widely. Planning institutions could advocate for such governance innovations at the local level (like recommending a citizen data ethics committee for city planning offices).

7.4. Transatlantic Lessons and Opportunities for Mutual Learning

Comparing the U.S. and European experiences with AI in planning is not just an academic exercise; it highlights opportunities where each side can learn from the other. In many ways, these differing approaches are complementary and suggest a more balanced strategy could be achieved by borrowing the best of both worlds.

The most obvious thing the U.S. can learn from Europe is the value of a proactive regulatory stance. While the U.S. planning community has moved swiftly in producing ethical guidance, there remains the possibility that without enforceable standards, some jurisdictions might ignore best practices, leading to harm. The EU's AI Act demonstrates that it is possible to create overarching rules for AI; the U.S. might consider at least some regulatory or legislative frameworks for particularly sensitive planning applications (e.g., laws that ensure transparency for algorithms used in public decisions). The U.S. could also learn from Europe's integration of AI goals with broader societal goals. For instance, the idea of explicitly tying AI use to equity outcomes and measuring that (common in EU-funded projects) could be more emphasized in the U.S., where sometimes innovation is pursued for efficiency without always quantifying equity impacts. In short, the EU approach reminds the U.S. that principles sometimes need teeth and alignment with public policy can drive more consistent and accountable outcomes.

Additionally, European initiatives show the value of structured funding for innovation. The APA or state/federal agencies in the U.S. might push for something akin to Horizon Europe's city pilots – perhaps federal grants that encourage cities to experiment with AI for planning under certain ethical guidelines. The U.S. has started

doing this in the smart cities realm, but a more explicit program focusing on planning and AI could accelerate learning and provide models.

Conversely, Europe can learn from the bottom-up, practitioner-centric approach that has emerged in the United States, where organizations such as the APA and ACSP have cultivated a rich ecosystem of practical knowledge, experimentation, and professional dialogue around AI. While Europe's regulatory leadership provides important safeguards, there is a risk that top-down mandates may be experienced primarily as compliance exercises unless they are accompanied by sustained professional engagement. In this sense, European planning bodies and professional associations could complement regulatory frameworks by fostering more practice-oriented conversations through webinars, applied guidance, and informal knowledge-sharing among planners. Regulations may set the floor for acceptable practice, but professional enthusiasm and collective learning raise the ceiling of what is possible. The U.S. experience suggests that planners are willing to take initiative in adopting and critically evaluating new tools when they are supported by accessible resources and professional communities. European institutions could strengthen AI governance not only by setting rules, but by actively convening practitioners to interpret, apply, and move beyond those rules in practice.

The U.S. also demonstrates a relatively strong integration between professional practice and academia through the relationship between APA and ACSP. Europe has a comparable structure in AESOP and ECTP-CEU, but the interaction between academic research and professional guidance could be further strengthened. Existing collaborations, such as joint ACSP–AESOP conference sessions, suggest a foundation for deeper exchange, particularly around shared challenges related to AI in planning education and practice.

More broadly, both contexts would benefit from stronger mechanisms for international knowledge sharing. Issues such as data bias, transparency, and accountability in AI systems are not geographically bounded, and lessons from one planning context are often transferable to others. Institutions such as UN-Habitat already provide a platform for this kind of cross-border learning. Beyond that platform role, Chapter 6 showed that UN-Habitat offers a third pole worth borrowing from: a rights-based framing of AI governance, and a deliberate focus on building capacity in under-resourced cities. Both are things the U.S. and European approaches tend to assume rather than build. Closer engagement between global, European, and U.S. planning bodies could help translate broad ethical principles into more practice-oriented guidance.

One area where this cooperation could be especially valuable is the systematic documentation of AI use in planning practice. A shared repository of case studies—focused on both successes and limitations—would help move discussions of "responsible AI" beyond abstract principles toward concrete, evidence-based

understanding. Finally, differences in planning culture suggest additional opportunities for mutual learning. U.S. planning's emphasis on public engagement offers useful lessons for making AI-related decisions more transparent and publicly accountable, while European experiences with institutionalized oversight and ethics mechanisms could inform more structured approaches in the U.S. Together, these perspectives point toward a more balanced model of AI governance that combines regulatory clarity with professional and public engagement.

8 Conclusion

This thesis examined how planning institutions in the United States, Europe, and the international arena are responding to the rapid expansion of artificial intelligence, focusing on governance rather than technology itself.

Chapter 1 established the comparative framework, centered on a transatlantic comparison between United States professional and academic bodies (APA and ACSP) and European regulatory and professional networks (the European Commission, ECTP-CEU, and AESOP), with UN-Habitat examined as the global normative context. Chapter 2 reviewed the literature across four thematic areas — policy and regulatory frameworks, AI applications in planning, ethics and equity concerns, and planning association publications — finding broad consensus that AI holds potential for planning but demands critical oversight. Chapter 3 described the qualitative, document-based research design used to analyze institutional responses. Chapter 4 examined the U.S. context, where APA and ACSP have responded through knowledge dissemination, ethical guidance, professional development, and academic research promotion — a largely voluntary, practitioner-driven approach. Chapter 5 turned to Europe, where the EU AI Act introduces binding requirements alongside softer guidance from ECTP-CEU and AESOP, creating a denser but more prescriptive governance environment. Chapter 6 considered UN-Habitat's global role in advocating for people-centered smart cities and building governance capacity in municipalities worldwide. Chapter 7 compared these approaches directly, identifying shared ethical commitments alongside significant differences in enforcement mechanisms, institutional scale, and the relationship between professional norms and legal mandates.

The institutional responses examined in these chapters do not reflect neutral assessments of what AI can do. They reflect deeply rooted regulatory cultures and professional traditions. Planning organizations on both sides of the Atlantic acknowledge AI's transformative capacity, but their approaches expose real tensions between innovation, accountability, and the public interest.

In the United States, planning institutions mostly treat AI as a professional responsibility, something to be managed through ethical guidance, education, and the discretion of individual planners. APA and ACSP put their weight behind capacity-building and interpretive ethics, casting planners as the responsible mediator between complex technologies and public decisions. This reliance on voluntary norms and professional self-regulation creates a real problem. Without binding requirements, AI governance ends up depending on local institutional capacity, individual expertise, and political context — and that dependency is uneven, as Section 7.3 explored. Well-

resourced agencies can build robust safeguards. Smaller or underfunded jurisdictions can adopt AI tools without much scrutiny, transparency, or community oversight at all. Where that happens, ethical guidance is functioning more as aspiration than enforcement (American Institute of Certified Planners, 2021; OSTP, 2022).

European and international institutions take a more interventionist line, building AI governance into regulatory and policy frameworks that put fundamental rights, transparency, and human oversight in writing. The EU's risk-based approach sets clearer limits on how public authorities can deploy AI in areas affecting access to services, infrastructure, or civic participation (European Commission, 2024). AESOP and ECTP-CEU reinforce these norms through professional discourse and knowledge exchange, and UN-Habitat ties AI governance to questions of equity and sustainable development. Even so, the model has its own weaknesses. Compliance with a regulation does not automatically produce meaningful accountability, and rules tend to harden into checklists if planning institutions cannot translate them into everyday practice. Regulation can draw a boundary; it cannot, on its own, resolve who decides, who interprets, and who answers politically for the decisions that get made inside that boundary.

Both contexts treat AI as a tool that should support human judgment rather than displace it, and on the surface this sounds like clear agreement. In practice it papers over a contradiction. Planning organizations promote AI for its efficiency, foresight, and capacity to broaden participation, while at the same time conceding that it can reproduce bias, hide its reasoning, and concentrate expertise in fewer hands. Telling planners to "remain in the loop" sounds reassuring, but the loop is shaped by resource constraints, political pressure to move quickly, and growing dependence on outside vendors. Those pressures push in one direction, which is toward greater trust in whatever the algorithm produces, even where the ethical concerns have not been resolved (Kitchin, 2021; Sanchez et al., 2024).

What this study challenges is the assumption that ethical principles or professional norms can govern AI in urban planning on their own. Awareness and training matter. They cannot substitute for institutional arrangements that decide who bears responsibility for AI-assisted decisions, how accountability gets exercised, and how affected communities can contest or shape those decisions. Without those structures, planning agencies risk using AI as a shortcut, letting complex models drive outcomes that never received real deliberation or transparency. Instead of democratizing planning, AI then ends up reinforcing technocratic tendencies, shifting authority away from public debate and toward opaque systems controlled by experts and external vendors.

The governance of AI in urban planning is, in the end, a struggle over what planning is for. Institutions can stretch their existing norms to make room for increasingly powerful technologies, or they can reassert planning's democratic and redistributive

The result is a denser, more connected response than the fragmented U.S. landscape. The connections are concrete and institutionally acknowledged: AESOP and ECTP-CEU formally describe their partnership as a synergy, with AESOP bringing academic expertise into ongoing professional debates (ECTP-CEU, 2023), while ECTP-CEU translates the AI Act's requirements into practical guidance for national associations (ECTP-CEU, 2023). Research, professional guidance, and legal requirements feed into one another through these channels rather than running in parallel.

Bibliography

Allam, Z., & Jones, D. (2020). On the nexus of AI and urban planning: Outcomes and implications. *Journal of Urban Technology, 27*(1–2), 59–75.

American Institute of Certified Planners. (2021). *AICP code of ethics and professional conduct (2021 update).* American Planning Association.

American Planning Association. (2020). *Planning with artificial intelligence* (PAS Report No. 604, by T. Sanchez). APA Publications.

American Planning Association. (2024). *Artificial intelligence and planning: KnowledgeBase collection.* APA Publications.

American Planning Association. (2025). *2025 trend report for planners: Shape the future.* APA Foresight.

Association of Collegiate Schools of Planning. (2024). *Annual conference sessions and newsletters on AI and planning education.* Association of Collegiate Schools of Planning.

Association of European Schools of Planning. (2025). *AI in planning and planning education: Webinar series.* AESOP.

Barcelona City Council. (2019). *Barcelona digital data charter.* Barcelona City Council.

Batty, M. (2018). *Inventing future cities.* MIT Press.

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

Chen, T., & Gascó-Hernández, M. (2024). Uncovering the results of AI chatbot use in the public sector: Evidence from US state governments. *Public Performance & Management Review*, 1–26.

European Commission. (2019). *Ethics guidelines for trustworthy AI.* High-Level Expert Group on Artificial Intelligence.

European Commission. (2024). *Proposal for a regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).* European Commission.

European Commission. (2025). *AI tools for cities: Supporting urban digital twins and simulation platforms* [Fact sheet]. Directorate-General for Communications Networks, Content and Technology.

European Council of Town Planners (ECTP-CEU). (2023). *Position paper on smart cities and AI in planning.* ECTP-CEU.

European Council of Spatial Planners (ECTP-CEU). (n.d.). *Synergies with other associations: AESOP*. Retrieved from https://ectp-ceu.eu/biennial/synergies-with-other-associations/aesop/

European Parliament. (2024). *Artificial Intelligence Act: Legislative text adopted by the European Parliament.* Publications Office of the EU.

European Transport Safety Council. (2020). *AI and road safety: Guidelines for autonomous vehicle integration.* ETSC.

Illinois General Assembly. (2019). *Artificial Intelligence Video Interview Act, Public Act 101-0260.* State of Illinois.

Jasanoff, S. (2005). *Designs on nature: Science and democracy in Europe and the United States.* Princeton University Press.

Journal of Planning Education and Research. (2024). Call for papers: Urban planning AI – Navigating challenges, seizing opportunities. *Journal of Planning Education and Research.*

Kitchin, R. (2021). Ethics and governance of autonomous urban systems: Ensuring responsible deployment of AI in cities. *Urban Planning, 6*(2), 141–151.

Ligmann-Zielinska, A., & Jankowski, P. (2014). Spatially explicit integrated uncertainty and sensitivity analysis of criteria weights in multicriteria land suitability evaluation.
Environmental Modelling & Software, 57, 235–247.

Malczewski, J., & Rinner, C. (2015). *Multicriteria decision analysis in geographic information science.* Springer.

Mills, E. (2025). *An introduction to Planning and AI* [Advice Note]. Royal Town Planning Institute.

Moreno, C., Allam, Z., Chabaud, D., & Jones, D. (2020). Introducing the "15-minute city": Sustainability, resilience and place identity in future post-pandemic cities. *Smart Cities, 4*(1), 93–111.

Naik, N., Philipoom, J., Raskar, R., & Hidalgo, C. (2014). Streetscore: Predicting the perceived safety of one million streetscapes. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 793–799.

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework* (Version 1.0). U.S. Department of Commerce.

New York City Council. (2021). *Local Law 144 of 2021: Automated employment decision tools.* New York City Council.

Office of Science and Technology Policy. (2022). *Blueprint for an AI bill of rights: Making automated systems work for the American people.* The White House.

Organisation for Economic Co-operation and Development. (2019). *OECD principles on artificial intelligence.* OECD Publishing.

Organisation for Economic Co-operation and Development. (2020). *OECD AI policy observatory report.* OECD Publishing.

O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy.* Crown.

Reades, J., De Souza, J., & Hubbard, P. (2019). Understanding urban gentrification through machine learning: Predicting neighbourhood change in London. *Urban Studies, 56*(5), 922–942.

Sanchez, T. W. (2025). *Artificial intelligence for urban planning.* Routledge.

Sanchez, T. W., Brenman, M., & Ye, X. (2024). The ethical concerns of artificial intelligence in urban planning. *Journal of the American Planning Association, 91*(2), 294–307.

Sanchez, T. W., Shumway, H., Gordner, T., & Lim, T. (2023). The prospects of artificial intelligence in urban planning. *International Journal of Urban Sciences, 27*(2), 179–194.

Tonnarelli, F., & Mora, L. (2025). Responsible AI for cities: A case study of GeoAI in African informal settlements. *Journal of Urban Technology, 32*(3), 111–137.

UN-Habitat. (2022). *People-centered smart cities initiative.* United Nations Human Settlements Programme.

UN-Habitat. (2024). *Global assessment of responsible AI in cities: Research and recommendations to leverage AI for people-centered smart cities.* United Nations Human Settlements Programme.

United Nations Educational, Scientific and Cultural Organization. (2021). *Recommendation on the ethics of artificial intelligence.* UNESCO Publishing.

United States Department of Transportation. (2020). *Preparing for the future of transportation: Automated vehicles 4.0.* U.S. Department of Transportation.

University of Florida Department of Urban and Regional Planning. (2024, November 25). *ACSP Seattle 2024: UF URP recap.* UF DCP News.

Veale, M., & Borgesius, F. Z. (2021). Demystifying the draft EU Artificial Intelligence Act. *Computer Law Review International, 22*(4), 97–112.

Walker, R. (2024). How planners are using AI in practice. *Planning, 90*(3), 22–27.

Wasserman, S., & Flaxman, S. (2023). Artificial intelligence and planning practice (PAS Memo). *Planning Advisory Service.* American Planning Association.

Xie, Y., & Batty, M. (2003). Calibrating cellular automata models for urban simulation. *Environment and Planning B: Planning and Design, 30*(5), 681–698.

Ye, X., Zhai, W., Du, J., Van Zandt, S., & Ye, Y. (2024). Authentic or artificial intelligence? Faculty's perspectives on ChatGPT's impact on U.S. urban planning Ph.D. programs. *Frontiers of Urban and Rural Planning, 2,* Article 23.

Yeh, A. G. O., & Shi, X. (2021). Spatial analysis for urban planning using AI and big data. *Computers, Environment and Urban Systems, 89*, Article 101681.

Zhang, F., Zhou, B., Liu, L., Liu, Y., Fung, H. H., Lin, H., & Ratti, C. (2018). Measuring human perceptions of a large-scale urban region using machine learning. *Landscape and Urban Planning, 180*, 148–160.

List of Figures

Figure 1: Comparative Framework of Planning Institutions by Governance Context. Source: Author.

Figure 2: American Planning Association (APA) Logo. Source: APA.

Figure 3: Association of Collegiate Schools of Planning (ACSP) Logo. Source: ACSP.

Figure 4: European Commission Logo. Source: European Commission.

Figure 5: ECTP-CEU Logo. Source: ECTP-CEU.

Figure 6: AESOP Logo. Source: AESOP.

Figure 7: UN-Habitat Logo. Source: UN-Habitat.

Figure 8: Generative AI Assistance in Thesis Composition. Source: Author.

Figure 9: EU AI Act Risk-Based Classification Framework. Source: European Parliament (2024).

Figure 10: UN-Habitat Global Assessment of Responsible AI in Cities. Source: UN-Habitat (2024).

Acknowledgements

I acknowledge the use of large language model tools, including ChatGPT and Claude, to support research and proofreading. All research design, analysis, interpretation, writing, and conclusions are my own.

I want to also thank Professor Davide Ponzini for his continued support throughout the the entire process.

Finally, I want to thank my parents and friends. Mom and Dad, I couldn't have embarked on this journey without you. Your unrelenting support and positivity were paramount for both the output of this thesis, as well as my success here in Italy.