

From Connected Infrastructure to Autonomous Governance: An AI-Driven Urban Decision Intelligence Framework for NEOM-Scale Smart Cities

CHOUDHRY BILAL MAZHAR

Digital Analytics and Artificial Intelligence Lead, ARASCO, Riyadh, Saudi Arabia

Abstract- Background: Smart cities are evolving from connected infrastructure models toward intelligent urban ecosystems capable of sensing, learning, predicting, simulating, and coordinating public-service decisions in near real time. NEOM-scale environments require more than 5G connectivity, IoT sensing, cloud platforms, and green data centers; they require a trusted decision intelligence layer that transforms continuous urban data into explainable governance actions. This paper proposes the NEOM Urban Decision Intelligence Framework (NUDIF), an AI-driven conceptual architecture that integrates connected infrastructure, interoperable data platforms, digital twins, multi-agent intelligence, governance dashboards, and sustainability optimization. The framework is organized into six interoperable layers spanning connected infrastructure, urban data management, analytics and digital twins, domain AI agents, decision intelligence, and autonomous governance execution. It explains how energy, mobility, water, environment, public safety, and public-service systems can be monitored and coordinated through specialized AI agents operating under transparent human oversight, where routine decisions may be automated, high-impact decisions require human approval, and strategic decisions remain under institutional control. The paper further outlines a four-stage implementation maturity path progressing from connected foundation and predictive intelligence to agentic coordination and controlled autonomy, and it defines governance controls, key performance indicators, and responsible AI safeguards covering data quality, model reliability, cybersecurity, privacy, explainability, and auditability. Representative indicators are mapped across energy, mobility, water, environment, public services, and governance domains to support measurable evaluation. The study contributes a practical model for future smart cities by linking digital infrastructure with accountable decision-making, urban resilience, sustainability alignment, and Vision 2030-oriented innovation, and it identifies validation directions for large-scale greenfield and brownfield urban environments.

Keywords: Artificial Intelligence, Autonomous Governance, Decision Intelligence, Digital Twin, NEOM, Smart Cities, Sustainability.

I. INTRODUCTION

The development of next-generation smart cities is moving beyond infrastructure connectivity toward cognitive, data-driven, and AI-enabled governance. NEOM and THE LINE represent a new model of urban development that prioritizes sustainability, renewable energy, advanced connectivity, livability, and technology-enabled services [1], [2]. NEOM's digital strategy also emphasizes trusted data, advanced digital infrastructure, cloud connectivity, and technology-enabled public services as foundations for future urban operating models [3]. However, the strategic value of these investments depends on the ability to convert continuous urban signals into decisions that are timely, explainable, auditable, and aligned with public objectives. This requires an integrated governance architecture in which sensing, data management, digital twins, analytics, AI agents, and human decision rights operate as one coordinated system.

This paper introduces an AI-driven Urban Decision Intelligence Framework for NEOM-scale smart cities. The objective is to show how 5G, IoT, edge computing, cloud services, urban data platforms, digital twins, AI analytics, and multi-agent systems can operate as one integrated governance architecture. Recent smart city research highlights AI-enabled decision making, data interoperability, transparency, human-AI collaboration, and smart urban governance as essential requirements for future urban ecosystems [4], [5], [17]–[20]. It also emphasizes IoT-enabled urban systems, fog computing, and edge computing as

foundational capabilities for real-time city operations [21]–[23]. The central research question is therefore: how can a NEOM-scale smart city progress from connected infrastructure to accountable autonomous governance without compromising privacy, trust, sustainability, or human oversight?

The contribution of this study is threefold. First, it defines NUDIF as a layered architecture that links physical infrastructure with decision intelligence and governance execution. Second, it explains the role of specialized urban AI agents in coordinating energy, mobility, water, environment, public safety, and citizen service domains. Third, it identifies responsible AI safeguards, interoperability requirements, and performance indicators that can guide future implementation and evaluation in large-scale smart city environments.

II. BACKGROUND AND THEORETICAL FRAMEWORK

Smart city scholarship has progressed through several distinct phases. Early contributions framed the smart city primarily as a technology-enabled environment in which networked infrastructure, human capital, and economic competitiveness reinforce one another [18]. Subsequent work broadened this view by identifying institutional, managerial, and policy factors as equally decisive success conditions, arguing that technology alone cannot explain why comparable investments produce very different urban outcomes [19]. A parallel stream examined the implications of continuous data capture for urban management and cautioned that real-time urbanism introduces new questions of accuracy, representation, and control that cities must address explicitly [17]. More recent governance-oriented research has consolidated these perspectives by treating smart urban governance, rather than smart technology, as the central unit of analysis [20]. Taken together, this trajectory indicates a steady movement from infrastructure deployment toward decision quality, accountability, and institutional capability.

The technological substrate that supports these ambitions has matured considerably. Internet of Things architectures established the sensing and communication foundations required for city-scale

observability [21], while fog and edge computing addressed the latency, bandwidth, and resilience constraints that arise when analytical workloads are concentrated exclusively in centralized cloud environments [22], [23]. These distributed paradigms are particularly relevant for NEOM-scale deployments, where mobility, safety, and utility functions require near-real-time responsiveness and cannot tolerate round-trip delays to remote data centers. In parallel, the convergence of artificial intelligence with connected infrastructure has given rise to artificial intelligence of things architectures and urban brain concepts that couple continuous sensing with learning and coordination capabilities [8].

Digital twin research provides the simulation foundation on which anticipatory governance depends. Foundational contributions established the digital twin as a mechanism for managing unpredictable and undesirable emergent behavior in complex systems [15], and later urban-focused work extended the concept from asset-level replication to city-scale representation [14]. Recent syntheses have mapped the convergence of the internet of things, artificial intelligence, and digital twins in smart city contexts, identifying thematic clusters and open research directions [7], while related studies have examined how digital twin and distributed ledger solutions can jointly advance sustainability objectives [6]. Evidence-synthesis reviews of data-driven sustainable cities further confirm that predictive and simulation capabilities are now central rather than peripheral to urban planning practice [16].

A newer and rapidly developing stream concerns agentic intelligence in urban settings. Research on agentic urban artificial intelligence argues that cities are moving beyond rule-based automation toward systems capable of goal-directed reasoning and adaptive coordination [11]. Work on language-model-powered urban agents describes emerging capabilities in perception, planning, and interaction across urban domains [12], and sensor-based urban computing studies demonstrate how distributed intelligence can be embedded within sustainable city administration [13]. These developments rest on established foundations in multi-agent systems, which formalize agent autonomy, negotiation, and coordination [24],

and on general artificial intelligence theory concerning rational agents operating in partially observable environments [25]. Applied reviews of artificial intelligence across smart city pillars confirm that domain-specific deployments are already widespread, though typically implemented as isolated solutions rather than as an integrated operating model [4].

Governance and accountability requirements have advanced in parallel with these technical capabilities. Global assessments of responsible artificial intelligence in cities document both the promise and the uneven readiness of municipal institutions [5], and complementary analyses emphasize inclusion as a precondition for advancing sustainable development objectives through urban artificial intelligence [10]. Normative frameworks provide the principles on which such practice rests, including ethical frameworks for beneficial artificial intelligence societies [26], trustworthy artificial intelligence guidelines [27], international ethics recommendations [28], and intergovernmental principles for trustworthy systems [29]. Operational translation of these principles is increasingly supported by risk management frameworks [30] and management system standards that define auditable organizational controls [31]. On the measurement side, established indicator sets for smart sustainable cities [32] and standardized city service and quality-of-life indicators [33] offer a basis for evaluating governance outcomes rather than technology outputs alone.

Three gaps emerge from this body of work. First, the technical literature on sensing, edge computing, digital twins, and agentic intelligence and the governance literature on accountability, ethics, and performance measurement have developed largely in parallel, with limited architectural integration between them. Second, existing contributions rarely specify how autonomy should be graduated across decision types, leaving the boundary between machine action and institutional judgment undefined. Third, few frameworks connect layered technical design to measurable governance indicators in a way that supports staged implementation in greenfield environments. The framework presented below addresses these gaps by treating infrastructure, data,

simulation, agents, decision logic, and governance execution as a single vertically integrated architecture. The proposed NEOM Urban Decision Intelligence Framework (NUDIF) is based on six integrated layers. The first layer is the connected infrastructure layer, consisting of 5G networks, IoT sensors, edge devices, cloud platforms, and green data centers. This layer creates the physical and digital nervous system of the city by collecting operational signals from buildings, transport systems, utilities, environmental sensors, logistics assets, and public-service channels. The second layer is the urban data platform, where structured, semi-structured, and unstructured data are ingested, classified, standardized, governed, and linked through common data models. Interoperability is critical because fragmented city data prevents cross-domain optimization and limits the ability of decision makers to understand system-wide impacts.

The third layer is the analytics and digital twin layer, where real-time monitoring, simulation, forecasting, scenario planning, and risk modeling are performed. Digital twin studies emphasize their value for predictive modeling, urban resilience, sustainability, evidence-based planning, and mitigation of complex-system uncertainty [6], [7], [14]–[16]. In NUDIF, the digital twin is not only a visualization tool; it is a governance instrument that allows leaders to test policy options, compare trade-offs, estimate service impacts, and validate interventions before execution. The fourth layer introduces specialized AI agents that operate across urban domains. These agents include an Energy Agent for power optimization, a Mobility Agent for traffic and autonomous transport coordination, a Water Agent for resource monitoring, an Environment Agent for emissions and climate indicators, a Public Safety Agent for risk detection, and a Governance Agent for policy recommendations. Research on IoT, smart city brains, agentic urban AI, and intelligent urban agents support the shift from reactive administration to proactive urban operations [8], [9], [11]–[13]. Multi-agent systems and broader AI foundations further support the design of specialized agents that coordinate sensing, reasoning, recommendation, and decision workflows [24], [25].

The fifth layer is the decision intelligence layer, which converts analytics outputs into ranked

recommendations, alerts, policy options, and operational workflows. This layer combines predictive models, rules, optimization logic, risk scoring, and explainability mechanisms so that decision makers can understand why an action is recommended. The sixth layer is the autonomous governance layer, where approved recommendations are translated into controlled workflows, escalations, service actions, compliance reporting, and performance dashboards. Autonomy in this framework does not imply unrestricted machine control; rather, it represents a graduated model in which routine decisions may be automated, high-impact decisions require human approval, and strategic decisions remain under institutional governance.

A practical NUDIF implementation should follow a staged maturity path. In the first stage, the city establishes sensor connectivity, data governance standards, cybersecurity controls, and core dashboards. In the second stage, it integrates domain-level digital twins and predictive analytics for energy, mobility, water, environment, and public services. In the third stage, it deploys AI agents that generate recommendations and coordinate cross-domain actions. In the fourth stage, it introduces controlled automation for low-risk operational decisions while maintaining human accountability for policy, safety, privacy, and resource-allocation choices. Figure 1 illustrates the overall layered architecture of the framework, Table 1 summarizes the six framework layers and their key components, and Table 2 outlines the specialized AI agents and their governance responsibilities. Figure 2 presents the four-stage implementation maturity path described above.

III. METHODOLOGY

This study adopts a conceptual, design-oriented research approach rather than an empirical or experimental one. The objective is to construct and justify an architectural artifact, the NEOM Urban Decision Intelligence Framework, that addresses the integration gaps identified in the preceding section. This methodological choice is appropriate because the target environment is a greenfield urban development in which fully operational city-scale autonomous governance systems are not yet available for

observational study. In such conditions, framework construction and analytical evaluation precede empirical validation, and the resulting artifact serves as a reference model against which future implementations can be designed and assessed.

The framework was developed from three categories of source material. The first category comprises peer-reviewed scholarship on smart cities, urban governance, distributed computing, digital twins, and multi-agent systems, which supplied the conceptual vocabulary and the structural logic of the layered design. The second category comprises institutional and strategic documentation describing the objectives, digital ambitions, and service models of NEOM and THE LINE, which established the contextual requirements the framework must satisfy [1]–[3]. The third category comprises normative and standards-based material, including responsible artificial intelligence guidance, risk management frameworks, management system standards, and smart sustainable city indicator sets, which supplied the governance controls and measurement constructs embedded in the upper layers [27]–[33].

Framework construction proceeded in four steps. In the first step, recurring functional capabilities reported across the source literature were extracted and grouped by their role in the urban decision chain, yielding the six-layer vertical structure summarized in Figure 1 and detailed in Table 1. In the second step, urban service domains with distinct data profiles, operating tempos, and accountability requirements were identified and assigned to dedicated agents, producing the agent specification presented in Table 2. In the third step, governance controls were mapped onto the architecture by determining, for each decision category, whether execution should be automated, subject to human approval, or reserved for institutional authority. In the fourth step, representative performance indicators were derived for each domain so that framework adoption can be evaluated through observable governance outcomes rather than implementation milestones alone, as set out in Table 3.

Evaluation at this stage is analytical rather than empirical. The framework was assessed against four

criteria derived from the reviewed literature: completeness, meaning that the architecture spans the full path from physical sensing to governance execution; interoperability, meaning that each layer exchanges information with adjacent layers through defined data and decision interfaces; accountability, meaning that every automated action can be traced to an authorizing control and an identifiable owner; and measurability, meaning that outcomes can be expressed through recognized urban performance indicators. Each proposed layer, agent, and control was tested against these criteria and retained only where it contributed to at least one of them. Empirical validation through simulation studies, pilot deployments, and comparative case analysis remains a task for subsequent research, and the limitations arising from this design choice are discussed in the closing section.




Figure 1. NUDIF layered architecture for AI-driven urban decision intelligence and autonomous governance.

Table 1. NUDIF Framework Layers and Key Components.

Layer	Purpose	Key Components
Connected Infrastructure	Enables real-time urban connectivity	5G, IoT, edge devices, green data centers
Urban Data Platform	Collects and standardizes city data	Data lakes, integration pipelines, master data
Analytics and Digital Twin	Supports simulation and prediction	Forecasting, scenario models, dashboards
AI Agent Layer	Coordinates domain-specific intelligence	Energy, mobility, water, environment, safety agents
Decision Intelligence	Converts analytics into recommendations	Alerts, ranked options, policy recommendations
Autonomous Governance	Implements approved decisions with accountability	Workflow automation, controls, human oversight

Table 2. AI Agents and Governance Responsibilities.

AI Agent	Governance Function	Expected Outcome
Energy Agent	Monitors demand, renewable supply and consumption patterns	Improved energy efficiency and lower carbon intensity
Mobility Agent	Analyzes traffic, autonomous transport and congestion indicators	Reduced delays and optimized mobility flows
Water Agent	Tracks water demand, leakage and distribution performance	Better conservation and resource reliability

Environment Agent	Monitors air quality, emissions and climate indicators	Enhanced environmental compliance and ESG reporting
Public Safety Agent	Detects risks, incidents and emergency patterns	Faster response and improved urban resilience
Governance Agent	Consolidates recommendations into policy and operational actions	More transparent and accountable decision-making




Figure 2. NUDIF four-stage implementation maturity path from connected foundation to controlled autonomy.

IV. RESULTS AND THEMATIC SYNTHESIS

The framework indicates that autonomous governance is not achieved by deploying isolated technologies, but by integrating infrastructure, data, analytics, AI agents, and governance workflows into one operating model. The first finding is that connected infrastructure becomes strategically valuable only when it is linked to decision intelligence capabilities that support action, accountability, and measurable public value. The second finding is that AI agents can reduce governance latency by continuously monitoring urban indicators, identifying anomalies, and recommending proactive actions before service failures become visible to residents. The third finding is that digital twins strengthen decision quality by enabling simulation before execution, especially where energy, mobility, water, safety, and

environmental outcomes are interdependent. The fourth finding is that responsible AI controls, transparency, privacy, cybersecurity, inclusion, and human oversight are essential for trust in autonomous urban governance [5], [10], [26]–[29]. Formal AI risk management and AI management system standards further strengthen governance accountability and operational control [30], [31]. The fifth finding is that NUDIF can support sustainability governance by linking operational data to carbon, resource-efficiency, resilience, service-quality indicators, and recognized smart sustainable city performance standards [32], [33]. Table 3 presents representative governance indicators supported by the framework across the main urban domains.

Table 3. Governance KPIs Supported by NUDIF.

Governance Domain	Representative KPI	Decision Intelligence Use
Energy	Energy efficiency, renewable utilization, peak demand	Forecast demand and recommend load-balancing actions
Mobility	Travel time, congestion rate, transport availability	Predict congestion and optimize mobility routing
Water	Water loss, consumption per zone, system reliability	Detect leakage and recommend conservation measures
Environment	Carbon emissions, air quality, waste performance	Monitor sustainability risks and trigger corrective actions
Public Services	Service response time, citizen satisfaction, incident closure	Prioritize service requests and improve delivery performance
Governance	Decision latency, compliance rate, policy effectiveness	Rank recommendations and support transparent governance

NUDIF also presents implementation challenges. Large-scale urban autonomy depends on data quality, model reliability, cybersecurity resilience, consent management, and governance clarity. Poorly integrated data can create misleading recommendations, while opaque AI models can reduce trust among decision makers and residents. For this reason, the framework requires a governance board, model-risk controls, audit trails, privacy-by-design practices, explainable AI outputs, and periodic human review. These safeguards ensure that AI supports institutional judgment rather than replacing public accountability.

Positioned against existing contributions, the framework differs in three respects. Compared with architectures that concentrate on sensing and connectivity, it treats infrastructure as an input to decision quality rather than as an outcome in itself [21]–[23]. Compared with digital twin models that emphasize representation and simulation fidelity, it assigns the twin an explicit governance function as an instrument for testing policy options before commitment [14]–[16]. Compared with agent-centered proposals that demonstrate domain-specific autonomy, it specifies the coordination and escalation logic through which multiple agents operate under a single accountability structure [11]–[13], [24]. The contribution therefore lies less in any individual component than in the vertical integration of components that the literature has generally treated separately.

Several institutional implications follow. Adoption requires a defined decision-rights model that states, in advance, which categories of action may be executed autonomously, and which must be escalated; without this clarity, automation tends to expand informally and accountability becomes ambiguous. It also requires investment in data stewardship roles, since the analytical layers depend on curated and semantically consistent data rather than on volume alone. Workforce implications are equally significant, as operators shift from monitoring dashboards toward reviewing, challenging, and approving machine-generated recommendations, a transition that demands new competencies in model interpretation and risk assessment. Finally, procurement and vendor management practices must accommodate continuous model updates, which differ materially from the fixed-specification purchasing models traditionally used for urban infrastructure. These organizational conditions align with the readiness gaps identified in international assessments of responsible urban artificial intelligence [5], [10].

The study is subject to clear limitations. The framework is conceptual and has not yet been instantiated in an operational environment, so the performance indicators presented here are proposed rather than measured. The architecture is derived primarily from a greenfield context and would require adaptation for brownfield cities with legacy systems, established institutional boundaries, and existing regulatory commitments. The agent specification is

deliberately domain-level and does not prescribe algorithms, model classes, or interface standards, which preserves generality but leaves substantial design work to implementers. Finally, questions of public acceptance, participatory oversight, and citizen consent are acknowledged within the governance layer but are not examined in depth, and they warrant dedicated investigation given their influence on the legitimacy of autonomous urban decision-making.

V. CONCLUSION

This paper proposed the NEOM Urban Decision Intelligence Framework as a practical model for moving from connected infrastructure to autonomous governance in smart cities. The framework connects 5G, IoT, edge computing, cloud platforms, digital twins, AI agents, decision intelligence, and governance workflows into a unified architecture. It shows how urban data can be transformed into proactive recommendations, policy simulations, operational workflows, and sustainability-focused actions. For NEOM-scale smart cities, the future of governance depends not only on advanced infrastructure, but on the ability to convert real-time data into trusted, explainable, secure, and responsible decisions. Future research may validate the framework through case-based simulations, domain-specific AI agents, measurable smart city performance indicators, responsible AI governance controls, and comparative studies across greenfield and brownfield smart city environments aligned with urban sustainability goals [2], [3], [5], [10], [26]–[33].

The wider contribution of this work is to reframe autonomy as a governance construct rather than a purely technical one. Cities that treat autonomy as an engineering milestone risk deploying capable systems without the institutional controls needed to sustain public confidence, while cities that treat it as a governance design problem can expand automation incrementally and defensibly. By linking six architectural layers to graduated decision rights, staged implementation, and recognized performance indicators, the framework offers urban authorities a structured path for that expansion. As NEOM-scale developments progress from construction toward sustained operation, this integration of technical

capability with accountable decision-making is likely to determine whether advanced infrastructure translates into measurable improvements in sustainability, resilience, and quality of urban life.

REFERENCES

- [1] Saudi Vision 2030, “NEOM Project,” Vision 2030, 2025. [Online]. Available: <https://www.vision2030.gov.sa/en/explore/projects/neom>. Accessed: Jul. 29, 2026.
- [2] Saudi Vision 2030, “THE LINE,” Vision 2030, 2025. [Online]. Available: <https://www.vision2030.gov.sa/en/explore/projects/the-line>. Accessed: Jul. 29, 2026.
- [3] NEOM, “Technology and Digital,” NEOM, 2026. [Online]. Available: <https://www.neom.com/en-us/our-business/sectors/technology-and-digital>. Accessed: Jul. 29, 2026.
- [4] J. John *et al.*, “Artificial intelligence for smart cities: A comprehensive review across six pillars and global case studies,” *Urban Science*, 2025.
- [5] UN-Habitat, “Global Assessment of Responsible Artificial Intelligence in Cities,” United Nations Human Settlements Programme, 2024. [Online]. Available: <https://unhabitat.org/global-assessment-of-responsible-ai-in-cities>. Accessed: Jul. 29, 2026.
- [6] I. Lukić *et al.*, “Advancing smart city sustainability through artificial intelligence, digital twin and blockchain solutions,” *Technologies*, 2025.
- [7] E. J. Sacoto-Cabrera, A. Perez-Torres, L. Tello-Oquendo, and M. Cerrada, “IoT, AI, and digital twins in smart cities: A systematic review for a thematic mapping and research agenda,” *Smart Cities*, vol. 8, no. 5, 2025.
- [8] S. E. Bibri and J. Huang, “Artificial intelligence of things for sustainable smart city brain and digital twin systems,” *Environmental Science and Ecotechnology*, vol. 26, 2025.
- [9] NVIDIA, “Into the Omniverse: How smart city AI agents transform urban operations,” *NVIDIA Blog*, 2025. [Online]. Available: <https://blogs.nvidia.com/blog/smart-city-ai-agents-urban-operations/>. Accessed: Jul. 29, 2026.

- [10] T. Isagah and S. Ben Dhaou, "Responsible and inclusive urban AI: Opportunities and challenges for advancing sustainable development goals," United Nations Sustainable Development, 2024.
- [11] A. Tiwari, "Beyond Automation: The Emergence of Agentic Urban AI," *Automation*, vol. 6, no. 3, 2025.
- [12] J. Han, Y. Ning, Z. Yuan, H. Ni, F. Liu, T. Lyu, and H. Liu, "Large Language Model Powered Intelligent Urban Agents: Concepts, Capabilities, and Applications," arXiv:2507.00914, 2025.
- [13] A. Matei and M. Coccoşatu, "Artificial Internet of Things, Sensor-Based Digital Twin Urban Computing Vision Algorithms, and Blockchain Cloud Networks in Sustainable Smart City Administration," *Sustainability*, vol. 16, no. 16, 2024.
- [14] M. Batty, "Digital twins," *Environment and Planning B: Urban Analytics and City Science*, vol. 45, no. 5, pp. 817–820, 2018.
- [15] M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in *Transdisciplinary Perspectives on Complex Systems*. Cham, Switzerland: Springer, 2017.
- [16] S. E. Bibri, "Data-driven smart sustainable cities of the future: An evidence synthesis approach to a comprehensive state-of-the-art literature review," *Sustainable Futures*, 2021.
- [17] R. Kitchin, "The real-time city? Big data and smart urbanism," *GeoJournal*, vol. 79, pp. 1–14, 2014.
- [18] A. Caragliu, C. Del Bo, and P. Nijkamp, "Smart cities in Europe," *Journal of Urban Technology*, vol. 18, no. 2, pp. 65–82, 2011.
- [19] H. Chourabi *et al.*, "Understanding smart cities: An integrative framework," in *Proc. 45th Hawaii Int. Conf. System Sciences*, 2012.
- [20] A. Meijer and M. P. R. Bolívar, "Governing the smart city: A review of the literature on smart urban governance," *International Review of Administrative Sciences*, vol. 82, no. 2, pp. 392–408, 2016.
- [21] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of things for smart cities," *IEEE Internet of Things Journal*, vol. 1, no. 1, pp. 22–32, 2014.
- [22] M. Aazam, S. Zeadally, and K. A. Harras, "Fog computing architecture, evaluation, and future research directions," *IEEE Communications Magazine*, vol. 56, no. 5, pp. 46–52, 2018.
- [23] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637–646, 2016.
- [24] M. Wooldridge, *An Introduction to MultiAgent Systems*, 2nd ed. Hoboken, NJ, USA: Wiley, 2009.
- [25] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. Hoboken, NJ, USA: Pearson, 2020.
- [26] L. Floridi *et al.*, "AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations," *Minds and Machines*, vol. 28, pp. 689–707, 2018.
- [27] High-Level Expert Group on Artificial Intelligence, "Ethics Guidelines for Trustworthy AI," European Commission, 2019. [Online]. Available: <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>. Accessed: Jul. 29, 2026.
- [28] UNESCO, "Recommendation on the ethics of artificial intelligence," United Nations Educational, Scientific and Cultural Organization, 2021. [Online]. Available: <https://unesdoc.unesco.org/ark:/48223/pf0000381137>. Accessed: Jul. 29, 2026.
- [29] OECD, "OECD principles on artificial intelligence," Organisation for Economic Co-operation and Development, 2019. [Online]. Available: <https://www.oecd.org/en/topics/ai-principles.html>. Accessed: Jul. 29, 2026.
- [30] National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, 2023. [Online]. Available: <https://doi.org/10.6028/NIST.AI.100-1>. Accessed: Jul. 29, 2026.
- [31] ISO/IEC, "ISO/IEC 42001:2023 Artificial intelligence management system," International

Organization for Standardization, 2023.
[Online]. Available:
<https://www.iso.org/standard/81230.html>.
Accessed: Jul. 29, 2026.

[32] ITU, “United for Smart Sustainable Cities: Key performance indicators for smart sustainable cities,” International Telecommunication Union, 2017. [Online]. Available: <https://www.itu.int/en/ITU-T/ssc/Pages/KPIs-on-SSC.aspx>. Accessed: Jul. 29, 2026.

[33] ISO, “ISO 37120: Sustainable cities and communities—Indicators for city services and quality of life,” International Organization for Standardization, 2018. [Online]. Available: <https://www.iso.org/standard/68498.html>. Accessed: Jul. 29, 2026.