



# Artificial Intelligence of Things and Digital Twins for Urban Resilience: A Critical Review of Architectures, Adaptive Intelligence and Explainable Decision Support

<sup>1</sup>Deepak Barkhane, <sup>2</sup>Md Imran Ansari

<sup>1</sup>M.Tech Scholar, <sup>2</sup>Assistant Professor

<sup>1,2</sup>Department of Computer Science and Engineering

<sup>1,2</sup>Patel College of Science and Technology, Bhopal, Madhya Pradesh, India

**Abstract:** Urban areas are increasingly exposed to compounding hazards—flooding, heatwaves, infrastructure strain, and climate-driven extremes—at a pace that outstrips the static, threshold-based monitoring systems most cities still rely on. Over the last five years, the convergence of the Internet of Things with artificial intelligence, popularly termed AIoT, together with the maturation of digital twin technology, has been proposed as the technical foundation for a new generation of urban resilience infrastructure capable of sensing, reasoning about, and intervening in city-scale hazards in near real time. This review critically synthesizes the literature on AIoT-enabled digital twins for urban resilience across four interdependent layers: sensing and multi-source data fusion, adaptive and self-learning prediction, explainable and causal decision support, and domain deployment in flood, water, and disaster-management contexts. Drawing on 46 peer-reviewed studies published predominantly between 2020 and 2026 in IEEE, Elsevier, Springer, ACM, MDPI, and Wiley venues, the review develops a layered taxonomy of the field, compares architectural and algorithmic approaches across scalability, adaptability, explainability, interoperability, sustainability, and security and trust, and identifies where the literature remains fragmented. The analysis shows that reliability-aware sensor fusion, concept-drift-aware learning, and counterfactual or causal reasoning have each matured as independent research threads, yet are rarely integrated within a single, continuously operating urban resilience system. The review closes by outlining the resulting research gaps and the directions they open for future work.

**Index Terms:** AIoT, Digital Twin, Urban Resilience, Smart City, Sensor Fusion, Concept Drift, Explainable AI, Causal Inference, Intelligent Decision Support Systems, Cybersecurity

## I. INTRODUCTION

Cities are absorbing hazards faster than their monitoring systems can characterize them. Rainfall intensities are exceeding the design assumptions embedded in drainage infrastructure built decades ago; heat, water stress, and infrastructure fatigue are compounding rather than arriving in isolation; and the sensor networks most municipalities operate remain, in practice, closer to instrumented thresholds than to reasoning systems. A gauge crossing a fixed level triggers an alert; the alert is rarely accompanied by a forecast of what happens next, a ranked account of which intervention would help most, or an honest signal that the gauge itself may be malfunctioning. Over roughly the last five years, two technological threads have been proposed, largely independently at first and increasingly together, as the foundation for closing this gap: the Artificial Intelligence of Things (AIoT), in which learning is embedded across the sensing-to-decision pipeline rather than bolted on as downstream analytics, and digital twin technology, in which a continuously synchronized

virtual counterpart of physical urban infrastructure supports simulation, forecasting, and eventually intervention planning.

The resulting literature is large, fast-moving, and fragmented across disciplinary silos that do not cite one another as often as their shared subject matter would suggest. Digital-twin architecture papers, sensor-fusion papers, concept-drift and continual-learning papers, explainable-AI papers, causal-inference papers, and flood- or disaster-specific application papers each form reasonably mature literatures in their own right. What is rarer is a synthesis that reads across these layers specifically through the lens of urban resilience, asking not "how good is this technique in isolation" but "what would it take for these techniques to compose into a system a city could actually operate."

This review undertakes that synthesis. It is deliberately not another general survey of digital twin definitions or another general AIoT taxonomy; several excellent examples of both already exist and are discussed critically within. Instead, it organizes 46 peer-reviewed studies, published predominantly between 2020 and 2026 in IEEE, Elsevier, Springer, ACM, MDPI, Wiley, and Nature-portfolio venues, around a four-layer lens purpose-built for urban resilience: sensing and multi-source data fusion; adaptive, self-learning prediction under non-stationary conditions; explainable and causal decision support; and domain deployment in flood, water, and disaster-management contexts. Foundational works published before 2020 are included only where the field has not produced a more recent equivalent treatment of the same foundational question.

The review's scope is bounded deliberately. It concentrates on civic and urban resilience applications of digital twins and AIoT rather than the broader industrial and manufacturing digital-twin literature, drawing on the latter only where a technique (for example, reliability-weighted sensor fusion developed for industrial safety) is directly transferable. It likewise treats "urban resilience" as a genuinely contested, evolving concept rather than a settled engineering target, a point taken up directly in Section II, because that conceptual looseness has concrete downstream consequences for how systems built in its name are evaluated.

The remainder of the paper is organized as follows. Section II traces the evolution from static IoT monitoring to the cognitive, AIoT-enabled digital twin, grounding the discussion in how the urban resilience concept itself has developed. Section III proposes the five-layer taxonomy used throughout the review. Sections IV through VII review, in turn, digital twin architectures, AIoT and sensor-level intelligence, adaptive and self-learning prediction, and explainable/causal decision support. Section VIII turns to domain applications, principally flood and water resilience, where these layers are meant to converge. Section IX considers the still-nascent role of generative AI and large language models. Section X synthesizes the review's findings across five evaluative dimensions—scalability, adaptability, explainability, interoperability, and sustainability—and Section XI sets out the resulting research gaps and future directions before Section XII concludes.

## II. BACKGROUND: FROM STATIC MONITORING TO COGNITIVE URBAN TWINS

Before reviewing the technology built in its name, it is worth pausing on what "urban resilience" is actually taken to mean, because the term has not settled into a single, stable engineering definition. Amirzadeh, Sobhaninia, and Sharifi [1] trace how the concept has evolved across disciplines—ecology, disaster risk reduction, urban planning—and argue that it functions less as a fixed target than as an evolving, multidimensional capacity encompassing absorption, adaptation, and recovery rather than resistance alone. This is not merely a semantic quibble for a technology review. A prediction system optimized to minimize forecast error is implicitly optimizing for a different objective than a decision-support system optimized to preserve a community's adaptive capacity after a shock, and the literature reviewed in this paper is not always explicit about which of these targets it is actually serving. This ambiguity resurfaces at several points below, particularly in Section VII, where the distinction between explaining a prediction and supporting an intervention decision becomes material.

Against this conceptual backdrop, the technical evolution of urban monitoring can be read as a series of four overlapping generations rather than a clean succession. The first generation is discrete-sensor, threshold-based monitoring: a gauge, a fixed alert level, a human responder. The second is the smart-city platform generation: dashboards, GIS overlays, and descriptive analytics that aggregate sensor feeds for human interpretation but do not themselves forecast or recommend. The third is the digital twin city, in which Deng, Zhang, and Shen [2] argue a qualitatively different pattern of urban governance becomes possible once the virtual model is bidirectionally synchronized with the physical city rather than serving as a static visualization—decisions can be rehearsed in simulation before being enacted in the physical system. The fourth, still largely confined to research prototypes rather than field deployments, is the AIoT-enabled cognitive twin, in which the twin does not only mirror physical state but continuously learns from it and participates in reasoning about intervention.

A critical reading of the literature surveyed in this paper suggests that most systems described as operational today sit at the second or third generation; the fourth is, at present, largely a research-paper phenomenon rather than a field-deployed one. This is not a criticism so much as a calibration: claims of "cognitive" or "intelligent" digital twins should be read against what is actually being synchronized, learned, and acted upon, a distinction Table I attempts to make explicit by tracing the chronological evolution of the field's representative capabilities alongside its intelligence level and characteristic data flow.

TABLE I. CHRONOLOGICAL EVOLUTION OF URBAN MONITORING PARADIGMS

| Era                      | Representative Capability                                | Data Flow                     | Intelligence Level       | Refs     |
|--------------------------|----------------------------------------------------------|-------------------------------|--------------------------|----------|
| 1. Static IoT monitoring | Discrete sensors, fixed-threshold alerts                 | Unidirectional (sensor→human) | None (rule-based)        | [14]     |
| 2. Smart-city platform   | Dashboards, GIS overlays, descriptive analytics          | Unidirectional, aggregated    | Descriptive only         | [4],[13] |
| 3. Digital twin city     | Bidirectional sync; decisions rehearsed in simulation    | Bidirectional                 | Predictive (model-based) | [2],[3]  |
| 4. AIIoT cognitive twin  | Continuous learning, adaptive prediction, recommendation | Bidirectional + learning loop | Adaptive / cognitive     | [7],[21] |

### III. A LAYERED TAXONOMY OF AIOT-ENABLED DIGITAL TWINS FOR URBAN RESILIENCE

The literature reviewed in this paper can be, and often is, organized chronologically or by underlying technology (machine learning method, communication protocol, hardware platform). Both organizing principles are useful but neither is well suited to a review whose purpose is evaluating readiness for urban resilience specifically, because resilience outcomes depend less on any single layer's sophistication than on how layers compose end to end. A city with an excellent flood-prediction model but no mechanism for translating a forecast into a ranked, explainable intervention has not meaningfully improved its resilience; neither has a city with excellent sensor coverage but no adaptive learning to keep pace with a changing hazard distribution.

This review therefore organizes its critical synthesis around five layers, illustrated in Figure 1: (1) Sensing and Perception, the physical and virtual sensor networks generating raw observations; (2) Fusion and Preprocessing, where heterogeneous, sometimes unreliable observations are combined into a coherent state estimate; (3) Adaptive Learning and Prediction, where models forecast future hazard states and must remain accurate as underlying conditions drift; (4) Reasoning and Decision Support, where predictions are made explainable and, ideally, causally actionable; and (5) Actuation and Feedback, where recommendations reach physical infrastructure or human responders and their outcomes feed back into the twin. Sections IV through VIII map the literature onto layers 1 through 5 respectively, with Section IV additionally covering the architectural scaffolding that hosts all five layers together.

A pattern that recurs throughout the review, and is made explicit in the critical synthesis of Section X, is that most individual studies instantiate only two or three of these five layers at once. A central contribution of this review is making that partial coverage visible across the literature as a whole, rather than within any one study, which is the vantage point a single-paper's related-work section cannot provide.




Fig. 1. Layered taxonomy of AIoT-enabled digital twins for urban resilience.

#### IV. DIGITAL TWIN ARCHITECTURES FOR SMART AND RESILIENT CITIES

The architectural literature on digital twins is comparatively mature and provides the scaffolding within which the sensing, learning, and reasoning layers discussed later must fit. Minerva, Lee, and Crespi [3], in one of the most widely used reference surveys in this space, set a technical bar that many applied papers invoke but do not always meet: a genuine digital twin requires bidirectional, continuously updated data flow between physical and virtual entities, not merely a detailed virtual model that is periodically refreshed. Read critically against the applied literature surveyed later in this paper, a non-trivial share of systems described as "digital twins" in domain-specific venues are closer to sophisticated simulation models than to twins in this stricter sense, a distinction worth carrying forward.

Fuller, Fan, Day, and Barlow [4] broaden the picture considerably, cataloguing the enabling technology stack—IoT, AI, extended reality, blockchain—alongside open research questions across application domains. The breadth is valuable for orientation but comes at a cost familiar to broad technology surveys generally: no single hazard domain receives sustained treatment, which is precisely the gap that domain-specific reviews, discussed in Section VIII, exist to fill. Qi et al. [5] approach the same technology from a more product-lifecycle-management-oriented angle, proposing a five-dimension model—physical entity, virtual entity, services, twin data, and connections—that has proven influential in manufacturing contexts. Applied to an urban setting, however, the "services" and "connections" dimensions are considerably harder to specify: a manufactured asset typically has one owner and one data pipeline, whereas a city's flood-relevant infrastructure is typically split across water boards, transport authorities, and municipal government, each with independent data governance. The five-dimension model does not itself resolve this multi-owner complication, and few of the urban-facing papers reviewed later explicitly address it either.

Mihai et al. [6], writing from a communications and networking perspective, foreground a concern that is easy to elide in higher-level architecture papers: the real-time synchronization problem, encompassing bandwidth constraints, latency budgets, and the placement of computation at the edge versus the cloud. This is a useful corrective, because a meaningful share of the urban resilience literature discussed in Sections VI through VIII implicitly assumes clean, timely data arrival and reasons about model accuracy without

reasoning about whether the data reached the model in time to matter for a fast-onset hazard such as flash flooding.

Rathore et al. [7] occupy a distinctive position in this section's literature by treating the twin and the learning system that runs on it as genuinely coupled, rather than as a passive replica onto which a separately developed machine learning model is bolted—their systematic review of AI, machine learning, and big data's role specifically within digital twinning is, among the architecture-focused papers surveyed, the closest in spirit to this review's own layered framing, and is returned to in the critical synthesis of Section X.

Three further reviews narrow the architectural discussion toward sustainability and urban planning rather than general-purpose technology cataloguing. Weil et al. [8] catalogue urban digital twin challenges with particular attention to data integration and governance; Bibri et al. [9] and Bibri, Huang, Omar, and Kenawy [10] examine, respectively, the synergy between AI and digital twins in environmental planning for sustainable smart cities and the integration of digital twins with zero-energy building design for climate mitigation. Read together, these three reviews are comparatively strong on data-governance and environmental-sustainability challenges but comparatively thin on the adaptive-intelligence question—concept drift, continual learning—that Section VI takes up; this is a genuine gap between literatures that do not cite each other as often as their shared subject matter would suggest. Ariyachandra and Wedawatta [11] narrow further still to disaster risk management specifically, tracing how the concept of a "digital twin smart city" for this purpose has evolved; notably, even within this narrower and more resilience-focused review, the authors find that deployed, operational systems remain rare relative to conceptual and prototype-stage proposals, a finding consistent with the fourth-generation calibration made in Section II.

Table II summarizes this section's comparative reading of the architectural literature across data-flow directionality, domain generality, reported maturity, and the principal limitation each study itself identifies or that this review identifies in it.

TABLE II. COMPARATIVE OVERVIEW OF DIGITAL TWIN ARCHITECTURAL APPROACHES

| Study | Architectural Focus                                             | Data Flow                          | Domain Generality              | Reported Maturity             | Key Limitation Noted                                    |
|-------|-----------------------------------------------------------------|------------------------------------|--------------------------------|-------------------------------|---------------------------------------------------------|
| [3]   | Technical features, scenarios, architectural models             | Bidirectional (defining criterion) | General                        | Conceptual / survey           | Many applied “twins” fail to meet the bidirectional bar |
| [4]   | Enabling technology stack (IoT, AI, XR, blockchain)             | Not domain-specified               | General                        | Conceptual / survey           | Breadth over depth; no single hazard domain treated     |
| [5]   | Five-dimension model (entity/virtual/services/data/connections) | Bidirectional                      | Manufacturing, ported to urban | Conceptual                    | Multi-agency ownership complexity underspecified        |
| [6]   | Real-time sync, communications/networking                       | Bidirectional, latency-aware       | General                        | Conceptual / survey           | Assumes clean/timely data other papers do not guarantee |
| [7]   | AI/ML/big-data role within twinning                             | Bidirectional, learning-coupled    | General                        | Conceptual / survey           | Closest to integrated framing, still survey-level       |
| [8]   | Urban challenges: DT data integration & governance              | Bidirectional                      | Urban planning                 | Conceptual / survey           | Thin on adaptive intelligence / concept drift           |
| [9]   | AI + DT synergy for sustainable planning                        | Bidirectional                      | Urban / environmental          | Systematic review             | Sustainability-focused; adaptive learning underexplored |
| [10]  | DT + zero-energy buildings, climate mitigation                  | Bidirectional                      | Building / urban               | Systematic review             | Narrow to energy/buildings; same gap as above           |
| [11]  | Disaster risk management DT concepts                            | Bidirectional (conceptual)         | Disaster resilience            | Mostly conceptual / prototype | Deployed systems rare even in this narrower scope       |

Read as a whole, the architectural literature surveyed in this section has matured considerably faster than the intelligence layers that are meant to run on top of it. Most of the reviewed architecture papers describe the data plumbing—how physical and virtual entities stay synchronized—in careful technical detail, while treating the cognition that will eventually run on that plumbing largely as future work. This asymmetry sets up the next three sections, which examine the sensing, learning, and reasoning layers in turn.

## V. AIOT AND SENSOR-LEVEL INTELLIGENCE: FROM STATIC IOT TO COGNITIVE SENSING

The term Artificial Intelligence of Things is often used loosely, but Siam et al. [12] draw a distinction worth preserving throughout this review: AIoT properly refers to intelligence embedded across the device, edge, and cloud tiers of the sensing stack, rather than AI as a downstream analytics layer applied after data has already been collected and centralized. The distinction matters for urban resilience specifically because latency-critical hazard response—a flash flood evolving over tens of minutes—cannot always tolerate a round trip to a centralized model before a warning is issued. Menon U et al. [13] broaden this picture with a wide-ranging survey of AI-IoT integration across security, efficiency, and application domains; read critically, the survey is considerably stronger on breadth of application inventory than on depth within any single safety-critical domain, a pattern that recurs across broad AIoT surveys generally and that motivates the domain-specific treatment in Section VIII.

Whether cognitive sensing can be trusted at all depends substantially on the fusion layer that combines heterogeneous, sometimes contradictory sensor observations into a single coherent state estimate. The foundational taxonomy here remains Khaleghi, Khamis, Karray, and Razavi's review [14], which organizes fusion approaches by the level at which combination occurs—raw data, extracted features, or independent decisions—and continues to structure how fusion is discussed more than a decade later. That durability is itself notable, but it also means the taxonomy predates the heterogeneous, failure-prone, officially-and-crowd-sourced mixed sensor networks that characterize contemporary municipal deployments, where the question is not only how to combine observations but which observations deserve to be trusted in the first place.

This latter, reliability-aware framing is precisely what Desikan et al. [15] contribute, in an industrial rather than urban context: their fault-tolerant fusion architecture for fire-risk detection uses Dempster-Shafer evidence combination to dynamically down-weight sensors exhibiting drift, noise, or outright failure rather than assuming all inputs are equally trustworthy. The technique is directly transferable to urban hazard sensing—a water-level gauge fouled by debris or a rain gauge with a clogged funnel is a common, if under-discussed, failure mode in municipal networks—yet reliability-weighted fusion of this kind remains far more common in the industrial safety literature than in published urban hazard-monitoring systems, which still largely assume sensors are equally trustworthy once installed.

Two further developments extend this sensing layer in directions the foundational fusion literature did not anticipate. Liu, Jin, and Dong [43] survey the architectures, algorithms, and applications of edge-computing-based intelligent IoT, arguing that pushing inference to the network edge rather than the cloud is what makes the latency budget of the fusion step compatible with the real-time synchronization concerns Mihai et al. [6] raise; read together, the two surveys suggest that reliability-weighted fusion and edge placement are complementary design choices, though few studies reviewed here combine both explicitly. A second, non-instrumented sensing modality is crowdsourced and social sensing: Yu and Wang [44] propose a framework for extracting spatio-temporal disaster information from multimodal social media content, demonstrating that citizen-generated data can correlate more strongly with ground conditions than text alone during a live event. This is a valuable complement to physical-sensor fusion—citizens report inundation in locations no gauge network reaches—but it introduces its own reliability problem, namely rumor, duplication, and sentiment noise, that the reliability-weighting techniques of [14] and [15] were not designed with social data in mind to handle, an integration point the literature has not yet closed.

A further, non-technical bottleneck deserves equal weight in any critical review of this layer: interoperability. Rita et al. [16] review the role of the IEEE 1451 smart transducer interface standard in resolving cross-vendor compatibility issues in smart cities and buildings, making a structural point easy to overlook in prediction-focused papers—municipal sensor networks typically span multiple agencies, vendors, and legacy protocols, so even a mathematically sound fusion algorithm cannot be deployed city-wide until the upstream heterogeneity it depends on is resolved. This is a constraint on deployment that purely algorithmic papers, including several reviewed in the sections that follow, tend not to address at all.

A related but distinct gap is security. Kim et al. [45], reviewing 154 smart-city cybersecurity studies published between 2015 and 2022, find that research attention has concentrated heavily on smart homes and general IoT security while cyber-forensics capability specific to city-scale infrastructure remains comparatively underdeveloped. This matters directly for the systems this review surveys: a sensing and fusion pipeline that is reliability-aware against ordinary sensor failure but not hardened against deliberate manipulation offers, at best, a partial guarantee, and almost none of the fusion- or learning-focused papers reviewed in this section discuss adversarial robustness at all—a blind spot revisited in the cross-cutting synthesis of Section X.

Table III compares the sensing and fusion literature discussed in this section along fusion level, reliability-awareness, real-time capability, and treatment of interoperability.

TABLE III. SENSOR FUSION AND AIOT APPROACHES COMPARED

| Study | Fusion Level                                   | Reliability-Awareness                | Real-Time Capability           | Domain                     | Interoperability          |
|-------|------------------------------------------------|--------------------------------------|--------------------------------|----------------------------|---------------------------|
| [12]  | N/A (stack-wide AIoT survey)                   | Not central focus                    | Emphasized (edge-cloud)        | General AIoT               | Not addressed             |
| [13]  | N/A (broad survey)                             | Not central focus                    | Varies by application          | Security / general         | Not addressed             |
| [14]  | Data / feature / decision-level (foundational) | Not addressed                        | Not addressed                  | General                    | Not addressed             |
| [15]  | Data-level, Dempster-Shafer combination        | Central focus (reliability-weighted) | Real-time (edge)               | Industrial IoT (fire risk) | Not addressed             |
| [16]  | N/A (standards review)                         | Not addressed                        | Enabled indirectly             | IoT / smart buildings      | Central focus (IEEE 1451) |
| [43]  | N/A (edge architecture survey)                 | Not central focus                    | Central focus (edge placement) | General intelligent IoT    | Not addressed             |
| [44]  | Multimodal (social physical) +                 | Not addressed (open problem)         | Near real-time                 | Disaster / flood sensing   | Not addressed             |
| [45]  | N/A (cybersecurity review)                     | Adversarial, not sensor-failure      | Not addressed                  | Smart city (general)       | Not addressed             |

Taken together, this section's literature shows that AIoT sensing has become genuinely intelligent at the algorithmic level, and has grown richer sensing modalities through edge placement and social/crowdsourced data; what remains comparatively rare, at least in urban-facing rather than industrial deployments, is the combination of reliability-awareness, standards-compliant interoperability, and adversarial security that would make that intelligence trustworthy and deployable at city scale. This is the first thread of the integration gap this review ultimately identifies in Section XI.

## VI. ADAPTIVE AND SELF-LEARNING INTELLIGENCE UNDER NON-STATIONARY URBAN CONDITIONS

A model trained once on historical hazard data and deployed statically will degrade, often silently, because the relationship it learned does not hold indefinitely: rainfall-runoff dynamics shift with land-use change and impervious-surface growth, sensor characteristics drift with age, and what counts as "normal" moves over time. This is the concept-drift problem, and Arora, Rani, and Saxena [17] provide a systematic categorization of detection and adaptation techniques—distinguishing real drift (the underlying relationship changes) from virtual drift (only the observed distribution changes) and gradual from sudden onset—while candidly noting that most detection methods in the surveyed literature are validated on synthetic or benchmark data streams rather than the noisy, low-sample-rate environmental streams typical of municipal sensor deployments. That

validation gap is worth carrying forward: a drift-detection method's benchmark performance is not strong evidence of its performance on the kind of data an actual city would generate.

Suárez-Cetrulo, Quintana, and Cervantes [18] focus specifically on recurring concept drift, in which a pattern that has disappeared reappears later—seasonal flood regimes are an obvious urban analogue. This recurring-pattern framing is important because it reframes adaptation from a purely forgetting-oriented problem (discard what is no longer true) to a recall-oriented opportunity (recognize and reuse a pattern the system has seen before). Notably, this recall-oriented framing is underused across the urban-resilience-specific literature reviewed elsewhere in this paper, which tends to treat model retraining as unidirectional adaptation to novelty rather than as an opportunity to recognize recurring seasonal or cyclical regimes.

Adaptivity in the urban setting is constrained by a second, non-technical factor: data-sharing barriers across the multiple agencies that jointly hold hazard-relevant data. Federated learning has been proposed as an architectural response, allowing models to adapt from distributed data without centralizing it. Alterkawi and Dib [19] synthesize this literature using a challenge-oriented taxonomy and explicitly flag reproducibility—the availability of datasets and deployment artifacts—as a weak point across the federated-learning-for-smart-cities corpus they review, a caution any future system claiming federated adaptivity should take seriously rather than assume away. Al-Huthaifi et al. [20] complement this with a privacy-and-security-focused survey cataloguing attack surfaces specific to federated deployments in smart cities—model inversion and poisoning among them—underscoring that a continuously adapting municipal system must be evaluated on robustness to these attacks, not accuracy alone.

Ramu et al. [21] bring federated learning and digital twins together directly, proposing that twins themselves be updated via federated rather than centralized learning across city zones or participating agencies. This is conceptually close to what a genuinely adaptive, cross-agency urban resilience twin would require, and it is worth noting explicitly that the proposal is framed at the level of concepts and a research roadmap rather than a validated field deployment—an honest characterization of where this particular sub-literature currently stands.

Table IV compares the adaptive and self-learning literature discussed in this section by adaptation mechanism, treatment of recurring drift, cross-agency data-sharing support, and validation setting.

TABLE IV. ADAPTIVE AND SELF-LEARNING STRATEGIES COMPARED

| Study | Adaptation Mechanism                   | Recurring-Drift Aware?              | Cross-Agency Sharing | Validation Setting               |
|-------|----------------------------------------|-------------------------------------|----------------------|----------------------------------|
| [17]  | Drift detection + adaptation taxonomy  | Partially (categorized, not solved) | No                   | Benchmark streams                |
| [18]  | Recurring-drift-aware ensemble methods | Yes (central focus)                 | No                   | Benchmark streams                |
| [19]  | Federated learning (thematic review)   | Not addressed                       | Yes (central focus)  | Mixed; flags reproducibility gap |
| [20]  | Federated learning (privacy/security)  | Not addressed                       | Yes                  | Survey-level (security focus)    |
| [21]  | Federated learning + digital twins     | Not addressed                       | Yes (central focus)  | Conceptual / roadmap             |

The central critical observation of this section is that adaptive learning techniques are individually mature, but are almost never coupled, in the reviewed literature, to the reliability-weighted fusion layer discussed in Section V. A model that adapts quickly to new patterns in unreliable, unweighted sensor data is adapting to noise as readily as to signal. The fusion-reliability literature and the adaptive-learning literature have developed largely in parallel, with limited cross-citation between them—itsself a symptom, visible at the level of citation graphs rather than any single paper, of the field's broader fragmentation.

## VII. EXPLAINABILITY, CAUSALITY, AND TRUSTWORTHY DECISION SUPPORT

Prediction alone is an incomplete foundation for resilience decision support. A flood-risk score without a stated reason does not tell an emergency manager why the model believes what it believes; more importantly, it does not indicate which intervention—opening a particular drainage gate, pre-positioning pumps at a particular junction, issuing an evacuation order for a particular zone—would actually change the outcome.

This gap motivates a progression, visible across the literature, from predictive modeling toward explainable modeling and, at the frontier, toward genuinely causal and counterfactual reasoning.

Barredo Arrieta et al. [22] provide one of the most widely cited framings of explainable AI (XAI), organizing the field around transparency and post-hoc explainability and introducing a "Responsible AI" framing that numerous later application-specific papers, including several reviewed elsewhere in this paper, explicitly inherit. Read critically from the vantage point of 2026, this influential taxonomy predates the recent shift toward foundation-model-based systems and, more importantly for this review's purposes, does not treat explanation stability over time as a first-class concern—a limitation that matters once a model is continually updated under concept drift (Section VI) rather than static. Yang et al. [23] provide a more recent survey extending the taxonomy toward limitations and application-specific aspects, and the contrast between the two is itself instructive: consensus has shifted toward evaluating explanation quality with dedicated metrics rather than treating the mere existence of an explanation as sufficient. Rawal et al. [24] extend the discussion further toward trustworthiness broadly construed—robustness, fairness, and accountability alongside explainability narrowly defined—which is directly relevant to urban resilience decision-making, since decisions about who is evacuated first or which infrastructure is prioritized for repair are equity-sensitive in ways that generic XAI benchmarks are rarely designed to test.

Narrowing specifically to counterfactual explanation, Stepin et al. [25] survey contrastive and counterfactual explanation-generation methods, and Karimi et al. [26] survey algorithmic recourse—both literatures organized around the question of what minimal change to the input would have altered a model's output. A distinction worth making explicit, because it is frequently blurred in applied papers, is that counterfactual explanation methods of this kind are typically model-relative rather than world-relative: they explain what the model would have predicted under a hypothetical input, which is not the same as explaining what would actually happen in the physical hazard process if a real-world intervention were made. The latter requires a structural causal model of the world itself, which is the proper subject of the causal-inference literature rather than the XAI literature narrowly defined. Yao et al. [27] survey causal-discovery and treatment-effect-estimation methods at a general machine-learning level, while Brand, Zhou, and Xie [28], writing from a social-science statistics perspective, highlight an ongoing and unresolved tension between the flexibility of machine-learning function approximation and the identification assumptions that genuine causal inference requires. Read against Sections IV through VI, it is notable that almost none of the urban-resilience-specific studies surveyed in this review actually operationalizes a structural causal model; where those studies mention explainability at all, it is typically post-hoc feature attribution of the kind surveyed by [22] and [23], which is a materially weaker guarantee than the counterfactual or causal reasoning surveyed by [25] through [28].

Kobayashi and Alam [29] offer the clearest bridge, among the reviewed literature, between digital twins and genuine explainability: their case study builds an intelligent digital twin explicitly for trustworthy, interpretable estimation of remaining useful life. It is a useful existence proof of what integrating Sections IV and VII could look like in practice, though its domain—industrial asset degradation—has cleaner failure semantics (a component fails or it does not) than multi-hazard urban resilience, where no single analogue of "remaining useful life" exists across flood, heat, and infrastructure-strain hazards simultaneously.

Table V compares the explainability and causal-reasoning literature discussed in this section by explanation type, whether the explanation is model-relative or world-relative, whether it has been evaluated on a real deployment, and its treatment of equity or fairness considerations.

TABLE V. EXPLAINABILITY AND CAUSAL REASONING APPROACHES COMPARED

| Study | Explanation Type                       | Model- or World-Relative | Real-Deployment Evaluation | Equity/Fairness |
|-------|----------------------------------------|--------------------------|----------------------------|-----------------|
| [22]  | Post-hoc / XAI taxonomy                | Model-relative           | No (conceptual/survey)     | Limited         |
| [23]  | Post-hoc, updated taxonomy             | Model-relative           | No (survey)                | Limited         |
| [24]  | Trustworthy XAI (robustness/fairness)  | Model-relative           | No (survey)                | Central focus   |
| [25]  | Contrastive counterfactual explanation | Model-relative           | No (survey)                | Not central     |

|      |                                         |                             |                                 |                             |
|------|-----------------------------------------|-----------------------------|---------------------------------|-----------------------------|
| [26] | Algorithmic recourse                    | Model-relative              | No (survey)                     | Partial (recourse fairness) |
| [27] | Causal discovery / effect estimation    | World-relative (structural) | No (survey)                     | Not central                 |
| [28] | Causal inference + ML methods           | World-relative (structural) | No (survey/methods)             | Not central                 |
| [29] | Interpretable digital twin (case study) | Model-relative, applied     | Yes (industrial RUL case study) | Not central                 |

This section's central critical finding is a three-way split: mature post-hoc explainability methods, mature counterfactual-explanation methods, and mature causal-inference theory constitute three literatures that are individually well developed but rarely integrated into a single decision-support pipeline, and are almost never integrated with the continually adapting models discussed in Section VI. This is the second, and arguably the central, thread of the integration gap this review identifies in Section XI.

## VIII. DOMAIN APPLICATIONS: FLOOD, WATER, AND DISASTER-RESILIENT URBAN SYSTEMS

The preceding sections review the technologies that domain applications are ultimately meant to employ; this section examines how far that employment has actually progressed, using flood and water resilience as the primary lens, consistent with the review's overall scope, while drawing in the broader disaster-decision-support literature where relevant.

Al-Rawas et al. [30] provide one of the broadest recent syntheses of flash-flood prediction technologies, examining artificial intelligence, machine learning, IoT, cloud computing, and robotics together; their own stated conclusion—that the field remains fragmented across technique families with little standardized benchmarking across studies—is direct evidence, from within the domain literature itself, for this review's broader fragmentation argument. Byaruhanga et al. [31] trace the evolution of flood forecasting models for early-warning systems from physically based hydrological models toward hybrid and data-driven approaches; read against Sections V through VII, this progression mirrors, with a discernible lag, the general AI/ML maturity curve documented earlier in this review—flood-specific applications tend to adopt drift-aware learning or explainability techniques only some time after those techniques mature in the general machine-learning literature. Hakim, Gernowo, and Nirwansyah [32] complement this with a systematic review focused specifically on time-series data-mining methods for flood prediction; critically, however sophisticated the forecasting method, a time-series-only framing cannot by itself answer the intervention question central to Section VII's discussion—predicting that a flood will occur is not the same as identifying which available action would most reduce its impact. Gude, Corns, and Long [33], an earlier study using deep learning with explicit predictive-uncertainty estimation, serves as a useful point of contrast: several years on, explicit uncertainty reporting—as opposed to a single point forecast—remains inconsistently adopted across the more recent studies surveyed in [30] through [32], a regression worth naming rather than assuming the field has uniformly progressed.

Broadening from prediction to deployed infrastructure, Alshami et al. [34] review IoT applications across sustainable water and wastewater management and water-quality monitoring, and Velayudhan et al. [35] provide a comparative technological review of IoT-enabled water distribution systems; both are valuable precisely because they catalogue what is operationally mature—leak detection, quality monitoring, distribution telemetry—as distinct from what remains largely conceptual, namely the predictive and causal decision-support layers discussed in Sections VI and VII. This operational-versus-conceptual distinction recurs throughout the review and is one of its more actionable findings: sensing and monitoring infrastructure for water systems is comparatively field-proven, while the intelligence layered on top of it is comparatively not.

Widening further to city-scale decision-making infrastructure, Sheikh and Chong [36] propose a multi-agent reinforcement learning framework for optimizing smart cities as systems-of-systems, coordinating agents across transportation, energy, public safety, and communication domains, while Wu et al. [37] examine deep-learning solutions across smart-city challenges more broadly. Both illustrate decision-support approaches operating at a coarse, multi-domain-optimization grain—balancing competing city-wide objectives—rather than at the fine-grained, explainable, single-hazard-intervention grain that Section VII shows remains comparatively underdeveloped. Gaur, Agarwal, and Chatterjee's edited volume on intelligent decision support systems for smart-city applications [38] represents a third, more classical decision-support

tradition—multi-criteria decision analysis, expert systems, dashboard-based support—that has evolved largely independently of, and with limited cross-pollination with, the causal and counterfactual XAI tradition reviewed in Section VII; the urban resilience literature specifically has yet to meaningfully bridge these two decision-support traditions. DeMedeiros, Hendawi, and Alvarez [39] round out this section with a survey of AI-based anomaly detection across IoT and sensor networks, addressing a cross-cutting reliability concern—detecting corrupted or anomalous sensor data before it propagates into downstream prediction or decision support—that recurs across every domain application discussed in this section.

Table VI summarizes this section's domain literature by hazard focus, sensing maturity, prediction maturity, decision-support maturity, and reported deployment status.

TABLE VI. DOMAIN APPLICATION MATURITY SUMMARY

| Study | Domain Focus                    | Sensing Maturity  | Prediction Maturity       | Decision-Support Maturity | Deployment Status          |
|-------|---------------------------------|-------------------|---------------------------|---------------------------|----------------------------|
| [30]  | Flash flood (multi-technology)  | Moderate-High     | Moderate                  | Low                       | Mixed / fragmented         |
| [31]  | Flood early warning             | Moderate          | Moderate-High             | Low                       | Mixed (scoping review)     |
| [32]  | Flood time-series prediction    | Moderate          | High (methods-rich)       | Low                       | Mostly research-stage      |
| [33]  | Flood DL with uncertainty       | Moderate          | Moderate (w/ uncertainty) | Low                       | Research-stage             |
| [34]  | Water/wastewater IoT            | High (deployed)   | Low-Moderate              | Low                       | Operational                |
| [35]  | Water distribution IoT          | High (deployed)   | Low                       | Low                       | Operational                |
| [36]  | Multi-domain smart city (SoS)   | N/A               | N/A                       | Moderate-High (coarse)    | Simulation-validated       |
| [37]  | Smart city challenges (broad)   | Varies            | Moderate                  | Moderate                  | Mixed                      |
| [38]  | IDSS for smart cities (general) | N/A               | N/A                       | Moderate (classical IDSS) | Conceptual (edited volume) |
| [39]  | Cross-cutting anomaly detection | High (focus area) | N/A                       | Low                       | Research-stage             |

The critical pattern this section surfaces is an inversion: domain-specific application literature is operationally strongest at the sensing and monitoring end and comparatively weakest at the explainable-decision end, which is the reverse emphasis from the general AI/XAI literature reviewed in Section VII, which is comparatively strong on decision-support theory and comparatively thin on domain deployment. The two halves of the field—general-purpose reasoning research and domain-specific deployment—have not yet substantially met in the middle.

## IX. EMERGING DIRECTIONS: GENERATIVE AI AND LARGE LANGUAGE MODELS FOR URBAN DIGITAL TWINS

The three studies discussed in this section are the most recently dated in the review (2024-2025) and represent its least settled sub-literature, included because they plausibly reshape several of the gaps identified in Sections V through VIII within a short horizon. Huang, Bibri, and Keel [40] propose a generative "large flow model" for spatial artificial intelligence within urban digital twins, aimed at synthesizing plausible urban scenario data in areas of sparse sensing coverage—a direct, if still conceptual, response to the sensing-sparsity concerns raised in Section V. Xu et al. [41] provide a wider scoping review of generative AI applied to urban

digital twins, covering the autonomous generation of urban data, scenarios, designs, and 3D city models across a broader set of early efforts than [40] alone; taken together, [40] and [41] indicate how nascent this sub-area remains relative to the comparatively mature architectural, sensing, and learning literatures reviewed in Sections IV through VI.

Jiang, Ma, and Jin [42] review large language models across the full urban-data-analytics pipeline—collection, preprocessing, modeling, and post-analysis—and are the most decision-support-relevant of the three, explicitly examining disaster management as one of several domains covered. Their review plausibly points toward large language models serving as a natural-language interface layered atop the causal and counterfactual reasoning discussed in Section VII, translating structured model output into the kind of actionable guidance an emergency manager could use directly. The review itself, however, flags continuing fragmentation across the disciplines it surveys and a lack of urban-specific evaluation benchmarks, which is worth stating plainly: an interface layer is arriving before the reasoning layer it would sit on top of has matured, and the two should not be conflated. On balance, generative AI and large language model integration are best read, at this stage, as plausible future accelerants for gaps already identified elsewhere in this review—sparse sensing, natural-language decision interfaces—rather than as a resolved capability in their own right.

## **X. CRITICAL SYNTHESIS: CROSS-CUTTING STRENGTHS, LIMITATIONS, AND UNRESOLVED CHALLENGES**

The section-level critical findings above can be pulled together against six evaluative dimensions applied consistently throughout this review: scalability, adaptability, explainability, interoperability, sustainability, and security and trust. Table VII summarizes this synthesis; the discussion below justifies each rating.

**Scalability.** The architectural literature reviewed in Section IV scales conceptually—the five-dimension and communications-oriented models in principle accommodate arbitrarily many physical entities—but few of the reviewed studies report field trials extending beyond a single site or zone. The multi-agent and federated-learning approaches reviewed in Sections VI and VIII are the clearer exceptions, being explicitly designed around cross-zone or cross-agency scale from the outset rather than scaling being an afterthought.

**Adaptability.** Adaptive and continual-learning techniques (Section VI) are strong at the algorithmic level, with well-developed drift-detection and federated-adaptation methods. However, this adaptability is rarely coupled, in the studies reviewed, to the reliability-weighted fusion layer of Section V that would determine whether the data being adapted to is trustworthy in the first place—a model can adapt precisely and confidently to sensor noise.

**Explainability.** Theoretically mature (Section VII), with well-developed post-hoc, contrastive, and causal-inference literatures, but comparatively thin in domain deployment (Section VIII): where domain-specific application papers mention explainability at all, it is typically the weaker, post-hoc-attribution variety rather than the counterfactual or causal reasoning that would support actionable intervention guidance.

**Interoperability.** Explicitly identified as a structural bottleneck by the standards-focused literature ([16]), but rarely treated as a first-class design constraint by prediction- or decision-focused papers elsewhere in the review, most of which implicitly assume clean, already-integrated data arrival.

**Sustainability.** Addressed thoroughly at the level of environmental and energy outcomes by the digital-twin sustainability literature ([9], [10]), but rarely connected to the computational sustainability of the intelligence layer itself. Barbierato and Gatti's methodological survey of the Green AI literature [46] confirms this is not merely an oversight specific to urban resilience: even in the general machine-learning literature, systematic reporting of energy and carbon cost remains the exception, a "Red AI" default the authors trace across the surveyed corpus rather than a problem this review's domain invented. The energy and compute cost of continuously fusing, retraining, and reasoning over urban-scale data streams is essentially unmeasured across the urban-resilience studies surveyed in this paper specifically, a blind spot that cuts across nearly every section of this review.

**Security and Trust.** A dimension this review's earlier sections surface but have not yet evaluated explicitly. Kim et al.'s literature review of smart-city cybersecurity [45] finds research attention concentrated on smart homes and general IoT security, with cyber-forensics capability specific to city-scale infrastructure comparatively underdeveloped; almost none of the sensing, fusion, learning, or decision-support papers reviewed in Sections IV through VIII discuss adversarial robustness as part of their design. This is a meaningful omission rather than a peripheral one: a reliability-aware fusion layer defends against sensor failure, not against a sensor deliberately fed false readings, and a causally reasoned decision-support recommendation is only as trustworthy as the data pipeline feeding it. Resilience and security are, in this sense, coupled properties that the reviewed literature treats separately.

TABLE VII. CROSS-CUTTING EVALUATION OF THE REVIEWED LITERATURE BY DIMENSION

| Dimension        | Maturity Rating                                       | Justification                                                                                                                                               | Refs           |
|------------------|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| Scalability      | Developing                                            | Conceptually scalable architectures; few multi-site field trials                                                                                            | [5],[6],[36]   |
| Adaptability     | Mature (algorithmic) / Emerging (integrated)          | Strong drift/FL methods; rarely coupled to reliability-aware fusion                                                                                         | [17]-[19],[21] |
| Explainability   | Mature (theory) / Emerging (deployment)               | Strong XAI/causal theory; domain papers mostly use weaker post-hoc methods                                                                                  | [22]-[29]      |
| Interoperability | Emerging                                              | Identified as bottleneck but rarely a first-class design constraint                                                                                         | [16]           |
| Sustainability   | Developing (environmental) / Emerging (computational) | Strong on environmental outcomes; AI-layer compute cost unmeasured even in the general ML literature                                                        | [9],[10],[46]  |
| Security & Trust | Emerging                                              | Smart-city cybersecurity research concentrated on homes/general IoT; adversarial robustness absent from nearly all reviewed fusion/learning/decision papers | [45]           |




Fig. 2. Convergence of the three core intelligence capabilities and the shared integration gap.

Figure 2 depicts this synthesis visually, showing the four intelligence layers reviewed in Sections V through VIII as overlapping domains of active research, with the comparatively sparse shared intersection—systems that are simultaneously reliability-aware, adaptive, causally explainable, and field-deployed—labeled as the integration gap that motivates the research directions set out in the next section.

## XI. RESEARCH GAPS AND FUTURE DIRECTIONS

The central finding of this review, evidenced independently at the level of each layer (Sections V through VII) and confirmed at the level of domain deployment (Section VIII) and cross-cutting evaluation (Section X), is that the literature has matured three capabilities largely in isolation from one another: reliability-aware, adaptive sensor fusion; continual, drift-aware self-learning under non-stationary conditions; and causal or counterfactual reasoning for actionable decision support. Almost no study reviewed in this paper integrates all three within a single, continuously operating urban resilience system. This is not an assumption offered without support; it follows directly from the parallel-but-rarely-cross-cited literatures documented in Sections V through VII, the operational-versus-conceptual mismatch documented in Section VIII, and the post-hoc-versus-structural confusion around "explainability" documented in Section VII.

Several future directions follow from this evidence at the level of generality a literature review can responsibly support, without prescribing any single implementation. First, sensor fusion architectures for urban hazard monitoring would benefit from explicitly propagating reliability estimates into downstream learning, rather than treating fusion as a preprocessing step whose output is handed to a learning model as if it were uniformly trustworthy. Second, self-learning prediction models need principled, triggered retraining

strategies validated on field environmental data streams rather than benchmark data streams alone, with particular attention to recurring rather than purely novel drift, following the recall-oriented framing discussed in Section VI. Third, decision-support layers for urban resilience need to move beyond post-hoc feature attribution toward genuine counterfactual and causal intervention analysis grounded in a structural model of the hazard process, rather than a model-relative explanation of a black-box predictor. Fourth, none of the above is deployable city-wide without parallel progress on cross-agency interoperability and federated data-sharing arrangements, which remain, per Sections V and VI, more often acknowledged as a challenge than actually resolved. Fifth, the computational and energy footprint of continuously adaptive urban AI systems deserves explicit measurement and reporting, an omission that cuts across nearly the entire reviewed corpus.

Taken together, these converging gaps—in sensor-level reliability awareness, continual adaptation under drift, and causal decision support, and the near-total absence of studies integrating all three for urban hazard resilience—motivate continued research at precisely this intersection, particularly within flood and water-related urban resilience contexts where the domain literature reviewed in Section VIII shows sensing infrastructure to be comparatively mature but adaptive, explainable decision support to remain comparatively underdeveloped.

## XII. CONCLUSION

This review has synthesized 46 peer-reviewed studies, published predominantly between 2020 and 2026, into a critical account of AIoT-enabled digital twins for urban resilience organized around a purpose-built, five-layer taxonomy: architecture, sensing and fusion, adaptive learning, explainable and causal decision support, and domain deployment. Rather than treating these as separate literatures to be summarized in turn, the review has traced how they relate—and, more often, fail to relate—to one another, evaluating the field consistently across scalability, adaptability, explainability, interoperability, sustainability, and security and trust.

The review's central finding is one of maturity without integration. Reliability-aware sensor fusion, continual and federated adaptive learning, and counterfactual or causal decision support have each, independently, become genuinely sophisticated research areas. What the literature has not yet produced, in any volume, is a single system that combines all three and operates continuously in a real urban hazard context. The bottleneck facing AIoT-enabled digital twins for urban resilience is therefore no longer, principally, a shortage of component techniques; it is the harder, less glamorous work of integrating those components, validating the result under field rather than benchmark conditions, and reporting honestly on the computational and organizational cost of doing so. The research directions set out in Section XI are offered as a starting map for that work.

## REFERENCES

- [1] Amirzadeh, M., Sobhaninia, S., Sharifi, A. 2022. Urban resilience: A vague or an evolutionary concept? *Sustainable Cities and Society*, 81: 103853.
- [2] Deng, T., Zhang, K., Shen, Z.-J.M. 2021. A systematic review of a digital twin city: A new pattern of urban governance toward smart cities. *Journal of Management Science and Engineering*, 6(2): 125-134.
- [3] Minerva, R., Lee, G.M., Crespi, N. 2020. Digital twin in the IoT context: A survey on technical features, scenarios, and architectural models. *Proceedings of the IEEE*, 108(10): 1785-1824.
- [4] Fuller, A., Fan, Z., Day, C., Barlow, C. 2020. Digital twin: Enabling technologies, challenges and open research. *IEEE Access*, 8: 108952-108971.
- [5] Qi, Q., Tao, F., Hu, T., Anwer, N., Liu, A., Wei, Y., Wang, L., Nee, A.Y.C. 2021. Enabling technologies and tools for digital twin. *Journal of Manufacturing Systems*, 58: 3-21.
- [6] Mihai, S., Yaqoob, M., Hung, D.V., Davis, W., Towakel, P., Raza, M., Karamanoglu, M., Barn, B., Shetve, D., Prasad, R.V., et al. 2022. Digital twins: A survey on enabling technologies, challenges, trends and future prospects. *IEEE Communications Surveys & Tutorials*, 24(4): 2255-2291.
- [7] Rathore, M.M., Shah, S.A., Shukla, D., Bentafat, E., Bakiras, S. 2021. The role of AI, machine learning, and big data in digital twinning: A systematic literature review, challenges, and opportunities. *IEEE Access*, 9: 32030-32052.
- [8] Weil, C., Bibri, S.E., Longchamp, R., Golay, F., Alahi, A. 2023. Urban digital twin challenges: A systematic review and perspectives for sustainable smart cities. *Sustainable Cities and Society*, 99: 104862.

- [9] Bibri, S.E., Huang, J., Jagatheesaperumal, S.K., Krogstie, J. 2024. The synergistic interplay of artificial intelligence and digital twin in environmentally planning sustainable smart cities: A comprehensive systematic review. *Environmental Science and Ecotechnology*, 20: 100433.
- [10] Bibri, S.E., Huang, J., Omar, O., Kenawy, I. 2025. Synergistic integration of digital twins and zero energy buildings for climate change mitigation in sustainable smart cities: A systematic review and novel framework. *Energy and Buildings*, 333: 115484.
- [11] Ariyachandra, M.R.M.F., Wedawatta, G. 2023. Digital twin smart cities for disaster risk management: A review of evolving concepts. *Sustainability*, 15(15): 11910.
- [12] Siam, M.A.I., Ahn, H., Liu, L., Alam, S.M.I., Shen, H., Cao, Z., Shroff, N., Krishnamachari, B., Srivastava, M., Zhang, M. 2024. Artificial intelligence of things: A survey. *ACM Transactions on Sensor Networks*, 21(1): Article 1.
- [13] Vivek Menon, U., Kumaravelu, V.B., Kumar C, V., Rammohan, A., Chinnadurai, S., Venkatesan, R., Hai, H., Selvaprabhu, P. 2025. AI-powered IoT: A survey on integrating artificial intelligence with IoT for enhanced security, efficiency, and smart applications. *IEEE Access*, 13: 50296-50339.
- [14] Khaleghi, B., Khamis, A., Karray, F.O., Razavi, S.N. 2013. Multisensor data fusion: A review of the state-of-the-art. *Information Fusion*, 14(1): 28-44.
- [15] Desikan, J., Singh, S.K., Jayanthiladevi, A., Bhushan, S., Rishiwal, V., Kumar, M. 2025. Hybrid machine learning-based fault-tolerant sensor data fusion and anomaly detection for fire risk mitigation in IIoT environment. *Sensors*, 25(7): 2146.
- [16] Rita, J., Salvado, J., Rocha, H.d., Espírito-Santo, A. 2025. A comprehensive review of IoT standards: The role of IEEE 1451 in smart cities and smart buildings. *Smart Cities*, 8(4): 108.
- [17] Arora, S., Rani, R., Saxena, N. 2024. A systematic review on detection and adaptation of concept drift in streaming data using machine learning techniques. *WIREs Data Mining and Knowledge Discovery*, 14(4): e1536.
- [18] Suárez-Cetrulo, A.L., Quintana, D., Cervantes, A. 2023. A survey on machine learning for recurring concept drifting data streams. *Expert Systems with Applications*, 213: 118934.
- [19] Alterkawi, L., Dib, F.K. 2025. Federated learning for smart cities: A thematic review of challenges and approaches. *Future Internet*, 17(12): 545.
- [20] Al-Huthaifi, R., Li, T., Huang, W., Gu, J., Li, C. 2023. Federated learning in smart cities: Privacy and security survey. *Information Sciences*, 632: 833-857.
- [21] Ramu, S.P., Boopalan, P., Pham, Q.-V., Maddikunta, P.K.R., Huynh-The, T., Alazab, M., Nguyen, T.T., Gadekallu, T.R. 2022. Federated learning enabled digital twins for smart cities: Concepts, recent advances, and future directions. *Sustainable Cities and Society*, 79: 103663.
- [22] Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., Herrera, F. 2020. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58: 82-115.
- [23] Yang, W., Wei, Y., Wei, H., Chen, Y., Huang, G., Li, X., Li, R., Yao, N., Wang, X., Gu, X., et al. 2023. Survey on explainable AI: From approaches, limitations and applications aspects. *Human-Centric Intelligent Systems*, 3: 161-188.
- [24] Rawal, A., McCoy, J., Rawat, D.B., Sadler, B.M., Amant, R.S. 2022. Recent advances in trustworthy explainable artificial intelligence: Status, challenges, and perspectives. *IEEE Transactions on Artificial Intelligence*, 3(6): 852-866.
- [25] Stepin, I., Alonso, J.M., Catala, A., Pereira-Fariña, M. 2021. A survey of contrastive and counterfactual explanation generation methods for explainable artificial intelligence. *IEEE Access*, 9: 11974-12001.
- [26] Karimi, A.-H., Barthe, G., Schölkopf, B., Valera, I. 2022. A survey of algorithmic recourse: Contrastive explanations and consequential recommendations. *ACM Computing Surveys*, 55(5): Article 95.
- [27] Yao, L., Chu, Z., Li, S., Li, Y., Gao, J., Zhang, A. 2021. A survey on causal inference. *ACM Transactions on Knowledge Discovery from Data*, 15(5): Article 74.
- [28] Brand, J.E., Zhou, X., Xie, Y. 2023. Recent developments in causal inference and machine learning. *Annual Review of Sociology*, 49: 81-110.
- [29] Kobayashi, K., Alam, S.B. 2024. Explainable, interpretable, and trustworthy AI for an intelligent digital twin: A case study on remaining useful life. *Engineering Applications of Artificial Intelligence*, 129: 107620.

- [30] Al-Rawas, G., Nikoo, M.R., Al-Wardy, M., Etri, T. 2024. A critical review of emerging technologies for flash flood prediction: Examining artificial intelligence, machine learning, internet of things, cloud computing, and robotics techniques. *Water*, 16(14): 2069.
- [31] Byaruhanga, N., Kibirige, D., Gokool, S., Mkhonta, G. 2024. Evolution of flood prediction and forecasting models for flood early warning systems: A scoping review. *Water*, 16(13): 1763.
- [32] Hakim, D.K., Gernowo, R., Nirwansyah, A.W. 2024. Flood prediction with time series data mining: Systematic review. *Natural Hazards Research*, 4: 194-220.
- [33] Gude, V., Corns, S., Long, S. 2020. Flood prediction and uncertainty estimation using deep learning. *Water*, 12(3): 884.
- [34] Alshami, A., Ali, E., Elsayed, M., Eltoukhy, A.E.E., Zayed, T. 2024. IoT innovations in sustainable water and wastewater management and water quality monitoring: A comprehensive review of advancements, implications, and future directions. *IEEE Access*, 12: 58427-58453.
- [35] Velayudhan, N.K., Pradeep, P., Rao, S.N., Devidas, A.R., Ramesh, M.V. 2022. IoT-enabled water distribution systems-a comparative technological review. *IEEE Access*, 10: 101042-101070.
- [36] Sheikh, A.A., Chong, E.K.P. 2026. Multi-agent reinforcement learning framework for optimizing smart cities as system of systems. *Systems Engineering*, 29(1): 3-19.
- [37] Wu, P., Zhang, Z., Peng, X., et al. 2024. Deep learning solutions for smart city challenges in urban development. *Scientific Reports*, 14: 5176.
- [38] Gaur, L., Agarwal, V., Chatterjee, P. (Eds.) 2023. *Intelligent Decision Support Systems for Smart City Applications*. Wiley-Scrivener.
- [39] DeMedeiros, K., Hendawi, A., Alvarez, M. 2023. A survey of AI-based anomaly detection in IoT and sensor networks. *Sensors*, 23(3): 1352.
- [40] Huang, J., Bibri, S.E., Keel, P. 2025. Generative spatial artificial intelligence for sustainable smart cities: A pioneering large flow model for urban digital twin. *Environmental Science and Ecotechnology*, 24: 100526.
- [41] Xu, H., Omitaomu, F., Sabri, S., Zlatanova, S., Li, X., Song, Y. 2024. Leveraging generative AI for urban digital twins: A scoping review on the autonomous generation of urban data, scenarios, designs, and 3D city models for smart city advancement. *Urban Informatics*, 3(1): 29.
- [42] Jiang, F., Ma, J., Jin, Y. 2025. Unleashing the potential of large language models in urban data analytics: A review of emerging innovations and future research. *Smart Cities*, 8(6): 201.
- [43] Liu, X., Jin, J., Dong, F. 2022. Edge-computing-based intelligent IoT: Architectures, algorithms and applications. *Sensors*, 22(12): 4464.
- [44] Yu, C., Wang, Z. 2024. Multimodal social sensing for the spatio-temporal evolution and assessment of nature disasters. *Sensors*, 24(18): 5889.
- [45] Kim, K., Alshenaifi, I.M., Ramachandran, S., Kim, J., Zia, T., Almorjan, A. 2023. Cybersecurity and cyber forensics for smart cities: A comprehensive literature review and survey. *Sensors*, 23(7): 3681.
- [46] Barbierato, E., Gatti, A. 2024. Toward Green AI: A methodological survey of the scientific literature. *IEEE Access*, 12: 23989-24013.