

Article

Governing Data Integrity Risks in Digital Construction: A Qualitative Liability Framework for Generative AI Across the AEC Lifecycle and Smart City Operations

Lu Jiang ¹, Xin Wang ² and Mingliang Li ^{2,*}

¹ Law School, Tianjin University, Tianjin 300072, China; lujiang@tju.edu.cn

² School of Architecture and Urban Planning, Jilin Jianzhu University, Changchun 130118, China

* Correspondence: limingliang@jlju.edu.cn; Tel.: +86-18270912656

Abstract

Generative artificial intelligence (GenAI) can accelerate Building Information Modeling (BIM) and digital-twin workflows, but probabilistic outputs and manipulated inputs create data-integrity risks that conventional construction-software liability rules do not fully explain. This study develops a reproducible doctrinal and comparative-law method that maps the GenAI value chain onto the Architecture, Engineering, and Construction (AEC) lifecycle. Its principal advantage over single-actor product or professional-negligence analyses is a qualitative proportional-liability assessment framework that evaluates technical control, foreseeability, verification capacity, causal contribution, and evidence preservation across developers, deployers, and professional users. The analysis identifies the following two recurrent pathways of harm: autonomous hallucination entering safety-relevant design information and adversarial or erroneous data entering digital-twin feedback loops. It also specifies syntax, semantic, cybersecurity, and human-verification controls linked to AEC information-management and quality standards. The framework is not an empirically calibrated formula and does not determine legally binding percentages. Instead, it provides a transparent decision aid for applying existing tort doctrines and regulatory duties to distributed GenAI-enabled construction workflows.

Keywords: generative AI; data integrity; qualitative liability assessment; digital twins; Building Information Modeling (BIM); AEC quality management; smart city governance

Academic Editor: Antonio Caggiano

Received: 6 June 2026

Revised: 18 June 2026

Accepted: 25 June 2026

Published: 27 June 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The integration of generative artificial intelligence (AI), particularly large language models (LLMs) and diffusion algorithms, has driven a significant transformation within the Architecture, Engineering, and Construction (AEC) sector toward data-driven construction and smart city governance [1]. By leveraging algorithmic capabilities in spatial synthesis, parametric modeling, and automated text generation, generative AI optimizes the workflows and decision-making processes of project stakeholders [2]. Within the domain of smart urban infrastructure, this technological trajectory has stimulated academic and industrial discourse regarding its impact on network digital twins [3] and Building Information Modeling (BIM) ecosystems [4]. Research indicates that integrating genera-

tive design tools into urban computing minimizes decision-making costs during the planning phase and enhances the operational efficiency of lifecycle management [5] in large-scale engineering projects [6].

However, the rapid adoption of these intelligent technologies within digital construction environments [7] has introduced notable data integrity risks, which necessitate re-evaluation from a cross-disciplinary socio-legal perspective. Currently, digital challenges within the construction industry have expanded from traditional software operational errors to algorithmic hallucinations [8] during autonomous system generation [9], as well as erroneous data injections induced by improper user prompts [10]. When generative AI is embedded into complex engineering workflows or smart city platforms, inaccurate technical parameters autonomously generated by algorithms [11] pose a substantial risk to structural integrity, engineering cost estimations, and overall property safety. For instance, if unverified compliance data are injected into the interactive virtual-real mapping layer of a smart city digital twin system, they can disrupt the automated closed-loop feedback mechanisms of urban planning. Furthermore, AI-generated design drawings containing incorrect fire safety codes or mechanical metrics may embed latent vulnerabilities into public buildings, potentially leading to severe engineering failures.

Conventional AEC software, referred to here as a general category rather than a specific product used in this study, generally operates through more predictable, rule-based functions than GenAI, although software defects and contractual disputes have long complicated liability. Existing project contracts, BIM protocols [12–14], professional duties, and software terms commonly allocate responsibility for input, validation, approval, and use. GenAI adds a distinct problem because the same prompt may produce variable outputs and because model behavior may depend on changing data, retrieval sources, and versions. The legal inquiry can therefore no longer stop at the assumption that software merely executes a professional's fixed instruction.

The transition from deterministic software to probabilistic GenAI exposes a domain-specific gap in current liability analysis. The existing AI-liability literature commonly examines autonomous vehicles, medical systems, financial services, or general-purpose platforms [15–17]. AEC workflows differ because safety-relevant information moves through fragmented contractual chains, licensed professionals, model-based design environments, and cyber-physical feedback loops. A single erroneous output may be transformed, approved, and reused across several project phases before physical damage appears. These features do not make existing product liability, platform liability, or professional-negligence doctrines irrelevant. They make a lifecycle-based method necessary to identify which actor controlled the failed risk node and which duty was breached.

Emerging AI regulation supplies relevant compliance duties [18–21] but does not itself resolve AEC tort allocation. The EU AI Act establishes risk-management, documentation, logging, and human-oversight obligations for systems that fall within its scope [22]. China's Interim Measures for the Management of Generative Artificial Intelligence Services primarily govern services offered to the public and emphasize provider and content-governance responsibilities. Neither instrument creates a complete AEC-specific civil-liability regime. Their requirements are therefore treated in this study as possible evidence of reasonable care, regulatory breach, or risk control, subject to the applicable private-law rules.

Liability analysis in this setting must distinguish model development, application integration, professional adoption, and physical execution. These stages distribute control over training data, system boundaries, prompts, verification gates, and final construction decisions. The relevant legal question is therefore not whether the AI system itself should bear liability. It is how existing duties and causal responsibility should be reconstructed

across the actors who could foresee, prevent, detect, document, or correct the integrity failure.

Accordingly, this study asks three questions. First, why are existing single-actor liability models incomplete for GenAI-enabled AEC workflows? Second, which technical and organizational controls identify the actor best positioned to prevent or document a data-integrity failure? Third, how should regulatory compliance, contractual duties, and professional verification inform qualitative responsibility allocation without replacing applicable tort doctrine? The study answers these questions through a structured doctrinal and comparative analysis, a dual-track lifecycle mapping method, and bounded scenario-based stress tests.

2. Methodology and Conceptual Framework

GenAI introduces probabilistic outputs into AEC systems that manage safety-relevant information and physical assets. The methodological challenge is to connect the technical origin and movement of an integrity failure with the legal duties of the actors who develop, integrate, verify, approve, and implement the affected information. This section therefore combines a reproducible source-selection protocol with a dual-track mapping and qualitative assessment method.

2.1. Research Design and Reproducible Source-Selection Protocol

This study uses a structured doctrinal and comparative-law design. The source base comprises peer-reviewed AEC and AI-risk research, primary legal instruments, official technical and quality-management standards, and bounded real-world or legal analogs. Searches and verification were completed on 18 June 2026. Academic materials were identified through Crossref, OpenAlex, and structured literature searches using combinations of “generative AI”, “construction”, “BIM”, “digital twin”, “data integrity”, “hallucination”, “prompt injection”, “liability”, “automation bias”, “software license”, “intervening cause”, and “superseding cause”. Primary legal materials were retrieved from official sources and reported legal authorities, including EUR-Lex, the Cyberspace Administration of China, reported case law, and the Restatements of Torts.

Materials were included when they directly addressed GenAI output integrity, AEC information exchange, digital-twin security, professional verification, multi-actor liability, or the regulatory scope of the compared instruments. Marketing materials, unverified incident reports, and sources that were only nominally related were excluded. Each retained source was coded against the following five analytical dimensions: the actor controlling the risk node, the foreseeable failure mode, the available preventive or verification control, the evidence needed to reconstruct causation, and the legal function of the relevant rule. Primary legal texts control over secondary descriptions when defining regulatory scope.

The framework was constructed in four reproducible steps. First, the GenAI value chain was mapped to the planning, design, construction, and operation phases of the AEC lifecycle. Second, the two recurring integrity-failure pathways were traced through these mapped nodes. Third, responsibility was assessed using the following five qualitative factors: technical control, foreseeability, verification capacity, causal contribution, and evidence preservation. Fourth, the resulting allocation logic was conceptually illustrated and normatively tested through two hypothetical AEC scenarios, with verified legal and engineering analogs used only to define the scenarios’ practical boundaries. This procedure produces analytical findings rather than empirical or statistical results.

No Delphi study, expert interview, AHP exercise, quantitative calibration, or damages dataset were used. The qualitative responsibility bands are therefore ordinal decision-support categories, not estimates of legally binding percentages. The scenarios test

whether the framework exposes relevant controls and evidence; they do not validate predictive accuracy, external validity, or compensation shares.

2.2. Dual-Track Mapping and Framework Construction

The dual-track mapping method aligns the GenAI value chain with the AEC project lifecycle. It treats model development and project delivery as connected but distinct chains. The first identifies who controls training, model behavior, integration settings, and interface safeguards. The second identifies when information is generated, reviewed, approved, transformed, and physically implemented. Linking the two chains makes the location of control and the point of verification visible.

For analytical clarity, actors are grouped as developers, deployers, and professional or organizational users. Developers control the foundation model, training and fine-tuning processes, model documentation, and known-limit disclosures. Deployers integrate the model into BIM, digital-twin, or review platforms and control access, retrieval sources, interface constraints, logging, and validation gates. Users provide prompts, interpret outputs, approve safety-relevant information, and authorize its use in project decisions. These categories describe functional control and may overlap in a vertically integrated organization.

Mapping these functions into planning, design, construction, and operation creates an auditable chain from model source to physical consequence. The method does not presume that one actor is always primarily responsible. It identifies which actor controlled the failed safeguard, which actor had the practical capacity to verify the output, and which records are needed to establish causation. Figure 1 illustrates this dual-track structure.




Figure 1. Dual-track mapping model aligning the generative AI value chain with the AEC project lifecycle.

2.3. Syntax, Semantic, and Algorithmic-Stability Controls

Heterogeneous AEC data cannot be placed into a common information space by aggregation alone. Syntax transformation first converts source files, sensor messages, schedules, and textual outputs into declared exchange structures. In BIM workflows, Industry Foundation Classes (IFCs) provide a machine-readable data schema, while the buildingSMART Information Delivery Specification (IDS) can express required properties and support automated compliance checks [23,24]. Files that fail schema, unit, identifier, or mandatory-field validation should be quarantined before they enter the shared model.

Semantic transformation then maps equivalent concepts across disciplines and project phases. A deployer should maintain controlled vocabularies, object classifications, unit mappings, provenance fields, and explicit links between a generated statement and the authoritative code, drawing, calculation, or sensor source it purports to represent. Ambiguous or unmapped concepts should be routed to professional review rather than silently normalized. These controls reduce the risk that syntactically valid information is treated as semantically correct.

Algorithmic convergence must also be separated from engineering correctness. Repeated sampling and self-consistency can improve reasoning performance by selecting a recurring answer among several generated paths [25], but agreement among probabilistic outputs does not prove compliance with physical laws or construction codes. For safety-relevant AEC use, output stability should therefore be tested through repeated generation, independent deterministic calculations, constraint checking, and professional approval. A stable but incorrect output remains a data-integrity failure.

2.4. Typology of Data Integrity Risks in Digital Construction

Having established the mapping model, this study deconstructs the technical mechanics of the following two core data-integrity risks within the AEC domain: algorithmic hallucinations and erroneous data injections. The AEC sector is not uniformly risk-intolerant. Construction practice accepts specified tolerances, uncertainty, optimization trade-offs, and managed nonconformities. Tolerance is substantially lower; however, where an error can compromise structural stability, fire safety, or another mandatory safety function. The relevant distinction is therefore between managed engineering tolerances and failures that cross legal, contractual, or professional safety thresholds.

Algorithmic Hallucinations in Generative Design: GenAI generates outputs from learned probability distributions and does not independently guarantee compliance with physical mechanics or statutory engineering codes. In a safety-relevant workflow, a fabricated clause, parameter, or calculation may appear credible and may remain syntactically valid. The principal safeguards are authoritative-source grounding, explicit citations, deterministic constraint checks, independent calculation, model-version records, and professional approval before adoption.

Erroneous Data Injection in Digital Twins: A digital twin combines data from physical assets, models, users, and connected services through a continuing virtual-physical feedback process. Natural-language interfaces and automated integrations expand the attack and error surface [26]. A malicious prompt, compromised sensor, mapping error, or unverified external record may be accepted as a true physical state unless the deployer applies access controls, provenance checks, schema and semantic validation, anomaly detection, and a separate approval gate before operational commands are issued.

The two risk types differ not only in their technical origins, but also in their attribution logic. Table 1 classifies these data integrity risks by source, workflow manifestation, engineering hazard, and legal attribution problem. Figure 2 then visualizes the two pathways as parallel mechanisms through which digital anomalies become built-environment risks.

Table 1. Typology of data integrity risks and legal attribution problems in GenAI-enabled digital construction.

Risk Category	Technical Origin	AEC Manifestation	Engineering Hazard	Legal Attribution Problem
Autonomous algorithmic hallucination	Feature conflicts in training data; weak alignment with engineering codes; probabilistic generation under insufficient source grounding	Fabricated structural parameters, non-existent fire-safety clauses, or unreliable compliance outputs in BIM-assisted design and review	Latent vulnerabilities may enter drawings, structural calculations, or code-compliance files	Requires separating developer source-data failure from professional verification failure through model provenance and adoption records
User-induced data injection	Adversarial prompts, malicious NLP manipulation, or inadequate input-validation safeguards	Fraudulent compliance reports, distorted lifecycle assessments, or contaminated operational data in smart-city digital twins	Digital twin databases may absorb false physical-state data, producing resource misallocation, subsidy misuse, or erroneous administrative decisions	Shifts liability toward intentional user misconduct and deployer cybersecurity omissions; developer exposure may reduce if the model was not the proximate source


**Figure 2.** Dual pathways of data integrity failure in GenAI-enabled digital construction: (A) autonomous or hallucinated output path; (B) adversarial or erroneous input path. (A) Autonomous algorithmic hallucination pathway from probabilistic generation to safety-relevant design information. (B) User-induced erroneous-input pathway from adversarial or invalid input to contaminated digital-twin records.

Figure 2 leads to a specific attribution conclusion. Autonomous hallucination directs attention toward model documentation, source grounding, platform guardrails, and professional verification. User-induced injection directs attention toward intent, interface security, access control, provenance, and operational approval. In both pathways, liability cannot be inferred from actor labels alone. It must be reconstructed from the failed control, the actor capable of operating that control, and the available audit evidence.

2.5. Theoretical Contribution and Boundary Conditions: Beyond Existing AI-Liability Pathways

The theoretical contribution is not a new freestanding tort regime. It is a domain-specific translation rule that connects technical control in the GenAI-AEC information chain to existing legal concepts of duty, causation, fault, product defect, and professional responsibility. This translation is needed because no single existing theory captures the complete pathway from model development to physical construction or infrastructure operation.

The first limitation is control. GenAI systems do not operate like deterministic construction software whose behavior is fully fixed by code at the moment of delivery. Training data, retrieval sources, prompts, deployment settings, feedback loops, and model updates may all affect later outputs. In an AEC project, a harmful parameter may emerge only after a generated statement has been transformed into a BIM object, reviewed by a professional, incorporated into a drawing, or reused by a digital twin. Traditional liability theories built around a single controlling actor therefore risk either over-concentrating responsibility on the upstream provider or losing the responsible actor within a fragmented value chain [27–29].

Product liability remains relevant, especially after the revised EU Product Liability Directive expressly includes software and AI-related products within its modernized framework [30]. It is nevertheless incomplete for GenAI-enabled AEC workflows. Product liability usually presupposes an identifiable product, a defect, a producer or seller, and a relatively linear link between defect and damage. In GenAI-enabled construction, the legally relevant failure may result from model design, training data, semantic mapping, platform integration, professional verification, or post-deployment use. A single product-defect inquiry cannot by itself explain how responsibility should be distributed across these stages.

Platform-liability theory is also insufficient. Platform duties are commonly organized around an intermediary role and a reactive obligation to manage third-party content, transactions, or user conduct. GenAI tools embedded in BIM and digital-twin systems do more than host or transmit third-party information. They actively generate technical statements, classify sensor data, suggest design parameters, and influence operational decisions. The central legal problem is therefore not only whether an intermediary failed to remove harmful content. It is how to attribute a safety-relevant output that has been generated, transformed, approved, and operationalized across several actors.

Risk-governance theory supplies useful ex-ante controls, but it does not independently answer ex-post private-law questions. Risk-management duties, documentation, logging, and human oversight can identify expected precautions and evidentiary trails. They do not determine whether a professional duty was non-delegable, whether an upstream defect and downstream verification failure were concurrent causes, how a claimant proves causation, or whether liability should be joint, several, or contractually allocated. Similarly, a broad strict-liability approach to all AI providers would be disproportionate where the core problem is not a single abnormal danger but information asymmetry, distributed control, and proof difficulty.

The dual-track mapping model advances existing theories by treating lifecycle position and functional control as the bridge between technical evidence and legal attribution. The qualitative proportional-liability assessment framework then organizes that bridge through the following five explicit factors: technical control, foreseeability, verification capacity, causal contribution, and evidence preservation. Its academic value lies in integrating doctrines that otherwise operate separately and in showing where each existing legal pathway stops explaining the AEC-specific risk chain.

The framework remains bounded. It does not override mandatory rules on joint and several liability, contractual allocation, professional licensing, product liability, or evidentiary burdens. It also does not assign numerical percentages. Courts, regulators, and contracting parties must apply the relevant jurisdictional rules to the established facts.

3. Doctrinal and Scenario-Based Findings

The structured analysis produces three doctrinal findings. First, GenAI-related AEC harm is usually distributed across several technical and professional control points rather than one autonomous machine act. Second, responsibility becomes more transparent when each failed safeguard is mapped to the actor able to implement, verify, or document it. Third, regulatory compliance and technical standards can inform the standard of care, but they do not mechanically determine private tort liability. These are analytical findings from doctrinal coding and scenario-based stress testing, not empirical or statistical results.

3.1. Systemic Rupture of Traditional Tort Elements Under Algorithmic Anomaly Scenarios

Traditional tort elements remain applicable, but their application becomes evidentially difficult when probabilistic outputs and cyber-physical feedback loops distribute decisions across several actors. A wrongful act may consist of defective model design, unsafe integration, inadequate verification, intentional manipulation, or a combination of these acts. Damage may remain latent, while causal proof may depend on model versions, prompts, transformations, approval records, and operational logs.

Product liability, platform responsibility, contract, and professional negligence each explain part of this chain. None should be presumed to fail categorically. The central difficulty is identifying the legal function of each actor after the same output has been generated, integrated, reviewed, transformed, and physically implemented. The four-element analysis in Table 2 therefore focuses on the evidence required to apply existing tort rules to a distributed workflow.

Doctrinal rupture can be expressed through the four elements of traditional tort liability. Table 2 shows why each element becomes unstable when probabilistic outputs and cyber-physical feedback loops are embedded into AEC workflows.

Table 2. Systemic rupture of traditional tort elements by generative AI in the built environment.

Traditional Tort Element	Traditional AEC Application	Disruption by Generative AI	Doctrinal Challenge
Wrongful act	Defective drawing, static software error, or identifiable professional drafting omission	AI outputs are dynamic, probabilistic, and affected by model updates, retrieval contexts, and prompt interactions	Difficult to classify a generated output as a static defective product or discrete wrongful act
Fact of damage	Physical defect, financial loss, code non-compliance, or project delay linked to a known actor	Digital anomalies may remain latent until construction, operation, or public-infrastructure feedback loops reveal the defect	Damage recognition may lag behind the original data corruption, complicating limitation periods and evidence
Causal link	A direct chain connects human decision, design document, construction action, and damage	BIM and digital twins distribute false data through multiple actors, databases, and feedback loops	Proximate cause requires prompt logs, model versions, audit trails, and data-origin verification
Subjective fault	Professional negligence, intent, or organizational omission is attributed to a recognized actor	Autonomous model behavior and automation bias blur the line between machine output, platform design, and human reliance	Fault must be reconstructed through control, foreseeability, auditability, and professional verification duties

3.2. Qualitative Proportional-Liability Assessment Framework

The proposed matrix is a qualitative proportional-liability assessment framework. It does not reject joint and several liability where applicable and does not calculate numerical shares. The high, medium, and low bands are ordinal descriptions of legal significance. They are not numerical scores and must not be added, averaged, or translated into compensation percentages.

The assessment proceeds actor by actor across five separate factors. Technical control concerns the actor's practical and legally relevant ability to prevent, modify, suspend, detect, or mitigate the failed process at the time when intervention remained possible. Foreseeability concerns whether the type of failure and harm was reasonably identifiable from the actor's knowledge, role, warnings, prior incidents, and intended use. Verification capacity concerns the actor's authority, information, competence, tools, and practical opportunity to test or review the relevant output or process.

Causal contribution concerns whether the actor's conduct was necessary, substantial, or risk-increasing in the sequence that produced the legally recognized harm, subject to the causation test of the governing jurisdiction. Evidence preservation concerns the actor's control over records needed to reconstruct the event and whether those records were reasonably created, retained, disclosed, altered, or lost. This factor supports forensic attribution but does not replace proof of duty, breach, causation, or damage.

A high band indicates strong and direct support for greater relative responsibility under the relevant factor. A medium band indicates material but shared, indirect, qualified, or evidentially uncertain support. A low band indicates limited practical capacity, remote foreseeability, attenuated causal involvement, or little responsibility for the relevant records. Every classification must identify the supporting facts, contrary evidence, and the applicable legal or professional rule.

Application follows three stages. First, the decision-maker separates established facts from disputed or unavailable evidence for each actor. Second, each factor is classified with a short reasoned statement. Third, the five classifications are synthesized under the governing doctrine. The factors are not presumed to carry equal weight because their legal relevance depends on the claim, jurisdiction, protected interest, and evidence.

The final assessment is a reasoned comparative judgment rather than arithmetic aggregation. Converging high classifications may support a stronger attribution, while mixed classifications require an explanation of which factors are legally decisive and why. No combination compels a predetermined result. The framework is intended to improve transparency, consistency, and reviewability without claiming empirical calibration or mathematical precision.

Table 3 identifies actor-specific control and fault indicators rather than predetermined liability shares. A developer's responsibility may increase where a known model limitation, deficient documentation, or avoidable source-grounding defect materially contributes to the output. A deployer's responsibility may increase where unsafe integration, missing validation gates, weak access control, or inadequate logging permits the output to enter the project workflow. A professional or organizational user's responsibility may increase where safety-relevant information is intentionally manipulated or adopted without the verification required by law, contract, or professional practice. The five operational definitions and the three-stage procedure above must be applied to the evidence in the individual case.

Table 3. Actor-specific controls and fault indicators for applying the qualitative proportional-liability assessment framework.

Stakeholder in AI Value Chain	Primary Scope of Liability (Risk Control Phase)	Standard of Care and Fault Attribution Criteria (Indicators of Negligence)
Foundation Model Developers (e.g., Algorithm Providers)	Source data integrity and algorithmic logic during the pre-training and fine-tuning phases.	Potential breach of model-development and documentation duties: - Failure to address known, material data-quality defects or disclose relevant model limitations; - Failure to support authoritative-source grounding or validated constraint integration for the intended safety-relevant use; - Inadequate testing, documentation, monitoring, or warning for foreseeable unsafe outputs.
Platform Deployers (e.g., BIM Cloud Services, Smart City Operators)	System integration, cybersecurity architecture, and administrative safeguards within automated workflows.	Potential breach of integration, cybersecurity, and workflow-control duties: - Inadequate access control, input validation, provenance checking, or prompt-injection mitigation; - Failure to implement required human approval and deterministic validation gates; - Inadequate logging, anomaly detection, incident response, or independent audit.
End Users (e.g., Registered Architects, Cost Engineers, Contractors)	Terminal prompt inputs, adoption of AI-generated outputs, and physical execution in engineering nodes.	Potential professional negligence or intentional misconduct: - Adopting safety-relevant AI outputs without the independent verification required by law, contract, or professional practice; - Approving or transmitting unverified information into construction or operational records. Intentional manipulation or malicious prompt injections to bypass compliance.

3.2.1. Regulatory Context and Its Evidentiary Role

Regulatory compliance may inform evidence of reasonable care or breach, but it does not mechanically allocate civil liability. The legal effect depends on the applicable private-law regime, the protected interest, causation, and the relationship between the breached rule and the harm. The EU and Chinese instruments considered below therefore operate as contextual inputs to the qualitative assessment rather than liability formulas.

The EU AI Act is principally a regulatory compliance framework. It imposes obligations according to the system's classification and the actor's role, including risk management, documentation, logging, data governance, and human oversight where applicable [22]. Not every GenAI tool used in AEC is automatically high-risk. A specific use must fall within the Act's scope and classification rules. The Act also preserves other Union and national remedies. Civil compensation must therefore be analyzed under the applicable tort, contract, and product-liability rules, including Directive (EU) 2024/2853 where its conditions are satisfied [30].

China's Interim Measures primarily govern GenAI services offered to the public in China and emphasize provider duties, lawful content, training-data governance, and service security [31]. The Measures expressly exclude research, development, and use that do not provide services to the domestic public. They do not establish a specific AEC tort-liability matrix. For a public-facing service within scope, compliance or non-compliance


may provide relevant evidence. For internal AEC systems outside scope, responsibility must instead be derived from other applicable civil, cybersecurity, data, construction, contract, and professional rules.

3.2.2. Contractual Risk Allocation and Forensic Evidentiary Challenges in AEC Workflows

Contractual allocation and forensic evidence constrain the qualitative assessment. Construction contracts, BIM execution plans, common-data-environment rules, professional appointments, and software-license terms can assign responsibilities for model use, validation, approval, record keeping, warranties, and remedies. Software-license terms may limit warranties or remedies between contracting parties, but they do not categorically eliminate civil liability. In *M.A. Mortenson Company, Inc. (Minneapolis, MN, USA) v. Timberline Software Corporation (Beaverton, OR, USA)*, the reported construction-estimating-software dispute illustrates that the effect of a license limitation depends on contract formation, wording, governing law, the type of loss, and applicable mandatory rules [32]. Such terms do not automatically determine duties owed to non-contracting third parties.

Forensic reconstruction should identify the model version, retrieval sources, prompts, transformations, validation results, approvals, and final project decision. Missing or unreliable records may prevent a claimant or defendant from establishing causation and fault. Tamper-evident logs and signed provenance can improve traceability, but they do not make the record complete or establish liability by themselves.

Figure 3 translates the five assessment factors into a qualitative screening pathway. It identifies a presumptive focus for doctrinal analysis and expressly avoids categorical or mechanical liability outcomes. Final responsibility remains subject to the applicable law, contract, causation, professional duties, and available evidence.



Screening aid only: liability remains subject to applicable tort, contract, product-liability, professional, causation, and evidentiary rules.

Figure 3. Qualitative fault-attribution screening pathway for identifying the presumptive focus of doctrinal analysis.

3.3. Doctrinal Anchors for Fault Attribution in Comparative Jurisprudence

The qualitative framework operationalizes existing doctrines rather than replacing them. Product liability, contractual allocation, professional negligence, network or platform duties, and general tort principles remain the legal anchors. The framework's function is to identify which doctrine and evidence correspond to each failed control point.

For developers, relevant questions include whether the software or AI system is legally treated as a product, whether a defect or inadequate instruction existed, whether the intended or reasonably foreseeable use was disclosed, and whether that defect caused compensable damage. The revised EU Product Liability Directive includes software and AI-related products, but liability still depends on the directive's conditions and national implementation [30]. Regulatory documentation or data-governance breach may be relevant evidence; it is not an automatic finding of tort liability.

For deployers, the relevant duty arises from their active control over system integration, access, retrieval sources, validation gates, logging, and operational interfaces. The legal basis may arise from contract, general negligence, cybersecurity and data obligations, sector regulation, or other applicable rules. Platform status alone does not establish liability. The inquiry concerns the deployer's actual control, the foreseeability of the failure, and whether reasonable safeguards were omitted.

For AEC professionals and organizational users, existing professional and contractual duties remain central. Automation bias is a human-factors explanation for inappropriate reliance on automated output [33,34]; it is not an independent legal wrong and does not itself establish negligence, gross negligence, or recklessness. A breach can arise only where the applicable law, contract, or professional standard required verification and the actor failed to satisfy that requirement.

The legal bridge requires four inquiries. First, was independent verification required by a statutory, contractual, project, or professional rule? Second, were the tool limitation and safety consequence reasonably foreseeable? Third, did the actor have authority, information, competence, and a practical opportunity to verify the output? Fourth, did the omission fall below the governing standard of reasonable professional care and materially contribute to the harm? General negligence is assessed against reasonable care under the circumstances, while any aggravated classification requires the additional jurisdiction-specific threshold for recklessness or gross negligence [35].

Expert verification is a key attribution point, but it is not the only one. A professional omission does not erase an upstream defect or unsafe deployment, and an upstream defect does not excuse a non-delegable professional duty. Repeated warnings, a safety-critical decision, an obvious contradiction, or deliberate bypass of a mandatory review may strengthen the inference of serious fault. A single understandable cognitive error does not. The qualitative framework permits concurrent responsibility where several failures materially contribute to the same harm.

3.4. Scenario Analysis: Tort Liability Attribution in Engineering Practice

Within the databases, search terms, and cutoff date stated in Section 2.1, this study identified no reported judicial decision adjudicating physical AEC damage caused by a large language model or other GenAI system. A recent taxonomy of generative-AI litigation identifies reported claims concentrated in copyright, unjust enrichment, negligence, and unfair competition rather than physical construction failure [36]. This is a bounded search result rather than a universal assertion that no such dispute exists. The two scenarios are therefore normative illustrations, not empirical validation. Their reasoning is bounded by the following two verified analogs: *M.A. Mortenson Co. v. Timberline Software Corp.*, reported at 140 Wn.2d 568, 998 P.2d 305 (2000), and a documented digital-

twin-supported bridge-monitoring case [32,37]. Neither analogue proves the proposed allocation.

3.4.1. Scenario A: Algorithmic Hallucination and Structural Design Failure in BIM Collaborative Workflows

Scenario A considers a structural engineering team using a GenAI parametric plugin within a cloud BIM platform. The model combines incompatible or outdated engineering sources and generates a reinforcement parameter below the applicable code threshold. The hypothetical output is not itself treated as proof of a product defect. The scenario asks which actor controlled authoritative-source grounding, rule-based constraint checking, professional verification, and final drawing approval.

The generated parameter becomes hazardous only after it passes several preventable control points. The platform does not flag the non-compliant value, and the responsible professional adopts it without the independent calculation and review required for safety-relevant design [38]. The error is then incorporated into approved construction information and later requires remedial work.

Applying the qualitative framework does not predetermine the legal result. Developer responsibility would increase only if evidence establishes a material model or documentation defect connected to the output. Deployer responsibility would increase if a reasonably required validation gate was omitted. Professional-user responsibility would increase if the applicable verification duty was breached. Contract terms, mandatory law, causation, and the available logs would determine the final legal allocation.

3.4.2. Scenario B: Adversarial Prompt Injection in Municipal Building-Compliance Administration

A parallel liability dilemma arises when a municipal digital-twin platform is used as an administrative decision-support system for building records, lifecycle assessments, and green-building incentives [39–41]. Consider a contractor seeking a tax rebate or public subsidy for a building that does not satisfy the applicable performance criteria. The contractor submits a manipulated natural-language report through an interface that uses GenAI to summarize supporting documents and recommend an eligibility classification.

The contractor uses adversarial instructions and fabricated supporting text to induce the system to disregard conflicting sensor records and classify the building as compliant. Because the deployer has not separated generated summaries from authoritative records or required independent approval, the inaccurate classification is written into the digital-twin administrative record and relied on in a payment or rebate decision. The scenario concerns erroneous administrative classification and financial loss; it does not assume that natural-language GenAI output directly controls a municipal power grid or other safety-critical infrastructure.

Scenario B shifts the dominant control failure toward intentional user manipulation and unsafe deployment. The contractor may occupy the highest responsibility band if intent, reliance, and causation are proved. The municipal operator or platform deployer may also bear responsibility if it exposed an operational interface without proportionate authentication, source validation, separation of duties, logging, or human approval. Developer responsibility may be lower where the dominant failure arose from deployment-specific controls, but it cannot be excluded without evidence concerning known vulnerabilities, documentation, and feasible mitigations.

Malicious input should not automatically be classified as a superseding cause. Under the Restatement (Second) of Torts, intervening and superseding causation turns on matters including foreseeability, the scope of the original risk, the independence and extraordinary character of the later act, and the relationship between the omitted safeguard and

the intervention [42]. If prompt injection or falsified input was a foreseeable reason for requiring authentication and validation, the malicious act may coexist with, rather than necessarily displace, upstream responsibility. The terminology and legal effect remain subject to the governing jurisdiction.

A comparison of the two scenarios clarifies how the center of liability changes when the dominant anomaly source shifts from autonomous model generation to intentional user manipulation. Table 4 summarizes the forensic distinctions, and Figure 4 visualizes the resulting allocation pattern.

Table 4. Comparative forensic analysis of scenario-based liability allocations.

Scenario Parameter	Scenario A: Algorithmic Hallucination in BIM	Scenario B: Adversarial Injection in Municipal Compliance Administration
Locus of control	Autonomous model generation followed by professional verification failure	Human-driven adversarial manipulation at the digital-twin interface
Technical defect	Polluted or contradictory training data; lack of hard alignment with structural code thresholds	Absence of authentication, prompt-injection defense, authoritative-record validation, separation of duties, and human approval before administrative classification
Primary liable actor	Foundation model developer and registered structural engineer, depending on proof of source defect and verification omission	Contractor or applicant if intentional manipulation and fraudulent submission are proved
Secondary liable actor	BIM platform deployer if rule-based guardrails or alert mechanisms were absent	Municipal digital-twin operator or platform deployer if access, validation, approval, and logging safeguards were inadequate
Reduced-liability actor	Deployer liability may reduce where safeguards, logs, and expert review gates were implemented	Developer responsibility may be lower where evidence shows that deployment-specific malicious manipulation was the dominant causal factor
Evidentiary requirement	Training-data records, model version, generated parameter logs, verification notes, and final drawing approval history	Prompt history, interface logs, source-document records, eligibility criteria, approval trail, payment record, and cybersecurity incident evidence




Figure 4. Scenario-based liability allocation under algorithmic hallucination and adversarial data-injection cases: (A) hallucination in BIM collaborative workflows; (B) adversarial prompt injection in smart city digital twins.

4. Discussion

The findings support a layered governance approach that combines private-law remedies, contractual duties, professional standards, and proportionate technical controls. Ex-post tort analysis remains necessary after harm, but prevention and evidence preservation determine whether a distributed GenAI-AEC failure can be detected, contained, and legally reconstructed.

4.1. Complementary Dynamics Between Private Remedy and Public Regulation

Private-law remedies and public-law controls perform different functions. Tort and contract rules allocate responsibility after a legally recognized loss, while regulation, standards, and project governance can reduce risk before harm occurs and preserve evidence if an incident arises. In AEC workflows, the most useful connection between these functions is a documented control chain that links each safety-relevant output to an accountable reviewer and an auditable decision.

4.2. Risk-Based Oversight of Automated Drawing Review

Oversight of automated drawing review should focus on the safety significance and function of the system rather than assume that every GenAI tool requires a new license. For systems that generate or assess safety-relevant construction information, regulators and project owners can require documented intended use, authoritative-source grounding, validated constraints, human approval, event logging, incident reporting, and periodic review.

The strongest controls should apply before an AI output can alter an approved drawing, compliance decision, or operational command. A non-bypassable human approval gate is important [43], but it must be paired with an identified competent reviewer, access to the underlying source and calculation, and a recorded decision. Human presence without authority, information, or time to verify the output does not provide effective oversight.

4.3. Layered Protection for Smart City Digital Twins

A digital twin should be protected as a set of connected resources rather than as a single trusted object [44,45]. Proportionate controls include least-privilege and zero-trust access, authenticated data sources, signed provenance, schema and semantic validation, independent anomaly detection, separation between analytical outputs and operational commands, and human approval for safety-relevant actions. NIST's AI risk-management resources and Zero Trust Architecture provide general control principles that can be adapted to this setting [46–48]. ISO 19650-5 provides a security-minded approach to information management in the built environment [49].

Auditability is equally important. Deployers should preserve model versions, retrieval sources, prompts, transformations, rejected inputs, approval records, and operational changes in tamper-evident logs. Incident response should permit the digital twin to enter a safe state and prevent unverified information from reaching physical control systems. Independent audits should test whether these controls operate in practice, while recognizing that no technical audit eliminates the need for professional judgment.

4.4. Alignment with Construction Quality and Information-Management Standards

The proposed controls can be linked to existing quality-management indicators rather than presented as an isolated digital solution. GB/T 50430-2017 supplies a construction-enterprise quality-management benchmark for documented responsibilities, process control, inspection, nonconformity handling, and continual improvement [50]. Within a

GenAI-enabled workflow, these indicators correspond to assigned approval authority, recorded verification, traceable corrective action, and periodic control review.

GB/T 51212-2016 and GB/T 51301-2018 provide examples of BIM application and design-delivery requirements [51,52]. Their information-management logic supports declared exchange requirements, consistent model objects and attributes, delivery checks, version control, and accountable review. The proposed syntax and semantic controls can therefore be incorporated into BIM execution plans, common data environments, design-delivery checks, and project quality records. Standards inform the standard of care where applicable or contractually incorporated, but they do not independently determine tort liability.

4.5. Regulatory Boundary Conditions for the Private Liability Framework

The EU and Chinese instruments are relevant only as boundary conditions for the private liability framework. The comparison asks whether an instrument applies to the use case, which actor controls the regulated function, and whether the resulting compliance record is relevant to duty, breach, causation, or evidence preservation. It is not an independent account of transnational AI governance.

Under the EU AI Act, any evidentiary effect depends on the system's function, classification, and actor role [22]. Applicable duties concerning risk management, documentation, logging, data governance, and human oversight may help identify expected controls. They do not determine damage allocation, and an AEC tool is not high-risk merely because it is used in construction.

China's Interim Measures apply to GenAI services offered to the public in China and exclude internal research, development, or use that does not provide such services [31]. For an in-scope public service, provider duties and service-security records may inform the control analysis. For an internal AEC system outside that scope, the Measures do not create liability by analogy.

Table 5 therefore treats both regimes as possible evidentiary modulators rather than liability formulas. Civil responsibility still requires the applicable private-law elements and sector-specific duties. This narrower use preserves the manuscript's focus on developer, deployer, and professional-user responsibility.

Table 5. Regulatory boundary conditions and their possible evidentiary relevance to private AEC liability.

Regulatory Framework	Core Legislative Philosophy	Key Administrative Mechanism	Modulation Effect on Private Tort Liability
European Union AI Act	Risk-based regulatory framework with heightened obligations only where a system falls within the applicable classification and use case	Ex-ante conformity assessment, continuous risk management, data-governance duties, event logging, and human oversight requirements	Compliance or non-compliance may be relevant evidence under the applicable private-law regime, but the Act does not mechanically allocate civil liability
China Interim Measures for Generative AI Services	Governance of public-facing GenAI services, centered on provider duties, content safety, data governance, and service supervision	Provider obligations, security assessment where required, lawful content, training-data governance, and service-level controls	Direct relevance depends on whether the service falls within scope; the Measures do not create an AEC-specific tort allocation regime
AEC-Specific Adaptation	Prospective translation of general AI governance into construction review, BIM protocols, and digital-twin safety requirements	Integration of applicable AI governance, automated drawing review, human approval, BIM information controls, and independent technical audit	Standards and regulatory compliance provide contextual evidence and do not substitute for private-law causation, fault, and damages analysis

4.6. Practical Recommendations for AEC Stakeholders

Developers should document intended and prohibited uses, known limitations, model and retrieval versions, evaluation results, and incident-reporting channels. Safety-relevant applications should support authoritative-source grounding, output citations, and machine-readable constraints rather than rely on free-form generation alone.

Deployers should define the system boundary in the BIM execution plan or digital-twin governance plan. They should implement least-privilege access, validated data mappings, prompt and input controls, deterministic code and physics checks, tamper-evident logs, change approval, fallback states, and periodic independent audits.

AEC professionals and project owners should identify which decisions require independent calculation or specialist review, prohibit unverified AI output from entering approved construction information, record acceptance or rejection decisions, and treat unexplained model changes as controlled nonconformities. Contracts should allocate documentation, verification, incident-response, and evidence-preservation duties without attempting to exclude mandatory liability.

Regulators and standard-setting bodies should prioritize function-specific requirements for safety-relevant uses, clarify how existing construction and professional duties apply to GenAI-assisted work, and support interoperable audit records. These recommendations are implementation options derived from the qualitative framework, not empirically validated performance claims.

5. Conclusions

This study contributes a reproducible doctrinal method for analyzing GenAI-related data-integrity risks across the AEC lifecycle. The method combines structured source selection, legal-document coding, dual-track lifecycle mapping, and scenario-based stress testing. Its central analytical finding is that responsibility should be reconstructed from failed controls and evidence across developers, deployers, and users rather than assigned solely from actor labels or the presence of an AI system.

The qualitative proportional-liability assessment framework connects existing product, platform, professional-negligence, and risk-governance theories through the following five operationally defined factors: technical control, foreseeability, verification capacity, causal contribution, and evidence preservation. Its three-stage procedure requires actor-specific fact identification, reasoned ordinal classification, and doctrine-sensitive synthesis. It also identifies syntax, semantic, cybersecurity, and professional-verification controls that can be embedded in BIM and digital-twin workflows and aligned with construction quality-management practices.

The framework is deliberately bounded. It is not empirically calibrated, does not calculate legally binding percentages, and does not replace jurisdiction-specific tort, contract, product-liability, causation, or professional rules. The scenarios normatively illustrate the framework and do not establish external validity. The legal search identified no reported GenAI-related physical AEC decision within its stated scope and cutoff date, but that bounded result does not prove universal absence. Future research should test the factor definitions through disclosed case datasets, expert elicitation, Delphi or multi-criteria methods, and comparative judicial outcomes before any quantitative allocation model is proposed.

Author Contributions: Conceptualization, L.J.; methodology, L.J.; formal analysis, L.J.; investigation, L.J. and X.W.; data curation, X.W.; writing—original draft preparation, X.W.; writing—review and editing, L.J., X.W. and M.L.; visualization, X.W. and L.J.; validation, L.J.; supervision, M.L.; project administration, M.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: No new empirical data were created or analyzed. This study analyzes publicly available legal instruments, standards, case materials, and the academic literature identified through the source-selection protocol described in Section 2.1.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
AEC	Architecture, Engineering, and Construction
BIM	Building Information Modeling
CAD	Computer-Aided Design
EULA	End-User License Agreement
EU	European Union
EU AI Act	European Union Artificial Intelligence Act
FEA	Finite Element Analysis
GenAI	Generative Artificial Intelligence
LCA	Life Cycle Assessment
LLM	Large Language Model
NLP	Natural Language Processing

References

1. Memon, S.A.; Shehata, W.; Rowlinson, S.; Sunindijo, R.Y. Generative Artificial Intelligence in Architecture, Engineering, Construction, and Operations: A Systematic Review. *Buildings* **2025**, *15*, 2270. <https://doi.org/10.3390/buildings15132270>.
2. BuHamdan, S.; Alwisy, A.; Bouferguene, A. Generative systems in the architecture, engineering and construction industry: A systematic review and analysis. *Int. J. Archit. Comput.* **2021**, *19*, 226–249. <https://doi.org/10.1177/1478077120934126>.
3. Deng, M.; Menassa, C.C.; Kamat, V.R. From BIM to digital twins: A systematic review of the evolution of intelligent building representations in the AEC-FM industry. *J. Inf. Technol. Constr.* **2021**, *26*, 58–83. <https://doi.org/10.36680/j.itcon.2021.005>.
4. Valdebenito, R.; Forcael, E. Integrating Artificial Intelligence and BIM in Construction: Systematic Review and Quantitative Comparative Analysis. *Appl. Sci.* **2025**, *15*, 12470. <https://doi.org/10.3390/app152312470>.
5. Bryde, D.; Broquetas, M.; Volm, J.M. The project benefits of Building Information Modelling (BIM). *Int. J. Proj. Manag.* **2013**, *31*, 971–980. <https://doi.org/10.1016/j.ijproman.2012.12.001>.
6. Pan, Y.; Zhang, L. Roles of artificial intelligence in construction engineering and management: A critical review and future trends. *Autom. Constr.* **2021**, *122*, 103517. <https://doi.org/10.1016/j.autcon.2020.103517>.
7. Onatayo, D.; Onososen, A.; Oyediran, A.O.; Oyediran, H.; Arowoia, V.A.; Onatayo, E. Generative AI Applications in Architecture, Engineering, and Construction: Trends, Implications for Practice, Education and Imperatives for Upskilling-A Review. *Architecture* **2024**, *4*, 877–902. <https://doi.org/10.3390/architecture4040046>.
8. Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.J.; Madotto, A.; Fung, P. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* **2023**, *55*, 1–38. <https://doi.org/10.1145/3571730>.
9. Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; et al. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *ACM Trans. Inf. Syst.* **2025**, *43*, 1–55. <https://doi.org/10.1145/3703155>.
10. Greshake, K.; Abdelnabi, S.; Mishra, S.; Endres, C.; Holz, T.; Fritz, M. Not What You have Signed Up For: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection. In Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security, Copenhagen, Denmark, 30 November 2023; pp. 79–90. <https://doi.org/10.1145/3605764.3623985>.
11. Maynez, J.; Narayan, S.; Bohnet, B.; McDonald, R. On Faithfulness and Factuality in Abstractive Summarization. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 5–10 July 2020; pp. 1906–1919. <https://doi.org/10.18653/v1/2020.acl-main.173>.
12. Volk, R.; Stengel, J.; Schultmann, F. Building Information Modeling (BIM) for existing buildings-Literature review and future needs. *Autom. Constr.* **2014**, *38*, 109–127. <https://doi.org/10.1016/j.autcon.2013.10.023>.

13. Azhar, S. Building Information Modeling (BIM): Trends, Benefits, Risks, and Challenges for the AEC Industry. *Leadersh. Manag. Eng.* **2011**, *11*, 241–252. [https://doi.org/10.1061/\(ASCE\)LM.1943-5630.0000127](https://doi.org/10.1061/(ASCE)LM.1943-5630.0000127).
14. Sacks, R.; Koskela, L.; Dave, B.A.; Owen, R. Interaction of Lean and Building Information Modeling in Construction. *J. Constr. Eng. Manag.* **2010**, *136*, 968–980. [https://doi.org/10.1061/\(ASCE\)CO.1943-7862.0000203](https://doi.org/10.1061/(ASCE)CO.1943-7862.0000203).
15. Kingston, J.K.C. Artificial Intelligence and Legal Liability. In *Research and Development in Intelligent Systems XXXIII*; Springer: Berlin/Heidelberg, Germany, 2016; pp. 269–279. https://doi.org/10.1007/978-3-319-47175-4_20.
16. Gless, S.; Silverman, E.; Weigend, T. If Robots cause harm, Who is to blame? Self-driving Cars and Criminal Liability. *New Crim. Law. Rev.* **2016**, *19*, 412–436. <https://doi.org/10.1525/nclr.2016.19.3.412>.
17. Soyer, B.; Tettenborn, A. Artificial intelligence and civil liability-do we need a new regime? *Int. J. Law. Inf. Technol.* **2022**, *30*, 385–397. <https://doi.org/10.1093/ijlit/eaad001>.
18. Chamberlain, J. The Risk-Based Approach of the European Union’s Proposed Artificial Intelligence Regulation: Some Comments from a Tort Law Perspective. *Eur. J. Risk Regul.* **2022**, *14*, 1–13. <https://doi.org/10.1017/err.2022.38>.
19. Jobin, A.; Ienca, M.; Vayena, E. The global landscape of AI ethics guidelines. *Nat. Mach. Intell.* **2019**, *1*, 389–399. <https://doi.org/10.1038/s42256-019-0088-2>.
20. Buiten, M.C. Towards Intelligent Regulation of Artificial Intelligence. *Eur. J. Risk Regul.* **2019**, *10*, 41–59. <https://doi.org/10.1017/err.2019.8>.
21. Smuha, N.A. From a race to AI to a race to AI regulation: Regulatory competition for artificial intelligence. *Law. Innov. Technol.* **2021**, *13*, 57–84. <https://doi.org/10.1080/17579961.2021.1898300>.
22. European Parliament and Council of the European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No. 300/2008, (EU) No. 167/2013, (EU) No. 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). *Off. J. Eur. Union* **2024**, L 2024/1689, 12 July 2024. Available online: <https://data.europa.eu/eli/reg/2024/1689/oj> (accessed on 26 June 2026).
23. buildingSMART International. Industry Foundation Classes (IFC) 4.3.2.0 Documentation. Available online: <https://ifc43-docs.standards.buildingsmart.org/> (accessed on 14 June 2026).
24. buildingSMART International. Information Delivery Specification (IDS), Version 1.0. Available online: <https://technical.buildingsmart.org/projects/information-delivery-specification-ids/> (accessed on 14 June 2026).
25. Wang, X.; Wei, J.; Schuurmans, D.; Le, Q.V.; Chi, E.H.; Zhou, D. Self-Consistency Improves Chain of Thought Reasoning in Language Models. In Proceedings of the International Conference on Learning Representations, Kigali, Rwanda, 1–5 May 2023; arXiv:2203.11171.
26. Papernot, N.; McDaniel, P.; Jha, S.; Fredrikson, M.; Celik, Z.B.; Swami, A. The Limitations of Deep Learning in Adversarial Settings. In Proceedings of the 2016 IEEE European Symposium on Security and Privacy, Saarbrücken, Germany, 21–24 March 2016; pp. 372–387. <https://doi.org/10.1109/EuroSP.2016.36>.
27. Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Giannotti, F.; Pedreschi, D. A Survey of Methods for Explaining Black Box Models. *ACM Comput. Surv.* **2018**, *51*, 1–42. <https://doi.org/10.1145/3236009>.
28. Attia, A.R. The impact of integrating artificial intelligence and Building information modeling (BIM) systems on the development of construction methodologies. *J. Umm Al-Qura Univ. Eng. Archit.* **2025**, *16*, 1537–1554. <https://doi.org/10.1007/s43995-025-00193-2>.
29. Doshi-Velez, F.; Kim, B. Towards A Rigorous Science of Interpretable Machine Learning. *arXiv* **2017**, arXiv:1702.08608. <https://doi.org/10.48550/arXiv.1702.08608>.
30. European Parliament and Council of the European Union. Directive (EU) 2024/2853 of 23 October 2024 on liability for defective products and repealing Council Directive 85/374/EEC. *Off. J. Eur. Union* **2024**, L 2024/2853, 18 November 2024. Available online: <https://data.europa.eu/eli/dir/2024/2853/oj> (accessed on 26 June 2026).
31. Cyberspace Administration of China; National Development and Reform Commission of the People’s Republic of China; Ministry of Education of the People’s Republic of China; Ministry of Science and Technology of the People’s Republic of China; Ministry of Industry and Information Technology of the People’s Republic of China; Ministry of Public Security of the People’s Republic of China; National Radio and Television Administration. Interim Measures for the Management of Generative Artificial Intelligence Services. 2023. Available online: https://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm (accessed on 14 June 2026).

32. M.A. Mortenson Company, Inc. v. Timberline Software Corporation and Softworks Data Systems, Inc., 140 Wn.2d 568, 998 P.2d 305 (Supreme Court of Washington, Olympia, WA, USA, 4 May 2000). Available online: <https://law.justia.com/cases/washington/supreme-court/2000/67796-4-1.html> (accessed on 14 June 2026).
33. Lee, J.D.; See, K.A. Trust in Automation: Designing for Appropriate Reliance. *Hum. Factors* **2004**, *46*, 50–80. <https://doi.org/10.1518/HFES.46.1.50.30392>.
34. Parasuraman, R.; Riley, V. Humans and Automation: Use, Misuse, Disuse, Abuse. *Hum. Factors* **1997**, *39*, 230–253. <https://doi.org/10.1518/001872097778543886>.
35. American Law Institute. *Restatement (Third) of Torts: Liability for Physical and Emotional Harm, Sections 2–3*; American Law Institute Publishers: St. Paul, MN, USA, 2010.
36. Atkinson, D.; Morrison, J. A Legal Risk Taxonomy for Generative Artificial Intelligence. *arXiv* **2024**, arXiv:2404.09479. <https://doi.org/10.48550/arXiv.2404.09479>.
37. Hagen, A.; Andersen, T. Asset management, condition monitoring and Digital Twins: Damage detection and virtual inspection on a reinforced concrete bridge. *Struct. Infrastruct. Eng.* **2024**, *20*, 1242–1273.
38. Bahner, J.E.; Huper, A.-D.; Manzey, D. Misuse of automated decision aids: Complacency, automation bias and the impact of training experience. *Int. J. Hum.-Comput. Stud.* **2008**, *66*, 688–699. <https://doi.org/10.1016/j.ijhcs.2008.06.001>.
39. Batty, M. Digital twins. *Environ. Plan. B Urban Anal. City Sci.* **2018**, *45*, 817–820. <https://doi.org/10.1177/2399808318796416>.
40. Boje, C.; Guerriero, A.; Kubicki, S.; Rezgui, Y. Towards a semantic Construction Digital Twin: Directions for future research. *Autom. Constr.* **2020**, *114*, 103179. <https://doi.org/10.1016/j.autcon.2020.103179>.
41. Dembski, F.; Wossner, U.; Letzgus, M.; Ruddat, M.; Yamu, C. Urban Digital Twins for Smart Cities and Citizens: The Case Study of Herrenberg, Germany. *Sustainability* **2020**, *12*, 2307. <https://doi.org/10.3390/su12062307>.
42. American Law Institute. *Restatement (Second) of Torts, Sections 440–453*; American Law Institute Publishers: St. Paul, MN, USA, 1965.
43. Parasuraman, R.; Sheridan, T.B.; Wickens, C.D. A model for types and levels of human interaction with automation. *IEEE Trans. Syst. Man Cybern. Part A Syst. Hum.* **2000**, *30*, 286–297. <https://doi.org/10.1109/3468.844354>.
44. Xia, H.; Liu, Z.; Efremochkina, M.; Liu, X.; Lin, C. Study on city digital twin technologies for sustainable smart city design: A review and bibliometric analysis of geographic information system and building information modeling integration. *Sustain. Cities Soc.* **2022**, *84*, 104009. <https://doi.org/10.1016/j.scs.2022.104009>.
45. Shishehgarkhaneh, M.B.; Keivani, A.; Moehler, R.; Jelodari, N.; Laleh, S.R. Internet of Things (IoT), Building Information Modeling (BIM), and Digital Twin (DT) in Construction Industry: A Review, Bibliometric, and Network Analysis. *Buildings* **2022**, *12*, 1503. <https://doi.org/10.3390/buildings12101503>.
46. National Institute of Standards and Technology. *NIST AI 100-1*; Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST: Gaithersburg, MD, USA, 2023. <https://doi.org/10.6028/NIST.AI.100-1>.
47. National Institute of Standards and Technology. *NIST AI 600-1*; Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST: Gaithersburg, MD, USA, 2024. <https://doi.org/10.6028/NIST.AI.600-1>.
48. Rose, S.; Borchert, O.; Mitchell, S.; Connelly, S. *NIST SP 800-207*; Zero Trust Architecture. NIST: Gaithersburg, MD, USA, 2020. <https://doi.org/10.6028/NIST.SP.800-207>.
49. International Organization for Standardization. *ISO 19650-5:2020*; Organization and Digitization of Information about Buildings and Civil Engineering Works, Including Building Information Modelling (BIM)—Information Management Using Building Information Modelling—Part 5: Security-Minded Approach to Information Management. ISO: Geneva, Switzerland, 2020.
50. Standardization Administration of China. *GB/T 50430-2017*; Code for Quality Management of Engineering Construction Enterprise. Standardization Administration of China: Beijing, China, 2017.
51. Standardization Administration of China. *GB/T 51212-2016*; Unified Standard for Building Information Modeling Application. Standardization Administration of China: Beijing, China, 2016.
52. Standardization Administration of China. *GB/T 51301-2018*; Standard for Building Information Modeling Design Delivery. Standardization Administration of China: Beijing, China, 2018.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.