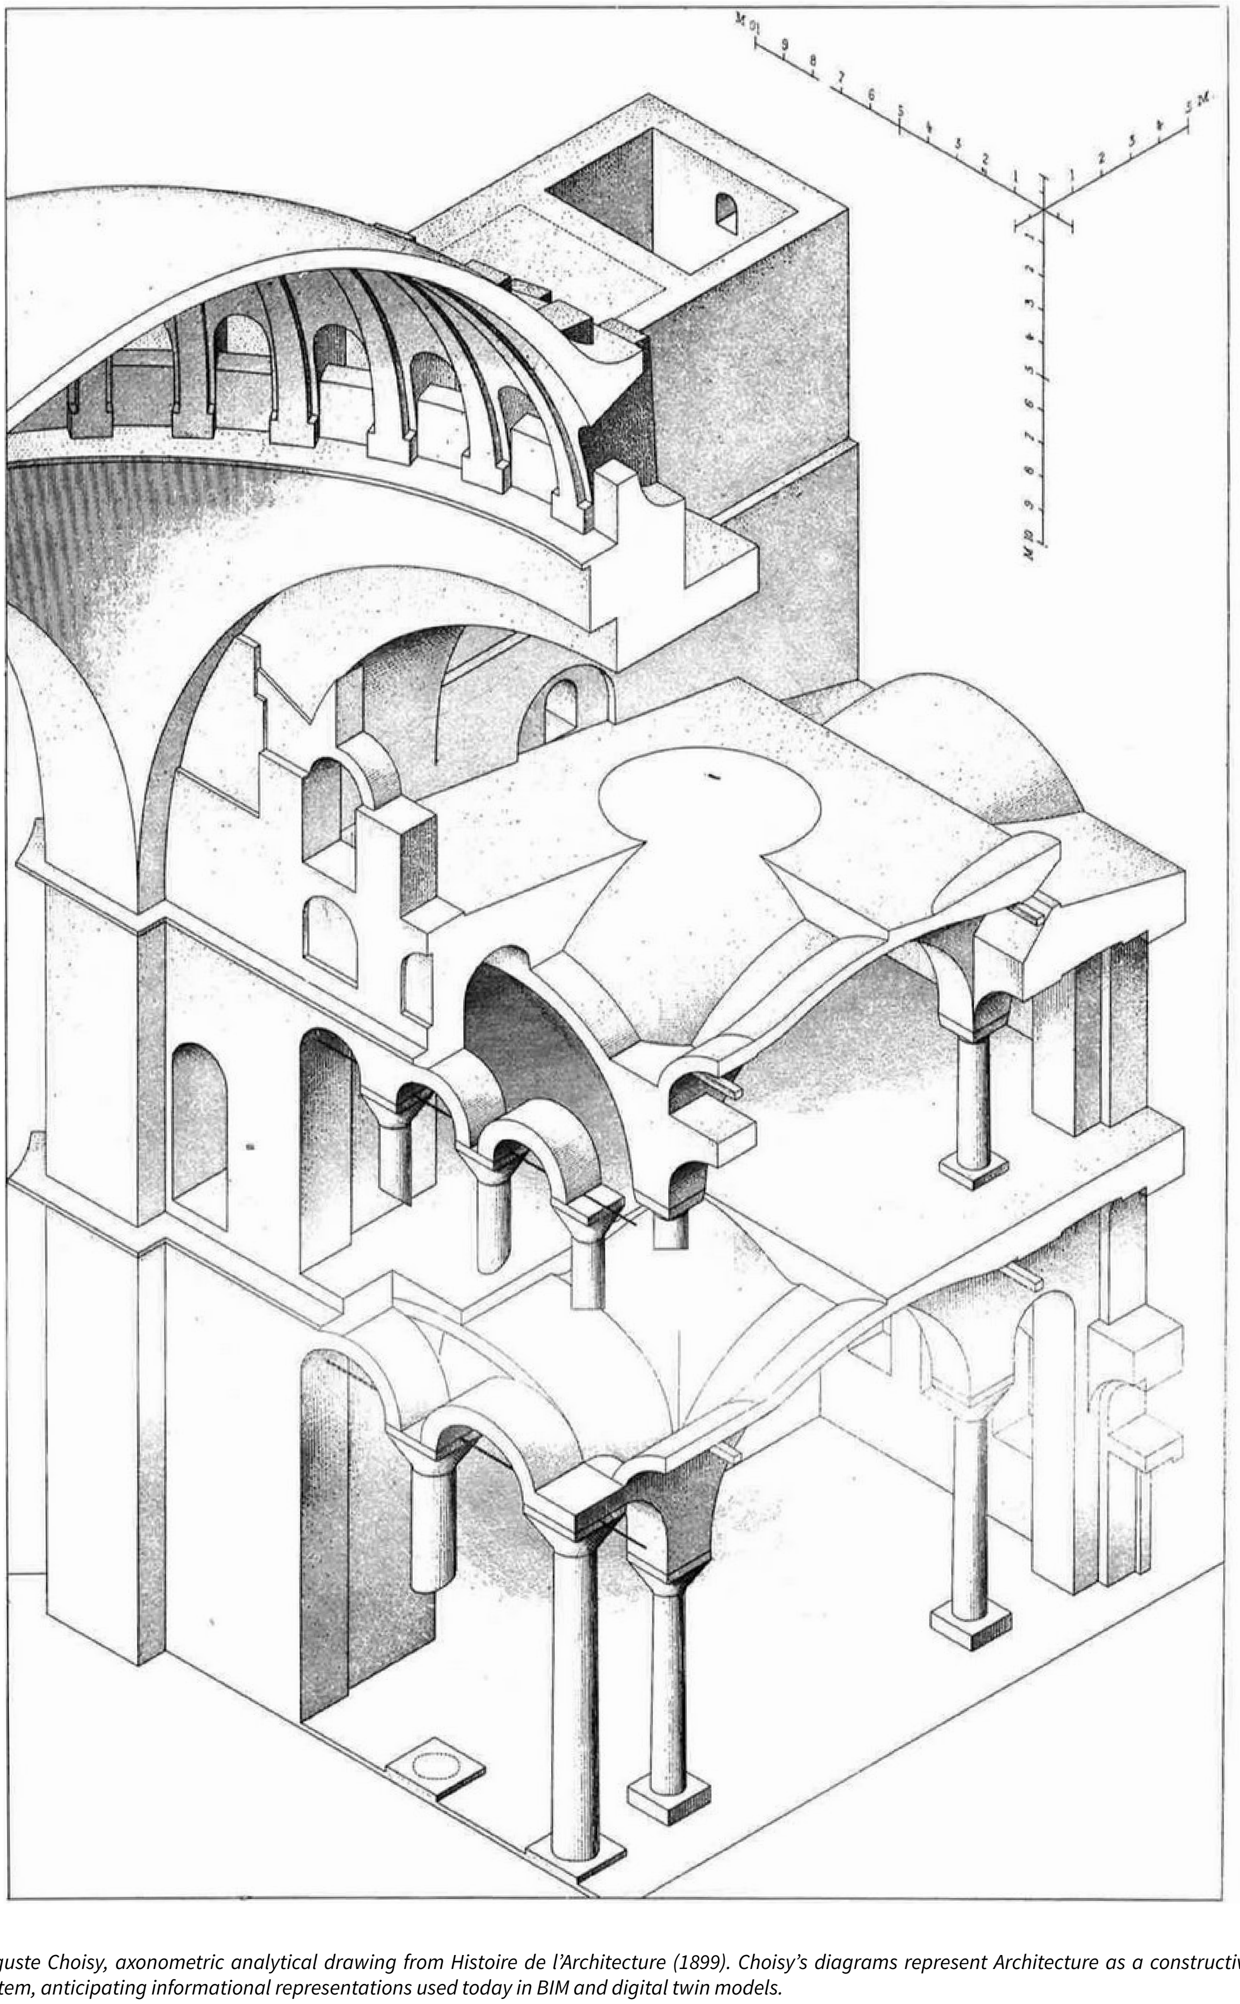

*Auguste Choisy, axonometric analytical drawing from Histoire de l'Architecture (1899). Choisy's diagrams represent Architecture as a constructive system, anticipating informational representations used today in BIM and digital twin models.*

# From BIM to Agentic AI: predictive management systems and phygital digital twins for built heritage

Daniele Fanzini *[1] , Giorgio Casoni,[1] Angelo de Cocinis[2]

[1] Department of Architecture, Built Environment and Construction Engineering, Politecnico di Milano, Italy.
[2] E-Making srl, San Giovanni in Persiceto, Italy.

**Email: daniele.fanzini@polimi.it*

**Abstract:** Information flow management is the backbone of digital transformation in the construction sector: the Data Sharing Environment (ACDat) enables integration and full interoperability, extending the potential of Building Information Modeling (BIM) and promoting sustainability in the management of the building life cycle. At the same time, the emergence of Agentic AI introduces a cognitive shift: autonomous, adaptive, and proactive agents take on design, cognitive, and institutional roles, helping redefine and even anticipate decision-making processes, organizational models, and market dynamics. In this context, the article presents the e-Building platform currently being tested, which combines ACDat, Agentic AI, and digital twins within a phygital logic. The platform automates document classification and building information handover through natural language and information constantly updated by the digital twin, with the result of evolving ACDat from a management infrastructure to a true cognitive ecosystem. The e-Building pilot (TRL 3–4) demonstrates that integrating Agentic AI with ACDat transforms predictive management from a technical tool to a paradigm of anticipatory governance.

**Keywords:** Agentic AI; ACDat; digital twin; anticipatory governance; phygital platforms.

**Received:** 21/10/2025; **Accepted:** 28/01/2026; **Available online:** 30/03/2026; **Published:** 30/06/2026.

**Cite as:** Fanzini, D., Casoni, G., De Cocinis, A. (2026). 'From BIM to Agentic AI: predictive management systems and phygital digital twins for built heritage'. *VITRUVIO - International Journal of Architectural Technology and Sustainability*, *11*(1), e24846. https://doi.org/10.4995/vitruvio-ijats.2026.24846

# 1. Introduction

The construction sector is undergoing a digital transformation that is redefining tools, processes, and skills. Building Information Modeling (BIM), which began as a primarily descriptive and prescriptive tool, has introduced collaborative approaches to design, but has shown limitations in managing complex and multidisciplinary contexts. This gap is being filled by the Data Sharing Environment (ACDat, internationally known as the Common Data Environment – CDE), which is evolving towards the role of cognitive infrastructure: a platform capable of integrating information flows, ensuring traceability, and feeding predictive systems throughout the entire life cycle of the project.

At the same time, the development of artificial intelligence (AI) is experiencing a radical discontinuity with the emergence of Agentic AI, characterized by autonomy, adaptability, and proactive capabilities. Agents do not merely process data but operate as cognitive and institutional actors, redefining decision-making processes, organizational models, and economic dynamics. This shift paves the way for an agentic economy in which knowledge, governance, and work are reorganized according to human-machine co-cognition logic.

This is the context for the experimentation of the e-Building platform, which integrates ACDat, Agentic AI, and digital twins in a phygital framework. The originality of the contribution lies in proposing a model in which ACDat evolves from a simple management archive to a cognitive ecosystem capable of supporting anticipatory and predictive governance practices for built heritage.

This contribution adopts a design science oriented approach, combining (i) a targeted literature review on BIM, common data environments and Agentic AI in the AEC domain, (ii) the design and prototyping of the e-Building platform as an information and decision-support artefact, and (iii) a formative evaluation based on a pilot deployment. The pilot is intended to verify functional feasibility and identify governance and skills requirements rather than to provide statistical generalisation. The evaluation protocol relies on structured walkthroughs of the main workflows (document upload and classification, level transitions across L0–L3, and PIM-to-AIM handover) and on an evidence log capturing inputs, feature paths, and outputs for replication. Where available, empirical indicators are reported qualitatively. Future work will extend validation with comparative benchmarks and quantitative performance measures.

The study addresses three questions: (RQ1) How can an ACDat evolve from a document repository to a cognitive infrastructure through Agentic AI and digital twins? (RQ2) Which functional roles and degrees of autonomy are required to automate classification, traceability, and handover workflows without compromising accountability? (RQ3) Which governance mechanisms and hybrid skills are necessary to operationalise predictive management systems for built heritage?

The article is structured into eight sections. Section 1 introduces the scope and methodological framing. Sections 2 and 3 outline the evolution from BIM to Agentic AI and map agent types and applications in the AEC sector. Section 4 details the e-Building architecture and the role of ACDat. Section 5 reports a formative evaluation protocol and the evidence log. Section 6 discusses multiscale governance and skills requirements. Section 7 develops evolutionary scenarios and governance prospects. Section 8 concludes and outlines limitations and future work.

# 2. From BIM to cognitive management of buildings

Building Information Modeling (BIM) has been instrumental in consolidating building data into unified models and improving coordination across AEC disciplines. However, fragmentation of information flows, uneven data quality, and document-centric practices still constrain lifecycle decision-making and the operationalisation of knowledge beyond project handover (Tran, 2024; Revolti et al., 2024; Miller et al., 2025; Radzi et al., 2025).

International and national standards have therefore reframed BIM within a broader information management paradigm. ISO 19650 and UNI 11337 foreground governance principles, including roles, responsibilities, information requirements, versioning, and interoperability, enabling more reliable information continuity across the asset lifecycle (ISO 19650; UNI 11337).

Accordingly, the transition is best described as a shift from document management to data-driven processes centred on a Common Data Environment and an Asset and Construction Data repository (ACDat). Here, the informational “heart” is not the model alone, but a governed data ecosystem supporting ingestion, validation, traceability, and auditable updates, while enabling shared metrics and KPIs that reduce rework and strengthen accountability mechanisms (Autodesk, 2024).

On this basis, Agentic AI marks a further discontinuity: cognitive management becomes feasible when AI systems generate content, use tools, navigate structured data, and execute bounded actions under governance constraints. This supports a human-machine co-cognitive regime where parts of operational reasoning are externalised into actionable workflows while preserving human responsibility for final decisions (Vepřek, 2024).

Recent literature distinguishes graded levels of "agency", from assistive utilities to tool-using and learning architectures, up to hierarchical and partially autonomous agents (Krishnan, 2025; Wissuchek & Zschech, 2025; Sapkota et al., 2025). Agency is therefore not binary, but configured by autonomy boundaries, access to tools and data, memory and planning capabilities, and the governance layer that constrains actions in safety-critical domains such as AEC.

Integrating these perspectives, we propose a taxonomy combining degrees of autonomy with application domains, clarifying how agentic systems can progressively reconfigure cognitive work in building lifecycle management. In this framing, cognitive management is not merely an efficiency upgrade, it is an infrastructural shift affecting coordination, role definitions, and institutional arrangements around accountability and information continuity (Figure 1).

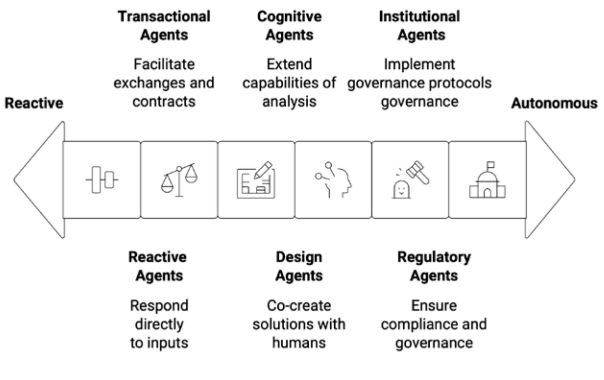


**Figure 1** | Spectrum of AI agent autonomy: from reactive to autonomous.

## 3. Types of agents and applications in the AEC sector

The agent typology above is reflected in AEC applications that translate agentic capabilities into operational value under explicit accountability constraints. Table 1 links agent type, core function, and expected impacts. In regulation, the UK Golden Thread exemplifies compliance-oriented configurations supporting traceability and lifecycle information continuity (Hackitt, 2018; UK Government, 2022). In site monitoring, systems such as Doxel (2025) combine sensing infrastructures with BIM-linked checks to detect deviations and support corrective actions. At the city scale, digital twins such as Virtual Singapore illustrate institutional implementations of integrated data environments for resilience and planning (National Research Foundation, 2018). Design-oriented agents and platforms, including ArkDesign AI and Spacemaker AI, show how generative and evaluative capabilities can support iterative option generation and assessment under constraints (Berger et al., 2022; Casoni, 2025) (Table 1).

Beyond tools, agentic systems affect how knowledge is codified and governed across projects and life cycles. In knowledge-intensive production contexts, expertise often remains tacit and distributed across hierarchical structures (Garicano, 2000). In the AEC sector, this tacit knowledge is amplified by project-based work, multi-disciplinary coordination, and craft-based practices (Addis, 2016; Carrillo et al., 2013). AI agents can partially externalise this knowledge into structured queries, checklists, and decision proposals, increasing procedural reliability and supporting more explicit operating rules.

Decision-making is similarly reshaped. Processes are slowed by fragmentation and incomplete information, while key constraints, especially aesthetic, social, and heritage-related, remain difficult to formalise. Tool-using agents can accelerate scenario exploration and consistency checks (performance, logistics, maintenance prioritisation), but inherit limitations from data quality and bias, and may induce overconfidence if treated as authoritative rather than advisory (Hasan, & Jha, 2024; Kang et al., 2024; Mousavi et al., 2024).

Market and organisational dynamics show analogous discontinuities. Agentic mediation can reduce coordination costs by standardising information exchange, identifying non-conformities earlier, and enabling more continuous model updates and cross-domain alignment (Agwa et al., 2025; Yan et al., 2025; Autodesk, 2024). These patterns align with the "agentic economy" framing, in which autonomous systems mediate interactions and reshape coordination mechanisms, including the informational role conventionally played by prices and contracts (Yang & Zhai, 2025; Radzi et al., 2025).

Labour impacts should be read as a rebalancing of substitution and augmentation, with hybrid roles emerging around digital twin operations, data quality governance, and validation of algorithmic outputs. Distributional effects reward complementary capabilities

**Table 1** | Types of agents and case studies in the agentic economy.

| Application | Type of agent | Function | Impacts |
|---|---|---|---|
| Golden Thread | Regulatory | Ensure information traceability and compliance | Increase safety, improve building governance |
| Doxel AI | Cognitive / Operational | Analyze data from drones and BIM, support decisions | Anticipate anomalies on construction sites, reduce construction times and costs |
| Virtual Singapore | Institutional | Simulate urban scenarios using digital twins | Enable sustainable and resilient planning |
| Monumental Robotics | Transactional | Manage processes with outcome-as-a-service logic | Introducing new economic models based on results |
| ArkDesign AI, Spacemaker AI | Design | Generate and optimize design solutions | Redefining the role of the designer, enabling co-creation |

and higher-level coordination work (Tran, 2024; Mäkelä & Stephany, 2024; Ganuthula & Balaraman, 2025). The central challenge remains twofold: to contain opacity, bias, and automation complacency, and to design accountability mechanisms and anticipatory governance practices that orient agentic infrastructures toward sustainability, equity, and social responsibility (Cabitza et al., 2017).

## 4. ACDat as the informational heart of the process

ACDat is now the key element of interoperability between the digital systems that accompany the life cycle of a structure. It no longer limits itself to collecting data, but guarantees its quality, consistency, and traceability, transforming it into useful knowledge for predictive systems and digital twins. Without a reliable and structured infrastructure, algorithms do not generate operational value, because classification errors and missing identifiers break traceability and auditability. This section describes the baseline ACDat configuration and introduces the agentic extensions implemented in e-Building. The evaluation protocol and evidence log are then presented in Section 5.

ACDat is a centralized and secure digital environment used to manage information models, documents, and project or asset data. It constitutes the infrastructure that ensures the consistency and quality of information flows throughout the entire life cycle of the project. It is divided into maturity levels: from Level 0, based on unstructured management, to advanced levels L1–L3, characterized by interdisciplinary collaboration and integration with digital twins.

ACDat organises content through four functional areas: WIP (Work in Progress) for content in progress, Shared for validated and shared data, Published for approved documents, and Archive for historical storage. This scheme ensures orderly flow, versioning, and traceability of changes.

ACDat is therefore not a static repository, but a cognitive platform that feeds artificial intelligence and digital twin systems. Through access rules and interoperability protocols, it enables reliable predictive models and new forms of anticipatory governance.

The platform embeds a semantic layer that transforms uploaded documents into machine-interpretable data. Files undergo text extraction, segmentation and vectorisation via embedding models, with vector representations co-located alongside governed metadata in a PostgreSQL repository. This architecture supports hybrid querying: structured filters (asset, phase, discipline, information state) can be combined with vector-based semantic retrieval, and both remain subject to the same access-control logic that governs the underlying ACDat repository, preventing unauthorised discovery through similarity search.

On this basis, the e-Building platform is being tested, introducing predictive management methods that are absent in traditional ACDat. Where ACDat requires manual activities, e-Building integrates AI agents capable of automating uploading, classification, and validation, making the transition between the information states required by UNI 11337 and ISO 19650 more fluid.

The innovations concern the automatic classification of documents: agents analyze content, extract metadata, and place files along the L0–L3 cycle; the management of transitions between levels, with support for the transition from L0 to L2, reducing errors and time; and handover management, where agents verify consistency and completeness, generate summary dashboards, and automatically feed the AIM, ensuring traceability.

Additional innovative elements include natural language interaction, which allows the platform to be queried to receive indicators and predictive scenarios, and the phygital approach, in which the digital twin is constantly updated and connected to real conditions.

The e-Building platform builds on ACDat by adding agentic services that operationalise information requirements and coordination tasks. Rather than redefining the repository itself, the platform augments ACDat with tool-using workflows for classification, semantic validation, traceable handovers across information states, and decision-oriented querying. In this way, AI acts as an accountable operational layer that reduces coordination frictions and increases the reliability of lifecycle updates (Table 2).

The table above highlights the transition from an ACDat based on manual activities to a system in which AI agents automate and enrich information flows. While ACDat only guarantees sharing, e-Building interprets, connects, and verifies data, directly feeding digital twins. Innovations affect governance and skills, requiring more robust protocols and new hybrid figures capable of guiding predictive scenarios. The result is a technical and institutional device that supports the governance of the built environment in an anticipatory and resilient manner.

Architecturally, e-Building is conceived as an extension of the ACDat (CDE) that adds an agentic layer for document understanding, semantic validation, and controlled handover between information states. The backbone remains a federated repository aligned with ISO 19650 information management principles and UNI 11337 information levels, while the platform introduces a service-oriented architecture that separates storage, orchestration, and verification concerns.

At system level, the platform comprises four macro-components. The user-facing layer provides a web interface for upload, query, and dashboarding, alongside APIs for integration with authoring tools. The data and knowledge layer combines the ACDat repository with structured metadata stores and a graph-oriented knowledge base linking assets, documents, requirements, and responsibilities. A workflow and governance layer implements versioning, access control, issue tracking, and audit logs, while the AI and agent layer orchestrates specialised agents and tools for classification, extraction, checking, and summarisation, with explicit human-in-the-loop gates before any authoritative update is committed to the information model (Figure 2).

From an information management perspective, the platform operationalises a progressive chain of information states. The Project Information Model (PIM) is populated during design and construction. It contains evolving models, temporary documents, and work-in-progress artefacts. At handover, the Asset Information Model (AIM) is created or updated as the authoritative record for operations. It consolidates validated deliverables, as-built documentation, and asset registers, and it becomes the substrate for maintenance, monitoring, and predictive services over the asset life cycle (Figure 3).

The platform distributes these functions across specialised agent roles. Ingestion agents manage upload and versioning. Semantic-indexing agents chunk

**Table 2 |** From traditional ACDat to the e-Building platform: innovations, governance, and skills.

| Dimension | Traditional ACDat | e-Building platform (with Agentic AI) | Impacts on governance and skills |
|---|---|---|---|
| Data management | Manual uploading, naming rules, possible errors | Automatic classification, semantic extraction of metadata | Need for advanced data governance, security and privacy protocols |
| Information flows | Transitions between levels (L0–L3) managed manually | AI agents govern transitions, versioning, and document consistency | Greater transparency and accountability in processes |
| Handover | Manual checks for completeness and compliance | Automated verification, summary dashboards, AIM feeding | Hybrid professional roles for AI governance and data quality |
| Interaction | Menu navigation, limited text search | Natural language interaction, predictive indicators | AI literacy skills and anticipatory capabilities |
| Digital Twin | Not integrated, sporadic updates | Updated in real time, digital-physical phygital logic | New professional skills for DT management and simulations |
| Role of the platform | Management infrastructure/centralized archive | Cognitive ecosystem, capable of interpreting, suggesting, and anticipating | Anticipatory leadership, systemic vision, and organizational resilience |

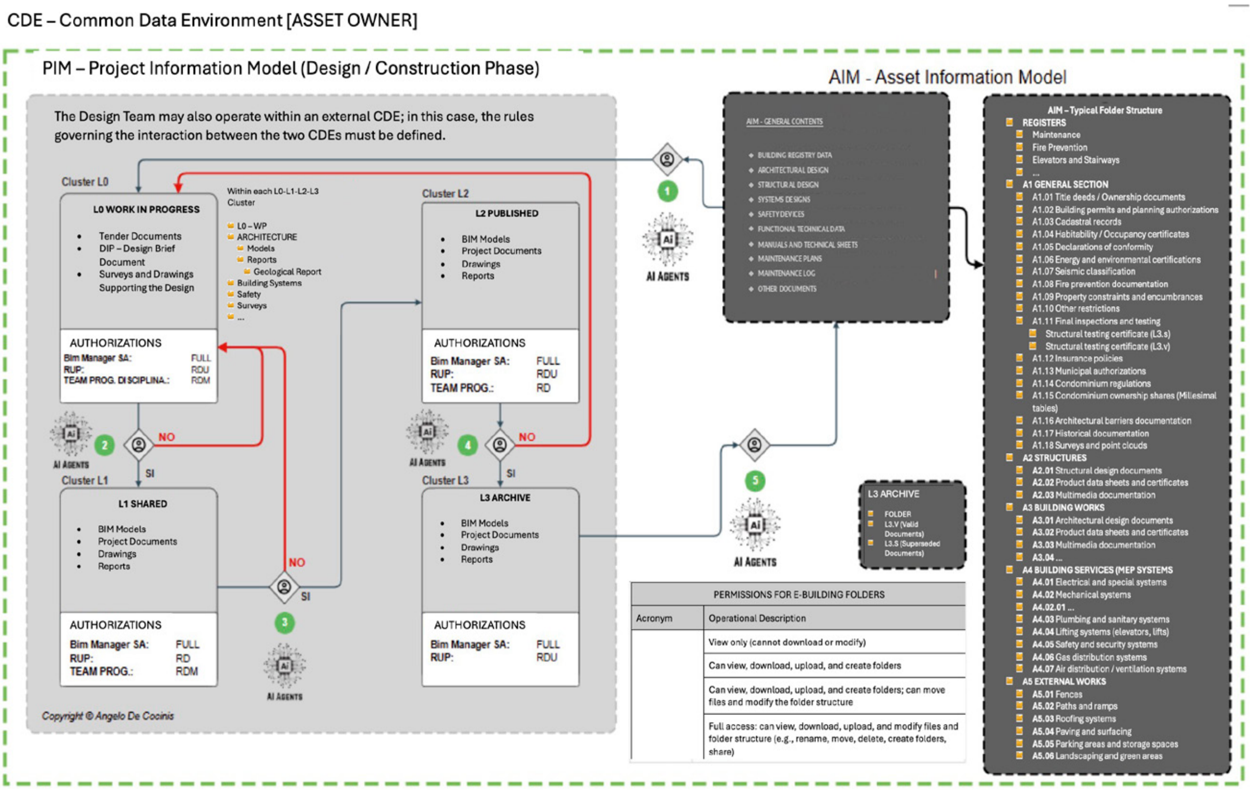


Figure 2 | ACDat architecture and information flows in a conventional CDE configuration.

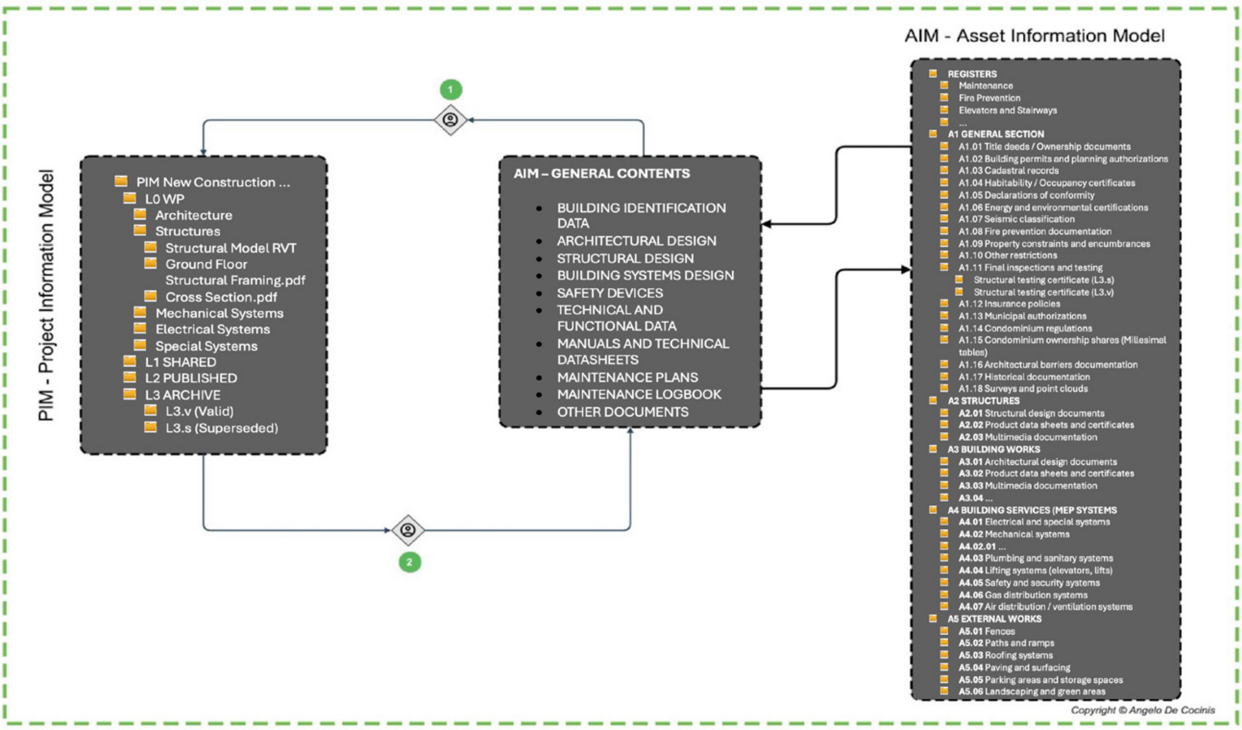


Figure 3 | e-Building platform architecture. Agentic extensions for classification, validation, and handover.

text and generate embeddings that are stored alongside metadata in the PostgreSQL-based relational database repository. Metadata agents propose document types, disciplines and asset links, while classification agents suggest placement within the L0–L3 maturity cycle. Governance agents check whether mandatory fields are complete and whether transitions between information states satisfy declared requirements. Handover agents assemble and validate information packages for transfer from PIM to AIM. Query agents combine semantic search with structured filters to support natural-language access, explaining which documents are retrieved and why they satisfy the query.

The agentic layer supports this chain by implementing three functions that are often weakly formalised in conventional CDE practice. First, automated classification and metadata enrichment. Incoming files (for example PDFs, DWGs, images, reports) are parsed and mapped to document classes, disciplines, and information containers, producing consistent metadata and linking them to the relevant asset instances. Second, semantic validation against explicit requirements. The platform can check completeness and conformity through machine-readable rules, such as information delivery specifications and model view constraints, and it can identify.

Non-conformities as issues that are traceable to roles and deadlines. Third, controlled handover: before promoting artefacts from PIM to AIM, the system enforces approval workflows and produces a traceable change set that can be inspected and signed off.

Operationally, the agentic pipeline unfolds as a tool-using orchestration workflow. Files are first uploaded or synchronised, and a fingerprint is created for versioning (ingestion). The system then applies parsers and language models to extract key entities, parameters, and references, normalising them to the project taxonomy (extraction). Extracted content is checked against rule sets and consistency constraints, producing a structured compliance report (verification). Based on this analysis, agents propose metadata assignments, links, and issue tickets, but do not automatically commit authoritative changes (action proposal). Appointed reviewers approve, reject, or request clarifications through a dedicated interface (human gate). Only after explicit approval are changes committed to the AIM and stored with full provenance information (commit and log).

To clarify the practical value beyond architectural statements, consider a representative use case in built heritage maintenance. A maintenance report is uploaded in PDF, together with photos and a short textual note from

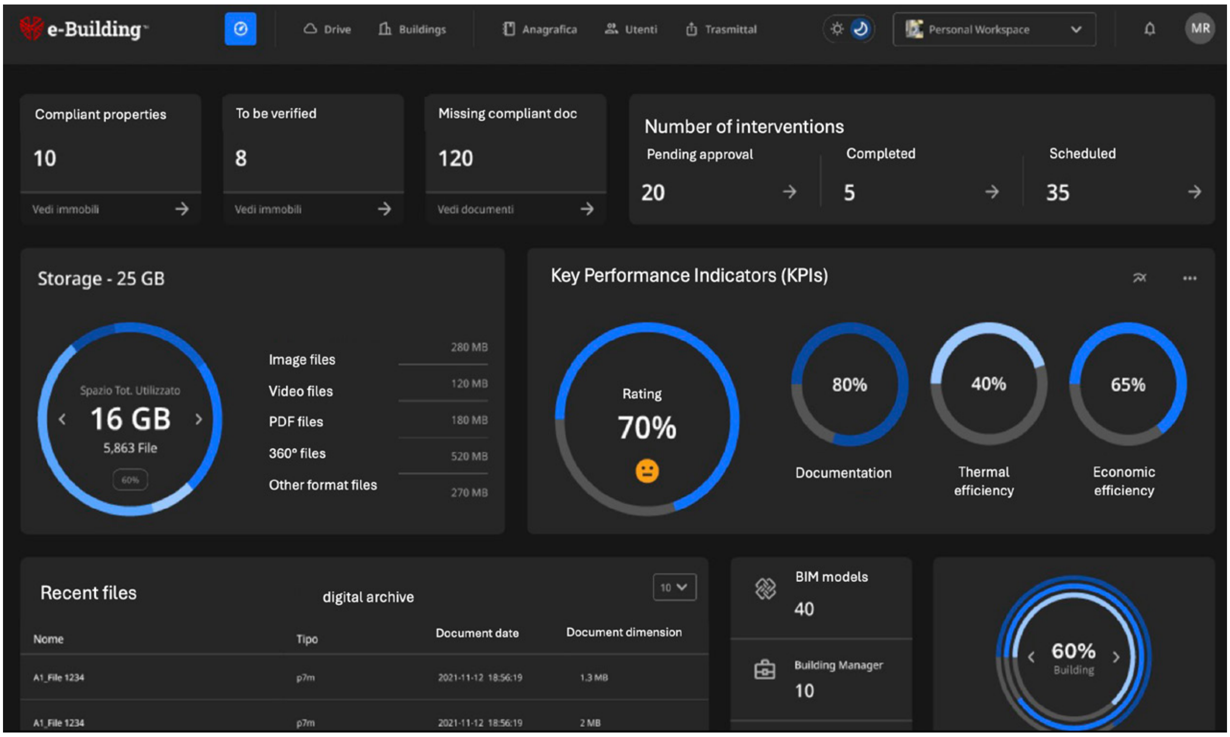


**Figure 4** | Phygital e-Building ecosystem.

the site. The ingestion agent classifies the document as an inspection deliverable, links it to the corresponding heritage element in the asset register, and extracts key signals. The verification agent checks whether the report contains mandatory fields and whether the described anomaly matches an existing issue history. The system then proposes an update to the AIM, creates a ticket for the responsible role, and updates a dashboard indicator. The human gate validates the proposal and, once approved, the platform records the update with full traceability. This is the minimal mechanism through which the phygital digital twin remains aligned with reality while preserving accountability.

The platform also supports natural language interaction as an additional access mode rather than as a replacement for formal models. Users can query the ACDat and the AIM through controlled prompts that return grounded answers based on the current information baseline and the evidence log. This design choice aims to increase accessibility for non-specialist stakeholders, while maintaining governance through traceability, role-based permissions, and explicit approval steps.

In operational terms, the platform's distinctive features are controlled data synchronisation, explicit role-based access rules, and versioned audit trails for agentic actions. These elements support continuous digital-twin updates while preserving accountability through traceable handovers and human validation gates (Figure 4).

## 5. Preliminary evaluation and evidence log

To provide empirical grounding, the work includes a formative, small-scale evaluation protocol aimed at assessing whether e-Building's agentic services improve information quality and reduce coordination frictions in ACDat workflows. The goal is not to claim statistical generalization at this stage, but to offer a replicable procedure and auditable traces of the platform's behaviour under realistic document and model conditions.

Evaluation setting and units of analysis. The unit of analysis is the ACDat transaction, defined as a traceable sequence from document ingestion (e.g., inspection report, maintenance log, survey outputs, IFC model revisions) to an actionable proposal (classification, metadata completion, linking to an asset, or handover between information states). Transactions are sampled across disciplines (architecture, structures, MEP) and across information states.

Protocol. Each transaction is evaluated along four dimensions reflecting ISO 19650-oriented information management requirements and Agentic AI accountability constraints: (D1) classification and metadata enrichment accuracy, (D2) semantic consistency and completeness against declared information requirements, (D3) traceability and auditability of the agent's actions, and (D4) reduction of human effort in validation and coordination.

Metrics and operationalization. For D1, accuracy is computed against a human-coded label set (e.g., document type, discipline, asset taxonomy, information container status). For D2, completeness is measured as the proportion of mandatory fields satisfied and the number of detected inconsistencies (e.g., missing identifiers, mismatched dates, inconsistent asset references). For D3, auditability is assessed through the presence of explainable rationales, reproducible feature paths, and versioned evidence links in the ACDat log. For D4, effort is proxied by time-on-task for validation, number of manual corrections, and coordination loops required before publication.

Evidence log. Each evaluated transaction is recorded in an evidence log containing access date, input artefacts and identifiers, agentic function invoked, outputs produced, human corrections (if any), and final publication decision. The log provides a minimal replication checklist and can be shared upon request, subject to data protection constraints typical of built heritage projects.

Early pilot runs demonstrate three key benefits: rapid triage and normalization of heterogeneous documentation; detection of missing or inconsistent information before publication; and reduction of latent coordination work through structured updates that identify requirement uncertainty. These observations motivate the future work, including broader pilots and controlled comparisons against baseline ACDat-only workflows.

## 6. New needs for multiscale governance and skills

This section extends the analysis from the building scale to the urban and territorial scale, where ACDat, digital twins, and agentic services intersect with multi-actor coordination processes. At this scale, information continuity and accountability depend not only on technical integration but also on governance arrangements that define access rights, validation procedures, and institutional responsibilities. The section examines these conditions and outlines implications for resilience, sustainability, and social responsibility in built environment management.

In light of the challenges outlined, a number of strategic needs emerge to make the adoption of agentic platforms in the construction sector sustainable. Advanced cybersecurity and traceability protocols must ensure data protection, algorithm transparency, and output reliability. At the same time, interoperability standards should enable integration between different systems and promote the widespread use of digital twins.

In terms of professional skills, hybrid figures must be trained, capable of combining technical skills (information modeling, data management, ISO 19650 and UNI 11337 standards), regulatory skills (privacy, security, compliance), and management skills. This is accompanied by the need to promote widespread AI literacy and anticipatory literacy, to enable decision-makers to understand, critically evaluate, and use algorithmic predictions in an informed manner.

Establishing clear mechanisms of accountability and decision-making responsibility is equally crucial, avoiding automation complacency and ensuring that the adoption of Agentic AI remains anchored to criteria of fairness, sustainability, and social responsibility.

In this sense, the governance of the built environment can no longer be limited to an administrative role, but must take on a proactive and anticipatory function, orchestrating human-AI cognitive ecologies oriented towards resilience and responsible innovation.

## 7. Evolutionary scenarios and governance prospects

Recent predictions on the development of artificial intelligence (Kokotajlo et al., 2025) suggest that the impact of agentic systems and, in perspective, forms of artificial general intelligence (AGI), could exceed the scope of the Industrial Revolution itself. These radical prospects are linked to the governance and skills requirements outlined in the previous section, but shift the focus to future scenarios in which the scope of Agentic AI extends beyond the individual building. When transposed to the construction sector, these projections outline scenarios in which cognitive agents not only automate design and maintenance processes, but also take on proactive roles in defining urban regeneration strategies, risk management, and the governance of built heritage.

This perspective, however, introduces new challenges in terms of security and governance. Just as international reports warn of the risks of loss of control or malicious use of advanced systems, in the AEC context, the integration of Digital Twin, Internet of Things, and Agentic AI raises crucial issues of cybersecurity, source reliability, and organizational resilience. These no longer concern only the level of individual projects, but the security of urban infrastructure and the sovereignty of territorial data. Information continuity cannot therefore be reduced to a technical problem of interoperability, but must be understood as an element of cognitive sovereignty of the built environment, which requires clear rules of access, control, and responsibility.

Furthermore, the growing predictive capacity of AI agents requires extending the notion of anticipatory governance from the building level to the urban and territorial level, including medium- to long-term scenarios (2030+), for example in predictive climate risk management, urban energy flow optimization, or resilient infrastructure planning. In this context, the AEC sector can be a privileged laboratory for experimenting with human-machine co-governance practices aimed at sustainability, reducing management costs, and increasing social and economic resilience.

As Table 3 shows, each dimension presents a delicate balance between transformative potential and emerging vulnerabilities, which necessitates integrated anticipatory governance.

**Table 3 |** Opportunities and risks of Agentic AI and digital twins in the AEC sector.

| Dimension | Opportunities | Risks |
|---|---|---|
| Cognitive automation | AI agents capable of managing maintenance, urban simulations, and energy optimization. | Loss of decision-making control, excessive dependence on agents. |
| Governance | Continuity of information and predictions throughout the entire life cycle, extended to the urban/territorial level. | Algorithmic opacity, difficulties with accountability and legal responsibility |
| Cybersecurity | Secure digital twins as cognitive infrastructures of built heritage | Vulnerability to cyber attacks and data manipulation. |
| Market and work | New professions related to digital twins and predictive management | Employment polarization, reduction of traditional roles |
| Geopolitics of data | Information sovereignty as a lever for economic resilience | International competition for control of cognitive infrastructure |

These scenarios confirm that Agentic AI is not only a technological evolution, but the beginning of a new institutional and cognitive regime, whose implications for the AEC sector will be discussed in the Conclusions.

## 8. Conclusions

The analysis has shown how built heritage management is undergoing a transformation that goes beyond the adoption of digital tools. The evolution from BIM to ACDat, assisted by Agentic AI, marks a shift from a documentary logic to a cognitive infrastructure capable of integrating data, actors, and processes into a continuous information flow. In this context, Agentic AI represents a cognitive and economic turning point that redefines functions and roles. Agents not only store or process information, but also interpret meanings, anticipate scenarios, and contribute to new technical and organisational institutions.

Experimentation with the e-Building platform has translated these dynamics into operational terms, showing how an ACDat enriched by AI agents can function as a phygital ecosystem that continuously updates the digital twin, automates complex steps, and returns predictive indicators in natural language. Compared with traditional ACDat, the platform indicates a qualitative leap from management infrastructure to cognitive platform, with AI acting as an operational and institutional partner.

These developments open opportunities as well as challenges. Technically, interoperability standards and security protocols must be consolidated. Organisationally, hybrid professionals combining digital, regulatory, and anticipatory skills must be developed. Institutionally, governance should nurture human-AI cognitive ecologies oriented toward resilience, sustainability, and social responsibility.

The agent types discussed (cognitive, regulatory, transactional, design, and institutional) provide a framework to interpret how Agentic AI can affect multiple layers, from knowledge production and interpretation to institutional regulation, and offer a conceptual basis for guiding future experimentation with integrated platforms such as e-Building.

Overall, the article highlights the transition from BIM to ACDat as a cognitive infrastructure, the discontinuity introduced by Agentic AI and its epistemic and economic implications, and the e-Building case as an innovative model, together with emerging governance and skill requirements. Taken together, these elements suggest that predictive management systems based on ACDat, Agentic AI, and digital twins are not only a technical evolution, but also contribute to a new regime of anticipatory governance aimed at transforming the built environment toward more intelligent, transparent, and inclusive outcomes.

From a technology readiness perspective, the e-Building pilot corresponds to TRL 3–4: a feasibility study for semantic-governed document lifecycle management within ACDat environments. The work demonstrates functional integration of embedding-based retrieval, agentic classification workflows, and traceable handover mechanisms in a controlled setting. Current pilot runs operate on a limited dataset to validate architectural choices and to log agent behaviour under realistic conditions. User-facing interfaces are not yet deployed, and proprietary configuration details are withheld at this stage. Advancing to TRL 5–6 will require quantitative validation on larger datasets, comparative benchmarking against conventional ACDat workflows, deployment with operational users, and broader organisational adoption across multiple projects.

Open issues remain, including implementation costs, technological maturity, and organisational acceptance, which require further empirical investigation. Future studies should examine long-term impacts of Agentic AI and data sharing through comparative analyses and shared metrics.

## References

Addis, M. (2016). ‘Tacit and explicit knowledge in construction management’. *Construction Management and Economics, 34*(7-8), 439-445. https://doi.org/10.1080/01446193.2016.1180416

Agwa, T.C., Oti, A.H., & Kurul, E. (2025). ‘From barriers to breakthroughs: A deep dive into BIM adoption’. *Buildings 15*(7), 1116. https://doi.org/10.3390/buildings15071116

ArkDesign AI. (2025). *AI-powered generative design platform* [WWW document]. URL https://www.arkdesign.ai (accessed 13 September 2025).

Autodesk. (2024). *Connected Data Environment for Digital Twins*. Autodesk Whitepaper [WWW document]. URL https://damassets.autodesk.net/content/dam/autodesk/www/pdfs/cde-whitepaper-en.pdf (accessed 13 September 2025).

Berger, H., Svensson, T., & Røstvik, H.N. (2022). Spacemaker.Ai: Using AI in developing urban block variations. In: M. Jabi (ed.), *Artificial Intelligence in Architecture*. Elsevier, Amsterdam. https://doi.org/10.1016/B978-0-12-823941-4.00010-X

Cabitza, F., Rasoini, R., & Gensini, G.F. (2017). 'Unintended consequences of machine learning in medicine'. *JAMA, 318*(6), 517–518. https://doi.org/10.1001/jama.2017.7797

Carrillo, P., Ruikar, K., & Fuller, P. (2013). 'When will we learn? Improving lessons learned practice in construction'. *International Journal of Project Management, 31*(4), 567-578. https://doi.org/10.1016/j.ijproman.2012.10.005

Casoni, G. (2025). Sperimentazione di sistemi di anticipazione autonomi. In: F. Celaschi, G. Formia, G. Casoni and V. De Matteo (eds.), *Design e anticipazione: Change studies, processi, metodi e tecnologie per l'innovazione*. Bologna University Press, Bologna. https://doi.org/10.30682/9791254776544

Doxel Inc. (2025). *AI-powered construction management* [WWW document]. URL https://www.doxel.ai (accessed 13 September 2025).

Ganuthula, V.R.R., & Balaraman, K. (2025). 'Skill-based labor market polarization in the age of AI: A comparative analysis of India and the United States'. *arXiv* [Preprint]. https://doi.org/10.48550/arXiv.2501.15809

Garicano, L. (2000). 'Hierarchies and the organization of knowledge in production'. *Journal of Political Economy, 108*(5), 874–904.

Hackitt, J. (2018). *Building a safer future: Independent review of building regulations and fire safety*. UK Government, London.

Hasan, A., & Jha, K. N. (2024). A holistic evaluation of BIM implementation barriers in the Indian construction industry: Pre- and post-adoption perspectives. *International Journal of Construction Education and Research, 20*(3). https://doi.org/10.1080/15578771.2024.2320108

Kang, J., Liu, Y., & Wang, Z. (2024). 'A comprehensive framework for digital twins of buildings: From BIM to real-time analytics'. *Journal of Building Engineering, 78*, 107654. https://doi.org/10.1016/j.jobe.2024.107654

Kokotajlo, D., Alexander, S., Larsen, T., Lifland, E., & Dean, R. (2025). *AI 2027*. AI Futures Project [WWW document]. URL https://ai-2027.com (accessed 13 September 2025).

Krishnan, N. (2025). 'AI agents: Evolution, architecture, and real-world applications'. *arXiv*.[Preprint]. https://doi.org/10.48550/arXiv.2503.12687

Mäkelä, E., & Stephany, F. (2024). 'Complement or substitute? How AI increases the demand for human skills'. *arXiv* [Preprint]. https://doi.org/10.48550/arXiv.2412.19754

Miller, J., Nguyen, T., & Sørensen, A. (2025). 'BIM interoperability: A systematic review and future research agenda'. *Applied Sciences, 15*(6), 3121. https://doi.org/10.3390/app15063121

Monumental Robotics. (2025). *Automating on-site construction with robotics and software.* [WWW document]. URL https://www.monumental.co (accessed 13 September 2025).

Mousavi, M., Hosseini, M.R., & Martek, I. (2024). 'Digital twins in the AEC industry: Challenges and opportunities'. *Digital, 4*(1), 23–41. https://doi.org/10.3390/digital4010004

National Research Foundation. (2018). *Virtual Singapore: A dynamic 3D city model and collaborative data platform*. Government of Singapore, Singapore.

Radzi, A. R., Azmi, N. F., Kamaruzzaman, S. N., Algahtany, M., & Rahman, R. A. (2025). 'Challenges in construction readiness for BIM-based building projects', *Journal of Asian Architecture and Building Engineering, 24*(3), pp. 1689–1704. https://doi.org/10.1080/13467581.2024.2343803

Revolti, A., Gualtieri, L., Pauwels, P., & Dallasega, P. (2024). 'From building information modeling to construction digital twin: A conceptual framework'. *Production & Manufacturing Research, 12*(1). https://doi.org/10.1080/21693277.2024.2387679

Sapkota, R., Roumeliotis, K.I., & Karkee, M. (2025). 'AI agents vs. Agentic AI: A conceptual taxonomy, applications and challenges'. *arXiv* [Preprint]. https://doi.org/10.48550/arXiv.2505.10468

Tran, H.V. (2024). 'A review of challenges and opportunities in BIM adoption for construction project management'. *Engineering Journal, 28*(2), 121–139. https://doi.org/10.4186/ej.2024.28.2.121

UK Government. (2022). *Building Safety Act 2022*. The Stationery Office, London.

Vepřek, L.H. (2024). *At the edge of AI*. transcript Verlag, Bielefeld. https://doi.org/10.14361/9783839472286

Wissuchek, C., & Zschech, P. (2025). 'Exploring Agentic AI systems: Towards a typological framework'. *arXiv* [Preprint]. https://doi.org/10.48550/arXiv.2508.00844

Yan, L., Chen, X., & Zhao, J. (2025). 'Common data environment for digital twins in the built environment: Enablers and challenges'. *Automation in Construction,* 162, 105612. https://doi.org/10.1016/j.autcon.2025.105612

Yang, K., & Zhai, C. (2025). 'Ten principles of AI agent economics'. *arXiv* [Preprint]. https://doi.org/10.48550/arXiv.2505.20273

Yitmen, I. (Ed.). (2023). *Cognitive digital twins for smart lifecycle management of built environment and infrastructure: Challenges, opportunities and practices*. Routledge. https://doi.org/10.1201/9781003230199