

Bridging BIM and operational sensor data through knowledge graphs: an interoperable digital twin for building energy maintenance

Anabelle Rahhal^{1*}, Pierre Leclercq¹,

¹ILUCID, University of Liege, Belgium

* Presenting author: a.rahhal@uliege.be

Abstract—The operational phase of buildings requires the integration of highly heterogeneous data streams, including Building Information Modelling (BIM) models, Internet of Things sensors, computerised maintenance management systems, and technical documentation, yet current BIM deliverables remain poorly suited to this task. This paper presents the JUMENGI project, which investigates how such heterogeneous data can be structured and articulated within an interoperable digital twin dedicated to building energy maintenance. An ontology-based information architecture, instantiated as a knowledge graph, is designed following a Competency Questions methodology and implemented in a graph-oriented database. The approach is tested on two contrasting real-world case studies, a residential project with available BIM models and sensor data, and a tertiary building initially lacking an exploitable BIM model, and operationalised through a digital platform combining three-dimensional visualisation, sensor monitoring, and alert management. The results demonstrate the technical feasibility of unifying multiple operational data sources within a single semantic model, while revealing persistent limitations of existing BIM deliverables for operation and maintenance. The contribution lies in a validated, industrially grounded information architecture that bridges design-oriented BIM practices and operational energy maintenance needs, with direct relevance to interoperability and digital twin research within the built environment.

Keywords—digital twin; Building Information Modelling (BIM); knowledge graph; ontology; energy maintenance; interoperability

I. INTRODUCTION

The development of digital twin-based approaches in the construction sector is part of a gradual transformation of information management practices. This transformation is marked in particular by the extension of Building Information Modeling (BIM) beyond the design phase toward operation and maintenance (Eastman & al., 2011; Succar, 2009). In this context, the digital twin is generally defined as a dynamic virtual representation of a physical asset, capable of integrating heterogeneous and evolving data from multiple sources in order to simulate, monitor, and optimise its behaviour throughout its life cycle (Grieves & Vickers, 2017).

The JUMENGI research project fits into this dynamic by proposing an integrated digital twin approach oriented toward the energy maintenance of buildings. The project aims to develop and validate an operational methodology based on the articulation of data derived from BIM models, business information systems, IoT sensors, and energy simulation models. Numerous recent studies confirm the potential of this approach: the integration of BIM and IoT provides complementary views of a building, with BIM offering a high-fidelity representation at component level, while IoT data provide near real-time performance statistics (Bjørnskov & Jradi, 2023).

This ambition raises significant challenges in terms of information structuring, quality, and interoperability. Indeed, data used in property and facilities management are produced according to distinct professional logics, at different temporal scales, and with varying levels of granularity (Succar, 2009; ISO 19650-1, 2018; ISO 19650-2, 2018). Interoperability between BIM systems and asset management systems remains one of the main obstacles to the operational use of digital models (Matarnah & al., 2019). The JUMENGI project thus provides a relevant analytical context in which to examine the ability of BIM-based systems to integrate these

different sources of information from an operational perspective.

While existing studies have explored the integration of BIM and IoT for building performance monitoring, they often remain either conceptually generic or focused on design-oriented use cases. In contrast, the JUMENGI project specifically addresses the operational phase by proposing and validating an ontology-based information architecture grounded in real-world data constraints. The originality of this work lies in its focus on energy maintenance practices, its validation through contrasting case studies, and its critical assessment of the limits of current BIM deliverables when confronted with operational requirements.

Several ontology-based frameworks have also been proposed to support semantic interoperability in the built environment, notably the Building Topology Ontology (BOT), Brick Schema, and SAREF for smart applications. These frameworks primarily target spatial topology and device-level semantics, but generally treat maintenance events, technical documentation, and organisational responsibilities as secondary concerns. JUMENGI differs from these approaches by explicitly integrating management-level entities, such as maintenance events, alerts, and work orders, alongside hierarchical and technical components within a single ontology, thereby directly targeting the operational needs of energy maintenance rather than general-purpose interoperability.

The JUMENGI project was carried out in the context of the Win2Wal programme, funded by the Walloon Region of Belgium, in partnership with the construction company Thomas & Piron and the research centre Cenaero. This industrial anchoring gave the project a distinctive character: rather than working from idealised or synthetic datasets, the research was confronted from the outset with the constraints and imperfections of real-world building data. This grounding in practice makes the findings particularly relevant for the

construction and facilities management sectors, where the gap between normative BIM prescriptions and actual operational data remains a persistent challenge.

II. RESEARCH QUESTIONS AND OBJECTIVES

Within this context, the central research question can be formulated as follows. Given its multifaceted nature, it is decomposed into four analytical sub-questions (RQ1–RQ4), which structure the remainder of this paper and are answered explicitly in Section 6:

RQ: *How can heterogeneous data derived from BIM, information systems, and sensors be structured and articulated to enable their operational use within a digital twin oriented toward the energy maintenance of buildings?*

RQ1: *What are the characteristics, formats, and limitations of the informational landscape of building energy maintenance, as observed across two contrasting real-world case studies?*

RQ2: *How can an ontology-based information architecture be designed to structure and interconnect this heterogeneous data within a unified knowledge graph?*

RQ3: *How can this architecture be operationalised into a digital platform that makes it usable by facility and energy managers in a real operational context?*

RQ4: *What are the limits of current BIM practices when confronted with operational requirements, and how can information management be improved across the building lifecycle?*

More broadly, this question relates to the capacity of information systems to support decision-making uses in an environment characterised by diversity of formats, temporalities, and stakeholders. It extends research on the management of distributed knowledge among professional actors in knowledge-intensive organisations (Boland & Tenkasi, 1995), applying it to the specificities of information systems used during the operational phase of buildings.

This central question is operationalised through four specific objectives. The first objective consists in characterising the informational landscape of building energy maintenance, by identifying the data sources, formats, and actors involved in two contrasting real-world case studies. The second objective is to design and implement an ontology-based information architecture capable of structuring and interconnecting heterogeneous data within a unified knowledge graph. The third objective is to develop an operational digital platform, integrating BIM visualisation, sensor data, and alert management, that makes this architecture usable by facility and energy managers in a real operational context. The fourth objective is to critically assess the limits of current BIM practices when confronted with operational requirements, and to formulate recommendations for improving information management across the building lifecycle. Together, these objectives position the JUMENGI project at the intersection of knowledge engineering, BIM research, and facilities management, contributing to the broader agenda of Asset Management 5.0.

III. METHODOLOGY

The methodology adopted in the JUMENGI project is based on an experimental and iterative approach, combining case study analysis with the progressive development of an information architecture. Two contrasting case studies were selected in order to confront the approach with real-world situations: a residential project with available BIM models and sensor data, and a tertiary building characterised by fragmented documentation and initially lacking an exploitable BIM model (Figures 1 & 2). This dual approach makes it possible to highlight the gap between the theoretical prescriptions of BIM standards and the actual practices observed in the field, in line with the recommendations of the ISO 19650 standards (ISO 19650-1, 2018; ISO 19650-2, 2018).




Fig. 1. BIM model of the Nivelles/Renard residential project (43 housing units), illustrating a case study with available BIM models and sensor data.



Fig. 2. BIM model of the Wierde tertiary building (3 stories, 3140 m²), illustrating a case study characterised by fragmented documentation and initially lacking an exploitable BIM model.

The analysis of the available data made it possible to identify the main limitations of existing systems, notably the absence of energy-related parameters in BIM models, difficulties in accessing structured data, and interoperability issues between systems. These limitations are consistent with findings in the literature, which emphasise that existing BIM models are generally insufficient to meet the needs of the operational phase (Eastman & al., 2011; Matarnah & al., 2019).

Ontology Design Methodology

To address these challenges, an ontology dedicated to energy maintenance was developed, defined as a "formal, explicit specification of a shared conceptualisation" (Studer & al., 1998, p. 185). Inspired by the IFC standard but adapted to the specific needs of the project, this ontology structures the relationships between building elements, technical equipment, maintenance events, and data sources. Similar approaches have demonstrated their effectiveness in linking BIM/IFC data with IoT data streams for operational building management (Tang & al., 2025).

The ontology development followed a structured and iterative process articulated around the Competency Questions (CQs) approach (Gruninger & Fox, 1995). Rather than defining classes and relationships in an abstract manner,

the CQ approach anchors the ontology design process in concrete informational needs expressed by its future users. Each question acts as a functional requirement: the ontology is considered valid only if it is capable of providing a correct and complete answer. In the JUMENGI project, a set of eighteen competency questions was formulated in collaboration with operational stakeholders, including facility managers and energy managers, derived from real maintenance scenarios encountered on the case study sites. The full validation results of this approach are presented in Section 4.1.

This ontology forms the foundation of a knowledge graph implemented in a graph-oriented database, presented in section 3.2, enabling the integration of data from multiple sources. In parallel, a digital platform was developed, based on a REST API, a graph database, and a BIM viewer built on BIMData.io, in order to make this data usable in an operational context.

This experimental setup also provides insights into the gap between standardised BIM prescriptions and actual information management practices observed during building operation.




Fig. 3. Conceptual diagram of the JUMENGI Building Maintenance Ontology, illustrating the four main class categories (hierarchical components, IFC components, technical elements, and management elements) and their explicit relationships within the knowledge graph.

The ontology was designed to be compatible with the IFC standard while extending it with concepts specific to energy maintenance, a scope that IFC does not natively cover in sufficient detail. Four main conceptual categories were defined: hierarchical building components (site, building, floor, space, equipment), IFC components, technical elements (sensors, actuators, controllers, distribution systems), and management elements (maintenance events, alerts, work orders, technical documentation). Relationships between these categories were explicitly formalised to enable multi-

level queries spanning spatial, systemic, and temporal dimensions (Figure 3).

Knowledge Graph Implementation

The ontology was instantiated as a knowledge graph hosted in a graph-oriented database (Neo4j). This architectural choice was motivated by the intrinsic limitations of relational databases when dealing with highly interconnected data structures, where queries across multiple relationship levels generate significant performance costs. The graph model, by contrast, natively represents entities and their relationships as nodes and edges, enabling efficient traversal and semantic querying (Figure 4). The resulting graph integrates data from four heterogeneous sources: BIM models (exported via the IFC format), CMMS (Computerised Maintenance Management Systems), IoT sensor streams, and static technical documentation. Each data source required specific preprocessing and semantic alignment, mapping proprietary identifiers, normalising units, and resolving spatial referencing inconsistencies, before integration into the unified knowledge graph.

As shown in Figure 4, a fifth node category, User, is also represented. Unlike the other four categories, this node is intentionally kept peripheral to the domain graph: it governs access control and traceability (authentication and action logging) rather than participating in the building-domain semantic relationships, and is therefore not connected to the hierarchical, IFC, technical, or management components.




Fig. 4. Simplified diagram of the DB-Graph implementation, showing the five node categories (hierarchical components in blue, IFC components in grey, technical elements in orange, management elements in yellow, and users in green) and their typed relationships within the Neo4j graph database. The User node is intentionally left unconnected to the domain graph, as it governs access and traceability rather than building-domain semantics.

IV. RESULTS

The results demonstrate the feasibility of integrating heterogeneous data within a unified environment based on a knowledge graph. This integration makes it possible to cross-reference information from BIM models, sensor data, and

technical documentation, thereby offering a coherent view of the building and its systems. These observations are consistent with recent literature reviews on building digital twins, which confirm the relevance of architectures combining BIM, IoT, and simulation models for energy management (Hosamo & al., 2022; Bjørnskov & Jradi, 2023).

However, the results also highlight the limitations of current approaches. Existing BIM models generally prove insufficient to meet the needs of the operational phase, due to a lack of relevant maintenance information and inadequate structuring. This finding echoes the persistent interoperability challenges between BIM systems and asset management systems (Matarneh & al., 2019). While the developed ontology makes it possible to structure this data, it requires continuous adjustment in order to adapt to real-world data constraints, particularly regarding equipment localisation and relationship consistency.

Furthermore, the development of a platform integrating visualisation, sensor data, and alert management constitutes a first operational demonstration of the digital twin (Figures 5 & 6). The integration of advanced analytical methods, particularly for fault detection, opens up promising perspectives for predictive maintenance. In this field, AI-driven models have shown significant reductions in unplanned downtime (Semeraro & al., 2025), although their effectiveness remains highly dependent on data quality.




Fig. 5. Electrical meter sensor data visualisation on the JUMENGI dashboard, illustrating how real-time sensor data are contextualised through BIM-based building elements within the digital twin.




Fig. 6. Alert detail screen for a technical room, showing the linkage between detected anomalies, building spaces, and associated equipment within the knowledge graph-based digital twin.

Ontology Validation: Competency Questions Approach

The use of Competency Questions as a validation methodology for ontology development has been widely recognised in the knowledge engineering literature (Gruninger & Fox, 1995). This user-centred perspective is particularly critical in the context of building operation and maintenance, where information needs are highly heterogeneous, spanning spatial queries (localisation of equipment), temporal queries (maintenance history, sensor records), systemic queries (equipment chains, distribution networks), and organisational queries (responsibilities, intervention workflows). In the JUMENGI project, the eighteen competency questions formulated with operational stakeholders were structured across these four information domains (Table 1). They were subsequently translated into graph queries (Cypher, the Neo4j query language) and executed against the populated knowledge graph. The ability of the system to return accurate, complete, and contextually meaningful answers constituted the primary validation criterion for the ontology and its implementation. This validation process was conducted through a series of collaborative expert sessions held at LUCID (ULiège), bringing together facility managers, energy managers, and project partners to collectively review the ontology against each question (Figure 7).

Table 1. Representative competency questions structured by information domain, used as design requirements and validation criteria for the JUMENGI Building Maintenance Ontology.

Domain	Example competency questions
Spatial	In which building, floor and space is the boiler of site X located? Where are the ventilation sensors located?
Temporal	What data have been recorded over the past two months by the sensors in apartment 23? When was the last intervention on the air conditioning system?
Systemic	How many radiators are supplied by the heating circuit of apartment X? How many distribution networks are connected to this riser?
Organisational	Who programmed this controller and when? Where is the documentation of the most recent repair located?



Fig. 7. Co-design validation session at LUCID (ULiège): stakeholders collectively reviewing the JUMENGI Building Maintenance Ontology against the eighteen competency questions, with the ontology diagram projected for group discussion.

Beyond their role as design requirements, the competency questions also served as a continuous diagnostic instrument: several questions could not initially be answered due to missing relationships in the BIM models, absent sensor metadata, or incomplete CMMS records. Rather than being treated as failures, these gaps were systematically documented and fed back into the data enrichment process, driving iterative improvements to both the ontology structure and the data integration pipelines.

Dashboard Design Methodology

The development of the JUMENGI dashboard followed an Agile Scrum methodology, structured around iterative design-and-validation loops (Figure 8). The process began with a systematic literature review to establish a functional baseline, a catalogue of expected features for a BIM-integrated operational dashboard. This baseline covered four functional domains: model management (3D visualisation, IFC import, model update mechanisms, data exchange), project management (multi-site and multi-building access, user profile management, electronic document management), alert management (contextualised display, filtering, status assignment, export), and interface management (3D navigation tools, customisable views).

From this baseline, a functional backlog was defined and iteratively refined through a series of brainstorming and co-design sessions involving the project partners, including the industrial partner Thomas & Piron and technical teams from Cenaero. Each sprint produced validated mockups, which were reviewed against two main operational usage scenarios before consolidation into a final validated design. This iterative process ensured that the resulting interface genuinely responds to the information needs of building operation practitioners, rather than reflecting design-phase assumptions.



Iterative Design Loop



↳ Iterative loop → returns to next sprint if validation is not reached



Fig. 8. Agile Scrum iterative loop applied to the JUMENGI dashboard design process, structured in five sequential stages: literature review, initial brainstorming, mockup design, brainstorming validation, and integration into the next sprint. The loop returns to the previous stage if validation criteria are not reached.

Dashboard Features and Operational Scenarios

The resulting JUMENGI dashboard is structured around two main usage scenarios (Figure 9). The first concerns *routine monitoring*: the operator accesses a consolidated view of sensor data, contextualised within the 3D BIM model, allowing spatial localisation of each measurement point and navigation across building systems (Figure 10). The second concerns *anomaly management*: when a fault detection algorithm triggers an alert, the dashboard enables the operator to access a structured detail screen showing the anomaly type, the affected building space, the associated equipment chain,

and the historical maintenance records linked to the relevant components within the knowledge graph.

Four alert categories were defined and implemented, each corresponding to a distinct data source and detection mechanism (Figure 11). Sensor alerts (CA) are triggered directly by IoT data streams when measured values exceed predefined thresholds. Simulation alerts (SI) are generated by the CENAERO energy simulation module (OpenStudio), which compares measured building performance against simulated reference models derived from the BIM-to-BEM chain. Fault Detection alerts (FD) are produced by an AI-based fault detection and diagnosis module developed by CENAERO, applying unsupervised learning methods to identify anomalous behavioural patterns in the heating and domestic hot water systems. Finally, Client alerts (CL) originate from after-sales service (SAV) reports submitted directly by building occupants or facility managers. All four alert types are routed through the JUMENGI back-end and made accessible via the dashboard interface, where each is associated with a configurable severity level and a recommended response workflow, enabling the facility manager to prioritise interventions effectively.




Fig. 9. Two main operational usage scenarios of the JUMENGI dashboard: managing alerts across multiple buildings (left) and energy monitoring of a specific site (right), each structured as a sequential user workflow from login to decision-making or analysis.




Fig. 10. Screenshot of the JUMENGI dashboard interface showing the integrated 3D BIM viewer (left panel) alongside the structured building hierarchy navigator (right panel), enabling multi-level navigation across floors, zones, systems, and associated technical resources for the Wierde case study.




Fig. 11. Alert management interface of the JUMENGI dashboard, displaying the four alert types (Sensor Alert, Simulation Alert, Fault Detection Alert, and Client Alert) with their associated details: alert date, affected site, description, and diagnostic results, enabling facility managers to prioritise and track maintenance interventions.

The platform reached TRL 3 (Technology Readiness Level), corresponding to an experimental proof of concept validated under laboratory conditions. Three operational digital twin instances were deployed, each corresponding to one of the real-world case studies integrated into the project. Although full-scale validation in production conditions remains as future work, this first deployment demonstrates the technical feasibility of the proposed architecture and highlights the practical challenges of aligning heterogeneous real-world data with a unified semantic model.

V. DISCUSSION

The results presented above confirm the feasibility of integrating heterogeneous building data sources within a unified knowledge graph, while also revealing structural limitations of current BIM deliverables for the operational phase. This section discusses these findings in relation to the existing literature, the trade-offs inherent to the chosen architecture, and their broader implications for Asset Management 5.0.

Compared with prior work combining BIM and IoT for building performance monitoring (Bjørnskov & Jradi, 2023; Tang & al., 2025), the JUMENGI approach extends the scope of integration to CMMS records, technical documentation, and organisational responsibilities, rather than restricting itself to sensor-to-model mapping. This broader scope reflects the operational orientation of the project, but also increases the complexity of the ontology and the effort required to maintain its consistency as new data sources are onboarded. This observation aligns with the interoperability challenges identified by Matarneh & al. (2019), and confirms that these challenges persist beyond the design phase into building operation.

The choice of a graph-oriented database (Neo4j) over a relational architecture involves a trade-off between semantic richness and long-term maintainability. The graph model natively supports the multi-level, cross-domain queries required to answer the competency questions (Section 4.1), and proved adequate for the two case studies considered. However, this flexibility comes at the cost of a less standardised schema, which requires clear governance to prevent drift as the ontology is extended, for instance for reuse in the CARMEN project. A systematic, quantified comparison of query performance against relational alternatives was outside the scope of this proof-of-concept

stage, and constitutes a natural direction for the larger-scale validation planned as future work.

From an Asset Management 5.0 perspective, the JUMENGI results emphasise that technical integration alone is insufficient: the competency questions process showed that a substantial share of information gaps originated not from the graph architecture itself, but from missing or poorly structured data at the source, in the BIM models, sensor metadata, or CMMS records. This suggests that progress toward the human-centred, resilient, and sustainable objectives of Asset Management 5.0 depends as much on upstream data governance and collaboration between designers, BIM managers, and facility managers, as on the downstream semantic architecture.

Overall, these results should be interpreted as evidence of technical feasibility (TRL 3) rather than of operational performance at scale. The absence of systematic quantitative metrics, such as query execution times or data alignment accuracy rates, is a limitation of the present study; the collaborative validation sessions conducted at LUCID nonetheless provided a rigorous, stakeholder-grounded qualitative assessment of the ontology's adequacy. The collection of such quantitative performance indicators is planned as part of the larger-scale validation foreseen in future work.

VI. CONCLUSIONS

In relation to the four sub-questions formulated in Section 2, the JUMENGI project shows that: (RQ1) the informational landscape of building energy maintenance is characterised by fragmented, heterogeneous data of variable quality and granularity, even under favourable BIM conditions; (RQ2) an ontology-based information architecture, validated through competency questions, can structure and interconnect this data within a unified knowledge graph across contrasting case studies; (RQ3) this architecture can be operationalised into a digital platform, reaching a first proof-of-concept demonstration (TRL 3) usable by facility and energy managers; and (RQ4) current BIM models remain insufficient to support energy maintenance requirements, calling for the earlier integration of operational needs into BIM modelling practices.

The JUMENGI project demonstrates that the implementation of an operational digital twin primarily depends on the ability to structure, contextualise, and articulate heterogeneous data. It highlights the importance of an integrated information approach combining modelling, semantic structuring, and the development of technical tools adapted to the needs of the operational phase. The results call for a rethinking of BIM modelling practices by integrating energy maintenance requirements from the earliest project stages, in line with the recommendations of the ISO 19650 standards (2018), while also developing systems capable of managing data complexity and variability.

From a professional perspective, the results underline the need to better anticipate operational and maintenance requirements during the early stages of BIM modelling. They highlight the importance of clearer information requirements, improved data structuring, and closer collaboration between designers, BIM managers, and facility and energy managers. In this sense, the proposed approach contributes to bridging the gap between design-oriented BIM practices and the practical needs of building operation.

The scientific and technical contributions of the JUMENGI project extend beyond the project itself. The ontology codes, graph database architecture, and data integration pipelines developed within JUMENGI are currently being reused and extended within the CARMEN project (Carbon, Maintenance, Energy and Finance), funded under the Win2Wal 2025-2028 programme. CARMEN aims to develop a decision-support platform for sequencing renovation and maintenance operations throughout the building lifecycle, adopting a Dynamics of Complex Systems (DSC) approach. This continuity demonstrates the scalability and transferability of the semantic architecture proposed in JUMENGI, and confirms the relevance of ontology-based knowledge graphs as a foundation for more advanced building lifecycle management tools. It also underlines the importance of designing information architectures that are not project-specific but sufficiently generic to support reuse across different operational contexts and building typologies.

From this perspective, the digital twin appears as an evolving system, constructed through the interaction of stakeholders, tools, and data, and oriented toward the production of operational knowledge. Future work will focus in particular on automating the semantic enrichment of BIM models, improving the robustness of ontologies when confronted with real-world data, and validating predictive maintenance approaches on a larger scale.

ACKNOWLEDGEMENTS

The authors acknowledge the Walloon Region for its financial support of the JUMENGI project. They also acknowledge Thomas & Piron for providing real-world case studies, Cenaero for its expertise in energy modelling and simulation, and Aurélie Jeunejean for the development of the digital platform.

REFERENCES

- Bjørnskov, J. & Jradi, M. 2023. An ontology-based innovative energy modeling framework for scalable and adaptable building digital twins. *Energy and Buildings* 292: 113146. <https://doi.org/10.1016/j.enbuild.2023.113146>
- Boland, R.J. Jr. & Tenkasi, R.V. 1995. Perspective making and perspective taking in communities of knowing. *Organization Science* 6(4): 350–372.
- Eastman, C., Teicholz, P., Sacks, R. & Liston, K. 2011. *BIM Handbook: A Guide to Building Information Modeling for Owners, Managers, Designers, Engineers and Contractors*. Hoboken, NJ: Wiley.
- Grieves, M. & Vickers, J. 2017. Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In Kahlen, F.-J., Flumerfelt, S. & Alves, A. (eds.), *Transdisciplinary Perspectives on Complex Systems*: 85–113. Cham: Springer.
- Gruninger, M. (1995). Methodology for the design and evaluation of ontologies. In *Proc. IJCAI'95, Workshop on Basic Ontological Issues in Knowledge Sharing*.
- Hosamo, H.H., Svennevig, P.R., Svidt, K., Han, D. & Nielsen, H.K. 2022. A digital twin predictive maintenance framework of air handling units based on automatic fault detection and diagnostics. *Energy and Buildings* 261: 111988. <https://doi.org/10.1016/j.enbuild.2022.111988>
- ISO 19650-1. 2018. *Organisation and Digitisation of Information About Buildings and Civil Engineering Works, Including BIM – Information Management Using BIM – Part 1: Concepts and Principles*. Geneva: International Organization for Standardization.
- ISO 19650-2. 2018. *Organisation and Digitisation of Information About Buildings and Civil Engineering Works, Including BIM – Information Management Using BIM – Part 2: Delivery Phase of the Assets*. Geneva: International Organization for Standardization.
- Matarnah, S.T., Danso-Amoako, M., Al-Bizri, S., Gaterell, M. & Matarnah, R. 2019. Building information modeling for facilities management: A literature review and future research directions. *Journal of Building Engineering* 24: 100755. <https://doi.org/10.1016/j.jobbe.2019.100755>
- Semeraro, C., Lezoche, M., Panetto, H. & Dassisti, M. 2025. Analysis of digital twin applications in energy efficiency: A systematic review. *Sustainability* 17(8): 3560. <https://doi.org/10.3390/su17083560>
- Studer, R., Benjamins, V.R. & Fensel, D. 1998. Knowledge engineering: Principles and methods. *Data & Knowledge Engineering* 25(1–2): 161–197. [https://doi.org/10.1016/S0169-023X\(97\)00056-6](https://doi.org/10.1016/S0169-023X(97)00056-6)
- Succar, B. 2009. Building information modelling framework: A research and delivery foundation for industry stakeholders. *Automation in Construction* 18(3): 357–375. <https://doi.org/10.1016/j.autcon.2008.10.003>
- Tang, S., Shelden, D.R., Gao, X., Zhang, C., Hao, J., Guo, F. & Zhang, H. 2025. A framework for integrating BIM, building automation systems, and IoT: An application in design and operation using open standards. *Journal of Building Engineering* 113: 114092. <https://doi.org/10.1016/j.jobbe.2025.114092>