2026 European Conference on Computing in Construction
Corfu, Greece
July 12–15, 2026

Structuring Knowledge for AI-Driven Schedule Management in Digital Construction Twins

Tomas Lay Herrera[1], Matthieu Gillet[2], Mehrzad Shahinmoghadam[3] and Simon Bee[4]
[1]Autodesk GmbH, Germany
[2]Autodesk France, France
[3]Autodesk Research, Canada
[4]Autodesk Ltd., United Kingdom

Abstract

Construction projects, which are prone to delays and budget overruns, struggle to transform fragmented data streams into coherent intelligence. We propose a novel approach to monitoring that combines onsite data capture, agentic workflows and knowledge graph architecture to provide a foundation for evidence-based progress tracking. This approach is built on an ontology for structured schedule management that unifies concepts from domain-specific ontologies and Autodesk Construction Cloud. Autonomous AI-driven agents enable knowledge evolution by interacting with a knowledge graph. This lays the foundation for a living digital twin that adapts as new data emerges, supporting consistency checking and cross-system integration.

Introduction

Construction project execution involves large volumes of heterogeneous data, including schedules, building information models, and field reports, yet this information is rarely transformed into timely and coherent project intelligence. This gap continues to contribute to schedule delays and cost overruns, particularly during the execution phase where coordination and interpretation of evolving site conditions are critical (Gómez-Cabrera et al., 2023). In this context, effective schedule management depends on continuously reconciling plan and actual status, identifying emerging deviations, and relating observed site conditions to the activities and building elements they affect. Digital twins have been proposed as a means to address these challenges by maintaining synchronized digital representations of physical assets and construction processes. In construction, digital twin information systems have been framed as closed-loop mechanisms that explicitly represent and compare project intent and project status to support production control and decision making (Sacks et al., 2020). Central to this perspective is the interpretation of observations in relation to planned work. Prior studies show that progress understanding and predictive analyses become possible when site observations are grounded in a formalized project representation capturing spatial and temporal structure (Lilis et al., 2017). In practice, however, sustaining a “living” construction digital twin remains difficult, as information is heterogeneous, incomplete, and distributed across tools and stakeholders, complicating the maintenance of a coherent and queryable project state. Recent work has increasingly explored semantic and ontology-based approaches to structure construction information for digital twin applications, highlighting the role of explicit representations in integrating planning, monitoring, and derived knowledge across systems. In parallel, knowledge-graph-based approaches demonstrate how BIM data, schedules, and site information can be fused into machine-interpretable representations supporting analysis and downstream services (Katsigarakis et al., 2022). Despite these advances, many approaches either remain primarily representational without clear operational workflows, or focus on analytics without maintaining an explicit, inspectable project state over time. Building on these directions, this paper proposes an ontology-grounded, agentic framework for AI-driven schedule management in digital construction twins. The approach introduces the Intelligent Construction (INTC) ontology, which provides a unified semantic model integrating building topology, construction planning, on-site observations, and construction management concepts drawn from BOT, Digital Twin Construction (DTC), and Autodesk Construction Cloud (ACC). The ontology governs a construction knowledge graph that serves as persistent project state and supports explicit representation of observations, evidence, and state assertions. Over this shared representation, we define an agentic knowledge workflow in which specialized agents ingest project data, align observations with planned activities and building elements, and derive evidence-backed progress and deviation assessments. The framework is evaluated through an observation-to-insight pipeline instantiated in a controlled construction setting and validated using competency-question-driven queries over the resulting knowledge graph.

Related Work

Construction Site Monitoring

Construction site monitoring research has largely focused on two complementary problems: capturing as-built conditions and interpreting those conditions relative to planned scope and sequence. Prior work shows that integrating multiple monitoring modalities with BIM improves the completeness and reliability of project status acquisition compared to any single technology alone (Hasan

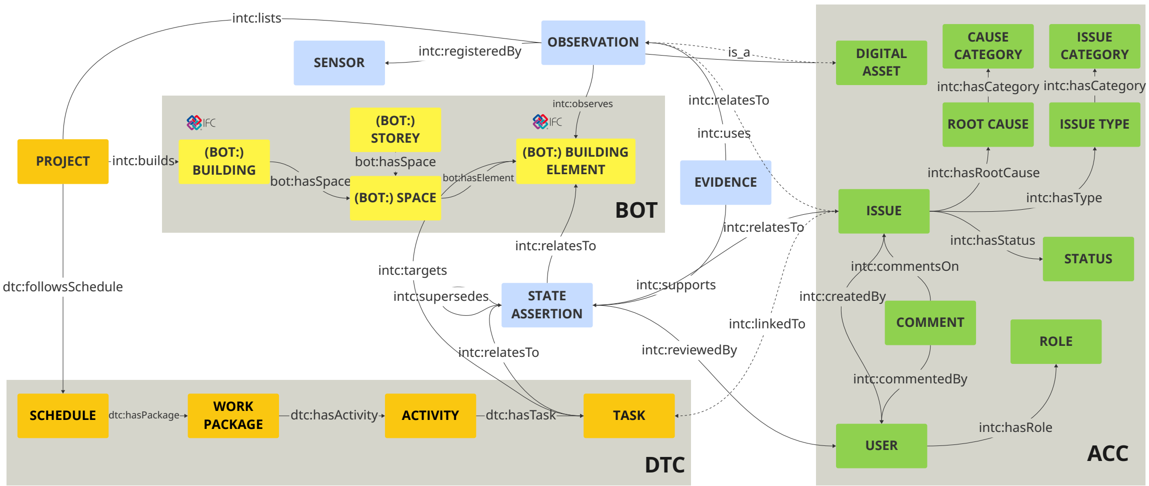

Figure 1: General overview of the INTC ontology

and Sacks, 2023). However, the central challenge remains not sensing per se, but interpreting heterogeneous observations in relation to formal project representations. Vision-based progress monitoring has received particular attention due to the practical availability of site imagery and time-lapse feeds. While recent systems have progressed beyond coarse, binary reporting toward richer state descriptions, they continue to face challenges related to occlusion, viewpoint variability, and robust mapping from detected signals to the project's formal objects and activities (Pal et al., 2023). As a result, visual signals alone often fail to produce reliable, attributable progress assessments. More recent work has explored lifting monitoring outputs into structured project knowledge rather than treating them as disconnected detections. Knowledge-graph based approaches demonstrate how BIM data, schedules, and site observations can be fused into machine-interpretable representations that support analysis and downstream services (Katsigarakis et al., 2022; Pfitzner et al., 2024). However, many such approaches either remain primarily representational, without clear operational workflows, or focus on analytics without maintaining an explicit, inspectable project state over time. This gap motivates our focus on persistent, evidence-linked representations that support plan–vs–actual reconciliation and traceable progress interpretation.

Ontologies for Construction

Ontologies provide a standardized machine-readable terminology to describe complex data and processes present in the construction process, while also enabling semantic interoperability across software capabilities involved in a construction project. Construction ontologies are an active research area, with multiple proposals that go from mid-level ontologies, such as the Digital Construction Ontology Suite (DiCon) (Zheng et al., 2021) or the Digital Twin Construction Ontology (DTC) (Schlenger et al., 2025), which attempt to cover all relevant aspects in the construction process, to domain-specific ontologies for scheduling, BIM data management, compliance and real-time site observations. Getuli (2020) presents a comprehensive study on knowledge modeling in construction planning, while Wu and Ma (2023) automatically generate construction schedules by combining constraint rules and genetic algorithms and Avogaro et al. (2024) proposed an IFC-based application for process scheduling. In the BIM data management field, Balaji et al. (2016) presented a unified schema for building elements, while Pauwels and Terkaj (2016) proposed ifcOWL, an attempt to convert IFC EXPRESS schemas into a usable ontology. In an attempt to address the practicality issues of ifcOWL, Rasmussen et al. (2020) published BOT, a shallower ontology focused on the topological relationships between building spaces. Ontologies such as OntoBPR from Hagedorn et al. (2025), or the Architecture, Engineering, Construction Compliance Checking and Permitting Ontology (AEC3PO) from Dridi et al. (2024) allow the modeling of compliance and permitting process using semantic web tools. Sensor-related ontologies like the Semantic Sensor Network Ontology (SSN) from Haller et al. (2018) or the lightweight ontology for sensors, observations, samples, and actuators (SOSA) from Janowicz et al. (2019) define a comprehensive vocabulary that enables the modeling of site monitoring activities.

AI-driven Agentic Frameworks

Recent work has increasingly framed large language models as agents, systems that interleave reasoning with tool use and state updates to support structured decision-making. ReAct (Yao et al., 2022) formalizes this paradigm by combining reasoning and acting in multi-step workflows, improving robustness and interpretability. While subsequent frameworks extend these ideas with modular orchestration, memory, and tool connectivity, most agent stacks remain domain-general. In construction, effective agent behavior requires grounding actions and inferences in a formal domain model and a continuously maintained project state, motivating our ontology- and knowledge-graph-centered approach.

Proposed Approach

The INTC ontology

Knowledge management in our approach is centered on a construction knowledge graph whose topology is governed by the Intelligent Construction (INTC) ontology. The INTC ontology acts as a connective tissue for four main areas of interest in construction: building information model (BIM), construction scheduling, on-site observations and construction management, the last one covered by Autodesk Construction Services (ACS), linking all these fields in a simple yet robust way. See Figure 1 for an overview of the INTC vocabulary and check (GitHub, 2026) for a version of the ontology, serialized in Turtle format. The BIM-related terminology is built upon an integration of the BOT and ifcOWL ontologies, to increase the coverage of building elements while considering IFC as a common input format for BIM data. The integration with the DTC ontology provides the required scheduling terms, while a one-to-one mapping with ACS concepts allows the definition of construction management concepts at the ontological level. The vocabulary for on-site observations is kept simple, focused less on sensor-related aspects and more on the modeling of insights derived from captured sensor data (see Figure 1 for an overview of the INTC vocabulary). The integration with ACS and the insights-driven on-site monitoring are the main reasons for the design on the INTC ontology.

INTC – ACS connection. ACS provides a comprehensive set of tools to interact with construction data stored in the Autodesk Construction Cloud (ACC). The ACC data model is intrinsically linked to the ACS services; however, there is no ontology to properly describe it. The INTC fills that gap, bringing ACC concepts like roles, users, and issues, among others, and adding semantic relationships between them. In order to avoid replicating the ACC data in the INTC knowledge graph, all entities related to ACC define an URN parameter that links to the corresponding resource for retrieval.

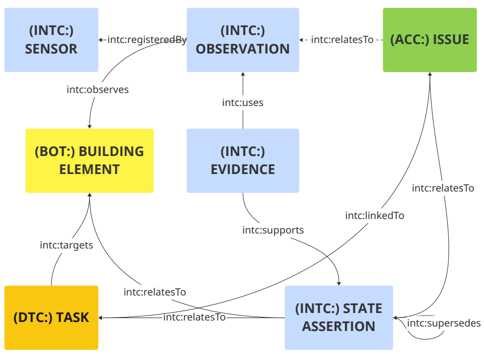

Figure 2: Observations, evidences, and state assertions in the INTC ontology

On-site observations in INTC. The sensor-related module of the INTC ontology aims to describe insights from on-site observations and how they connect with building elements, schedule activities, and potential ACC issues (see Figure 2). An important concept is the "Observation", captured by a "Sensor" (a camera, a microphone or any other sensing device) and used by a "Evidence". An "Evidence" leverages multiple "Observation" to support a "State Assertion" with a certain level of confidence. "State Assertion" is another central concept in this module and can be seen as a sort of update on the building element and the corresponding task in the schedule that targets it. An assertion might also eventually be related to an ACC issue, thus connecting the on-site monitoring with the ACC world. In summary, the core idea behind this design is that the progress of a task, associated with a building element, can be tracked through the associated state assertions, which are supported by evidence built upon observations. The "Ontology Validation" chapter describes a real-life scenario in which all these concepts are leveraged.

Agentic Knowledge Workflow

To operationalize AI-driven schedule management within a digital construction twin, we define an agentic knowledge workflow in which specialized agents operate on a shared construction knowledge graph governed by the INTC ontology (see Figure 3). The knowledge graph serves as the persistent project state and as the main coordination structure: agents record extracted information as typed entities and relations with provenance, timestamps, and confidence. Coordination is event-driven rather than purely sequential, so updates to graph state can activate subsequent processing steps while preserving an auditable record of how project understanding develops over time. The agentic workflow is decomposed into the following roles:

ACC Ingestion Agents (read-only). Retrieve authoritative project artifacts from Autodesk Construction Cloud (ACC) via Autodesk Platform Services (APS) and represent BIM/model artifacts, schedule structures, and documentation/specifications as INTC-aligned entities.

Semantic Linking Agent. Establishes and maintains the links between building elements, tasks, and relevant documentation/specifications required for monitoring, while preserving the provenance of these links.

Visual Observation Agent. Processes site images, registers each image as a DigitalAsset, and creates structured Observation entities linked to the most likely building element and task, while retaining confidence for perception and localization.

Progress and Deviation Agent. Uses newly added or updated observations to build Evidence supporting StateAssertion entities describing task progress and plan vs actual deviation. These outputs should be interpreted as evidence-based assessments rather than as ground truth. Conflicting evidence or assertions may coexist in the graph with their provenance and confidence, allowing later evidence or human review to refine the effective project state without removing prior interpretations.

ACC Sync Agent. Propagates selected progress updates back to ACC via APS. Synchronization is gated: low-confidence or conflicting interpretations are held for hu-

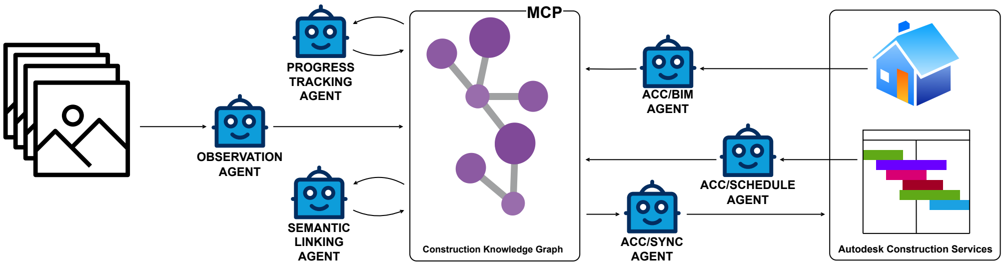

Figure 3: Simple, single-purpose agents operating on the construction graph

man review before write-back.

Overall, ACC provides the authoritative project context, while the knowledge graph provides the semantic structure through which site imagery is connected to observations, evidence, and state assertions for progress tracking and deviation assessment.

Observation-to-Insight Pipeline

The observation-to-insight pipeline converts raw observations into evidence-grounded project knowledge that can be compared against the plan to support activity-level progress monitoring and issue surfacing. The pipeline assumes an initialized project graph (from ACC ingestion and semantic linking) containing Task (planned activities) and BuildingElement entities, plus links to relevant requirements (e.g., execution specifications). This aligns with the DTC architecture that emphasize explicitly managing planning information, monitoring inputs, and derived knowledge in a unified structure (Schlenger et al., 2025).

Figure 4 instantiates the pipeline in the context of image-based progress monitoring as follows: For each new site image, the system (i) registers the file as a DigitalAsset (URI/URN, timestamp, capture metadata), (ii) applies a computer vision toolchain to extract candidate signals (e.g., detected components, stage indicators, potential defects), consistent with vision-based progress monitoring trends in construction (Pal et al., 2023). Next, (iii) it instantiates one or more Observation entities that are explicitly linked to the originating DigitalAsset and grounded to the most likely BuildingElement and Task (preserving uncertainty via confidence scores and, if needed, multiple candidate links). This "data → structured knowledge → metrics" flow is consistent with approaches that lift continuous image capture and object detection outputs into a knowledge-graph representation for construction process analysis (Pfitzner et al., 2024). Finally, with this pipeline in place, (iv) we can consolidate observations over time to maintain a current plan-vs-actual view at the activity level and emits actionable exceptions as Issue entities (e.g., suspected non-compliance or defect), each traceable through an Issue → Observation → DigitalAsset evidence chain. Where appropriate, issues can be exported back to ACC (e.g., as ACC Issues) under guarded policies (confidence thresholds and/or human approval).

While the individual stages of this pipeline (image capture, visual interpretation and progress reporting) are well represented in prior work, the central challenge lies in the alignment step: reliably mapping site observations into the project's object-activity model in a way that remains inspectable over time. In many existing systems, this alignment is either implicit (hard-coded mappings from detections to BIM objects, or transient (performed ad hoc for a specific metric or dashboard), which limits both extensibility and auditability. As a result, downstream "insights" often collapse heterogeneous evidence into a single status signal, obscuring how conclusions were reached and making them difficult to contest or revise as new information arrives.

The pipeline described here treats alignment as a persistent knowledge modelling operation. Observations are explicitly instantiated as first-class entities that mediate between raw digital artefacts and higher-level project claims. This allows multiple, potentially competing interpretations of the same evidence to coexist, each grounded in the project graph with explicit confidence and provenance, rather than forcing early commitment to a single explanation.

Experiments

Building on the pipeline logic introduced above (see Figure 4), we conducted a set of experiments focused on progress tracking using time-indexed site imagery. We assess the effectiveness of the ontology in: (i) grounding visually observable state changes in a knowledge graph instance, and (ii) maintaining an evidence-linked and traceable representation of construction state updates. It should be noted that our evaluation is intentionally scoped to the pipeline's semantic backbone rather than the end-to-end performance of the agentic workflow.

Experimental Setup

Our experiments are based on the ongoing construction of a real-scale mock-up room conducted in one of the Autodesk Technology Center in Boston. Figure 5 shows the 3D BIM representation of the room. The workshop floor is instrumented with multiple fixed cameras positioned to capture time-lapse sequences of construction activities from different angles.

Our setup takes Neo4j, a graph database following the Labelled Property Graph (LPG) as the storage solution for the construction knowledge graph. The INTC ontology is represented in RDF and serialized using the Turtle syntax, and

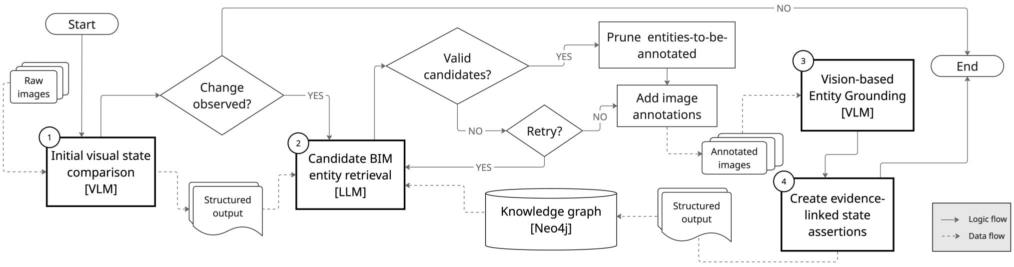

Figure 4: Logic and data flows of the observation-to-insight pipeline in the context of image-based progress monitoring; Steps 1–3 represent agentic enrichment of ingested imagery, while Step 4 establishes evidence and writes state assertions to the knowledge graph

Figure 5: 3D model of the mock-up room

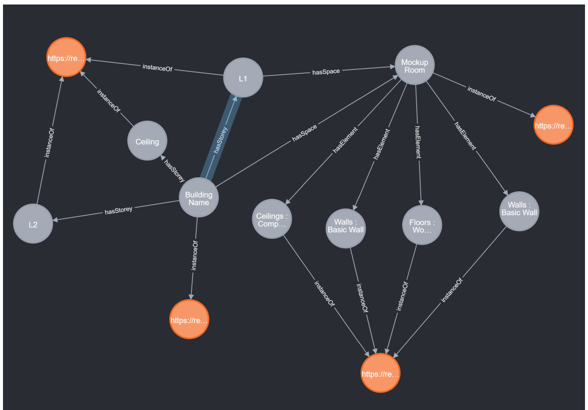

Figure 6: Subset of the building elements (gray nodes) of the mock-up room in Neo4j, connected to their corresponding ontology classes (orange nodes)

we leverage Neo4j's Neosemantics package to import it into the graph, forming the ''TBox'' (vocabulary) layer in the database. In order to link the instance data (''ABox'') with the vocabulary, we label the instance nodes with a single ''Instance'' label and link them with the corresponding class in the ''TBox'' with an ''instanceOf'' predicate, and ensure that all created relationships comply with the restrictions defined in the vocabulary, thereby leveraging ABox consistency while preserving key RDF semantics within an LPG environment (see Figure 6).

Figure 7 demonstrates instantiation of the main INTC ontology constructs, where the main subject of monitoring is the progress (stateType=''PercentComplete'') of one of the wall elements. An overview of instantiating the ontology artiacts for imagery observations is as follows:

Incremental ingestion: The camera captures from the the monitoring interval are represented as Observation instances (with stable external storage identifiers). Note that other sources of information in different formats, such as the site inspection report (text file), can be also represented as observation nodes (see top-left of Figure 7).

Agentic enrichment: The frames from the start and end of the current interval are fed to the image enrichment pipeline (see Figure 4), where a Vision Language Model (VLM) first compares the images to detect visible state changes. Then, the ACC/BIM agent (LLM-based) uses the generated structured change descriptions (in JSON format) and queries the BuildingElement nodes within the Neo4j instance to retrieve a set of likely BIM entity candidates that are set to be visible within the camera's field of view. Finally, it uses annotated overlays (pre-computed) to match each change to the correct BIM GUID by re-invoking the VLM for entity grounding.

Evidence establishment: The enriched observation artifacts become materialized as EvidenceItem nodes. Maintaining this explicit evidence layer is critical because multiple observations may support/contradict a single assertion (derived conclusion), and a single observation may be re-used to support multiple conclusions (e.g., element-level status and task-level completion). Evidence units can also carry indicators and references to intermediate outputs (e.g., confidence, image segmentation masks generated throughout enrichment steps).

State assertion: We use StateAssertion nodes to represent time-scoped and evidence-backed claim about the evolving condition of a subject entity. Each assertion node is associated with a set of key properties, e.g., StateType label (here ''PercentComplete''), state value (here percentage float), validity time (instant or interval), and provenance metadata (e.g., pipeline version). In the scenario depicted in Figure 7, we can see how multiple evidence units (originating from different sources at different times) can support a new assertion about the task's completion and supersede an existing one, while preserving the auditability chain from drawn conclusions back to raw observations.

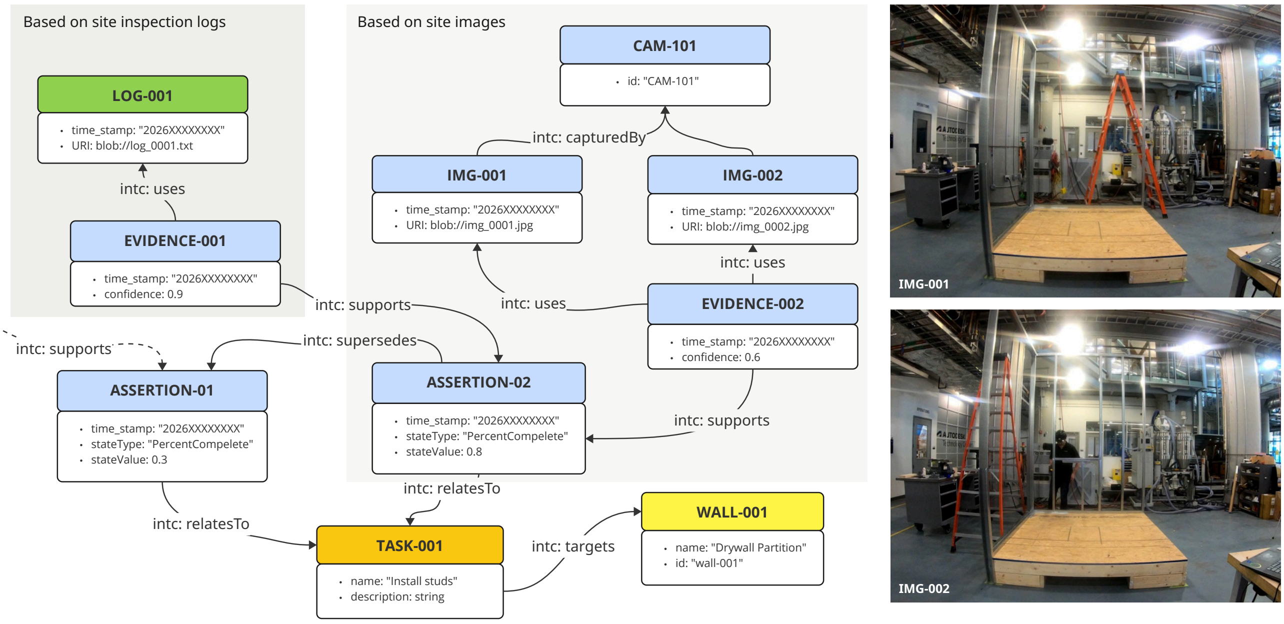

Figure 7: Example of ontology instantiation illustrating observation-to-insight pipeline in an experimental setting

Table 1: Summary of ontology validation competency questions

Category	Validation Purpose	Example Competency Question
Effective Progress State	Resolve multiple associated assertions into a single effective state (via validAt, confidence, etc.)	What is the current effective progress state of a given BuildingElement?
Temporal Reconstruction & Trajectories	Reconstruct progress as-of time t (via validAt); progress state evolution as a time series (trajectory, rates, stagnation)	What was the progress state and completion rate of a BuildingElement as of a time t?
Traceability & Evidence	Enable auditability: trace assertions to evidence, sensors, and raw data references	What evidence (observations) support a specific progress update?
Revisions & Updates	Track non-destructive revisions and state assertions history (when, what, who)	Was a progress state revised? what was the updated value? who initiated?
Spatial Context	Contextualize progress spatially (aggregate/filter by Zone/Space)	What is the progress of a given BuildingElement within a given Space?

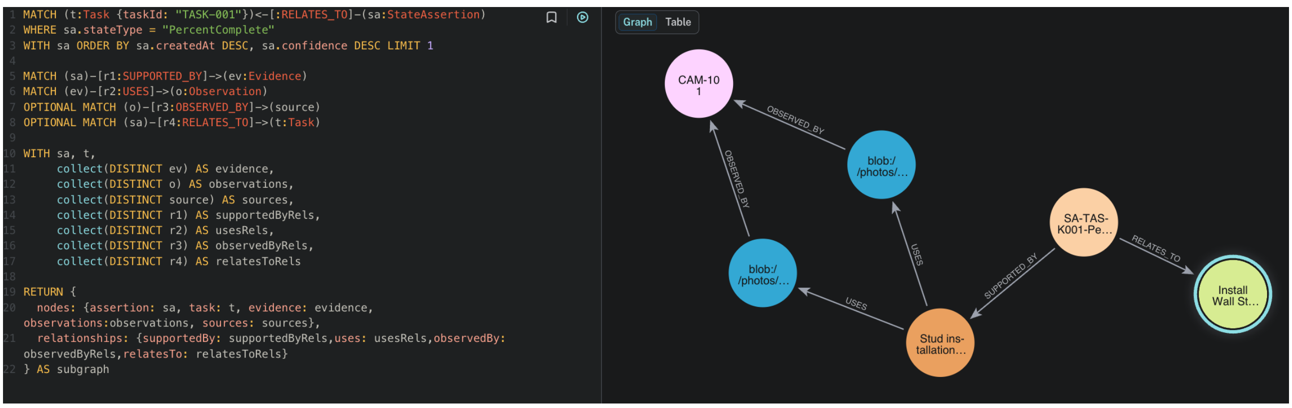

Figure 8: Example of CQ queries: retrieving latest state assertion and its chain of evidence for the data depicted in Figure 7

Validation

We adopt a Competency Question (CQ) methodology (Wiśniewski et al., 2019) to validate whether our proposed knowledge structuring can support the core representational requirements of the intended agentic workflow. The CQs were derived by first distilling the recurring information needs implied by the observation-to-insight pipeline (e.g., grounding, provenance) and then refining and prioritizing them through discussions with the domain researchers involved in the mock-up room deployment. A summary of the key patterns studied here is presented in Table 1. Each CQ is instantiated as Cypher query statements executed over the Neo4j instance, which contains data from the observation enrichment pipeline, as well as synthetic data needed to support the monitoring of the wall's construction progress. An example query associated with the scenario depicted in Figure 7 and it corresponding results are shown in Figure 8. To facilitate reproducibility of the validation results, QC evaluations are represented in a native LPG model (not RDF-based). The full list of questions studied in our experiments can be found in a designated public repository (GitHub, 2026).

Discussion

Our experiments indicate that our proposed ontology can support essential query patterns such as, "what changed," "why do we believe it changed," and "when did this belief become valid/invalid," using explicit representations. One of the key practical implications of our proposed ontology can be attributed to its support for traceable and auditable progress monitoring updates. First, it separates raw observations from interpreted claims about progress, which allows the system to retain historical assertions (including potentially conflicting ones) while updating the current estimate as further evidence becomes available. Second, the explicit linking of assertions to evidence and inference provenance makes each reported progress state inspectable: a user (or downstream analysis) can retrieve the originating sources (e.g., images, site logs), the intermediate evidence artifacts (enriched observation), and the method used to generate the claim (e.g., agent's model version). These aspects are crucial for construction practice, since field teams and project managers often have access to large volumes of signals (images, site reports, schedule logs), but the bottleneck lies in converting those signals into accountable estimates (state assertions) that can be traced back to evidence and acted upon. Moreover, progress updates and issue reports routinely require justification, especially when they trigger downstream actions such as schedule replanning, payment approvals, or rework decisions.

Our proposed representation is intended to be applicable across two complementary consumption modes. First, it can support deterministic workflows through formal queries embedded in conventional software. Second, it offers a structured representation for AI-enabled components and agentic workflows. Importantly, this dual use is enabled by a representation that remains relatively minimal: it captures the constructs required for traceability of evolving states without imposing an extensively granular design that can cause scalability issues for the broader system.

Finally, our evaluation primarily validates CQ answerability, however, does not directly quantify the advantage of the proposed model over existing data structures used in conventional schedule tracking systems. Although we demonstrate that the representation supports critical query patterns, we have not yet measured downstream impacts. In addition, our evaluation setup focuses on the construction of wall elements in a controlled environment. While this scenario can provide a clear sequence of visually observable states, other trade activities (e.g., mechanical/electrical rough-ins) may exhibit subtler visual indicators and greater ambiguity in mapping observations to relevant subjects or require additional sensing modalities. As such, the present contribution should be interpreted as a novel approach towards deviation assessment in construction progress.

Conclusions and Future Work

This paper introduced a semantic backbone for construction digital twins in the context of progress tracking. The proposed ontology provides a principled way to ground heterogeneous site observations into the planned project context by linking them to semantically-related entities (e.g., building elements or tasks). It enables the digital twin to maintain an evolving plan–vs–actual view as a set of time-scoped, evidence-backed assertions that remain queryable, interoperable, and inspectable as project knowledge is incrementally updated over time.

Future work should link the proposed representation to operational metrics and assess how effectively AI agents can exploit it, while clearly separating representational capability from agent performance. Larger-scale validation under more diverse sensing conditions remains necessary. The approach also opens avenues toward quality-aware progress monitoring and evidence-based causal analysis.

References

Avogaro, D., Berlato, M., and Zanchetta, C. (2024). Process scheduling application based on web ontology language representation of the ifc schema. In Construction Management Workshop, pages 125–135. Springer.

Balaji, B., Bhattacharya, A., Fierro, G., Gao, J., Gluck, J., Hong, D., Johansen, A., Koh, J., Ploennigs, J., Agarwal, Y., et al. (2016). Brick: Towards a unified metadata schema for buildings. In Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments, pages 41–50.

Dridi, A., Patlakas, P., Lefrancois, M., Beach, T., Petr, H., and Vakaj, E. (2024). AEC3PO: A knowledge model for machine executable construction regulations. Available at SSRN 4784123.

Getuli, V. (2020). Ontologies for Knowledge modeling in construction planning: Theory and application. Firenze University Press.

GitHub (2026). EC3-2026-intc. `https://github.com/tomaslay/EC3-2026`. Accessed: 2026-01-27.

Gómez-Cabrera, A., Gutierrez-Bucheli, L., and Muñoz, S. (2023). Causes of time and cost overruns in construction projects: a scoping review. International Journal of Construction Management.

Hagedorn, P., Fauth, J., Zentgraf, S., Seiß, S., König, M., and Brilakis, I. (2025). OntoBPR: An ontology-based framework for performing building permit reviews using standardized information containers. Advanced Engineering Informatics, 66:103369.

Haller, A., Janowicz, K., Cox, S. J., Lefrançois, M., Taylor, K., Le Phuoc, D., Lieberman, J., García-Castro, R., Atkinson, R., and Stadler, C. (2018). The modular ssn ontology: A joint W3C and OGC standard specifying the semantics of sensors, observations, sampling, and actuation. Semantic Web, 10(1):9–32.

Hasan, S. and Sacks, R. (2023). Integrating bim and multiple construction monitoring technologies for acquisition of project status information. 149(7). Publisher Copyright: © 2023 American Society of Civil Engineers.

Janowicz, K., Haller, A., Cox, S. J., Le Phuoc, D., and Lefrançois, M. (2019). SOSA: A lightweight ontology for sensors, observations, samples, and actuators. Journal of Web Semantics, 56:1–10.

Katsigarakis, K., Lilis, G., Rovas, D., González-Gerpe, S., Bernardos, S., Cimmino, A., Poveda-Villalón, M., and García-Castro, R. (2022). A digital twin platform generating knowledge graphs for construction projects. CEUR Workshop Proceedings.

Lilis, G. N., Giannakis, G. I., and Rovas, D. V. (2017). Automatic generation of second-level space boundary topology from ifc geometry inputs. Automation in Construction, 76:108–124.

Pal, A., Lin, J. J., Hsieh, S.-H., and Golparvar-Fard, M. (2023). Automated vision-based construction progress monitoring in built environment through digital twin. Developments in the Built Environment, 16:100247.

Pauwels, P. and Terkaj, W. (2016). EXPRESS to OWL for construction industry: Towards a recommendable and usable ifcowl ontology. Automation in construction, 63:100–133.

Pfitzner, F., Braun, A., and Borrmann, A. (2024). From data to knowledge: Construction process analysis through continuous image capturing, object detection, and knowledge graph creation. Automation in Construction, 164:105451.

Rasmussen, M. H., Lefrançois, M., Schneider, G. F., and Pauwels, P. (2020). Bot: The building topology ontology of the w3c linked building data group. Semantic Web, 12(1):143–161.

Sacks, R., Brilakis, I., Pikas, E., Xie, H. S., and Girolami, M. (2020). Construction with digital twin information systems. Data-centric engineering, 1:e14.

Schlenger, J., Pluta, K., Mathew, A., Yeung, T., Sacks, R., and Borrmann, A. (2025). Reference architecture and ontology framework for digital twin construction. Automation in Construction, 174:106111.

Wiśniewski, D., Potoniec, J., Ławrynowicz, A., and Keet, C. M. (2019). Analysis of ontology competency questions and their formalizations in SPARQL-OWL. Journal of Web Semantics, 59:100534.

Wu, Z. and Ma, G. (2023). Automatic generation of bim-based construction schedule: Combining an ontology constraint rule and a genetic algorithm. Engineering, Construction and Architectural Management, 30(10):5253–5279.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K. R., and Cao, Y. (2022). ReAct: Synergizing reasoning and acting in language models. In The eleventh international conference on learning representations.

Zheng, Y., Törmä, S., and Seppänen, O. (2021). A shared ontology suite for digital construction workflow. Automation in Construction, 132:103930.