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

TOWARDS SEMANTIC VOXEL-BASED REASONING FOR ROBOTIC FLEET MANAGEMENT IN CONSTRUCTION

Chu Han Wu[1, 2], Pranav Shevkar[1, 2], and Sigrid Brell-Cockan[1, 2]
[1]Cluster of Excellence CARE, TU Dresden and RWTH Aachen, Germany
[2]Chair for Individualized Production, RWTH Aachen University, 52074 Aachen, Germany

Abstract

Current robotic fleet management in construction primarily relies on 2D maps and geometry-based planners, with limited integration of process context and machine capabilities. This restricts scalability and adaptability in complex, dynamic sites. This paper proposes a semantic voxelisation framework that combines 3D voxel models with knowledge of construction processes, resources, and capabilities using IOC, BOT, and CaSkMan ontologies. The knowledge for capability-aware planning and task scheduling is queried using SPARQL and translated into domains and constraints for the planning framework. The contribution is a unified semantic-to-robotic pipeline enabling context-aware planning and scheduling for holistic fleet management.

Introduction

The increasing adoption of mobile robots and autonomous fleets in construction sites introduces new challenges for planning, scheduling, and coordination in highly dynamic, spatially complex environments. Unlike structured industrial settings, construction sites evolve continuously, involve heterogeneous resources, and impose strong dependencies between space, process, and machine capabilities. Current robotic fleet management systems predominantly rely on 2D geometric maps and local navigation constraints, which limit their ability to reason about volumetric space, task context, and resource suitability at a global planning level.

Parallel to advances in construction robotics, the construction informatics community has developed rich semantic representations of the built environment and construction processes, notably through ontologies such as the Building Topology Ontology (BOT), the Internet of Construction (IOC), and Capabilities and Skill of Machines, CaSkMan (CSS). However, these semantic models are typically used for information management, coordination, or analysis and are rarely integrated directly into task-planning and robotic fleet-management pipelines.

This paper argues that semantic voxelisation can serve as an effective intermediate representation between construction knowledge models and robotic systems. By discretising the construction site into a 3D voxel grid enriched with semantic information about space, processes, and resources, it becomes possible to perform capability-aware reasoning at multiple spatial and temporal scales. Tasks with explicit spatial and resource requirements can be matched against available voxel regions using semantic queries, while the resulting spatial constraints can be projected into representations consumable by robotic middleware such as ROS 2 and also industrial robotics.

Automated planning and scheduling are a branch of artificial intelligence that uses reasoning for generating a sequence of tasks that will lead from an initial state to a desired goal state.

We present a semantic voxelisation framework that integrates IOC, BOT, and CSS ontologies as a reasoning layer for task, process, and resource determination. The framework supports querying feasible workspaces and resource allocations via SPARQL, translating 3D semantic voxel information into 2D navigation costmaps, and providing higher-level inputs for fleet scheduling and task allocation. A prototype implementation demonstrates how semantic reasoning can inform navigation and fleet management decisions without replacing existing robotic planners.

The contribution of this work is a unified semantic-to-robotic pipeline that enables context-aware planning and holistic fleet management in construction environments, bridging the gap between construction knowledge representation and operational robotic systems.

State of the art

Fleet Management and Robotic Navigation in Construction

Current robotic fleet management systems in construction predominantly rely on 2D occupancy grids or topological maps derived from SLAM or BIM floor plans (Vega Torres et al. 2023). These representations are primarily geometry-centric and are decoupled from construction processes, task semantics, and machine capabilities. As a result, navigation and task allocation are typically treated as separate problems, handled by rule-based dispatchers or centralized schedulers (Fragapane et al. 2021).

Semantic Models and Digital Twins

Research on semantic digital twins and Linked Building Data (LBD) has introduced ontologies such as BOT, IFC-derived schemas, and construction process models to describe building topology, assets, and workflows (Boje et al. 2020). These models are largely used for documentation, monitoring, or offline analysis. While not unheard-of, there are attempts of using semantic data and IFC information for robotic planning such as works from (Crespo et al. 2020) and (Byers and RazaviAlavi 2022).

Limitations of Existing Approaches

Despite progress in robotics and semantic modelling, current approaches remain fragmented across geometric modelling, task planning, and robotic execution, limiting their effectiveness in dynamic construction environments. In particular, the lack of integration between spatial context, process knowledge, and execution frameworks restricts context-aware and adaptive decision-making:

- Spatial representations are largely 2D and static, ignoring volumetric constraints.
- Task planning is largely a manual process and does not explicitly account for site conditions, machine capabilities, or spatial affordances.
- Semantic reasoning is typically disconnected from robotic navigation and scheduling systems, such as ROS2.

This gap motivates a unified representation that bridges 3D space, semantics, and robotic execution.

Methodology

Overall System Architecture

The proposed approach introduces a semantic voxel layer, integrating ontologies from different domains with the Construction Site Voxel ontology as an intermediate representation between construction knowledge models and robotic fleet management systems. The framework integrates a voxelised spatial model of the construction site, semantic annotations using construction ontologies, and a reasoning pipeline to derive task-, resource-, and space-aware planning inputs.

The Construction Site Voxel Ontology, CVX

The **CVX Ontology** provides a semantic representation of construction sites based on voxelised volumetric space. Its primary purpose is to bridge geometric discretisation with construction semantics, enabling capability-aware planning, scheduling, and robotic fleet management. CVX acts as an intermediate semantic layer between building topology (BOT) (Rasmussen et al. 2020), construction processes (IOC) (Kirner et al. 2024), and capabilities and skills for machine capabilities (CSS) (Christian Diedrich et al. 2022), allowing volumetric regions of a construction site to be queried, reasoned over, and translated into robotic navigation and scheduling inputs.

Unlike traditional 2D grid maps or purely geometric voxel models, CVX treats each voxel as a first-class semantic entity, capable of representing spatial location, occupancy state, temporal validity, and process relevance. This enables integrated reasoning over where tasks can be performed, when space is available, and which resources are suitable. An overall concept of the ontology and its links to other ontologies are shown in Figure 1.

cvx:Voxel

The class *cvx:Voxel* represents an atomic volumetric unit of space within a construction site. Each voxel corresponds to a bounded 3D region and serves as the fundamental spatial abstraction in CVX. Voxels are uniquely identified and persist across time, allowing them to be referenced consistently in semantic queries, scheduling decisions, and robotic planning pipelines. The *cvx:voxelAttribute* represents non-tangible, contextual, and informational properties associated with a voxel that are relevant for planning, navigation, and fleet management. Unlike geometric or topological properties, voxel attributes capture dynamic, sensed, or inferred data

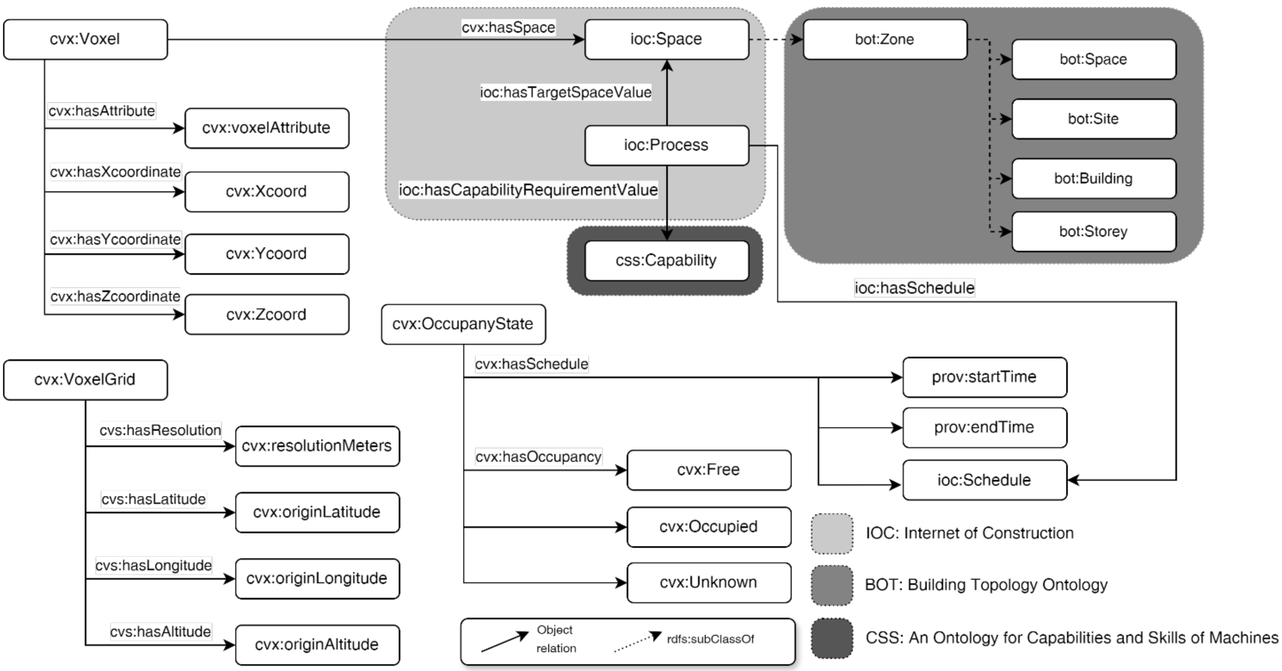

Figure 1: A first CVX ontology concept integrating IOC, CSS and BOT

that may influence robotic decision-making without constituting a physical obstruction.

cvx:VoxelGrid

The class *cvx:VoxelGrid* represents the global discretisation context within which voxels are defined. It captures parameters necessary to interpret voxel coordinates in real-world space. The key properties include *cvx:hasResolution*, specifying voxel size in meters, *cvx:originLatitude*, *cvx:originLongitude*, and *cvx:originAltitude*, defining the spatial reference origin of the grid. This design allows CVX to align voxelised construction sites with georeferenced building and site models, supporting integration with BIM, GIS, and digital twin systems.

cvx:OccupancyState

The class *cvx:OccupancyState* models the usage and availability of a voxel over time. It enables CVX to distinguish between static geometry and temporally evolving construction site conditions. A voxel's occupancy state is linked via *cvx:hasOccupancy* and can take one of three values:

- *cvx:Free* – the voxel is available for traversal or task execution
- *cvx:Occupied* – the voxel is currently in use or obstructed,
- *cvx:Unknown* – the voxel's state cannot be reliably determined.

This tri-state representation reflects real construction environments, where information may be incomplete or uncertain.

Data Storage

Voxel data and associated semantic annotations are persisted in an RDF triple store (e.g., GraphDB, Apache Jena, Fuseki, or Blazegraph) following Linked Building Data (LBD) modelling conventions. Voxels are represented as uniquely identified spatial entities and linked to building topology (BOT), construction processes (IOC), and machine capabilities (CSS). OWL DL axioms are used to capture class hierarchies, spatial affordances, and capability constraints, enabling lightweight semantic reasoning beyond purely syntactic SPARQL queries. SPARQL is subsequently employed for explicit selection and filtering of voxels that satisfy combined spatial, process, and resource conditions. Temporal aspects of voxel occupancy and usage are represented using time-indexed relations, allowing the semantic model to distinguish between current, planned, and historical states of the construction site.

AI-Based Reasoning

Planning Domain Definition Language (PDDL) (Haslum et al. 2019) is employed to model the current construction site state, the intended construction workflow, and the capabilities and status of the robots using the data from the CVX voxel ontology. This symbolic representation enables matching task requirements with available resources and reasoning over spatial feasibility to compute a valid task allocation. The planning output is then translated into an optimized execution schedule, which is consumed by the robot fleet manager. The fleet manager is responsible for low-level execution, including collision avoidance, navigation in dynamic environments, and handling unforeseen on-site changes. In the current implementation, this reasoning is limited to constraint extraction and does not yet include full symbolic plan generation.

Planning Methods

AI-based planning is formulated through two complementary specifications: a domain description and a problem definition. The domain captures the rules of the system, including available actions that are defined by preconditions and effects, objects, and logical predicates that represent world states. The problem specification provides the initial state and the desired goal state (Amer et al. 2021). A planning algorithm operates on both the domain and problem representations to synthesize a sequence of actions that transitions the system from the initial configuration to the goal (Sherkat et al. 2022).

Hierarchical Task Networks (HTN) is a technique that solves complex problems by decomposing high-level, abstract tasks into smaller and executable primitive actions (operators). It uses domain-specific methods to break down tasks, significantly reducing the search space for efficient planning in robotics (Höller et al. 2014). This enables the planner to exploit construction-specific structure and generate solutions that are aligned with established construction practices.

Temporal Planning methods generate action sequences while explicitly accounting for time-related constraints such as action durations, deadlines, and concurrency. These methods enable robots or other agents to (1) adhere to precedence relations such as "Action A must finish before Action B starts" or temporal requirements like "Complete Task before 3 PM", and (2) to coordinate concurrent activities if they do not conflict for optimizing resources. In this context, actions are modelled as temporally extended processes with defined start and end times, finite durations, and effects that may occur at specific time points or continuously over their execution. A combination of numeric and temporal planning can be used to model numeric changes over time, which is useful for resource management (e.g., the battery level) using planners such as OPTIC (Karpas and Magazzeni 2020). A key limitation of temporal planning is that action durations are typically modelled as estimates rather than guaranteed execution times. In real-world construction environments, the actual duration of robotic actions can deviate significantly due to factors such as environmental uncertainty, human intervention, dynamic obstacles, and resource conflicts. These discrepancies necessitate continuous execution monitoring and, when deviations exceed acceptable bounds, the ability to trigger online plan adaptation or re-planning to ensure timely and safe task completion (Staniaszek et al. 2025).

While this section introduces symbolic planning methods such as HTN decomposition and temporal planners (e.g., OPTIC), these are not fully integrated into the current prototype. The present work focuses on the semantic voxel reasoning layer and its translation into navigation-relevant representations (e.g., costmaps) for execution within ROS2 Nav2. The described planning methods represent a higher-level extension, where the derived spatial and resource constraints can be consumed by external planning frameworks for task allocation and scheduling.

Voxelisation of the construction site to fleet management

Semantic Voxel to Navigation Representation

In the proposed framework, voxels are not treated merely as discretised geometry, but as semantic mediation units between the Linked Building Data layer and fleet-level navigation and task allocation. Each voxel is modelled as a first-class semantic individual with explicit spatial extent and temporal validity, linked to its surrounding context through LBD-aligned relations such as *bot:containsElement* and *bot:intersectsZone*. Rather than encoding navigation logic directly, voxels act as semantic carriers of affordances and constraints. Through reasoning over CVX, BOT, IOC, and CSS, voxels inherit process- and resource-dependent semantics, such as work zones, safety buffers, or machine-restricted areas, with *ioc:process* instances associated with CSS machines and robots so that voxel states are interpreted relative to the capabilities and roles of specific fleet members, rather than as globally fixed occupancy values.

This semantic enrichment allows a voxel to represent not only whether space is free or occupied, but how it may be used by whom and when. For example, the same ground voxel may be traversable for a transport robot, temporarily restricted during a lifting process, or costly under degraded 5G conditions encoded via *cvx:voxelAttribute*. The ontology is queried using SPARQL to extract voxels satisfying fleet-relevant constraints, producing a temporally indexed set of 3D semantic voxels projected into a 2D costmap. Costs are derived from aggregated semantic attributes such as process activity, machine compatibility, or non-tangible signals, making the costmap directly usable by fleet planners while remaining grounded in the ontology. A Hybrid-A* algorithm then finds paths between all feasible start and goal positions for each capable robot, with metrics such as distance, travel time, cost, and battery consumption used to optimise the plan. Task allocation and scheduling are calculated from these estimates, with each robot assigned a priority reflecting task criticality, temporal dependencies, and resource constraints to resolve conflicts in shared or congested areas. Continuous monitoring detects deviations from estimated costs or travel times, enabling dynamic re-planning while ensuring collision-free paths that respect process constraints and overall mission objectives.

Use Case

To evaluate the proposed semantic voxel-based reasoning and planning framework, we consider a representative robotic construction scenario involving a heterogeneous multi-robot team executing a logistic workflow on a construction site for transportation. The use case is designed to simulate real construction environments. Sequential process dependencies, heterogeneous robot capabilities and coordinated multi-robot actions. The scenario considers a dynamic construction site where material handling and safety monitoring are performed by a heterogeneous robotic fleet, consisting of:

- 2 Unmanned Guided Vehicles (UGVs) dedicated to material transport within the site.
- One UGV dedicated to human presence monitoring, ensuring that restricted areas remain free of personnel during automated operations.

The site is represented as a voxelised three-dimensional spatial grid, where each voxel captures semantic, temporal, and environmental properties. The voxel grid forms the basis for navigation planning, task allocation, and safety enforcement. The use case serves as a concrete demonstration of how semantic voxel layer, *cvx* enables context aware task allocation and fleet coordination beyond conventional map-based approaches. In the present prototype, fleet coordination is demonstrated at the level of shared spatial constraints, task-dependent zone reservations, and robot-specific access conditions derived from the semantic voxel model. The validation therefore emphasizes coordination support through task feasibility and time-indexed navigation representations, rather than a full experimental benchmark of decentralized or reactive multi-robot coordination behaviours.

Voxel Grid Generation

The construction site is modelled and discretised into a 3D voxel grid using a custom C# script, in which each voxel represents a bounded volumetric region of space. In addition to geometric information, the script extracts and associates semantic metadata, including building elements derived from IFC models and machine-related information. The resulting A-Box is then translated to Turtle (.ttl) format and uploaded to the triplestore. Each voxel or voxel-based zone is subsequently assigned a contextual state, including:

- Occupancy state (Free, Occupied, Unknown)
- Process allocation (reserved by planned or executing tasks)
- Environmental metadata, i.e., 5G signal strength
- Temporal validity

Voxels are generated deterministically based on grid coordinates and world-space positions, ensuring stable identifiers across different systems. Figure 2 depicts the construction model, and Figure 3 shows the construction site after voxelisation. From this voxelised spatial–temporal representation, a binary navigation map is generated by querying voxels that are semantically classified as free at a specific time instant, as illustrated in Figure 4.

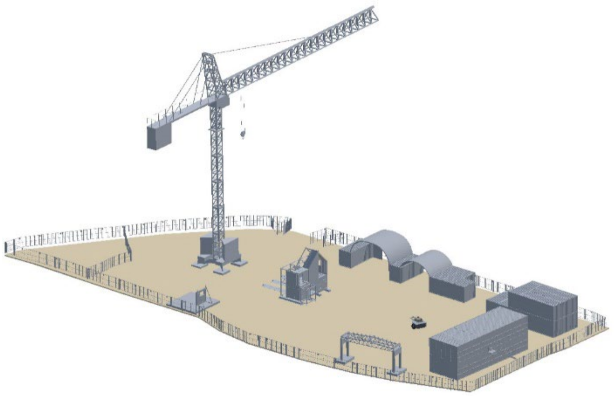

Figure 2: A model of the construction site

Figure 3: A voxelised construction site.

5G Signal Strength as a Spatial Cost Factor

A site-wide 5G signal strength map is integrated into the voxel grid as an environmental attribute. Signal strength varies spatially due to obstructions, distance to base stations, and site topology. This information is used to derive a communication-aware navigation cost map. The on-site signal strength is recorded using a Qualcomm

Figure 4: A 2D costmap queried from the knowledge graph.

SM8450 MTP (Mobile Test Platform), tested on Ericson NR Sub6 RU 3700MHz radio unit. Voxels with weak 5G coverage incur higher traversal cost. Transport UGVs prefer routes that maintain stable connectivity to ensure:

- reliable task updates
- safety monitoring
- remote intervention if required

By representing 5G signal strength as a spatially and temporally grounded attribute, the ontology enables context-aware path planning without hard-coding communication assumptions into the navigation stack. Figure 5, the top left of the figure shows a real test on the testbed using the MTP. On the right of the figure shows the distribution of the 5G signal strength. As tested in this paper (Behjati et al. 2022), it is shown that any Reference Signal Received Power (RSRP) above the value of -79dBM is able to provide a reliable connection for navigation and human observation, it is set that any signal strength less than the value of -80dBm has a high cost in the navigation map as shown in. The data is integrated into the knowledge graph and queried as voxels, the output is shown at the bottom left of Figure 5.

Robots, Capabilities, and Skills

At the semantic level, robots are modelled as resources that provide capabilities, which are realized through skills

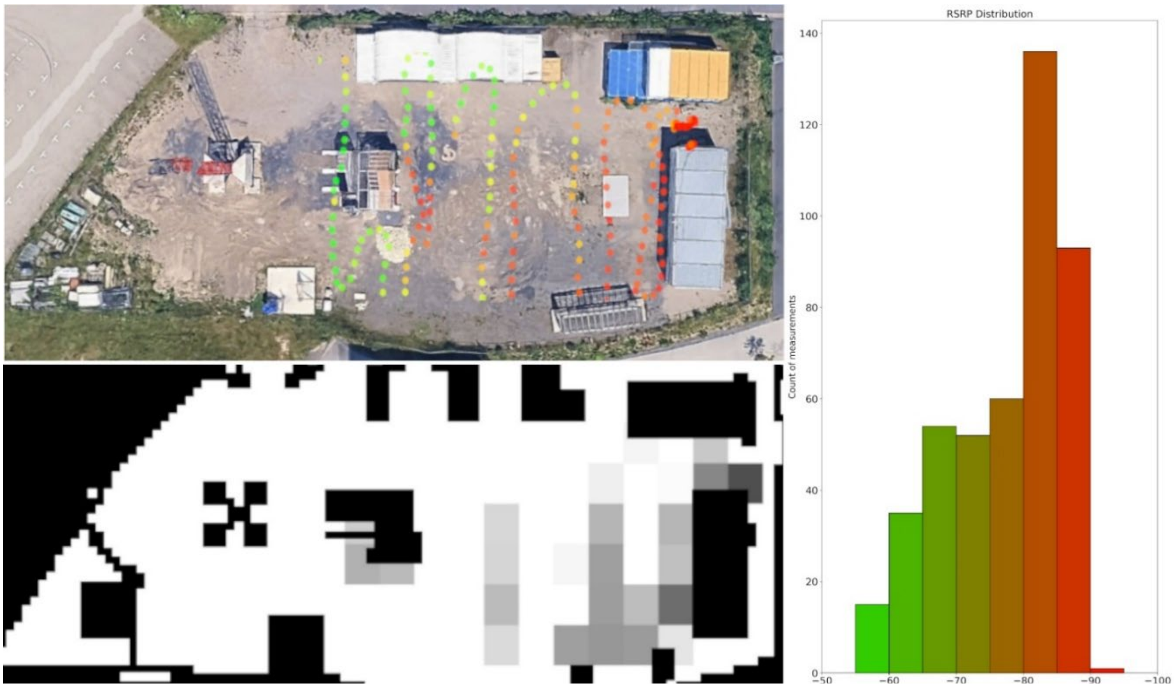

Figure 5: Top, Result of real test of the 5G signal strength on the test site. Bottom, a 2D costmap queried from the knowledge graph with 5G signal strength attribute. Right, a histogram of the data distribution.

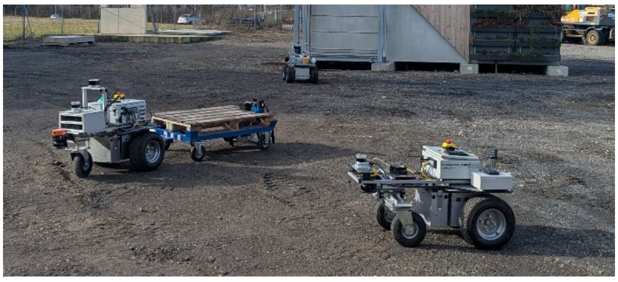

Figure 6: UGVs responsible for transport (left) and human observation

and accessed via skill interfaces. The two transport UGVs provide a Transport Capability, realised through:

- Locomotion skills (for navigation)
- Towing skills (for pallet dragging)
- Payload handling skills (for coupling and release)

The monitoring UGV provides a Safety Monitoring Capability, realised through:

- Locomotion skill (for navigation)
- Perception skill (for human detection)

enabling it to detect human presence within restricted zones. This separation allows tasks to be assigned based on capabilities, while execution selects appropriate skills depending on spatial and environmental conditions. Figure 6 shows the robots that are used for transportation and human observation. The robot on the left has a pallet that the robot can couple to and used to transport items while the robot on the right has an extra 3D LiDAR scanner and on-board computer with GPU for human detection algorithm.

Real-Time Updates and Consistency

The framework supports dynamic construction environments through incremental, event-driven updates of the semantic voxel model. Changes detected from sensors such as camera, LiDAR, or robot feedback are mapped to affected voxels, updating occupancy states, temporal validity (*cvx:hasOccupancy*), and contextual attributes without recomputing the full grid. These updates are propagated to the RDF store. The planning layer maintains consistency through triggered replanning and rolling-horizon execution, where only affected regions are reconsidered. At the execution level, the fleet manager continuously updates navigation costmaps derived from the semantic voxels, ensuring that robot behaviour remains aligned with the latest site conditions while preserving responsiveness.

Spatial-Temporal Map via contextualised working space

The proposed system maintains a spatial-temporal representation of the working space by contextualising spatial occupancy and 5G signal strength within a unified knowledge graph. Occupancy is updated from robotic perception and human-in-the-loop annotations, while communication data is integrated as a dynamic environmental layer. Unlike conventional layered costmaps, these attributes are modelled as first-class semantic entities with explicit temporal validity, enabling consistent querying of historical, current, and predicted site states.

Based on this contextualised representation, robotic task planning and execution are performed using standard frameworks such as Nav2 in ROS2. The knowledge graph acts as an intermediary semantic layer, where relevant spatial subsets are selected via SPARQL queries and lightweight semantic inference to identify regions satisfying task- and capability-specific requirements. These are then translated into navigation-compatible representations (e.g., occupancy grids or costmaps) for a given time horizon and task context, enabling capability-aware task allocation, temporal conflict avoidance, and context-aware navigation where traversal costs reflect both geometric and non-tangible factors.

Voxel-based reasoning refers to the process of evaluating semantically enriched voxels to derive planning-relevant constraints for task allocation and navigation. Each voxel encodes spatial, temporal, process, and contextual attributes, such as assigned processes, zones, machine capability requirements, occupancy state, and 5G signal strength. For example, a transport task requiring a UGV with a specific payload capacity will only be allocated to a voxel route where the associated capability constraints are satisfied, the occupancy state is unobstructed for the relevant time window, and signal strength exceeds the threshold required for teleoperation. This directly shapes both task allocation and the costmap weights used for navigation, meaning semantic voxel attributes are not merely annotations but active inputs for planning.

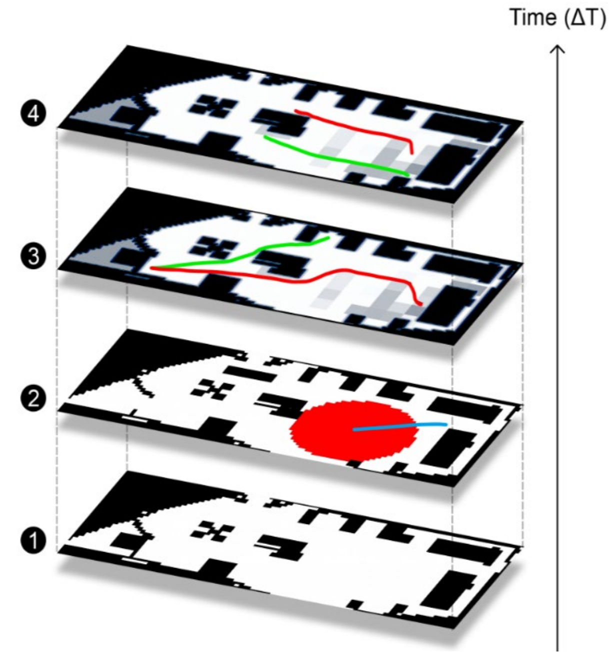

Figure 7: A voxel-based costmap for navigation queried from the knowledge graph at discrete time instance for initial navigation (1), human observation (2) and multiple and parallel processes (3,4).

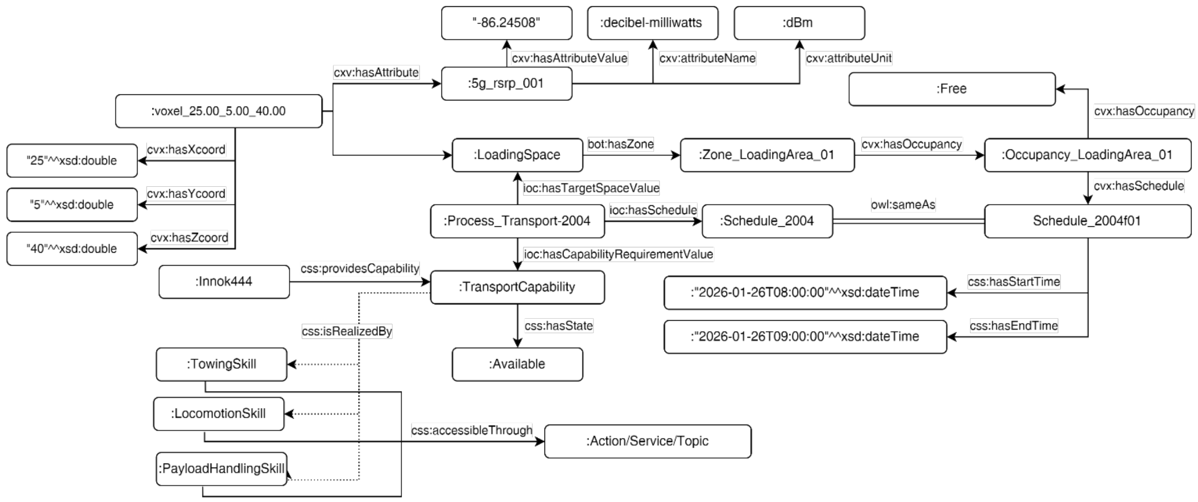

Figure 8: A excerpt of the instantiated voxel, linked to process, zones and machine capabilities.

As tasks are executed, routes, execution states, and temporal dependencies are written back into the knowledge graph, closing the perception–planning–execution loop. This supports continuous refinement of the spatial-temporal model and coordination among multiple agents operating in the same environment.

Figure 7 illustrates snapshots of navigation maps at different time periods. Map 1 shows a conventional binary occupancy map, while Map 2 incorporates human observation zones and a UGV navigation path. Maps 3 and 4 include 5G signal strength with two UGVs assigned transport tasks at different time intervals, where the planner avoids low-signal regions. While the resulting maps resemble layered costmaps (Lu et al. 2014), the key distinction is that this framework generates time-indexed spatial representations, rather than aggregating layers into a single static grid. Figure 8 shows an instance of a voxel linked to an IOC process and BOT zone, with the process associated with CSS-defined machine capabilities.

Conclusion

This paper presents an initial concept for a semantic, spatial-temporal framework for contextualising robotic working spaces in dynamic construction environments. By integrating occupancy states, task semantics, and attributes into a knowledge graph, the proposed approach enables robots to reason over space not only as a geometric entity but as a temporally evolving and context-dependent resource. This allows navigation and transport planning to move beyond static map assumptions and towards situation-aware decision-making.

The contextualised working space acts as a semantic intermediary between perception, planning, and execution. Standard robotic navigation stacks, such as Nav2 in ROS 2, can be leveraged without modification by exporting task- and time-specific spatial representations derived from the knowledge graph. In parallel, execution feedback and transport process updates are written back into the graph, enabling continuous synchronisation between the physical site and its digital representation. This bidirectional coupling supports coordination among heterogeneous agents, including UGVs and human observers, within a shared operational model. However, the current implementation focuses on generating semantically grounded constraints and navigation representations, while full integration with symbolic planning methods such as HTN or temporal planners (e.g., OPTIC) remains part of future work.

By treating spatial regions as first-class semantic entities rather than passive costmap layers, the framework supports traceability, explainability, and extensibility across planning horizons and task types, providing a foundation for planning under real-world constraints common to construction sites, such as dynamic obstacles, intermittent connectivity, and human presence.

Future Works

As a first approach, the proposed framework establishes the groundwork for a structured research roadmap. The following directions outline the planned extensions to mature the system toward real-world deployment. Several extensions can be pursued to further enhance the proposed spatial-temporal framework and broaden its applicability in real-world construction environments.

First, the framework can be extended through formal integration with symbolic planning models, such as Hierarchical Task Networks (HTN) and classical PDDL. Spatial-temporal regions represented in the knowledge graph can be abstracted into planning predicates, enabling high-level task decomposition and constraint-aware reasoning. HTN models are particularly promising for representing construction workflows, where transport, assembly, and inspection tasks follow structured yet adaptable procedures. Such an integration would allow semantic spatial constraints and communication conditions to directly influence task planning and execution at the symbolic level.

Second, the 3D nature of the voxel representation can be further exploited beyond conventional 2D costmap applications. Extending the framework to full 3D spatial–temporal planning would enable coordination of heterogeneous platforms operating across multiple elevation layers, including quadrupedal mobile robots capable of stair traversal, automated cranes transporting materials across building stories, and aerial robots operating in the airspace above the site.

Third, with sufficient data, learning-based occupancy prediction becomes feasible. Graph Neural Networks (GNNs) can be leveraged to model spatial-temporal dependencies between adjacent regions, robotic activities, and construction processes. By conditioning predictions on process states and task progress, the system could anticipate future occupancy changes and proactively adjust navigation and scheduling strategies. This would shift the framework from reactive planning towards predictive and anticipatory behaviour.

Finally, the proposed approach can be generalised into a knowledge-driven AI fleet management framework. By unifying semantic spatial representations, communication context, task models, and execution feedback, the system can serve as a coordination layer for heterogeneous robotic fleets. Such a framework would support dynamic task and resource allocation, and human-robot coexistence across large-scale sites, positioning the knowledge graph as the central decision backbone for adaptive, multi-agent construction robotics. Future evaluation would include explicit multi-robot coordination experiments, such as shared corridor negotiation, synchronized task execution, and conflict resolution in overlapping work zones, to quantify the impact of semantic voxel framework on coordination performance.

Acknowledgement

This project is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 3115 – 533767731.

References

Amer, Fouad; Koh, Hui; Golparvar-Fard, Mani (2021): Automated Methods and Systems for Construction Planning and Scheduling: Critical Review of Three Decades of Research. In Journal of Construction Engineering and Management 147. DOI: 10.1061/(ASCE)CO.1943-7862.0002093.

Behjati, Mehran; Zulkifley, Muhammad Aidiel; Alobaidy, Haider A. H.; Nordin, Rosdiadee; Abdullah, Nor Fadzilah (2022): Reliable Aerial Mobile Communications with RSRP & RSRQ Prediction Models for the Internet of Drones: A Machine Learning Approach. In Sensors (Basel, Switzerland) 22 (15). DOI: 10.3390/s22155522.

Boje, Calin; Guerriero, Annie; Kubicki, Sylvain; Rezgui, Yacine (2020): Towards a semantic Construction Digital Twin: Directions for future research. In Automation in Construction 114, p. 103179. DOI: 10.1016/j.autcon.2020.103179.

Byers, Gareth; RazaviAlavi, SeyedReza (2022): Layout modelling of the built environment for autonomous mobile robots using Building Information Modelling (BIM) and simulation. In mocs, pp. 201–208. DOI: 10.29173/mocs283.

Christian Diedrich; Alexander Belyaev; Rafi Blumenfeld; Juergen Bock; Stephan Grimm; Jesko Hermann et al. (2022): Information Model for Capabilities, Skills & Services. DOI: 10.13140/RG.2.2.30098.53440.

Crespo, Jonathan; Castillo, Jose Carlos; Mozos, Oscar Martinez; Barber, Ramon (2020): Semantic Information for Robot Navigation: A Survey. In Applied Sciences 10 (2), p. 497. DOI: 10.3390/app10020497.

Fragapane, Giuseppe; Koster, René de; Sgarbossa, Fabio; Strandhagen, Jan Ola (2021): Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda. In European Journal of Operational Research 294 (2), pp. 405–426. DOI: 10.1016/j.ejor.2021.01.019.

Haslum, Patrik; Lipovetzky, Nir; Magazzeni, Daniele; Muise, Christian (2019): An Introduction to the Planning Domain Definition Language. Cham: Springer International Publishing.

Höller; Daniel, ller; Gregor, Behnke; Pascal, Bercher; Susanne, Biundo (2014): Language Classification of Hierarchical Planning Problems. In : ECAI 2014: IOS Press (Frontiers in Artificial Intelligence and Applications).

Karpas, Erez; Magazzeni, Daniele (2020): Automated Planning for Robotics. In Annu. Rev. Control Robot. Auton. Syst. 3 (1), pp. 417–439. DOI: 10.1146/annurev-control-082619-100135.

Kirner, Lukas; Oraskari, Jyrki; Wildemann, Peter; Brell-Cokcan, Sigrid (2024): Internet of Construction Process Ontology (ioc) v 0.5: Zenodo.

Lu, David V.; Hershberger, Dave; Smart, William D. (2014): Layered costmaps for context-sensitive navigation, pp. 709–715. DOI: 10.1109/IROS.2014.6942636.

Rasmussen, Mads Holten; Lefrançois, Maxime; Schneider, Georg Ferdinand; Pauwels, Pieter (2020): BOT: The building topology ontology of the W3C linked building data group. In Semantic Web 12 (1), pp. 143–161. DOI: 10.3233/SW-200385.

Sherkat, Shermin; Wortmann, Andreas; Wortmann, Thomas (2022): Potentials of Symbolic AI Planning for Construction.

Staniaszek, Michal; Brudermüller, Lara; You, Yang; Bhattacharyya, Raunak; Lacerda, Bruno; Hawes, Nick (2025): Time-bounded planning with uncertain task duration distributions. In Robotics and Autonomous Systems 186, p. 104926. DOI: 10.1016/j.robot.2025.104926.

Vega Torres, M. A.; Braun, A.; Borrmann, A. (2023): OGM2PGBM: Robust BIM-based 2D-LiDAR localization for lifelong indoor navigation. Available online at https://arxiv.org/pdf/2308.05443.