A Hybrid Semantic Web and Data-Driven Framework for Spatial Constraint-Aware Construction Planning

Zeyu Mao[1], Eyob Mengiste[2], Borja García de Soto[2], Vicente A. Gonzalez[3]

[1] Infrastructure and Human Tech (IHT) Lab, Department of Civil and Environmental Engineering, University of Alberta, Canada
[2] S.M.A.R.T. Construction Research Group, Division of Engineering, New York University Abu Dhabi (NYUAD), United Arab Emirates
[3] Infrastructure and Human Tech (IHT) Lab, Strategic Projects Insight Centre in Engineering (SPICE), Department of Civil and Environmental Engineering, University of Alberta, Canada
zmao5@ualberta.ca, eyob.mengiste@nyu.edu, garcia.de.soto@nyu.edu, vagonzal@ualberta.ca

Abstract –

Construction workflows are restricted by multiple constraints, factors that limit when and how tasks can be performed. Among them, spatial constraints are particularly important as workers and equipment often compete for limited areas, and insufficient clearance can make tasks unsafe or reduce productivity. However, existing planning tools, such as the Last Planner System (LPS), a Lean-based production planning and control tool for construction, typically treat spatial constraints in binary terms, i.e., the relationship between space constraints and task performance, in decision-making regarding space constraint identification and removal in LPS. To investigate such a grey area and provide data-informed planning, this paper introduces a hybrid semantic web and data-driven approach to capture the relationship between task performance and spatial variability. Among them, the semantic web models the relationships between spatial variabilities and task requirements as ontology and reasoning rules, while space conditions are captured and mapped into a semantic model through a point cloud to an Ontology pipeline. By combining semantic modelling, data-driven decision-making, and AI knowledge capture, the proposed approach contributes to the automation of construction digital planning tools and offers a pathway towards a self-learning system in construction.

Keywords –

The Last Planner System, Semantic Web, Space Constraint Management, Point Cloud

1 Introduction and Related Work

Construction production is restricted by constraints that limit when, where, and how work can be executed. Among these, spatial constraints are particularly influential because construction activities are performed in physically bounded and dynamically changing environments where crews, equipment, and materials compete for limited workspace [1]. Insufficient clearance, congestion, and temporary obstructions can compromise safety and degrade productivity, making space a critical determinant of workflow reliability.

Early digital approaches to space constraint management primarily relied on geometric representations and binary logic. Studies such as Burcu et al. [2] and Guo [3] modeled task spaces as static volumes and assessed feasibility through overlap or scheduling checks. However, these approaches oversimplified spatial dynamics by assuming static environments and treating space as either available or unavailable. With advances in sensing and reality capture, later work incorporated laser scanning and BIM to achieve more timely geometric awareness of site conditions [4]. However, despite improved geometric fidelity, these systems largely remain descriptive, offering limited semantic interpretation of how spatial conditions influence task feasibility or performance.

Ontology-based models have been proposed to encode spatial semantics and relationships, enabling machines to reason about task requirements, hazards, and workspace zones [5, 6]. These approaches enhance transparency and knowledge structuring, but are typically static, manually updated, and loosely connected to real-time field data. As noted in a recent review [7], most construction ontologies fail to explain how evolving spatial constraints impact task performance, thereby limiting their usefulness for operational decision-making.

More recent research has explored hybrid semantic and AI-enabled systems that integrate ontological reasoning with data-driven methods and sensor inputs [8, 9]. While these approaches demonstrate the potential of self-updating knowledge models, they rarely quantify the productivity implications of spatial constraints or provide mechanisms to refine reasoning logic using empirical performance outcomes.

In parallel, production planning and control methods, such as the Last Planner System (LPS) [10], typically treat spatial constraints in binary terms (space is ready / not ready). While this abstraction supports the identification of severe conflicts, it provides limited diagnostic value during lookahead and weekly planning. A binary label does not offer planners or field teams a measurable, visible, or traceable explanation of why a space condition is considered a space constraint, nor does it quantify the expected performance consequences of proceeding under partially constrained conditions. In practice, readiness assessments are often informal, experience-driven, or based on qualitative judgment, which can lead to inconsistent interpretations across planners and trades. As a result, work may still be committed under partially constrained spatial conditions due to vague space constraint clarification: situations that Lean theory characterizes as making-do and identifies as waste. To address these limitations, this paper proposes a hybrid semantic web and data-driven framework that represents space as a graded constraint and explicitly links observed spatial conditions to their expected impact on task productivity. Rather than legitimizing execution under constrained conditions, the framework provides a structured, measurable, and traceable diagnosis of spatial readiness, making the performance consequences of spatial constraints visible during lookahead and weekly planning.

An ontology-based representation was adopted in this study as a proof-of-concept because it provides an interpretable and structured way to formalize task-space relationships, connect measured site conditions to reasoning rules, and generate traceable planning-relevant diagnostics from observed WorkZone conditions. By integrating ontology-based reasoning with point-cloud-derived spatial metrics and empirical productivity feedback, the proposed approach supports evidence-based decision-making for constraint removal and commitment planning, thereby strengthening spatial constraint management within LPS-based workflows.

2 Proposed Framework

This section presents the proposed framework, which consists of three interconnected phases (Figure 1): (1) extending the base ontology to incorporate computable spatial concepts, (2) generating a Space Condition Profile (SCP) from point cloud–derived geometric attributes, and (3) applying productivity reasoning rules that translate spatial conditions into quantitative productivity adjustments. Each phase is described below.

2.1 Phase 1 – Ontology Extension

The proposed ontology model builds upon the authors' previously developed Space Constraint Management Ontology (SCMO) [1] and extends its foundation to provide automated space condition evaluation and productivity reasoning.

The extension, as shown in Figure 2, focuses on enriching existing concepts, particularly WorkZone and Activity, with measurable WorkZone attributes and task-level productivity descriptors. Spatial measurements obtained from sensing processes are incorporated as data properties. For example, we added data properties "hasOccupiedAreaRatio," "hasMaxObjectiveHeight," "hasHumanOccupancy," and "hasFreeFloorArea" to the WorkZone class to represent spatial clutter, vertical obstruction, and usable area. On the activity side, we added three new productivity-related properties, "hasNominalProductivity," "hasProductivityAdjustmentFactor," and "hasExpectedProductivity," to allow tasks to store both baseline performance and reasoning outputs.

2.2 Phase 2 – Space Condition Profile (SCP) Generation

The second phase generates an SCP for each WorkZone instance by integrating geometric measurements and contextual information extracted from daily reality-capture data. In this study, spatial information is captured through a combination of daily 360-degree panoramic images and laser-scanning point clouds.

The SCP generation process relies on two complementary information streams, which were derived from the author's previous research [11]. First, visual context from daily 360-degree images is used to capture execution-related information that cannot be inferred from geometry alone. Each panorama is analyzed to identify active construction tasks, estimate the number of workers present in the WorkZone, and assess task progression relative to the plan. A Vision-Language Model (VLM) is employed to synthesize these observations. The resulting contextual interpretation is verified by a human to ensure reliability.

Second, geometric and occupancy measurements are derived from point cloud data to quantify the physical utilization of the WorkZone. Laser scans reconstructed are analyzed to identify temporary obstructions and human presence within the WorkZone. From these data, spatial parameters such as clutter density, average clutter height, human occupancy, and free floor area are computed. Each spatial metric used in the SCP is defined as shown in Table 1.

While the 360-degree imagery and VLM analysis provide semantic understanding of construction execution, the point cloud-derived metrics enrich this understanding with quantitative measures of spatial utilization. The resulting spatial attributes are summarized in tabular form, with each row corresponding to a WorkZone instance and each column representing a derived spatial parameter. This structure enables direct mapping of observed spatial conditions to semantic data properties, where each row corresponds to a WorkZone instance, and each column stores a derived spatial attribute, allowing it to map into and populate OWL data properties for each WorkZone instance within the ontology-based reasoning environment, such as Protégé. The derived spatial metrics are exported as tabular data and automatically mapped into ontology-aligned RDF instances through a CSV-to-RDF conversion process.

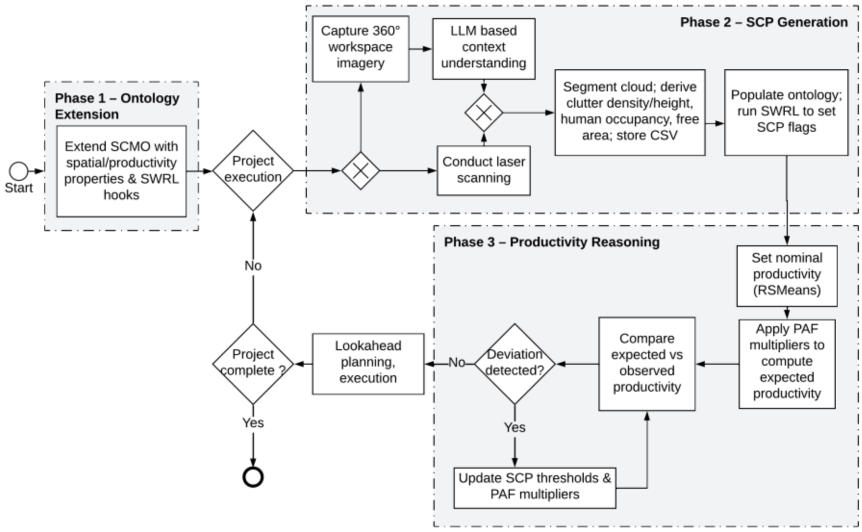

Figure 1. Overview of the proposed hybrid semantic web and data-driven framework for spatial constraint-aware construction planning

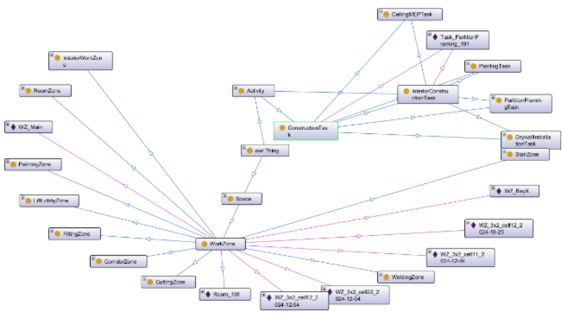

Figure 2. Extension of the Space Constraint Management Ontology (SCMO) to support spatial condition evaluation and productivity reasoning

For SWRL rule design, we draw on prior research that characterizes productivity impacts in terms of qualitative transitions between unconstrained, moderately constrained, and severely constrained WorkZone conditions. Empirical studies on space conflict report productivity reductions of approximately 10–40% under moderate WorkZone restrictions, with losses reaching 60–65% under severe congestion and access obstruction, particularly when working clearances and movement paths are constrained [12]. Grounding on prior research, this study introduces the initial operational discretizations of these qualitative transitions as a starting point and refines them through iterative comparison with project-specific productivity observations.

Several SWRL rules have been developed to transform the numeric observation from data properties into interpretable spatial conditions.

- SCP-1 (Low Clearance): WorkZone(?z) ^ minClearHeight_m(?z, ?h) ^ swrlb:lessThan(?h, 6.5) → hasLowClearance(?z, true)
- SCP-2 (Congested WorkZone): WorkZone(?z) ^ blockageRatio(?z, ?br) ^ swrlb:greaterThan(?br, 0.35)→ hasCongestedWorkspace(?z, true)
- SCP-3 (Human Crowding): WorkZone(?z) ^ humanOccupancy_points(?z, ?hp) ^ swrlb:greaterThan(?hp, 3) → hasHumanCrowding(?z, true)

The SCP for each WorkZone is defined as the set of all condition flags generated through SWRL rules:

- SCP (Zone(i))={LowClearance?, Congested WorkZone?, HumanPresence?,…}

2.3 Phase 3 – Productivity Reasoning Rules

The third phase assesses how the WorkZone's spatial conditions impact the expected productivity of construction tasks. While SCP captures space conditions, the productivity reasoning rules translate these conditions into quantitative adjustments that reflect their operational impact, which move beyond binary feasibility assessments in traditional space constraint management practices.

Table 1. Definition and aggregation of spatial metrics used for SCP inference

METRIC	DESCRIPTION	UNIT	SPATIAL AGGREGATION	TEMPORAL AGGREGATION
MINCLEARHEIGHT_M	Minimum vertical clearance between the floor plane and the lowest overhead obstruction within a grid-based WorkZone	m	Per WorkZone (grid cell)	Per day (single scan)
BLOCKAGERATIO	Proportion of the WorkZone floor area occupied by non-structural clutter	– (0–1)	Per WorkZone (grid cell)	Per day (single scan)
HUMANOCCUPANCY_POINTS	Number of point cloud points classified as human presence within the WorkZone	count	Per WorkZone (grid cell)	Per day (single scan)

To begin with, each construction task is assigned a nominal productivity rate, representing its expected output under ideal WorkZone conditions. The baseline productivity values in this study are informed by RSMeans [13], which provides standardized industry expectations for typical construction task performance. Afterwards, the nominal productivity is adjusted using a set of Productivity Adjustment Factors (PAFs). PAFs are initially specified a priori based on domain knowledge and prior studies, and are subsequently calibrated using empirical productivity observations. Calibration is performed by comparing predicted productivity (from PAF-adjusted nominal rates) with observed productivity proxies derived from daily progress logs. For each SCP, deviations are evaluated, and PAF values are iteratively adjusted when consistent over- or under-estimation is observed. In the current implementation, this process is conducted through structured empirical comparison and manual refinement.

Each PAF represents the expected proportional impact of a specific spatial condition, such as space congestion, on task performance. For example:

- PAF-1 (Low Clearance): Activity(?a) ^ locatedAt(?a, ?z) ^hasLowClearance(?z, true) ^hasCongestedWorkspace(?z, false) ^hasHumanCrowding(?z, false) → productivityReductionFactor(?a, 0.85)
- PAF-2 (Congested WorkZone): Activity(?a) ^ locatedAt(?a, ?z) ^hasCongestedWorkspace(?z, true) ^hasHumanCrowding(?z, false) ^hasLowClearance(?z, false) → productivityReductionFactor(?a, 0.80)
- PAF-3 (Human Presence/Crowding): Activity(?a) ^ locatedAt(?a, ?z) ^hasHumanCrowding(?z, true) ^hasCongestedWorkspace(?z, false) ^hasLowClearance(?z, false) → productivityReductionFactor(?a, 0.90)

Cumulative productivity impacts are modeled using a multiplicative interaction assumption, where each spatial constraint contributes an independent proportional reduction to task productivity. This formulation reflects the compounding nature of concurrent workspace limitations, as multiple constraints (e.g., crowding and restricted clearance) simultaneously affect different aspects of task execution, such as movement efficiency and working posture. Under this assumption, combined effects are computed as the product of individual PAFs (e.g., $0.90 \times 0.85 = 0.765$). For example:

- PAF-4 (Human Crowding + Low Clearance): Activity(?a) ^ locatedAt(?a, ?z) ^ hasHumanCrowding(?z, true) ^ hasLowClearance(?z, true) ^ hasCongestedWorkspace(?z, false) → productivityReductionFactor(?a, 0.765)
- PAF-5 (Congestion + Human Crowding + Low Clearance): Activity(?a) ^ locatedAt(?a, ?z) ^ hasCongestedWorkspace(?z, true) ^ hasHumanCrowding(?z, true) ^ hasLowClearance(?z, true) → productivityReductionFactor(?a, 0.612)

3 Implementation

3.1 Project Background and Data Processing

The proposed framework was implemented and evaluated on an interior fit-out construction project conducted on the New York University Abu Dhabi (NYUAD) campus. The project involved multiple overlapping interior construction activities, including ceiling framing and services installation, partition wall construction, stepped (amphitheater-style) flooring installation, electrical works, and furniture installation. Reality-capture data were collected on a near-daily basis over 51 days using a Leica BLK360 laser scanner and an Insta360 camera, capturing both geometric workspace conditions and execution context.

Phase 2 is implemented through a data-processing pipeline that operationalizes the SCP methodology using daily reality-capture inputs. Each daily laser scan is registered to the BIM using a rigid alignment formulation (Umeyama/Kabsch solution) [14] and clipped to BIM-aligned WorkZone regions to ensure spatial consistency across observations. The aligned point cloud is rendered into a synthetic equirectangular panorama with an associated depth map, enabling correspondence between image pixels and 3D geometry.

Human presence is detected by applying a pretrained semantic segmentation network (LR-ASPP with MobileNetV3-Large backbone [15], trained on the COCO dataset [16] to the synthetic panoramas. The resulting human masks are back-projected into 3D using the depth map and camera model. Human labels are propagated to the full dense point cloud through voxel-based association, producing an explicit 3D representation of human occupancy. The remaining scan points are classified as either structural or non-structural, using the BIM model as a geometric prior. Non-structural points represent temporary clutter and materials.

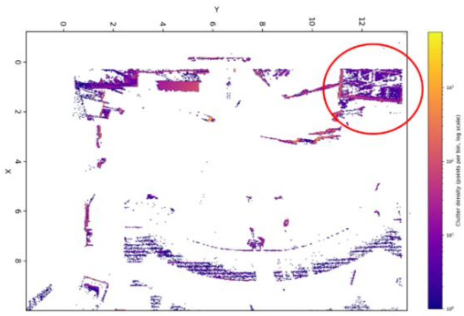

Figure 3. Clutter points density map on the xy plane

Spatial metrics used for SCP generation are computed from non-structural points only, including grid-based measures of clutter density, average clutter height, and human occupancy. The WorkZone footprint is discretized into BIM-aligned grid cells, and spatial attributes are computed per cell and per day. All derived metrics are exported in tabular form, where each row corresponds to a WorkZone instance, and each column represents a spatial attribute. These CSV files serve as the direct interface between the sensing pipeline and ontology-based semantic reasoning. Figure 3 and Figure 4 present example clutter density and height distribution maps that are used for visual validation of spatial measurements. Figure 5 summarizes the data exported into the CSV files, including grid resolution, cell identifiers, and associated spatial metrics. These CSV tables serve as the direct interface between the sensing pipeline and semantic reasoning.

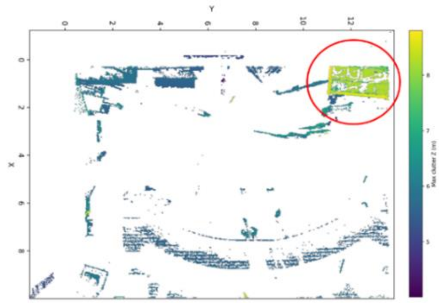

Figure 4. Clutter points height distribution

In parallel, real daily 360° panoramic images are analyzed to extract execution context using a VLM. Visual cues such as workers, tools, materials, and spatial groupings, which are selected based on the project ontology and represent the minimal set of observable entities required to characterize construction execution at the WorkZone level, are combined with zone and temporal metadata and provided to the model to synthesize structured descriptions of worker distribution and active tasks within the WorkZone. A commercially available model [17] was employed within a VLM-based contextual interpretation framework to analyze a dataset of 360 images. Example VLM-generated contextual interpretations are presented in Figure 7 and Table 2. The sample image shown in Figure 6 is provided as input to the VLM, which estimates the set of construction tasks visible in the scene and their corresponding percentage completion (Figure 7). In addition, the model extracts detailed contextual descriptions of ongoing activities, material staging, equipment usage, and worker presence, summarized in Table 2.

To ensure reliability, all VLM outputs were subsequently reviewed and validated by a human annotator prior to ingestion. The verified VLM-derived contextual descriptors were then integrated with geometric spatial metrics and jointly associated with their corresponding WorkZone instances. These combined representations served as input for SCP inference and downstream productivity analysis.

3.2 Ontology-Based Instance Population and Semantic Reasoning Execution

The populated CSV file was subsequently used as the input to the automated CSV-to-RDF script. After generating the RDF instances, the resulting data were imported into the extended SCMO Ontology using Protégé's ontology import mechanism. Semantic reasoning was then performed using SWRL rules executed within Protégé, with Pellet serving as the underlying inference engine.

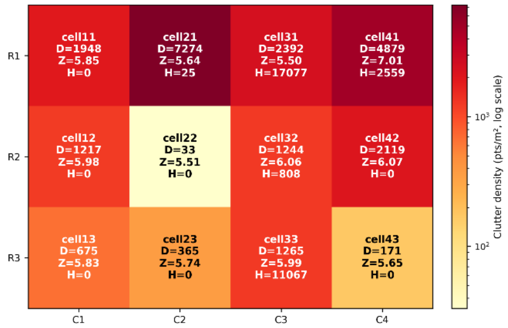

Figure 5. Cell level SCP metrics (assuming 3 x4 grid) D = Point cloud clutter density, Z = average clutter height, H = human occupancy

Figure 6. Sample 360 image.

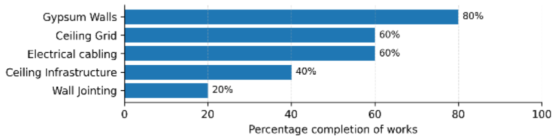

Figure 7. VLM generated percentage completion estimation

Table 2. Context details of the example 360 photo

Category	Extracted Context
Date	October 14 2024
Workers present	5 total workers - 1 kneeling near rear-center on metal stacks; 2 near front doors inspecting; 2 moving materials in the central area
Active tasks	Ceiling framing and services installation (center area, Side 1)
Material storage	Metal deck/corrugated sheets stacked across center toward rear and Side 2; boxes near front wall; boards and ladder along Side 1
Temporary equipment	Temporary site lighting on front and side walls; extension cables; ladders and small platforms along Side 1 and front
Spatial implications	Central circulation partially obstructed; Side 1 constrained by ladders and platforms; moderate movement interference observed

In the first reasoning stage, SCP rules were applied to interpret raw spatial measurements into qualitative WorkZone conditions. Three numeric data properties, humanOccupancy_points, blockageRatio, and minClearHeight_m, were evaluated against predefined thresholds to infer spatial conditions. In the illustrated case in Figure 8, elevated human occupancy values triggered the inference of "hasHumanCrowding = true," while a high blockage ratio resulted in the inference of "hasCongestedWorkspace = true." In parallel, clearance-related rules evaluated vertical space conditions and inferred "hasLowClearance = true" where applicable. These inferred SCPs serve as the semantic basis for downstream productivity reasoning.

After the SCPs were inferred at the WorkZone level, the reasoning process proceeded to activity-level productivity assessment by linking activity instances to their execution space segmentations. Each activity was associated with a corresponding *WorkZone* instance using the *locatedAt* object property to enable the propagation of inferred spatial conditions to the activity context.

Figure 8. Example of SCP inference through SWRL reasoning

To illustrate the reasoning process, Figure 9 presents an example involving the activity instance *Act_Gypsuminstall_2024-10-25*. Prior to reasoning, the baseline productivity value of this activity instance was calculated from RSMeans and represented in the ontology as *nominalProductivity_unitsPerHour = 83.3 S.F./hour*. This value reflects expected productivity under unconstrained spatial conditions and serves as a reference benchmark.

The activity instance Act_Gypsuminstall_2024-10-25 was located in WorkZone WZ_3x2_cell11_2024-10-25, where prior SCP reasoning inferred the simultaneous presence of human crowding and low-clearance conditions (hasHumanCrowding = true and hasLowClearance = true), while material congestion was not detected (hasCongestedWorkspace = false). Given this combination of spatial conditions, a composite productivity adjustment rule was triggered, resulting in the inference of a PAF of 0.765. This reduction factor was subsequently applied to the baseline productivity value, leading to an inferred effectiveProductivity_unitsPerHour (expected productivity) of 63.72 SF/hour.

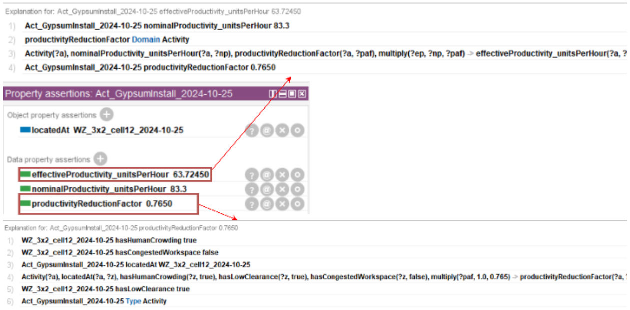

Figure 9. Activity-level productivity reasoning example

Table 3. Comparison of SCP-based productivity predictions with observed progress-derived productivity proxies for the Gypsum_Walls activity

DATE FOR GYPSUM INSTALL	SCP FLAGS (LC/CW/HC)	PREDICTED PAF	WORKERS	OBSERVED PROXY (SF/WORKER-DAY)	OBSERVED INDEX	ERROR (PAF − INDEX)
16/10/2024	T / F / T	0.765	3	53.82	1	−0.235
18/10/2024	T / F / T	0.765	4	40.36	0.75	0.015
21/10/2024	F / F / T	0.9	6	26.91	0.5	0.4
25/10/2024	T / F / T	0.765	4	40.36	0.75	0.015

To empirically ground the inferred productivity adjustment, beyond a single illustrative instance, activity progress logs for the Gypsum_Walls task were examined across multiple observation days. Table 3 summarizes the inferred SCP flags, predicted productivity adjustment factors, and observed productivity indices from October 16, 2024, to October 25, 2024. For example, the daily progress log for Gypsum_Walls on 2024-10-25 progressed from 95% to 100% completion with a crew of four workers. Given the known wall geometry (four walls, each 5 m × 15 m), this 5% progress increment corresponds to approximately 161.45 SF of completed work for the day. Dividing this quantity by the crew size yields an observed productivity proxy of approximately 40.36 SF per worker-day. However, WorkZone snapshot records indicate that multiple spatial events and task interactions were present on the same observation day, demonstrating that gypsum wall installation was executed alongside other concurrent activities rather than as a continuous, uninterrupted process. As a result, only a portion of the available labor time contributed directly to measurable gypsum-wall progress. This observed productivity proxy is therefore not interpreted as a direct measure of continuous installation rate, but rather as a relative indicator of net progress achieved within the reporting period. For comparative purposes, the proxy was normalized relative to the reference day (2024-10-16) to derive a dimensionless observed productivity index.

While the initial productivity reduction values setting in the ontology show a discrepancy (maximum 0.4), the proposed framework establishes a structured and traceable mechanism through which empirical observations can be systematically fed back into the ontology to iteratively update rule thresholds and adjustment factors. This provides a reusable semantic backbone that enables continuous learning and refinement of space-constraint-related productivity knowledge over time.

4 Conclusion

This paper presented a hybrid semantic web and data-driven framework for spatial constraint-aware construction planning that moves beyond binary space readiness assessments. This paper contributes to the development of production planning and control tools, such as LPS, in regard to space constraint-related issues by providing a transparent and explainable decision support framework for planners operating under spatially constrained conditions.

However, several limitations of the current implementation point to important directions for future research. First, the productivity adjustment factors were stochastically data-driven, rather than through a fully automated calibration mechanism. Future work should integrate statistical learning or optimization techniques to continuously update PAF values and SCP thresholds from accumulating project data. Second, the current implementation relies on Protégé as a static reasoning environment. Although automation scripts were introduced, the workflow remains semi-manual. Future development will migrate the ontology to a graph database to enable continuous data ingestion, automated reasoning execution, and real-time feedback. Third, the current rule set models spatial conditions primarily as productivity-reducing constraints and assumes multiplicative interactions when estimating their combined effects. In practice, however, some spatial configurations (e.g., temporary installations or support structures) may simultaneously hinder certain activities while enabling others, and interactions between constraints may not be fully independent. These task-dependent and potentially interdependent effects are not yet operationalized in the present implementation and warrant further investigation. Future research will extend the ontology and reasoning rules to encode task–space interaction semantics explicitly. Fourth, the proposed workflow may become computationally demanding when applied to large-scale point cloud datasets and complex construction scenarios involving multiple concurrent activities. Future research should therefore examine its computational efficiency relative to comparable workflows and explore alternative technological pathways that reduce the burden of data collection. For example, lighter multimodal approaches that leverage image-based observations and NLP/VLM-driven scene interpretation to encode spatial context into semantic representations, further enhanced through ontology-based reasoning, may offer a less data-intensive alternative while preserving planning-relevant diagnostic capability.

5 Acknowledgement

This research was partially supported by SHORES and CITIES (Centers funded by Tamkeen under NYUAD Research Institute Awards CG013 and CG001, respectively) and the NYUAD Environmental Health and Safety Department. The authors also acknowledge support from the Alliance Research Grant (ALLRP 572086), Federated Platform for Construction Simulation, from the Natural Sciences and Engineering Research Council of Canada (NSERC). The authors also thank William Fulton and Mohamed Rasheed (NYUAD Campus Planning and Projects Office) for facilitating access to campus projects.

References

[1] Z. Mao, R. Ahmad, F. Hamzeh, and V. A. Gonzalez, "Development of a Space Constraint Management Ontology for Lean Construction: Enhancing Constraint Management in Steel Fabrication Projects," in Proceedings of the 33rd Annual Conference of the International Group for Lean Construction (IGLC 33), O. Seppänen, L. Koskela, and M. Koichi, Eds., Osaka and Kyoto, Japan, 2025, pp. 1035–1045. doi: 10.24928/2025/0196.

[2] A. Burcu, F. Martin, and K. John, "Automated Generation of Work Spaces Required by Construction Activities," J. Constr. Eng. Manag., vol. 128, no. 4, pp. 306–315, Aug. 2002. doi: 10.1061/(ASCE)0733-9364(2002)128:4(306).

[3] S.-J. Guo, "Identification and Resolution of Work Space Conflicts in Building Construction," J. Constr. Eng. Manag., vol. 128, no. 4, pp. 287–295, 2002. doi: 10.1061/(ASCE)0733-9364(2002)128:4(287).

[4] F. Bosché and E. Guenet, "Automating surface flatness control using terrestrial laser scanning and building information models," Autom. Constr., vol. 44, pp. 212–226, 2014. doi: 10.1016/j.autcon.2014.03.028.

[5] S. Zhang, F. Boukamp, and J. Teizer, "Ontology-based semantic modeling of construction safety knowledge: Towards automated safety planning for job hazard analysis (JHA)," Autom. Constr., vol. 52, pp. 29–41, 2015. doi: 10.1016/j.autcon.2015.02.005.

[6] K. Farghaly, R. K. Soman, W. Collinge, M. H. Mosleh, P. Manu, and C. M. Cheung, "Construction safety ontology development and alignment with industry foundation classes (IFC)," ITcon, vol. 27, no. 5, pp. 94–108, Jan. 2022. doi: 10.36680/J.ITCON.2022.005.

[7] W. T. Chen and T. A. Bria, "A Review of Ontology-Based Safety Management in Construction," Sustainability, vol. 15, no. 1, 2023. doi: 10.3390/su15010413.

[8] A. Schlachter, M. H. Rasmussen, and J. Karlshøj, "Using Linked Building Data for managing temporary construction items," Autom. Constr., vol. 139, p. 104258, 2022. doi: 10.1016/j.autcon.2022.104258.

[9] V. Getuli et al., "From BIM to digital twin: IoT data integration in asset management platform," ITcon, vol. 29, no. 49, pp. 1103–1127, Dec. 2024. doi: 10.36680/J.ITCON.2024.049.

[10] H. G. Ballard, The Last Planner System of Production Control. Birmingham, UK: The University of Birmingham, 2000.

[11] E. T. Mengiste and B. Garcia de Soto, "Harmonizing Construction Flow: Integrating Takt Planning, Workspace Management, and Lookahead Strategies," in Proceedings of the 33rd Annual Conference of the International Group for Lean Construction (IGLC 33), O. Seppänen, L. Koskela, and M. Koichi, Eds., Osaka and Kyoto, Japan, 2025, pp. 856–868. doi: 10.24928/2025/0226.

[12] M. S. Mhd Azmin, H. Mohd Hanizun, and A. Shardy, "Main Factors Lack of Workspace Planning That Causes Workspace Conflict on Project Environment: Industrialised Building System In Malaysia," Aust. J. Basic Appl. Sci., vol. 7, no. 6, pp. 408–419, 2013.

[13] RSMeans, "RSMeans online cost and productivity data." Gordian.

[14] S. Umeyama, "Least-squares estimation of transformation parameters between two point patterns," IEEE Trans. Pattern Anal. Mach. Intell., vol. 13, no. 4, pp. 376–380, 1991. doi: 10.1109/34.88573.

[15] A. Howard et al., "Searching for MobileNetV3," in 2019 IEEE/CVF International Conference on Computer Vision (ICCV), 2019, pp. 1314–1324. doi: 10.1109/ICCV.2019.00140.

[16] T.-Y. Lin et al., "Microsoft COCO: Common Objects in Context," in Computer Vision – ECCV 2014, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds., Cham: Springer International Publishing, 2014, pp. 740–755.

[17] OpenAI, "ChatGPT." [Online]. Available: https://chatgpt.com/. Accessed: Jan. 18, 2026.