

OPERATIONALIZING LEAN PROJECT CONTROL THROUGH DATA-DRIVEN WEEKLY REASONING

Eyob T. Mengiste¹ & Borja Garcia de Soto²

ABSTRACT

Lean construction emphasizes flow, pull-based planning, and continuous risk management; however, its practical implementation still is very labor-intense and relies heavily on periodic meetings, manual interpretation of fragmented data, and lagging indicators of project performance. This paper presents an automated operationalization framework for Lean project management decision-making. The framework fuses standard project documentations such as periodic site observations, structured textual reports, and approvals into a continuously updated temporal representation of the project state. This fusion of information is used to evaluate constraints and estimate the execution readiness level of tasks. Moreover, this framework helps project managers detect workflow disruptions and constraint-related risks and generate Lean-consistent execution alternatives. The paper provides the system architecture, data integration logic, and reasoning mechanisms, and validates the approach through an automated deployment over 16 consecutive weekly planning cycles of a real construction project, demonstrating its feasibility as an operational tool for Lean project control.

KEYWORDS

MCTS, multi-armed bandits strategy, constraint management, LPS, lookahead.

INTRODUCTION

Construction projects are complex production systems characterized by spatial constraints, temporal dependencies, and high levels of uncertainty. Lean construction theory reframes project delivery as a problem of production flow and control, rather than a mere aggregation of contractual tasks (Koskela 1992). Central to Lean thinking is the idea that reliable production emerges not from detailed long-term schedules, but from continuous management of constraints, variability, and commitments at the operational level. The Last Planner System (LPS) operationalizes this perspective by emphasizing make-ready planning, pull-based task release, and collaborative decision-making (Ballard 2000; Ballard and Howell 1998). Despite its demonstrated benefits, Lean project management in practice still relies heavily on decentralized, manually processed data, which is subject to bias and delays. Moreover, the project performance indicators, such as Percent Plan Complete (PPC), do not reflect active constraints or the level of risk they pose to the construction project. Lookahead plans and constraint logs are typically updated on a weekly basis and depend on human judgment to synthesize information from site observations, reports, schedules, and models. As projects

¹ Instructor, Civil and Urban Engineering Program, Researcher, S.M.A.R.T Construction Research Group, New York University Abu Dhabi, Abu Dhabi, UAE, eyob.mengiste@nyu.edu, orcid.org/0000-0002-4841-7476

² Associate Professor, Civil and Urban Engineering Program, Director, S.M.A.R.T Construction Research Group, New York University Abu Dhabi, Abu Dhabi, UAE, garcia.de.soto@nyu.edu, orcid.org/0000-0002-9613-8105

increase in size and complexity, this manual synthesis becomes cognitively demanding and increasingly prone to delayed or incomplete situational awareness.

Existing efforts to strengthen Lean project control have largely concentrated on organizing planning and execution information to support coordination, reporting, and managerial oversight. Frameworks for Lean-oriented project management articulate how commitments, plans, and execution records can be structured to improve control during construction, often through formalized management processes and information flows that emphasize transparency and compliance (Dumrak and Zarghami 2023). In parallel, automation initiatives within planning and production management have typically focused on discrete functions, such as scheduling support, progress measurement, or workflow monitoring, rather than on the decision logic that governs task release and execution (Kolberg et al. 2017; Singh and Kumar 2020). Systematic analyses of Lean planning automation further indicate that most implementations emphasize stage-specific support, particularly for lookahead planning and performance tracking, while relying on predefined workflows and fragmented information inputs (Agrawal et al. 2024). Although these approaches improve the availability and structure of project information, they leave a critical aspect of Lean control under-specified: the reasoning process by which information is transformed into actionable production decisions. In practice, assessments of whether work can proceed, considering unresolved dependencies, approvals, and resource conditions, remain implicit and depend on practitioner judgment. As a result, Lean control continues to rely on lagging indicators and irregular interpretation rather than systematic, repeatable evaluation of production feasibility. This gap points to the need for an explicit operational layer that formalizes how routinely available project artifacts are interpreted to evaluate constraints, assess execution readiness, and support near-term planning decisions in a continuous manner.

Moreover, commonly used project performance indicators such as PPC are valuable for assessing planning reliability after execution, but they do not explicitly represent the type, severity, and timing of active constraints at the activity level, nor do they express the level of execution risk that these constraints create for upcoming work (Ballard 2000; Jang and Kim 2007; Masoetsa et al. 2022). As a result, existing planning indicators are useful for retrospective control, yet they offer limited support for judging whether specific tasks are sufficiently ready for short-horizon commitment under evolving site conditions. This creates an important gap between reporting planning performance and reasoning about execution readiness. In response, the present study introduces the Constraint Readiness Index (CRI) as a forward-looking and activity-level measure that translates fragmented project evidence into an explicit readiness score. The added value of the CRI lies not only in making constraint status more visible, but also in enabling traceable make-ready screening, comparison of alternative weekly plans, and learning across successive planning cycles.

This paper proposes an operational framework of generating a decision support mechanism where routinely produced project information is consolidated and interpreted during construction. The framework specifies a reasoning process for evaluating constraints and estimating task execution readiness based on a continuously updated representation of the project state. Using these evaluations, the framework supports the generation of Lean-consistent execution alternatives for near-term planning. The case implementation reported later in the paper demonstrates that the framework can be executed on routinely produced project records to generate auditable readiness scores and alternative weekly planning policies, thereby providing an operational rather than purely conceptual contribution.

LITERATURE REVIEW

Prior research has explored multiple avenues for operationalizing Lean construction principles through formalized planning routines, structured information flows, and automation of planning

and control activities. Several studies emphasize improving transparency and coordination by organizing planning information into structured management representations, often through visual management or standardized reporting mechanisms, to enhance planning consistency and managerial oversight (Brady et al. 2018). Researchers have also investigated the interaction between Lean practices and BIM-enabled environments (Heigermoser et al. 2019), focusing on how digital representations can support coordination and information exchange across planning and production activities (Bayhan et al. 2023).

Moreover, research on workflow automation (Mengiste et al. 2025) and production control, often inspired by manufacturing and cyber-physical systems, has introduced automation concepts applicable to construction planning contexts (Kolberg et al. 2017; Sobieraj and Metelski 2022). These studies propose a method of automating specific functions such as sequencing, monitoring, or performance estimation, rather than the planning logic governing commitment-making. A systematic mapping of automation within the LPS confirms this tendency, showing that automation efforts are concentrated in individual LPS stages (Mengiste et al. 2023; Mengiste and Garcia de Soto 2025), particularly lookahead planning and performance tracking, with limited integration across planning cycles and data sources (Agrawal et al. 2024). Within most of the literature, Lean planning is supported through structured processes or isolated automation components, but the reasoning process underlying planning decisions is not explicitly represented. The basis for classifying activities as executable, deferred, or excluded remains implicit, limiting traceability, repeatability, and systematic learning across planning cycles.

Within the Lean project control, constraint management is widely recognized as central to production performance and planning reliability. Early system dynamics and constraint-centric perspectives establish that system behavior is governed by bottlenecks and limiting conditions rather than by isolated task efficiencies, motivating explicit attention to constraints in planning and control (Gingras and Schweber 1986). This perspective is reflected in make-ready planning, where constraint removal is treated as a prerequisite for reliable task execution. One of the earliest quantitative approaches to constraint management in Lean planning is the Percent of Constraint Removal (PCR) metric, which measures the proportion of identified constraints resolved during the make-ready process (Jang and Kim 2007). Subsequent studies examine constraints as contributors to workflow disruption, schedule instability, and planning quality through empirical analysis and analytical modelling (Dallasega et al. 2018; Masoetsa et al. 2022; Wu et al. 2023). Additional research adopts multi-model decision-support and engineering control perspectives, representing constraints within analytical or simulation-based frameworks to evaluate system performance under uncertainty (Scherer and Schapke 2011).

These approaches contribute formal representations of constraint effects at the system or project level. While these studies establish that constraints are measurable and influential, constraint evaluation is typically performed at an aggregate or analytical level. Constraints are rarely localized to individual activities in a consistent manner, their persistence over time is not systematically tracked, and their role in determining near-term executability is often indirect. Consequently, constraint metrics and models are weakly connected to day-to-day planning decisions and commitment formation.

METHODOLOGY

The proposed approach follows a weekly lookahead reasoning workflow illustrated in Figure 1. Each planning week is treated as a self-contained reasoning cycle. At the beginning of the week, all available project information (e.g., project reports, approvals, submittals, procurement logs, schedule updates, and financial records) is consolidated into a unified state representation. Based on this state, feasible execution plans are generated and evaluated. Execution outcomes observed in the subsequent week are used exclusively for evaluation and learning and do not

retroactively influence planning decisions. This cycle is repeated for each reporting week, with outputs from one cycle informing subsequent cycles.

DATA PREPROCESSING

The workflow input includes a baseline CPM schedule, weekly project reports, and supporting project records, such as approvals, submittals, procurement logs, schedule updates, and financial records. The baseline schedule establishes the execution logic of the project and is represented as an Activity-on-Node (AON) network (Neumann and Schwindt 1997), where nodes correspond to activities and edges represent task dependency relationships.

Each incoming weekly report is parsed to extract structured information, including progress updates, planned work for the upcoming period, constraints, risk and safety notes, and financial remarks. All extracted items are tagged with the reporting period. Items that directly reference activities are mapped to activities in the baseline schedule. Because report entries may span multiple activities, the mapping step allows one-to-many associations when appropriate, and the resulting evidence is assigned to activity nodes in the AON network. Constraints are extracted from the reports, but they are often stated in free-text form often without explicit links to activities or work packages. To attach constraints to the schedule representation, a localization step is applied during preprocessing to associate each constraint statement with the affected AON activities. In this study, localization relies on fuzzy matching and bounded heuristic rules, allowing a constraint to be linked to one or more activities when appropriate. Localized constraints are then categorized by type (e.g., approval-related, material-related, interface-related, or scope-related) and stored as per-activity constraint sets for downstream readiness assessment.

To construct the weekly project state, report statements were first transformed into structured fields and then matched to baseline activities through a two-stage fuzzy localization procedure. Candidate activities were shortlisted using sequence similarity between report phrases and activity names, after which token overlap and domain-specific keywords related to approvals, materials, areas, and trades were used to refine the match score. When multiple activities appeared to yield a high probability of matching, the highest-scoring candidates were retained as reviewable evidence to keep the localization process auditable.

Although the workflow automates recurring parsing, matching, scoring, and plan generation steps, the current implementation should be understood as semi-automated decision support. This is due to the need for human intervention at points where the report language is ambiguous, documentation is incomplete, or multiple activity matches remain plausible.

The goal of the extraction and localization steps is to establish consistent activity-level links from routinely produced reports; advancing report-parsing methods is outside the scope of this study.

CONSTRAINT ANALYSIS AND READINESS ASSESSMENT

Readiness is quantified using a CRI, which indicates the extent to which an activity can be considered make-ready for execution in the upcoming planning period. CRI is computed using a predefined rule set and is bounded between 0 (not ready) and 1 (fully ready). For each activity and reporting period, any reported constraint linked to that activity reduces its CRI, such that readiness decreases as unresolved constraints accumulate. Based on the resulting score, activities are categorized as ready, near-ready, at-risk, or not ready. CRI is used solely as an eligibility signal for make-ready screening and planning; it is not intended as a measure of productivity, performance, or optimization. CRI is embedded within a recurring weekly planning cycle. Each week, candidate execution plans are formed from activities that meet the CRI eligibility threshold and are then evaluated against subsequent execution evidence. The observed execution outcomes provide feedback on the adequacy of the current CRI rule set,

and the readiness scoring rules are iteratively adjusted over successive cycles to better reflect project-specific constraint behavior and reporting patterns, while remaining transparent and interpretable. In parallel with readiness assessment, activity-level risk indicators are constructed. Risk metrics are evaluated only for activities that satisfy the CRI threshold and are used to compare feasible execution plan alternatives during the planning stage rather than to determine eligibility.




Figure 1: Overview of key components and process of the research methodology.

The CRI is defined in this paper as a bounded readiness indicator for short-horizon construction planning. It should not be confused with the coefficient of variation (CV). It summarizes the assumed executability of an activity during weekly planning. Its calibration is case-oriented and iterative: initial values are assigned based on expected readiness and then refined against observed execution outcomes. Accordingly, the thresholds adopted in this study should be understood as case-calibrated planning settings rather than universal industry constants.

PLANNING

The execution planning is formulated as a short-horizon selection problem, with the objective of assembling a weekly plan from activities that satisfy the CRI readiness threshold. The eligible set is defined from the baseline schedule and updated readiness status: only activities whose required predecessors are complete and whose remaining durations can be accommodated within the weekly capacity are considered. Using this eligible set, the method generates alternative plans through a Monte Carlo Tree Search (MCTS) procedure (Zhao et al. 2025) that runs on the AON network. MCTS incrementally constructs candidate weekly plans by adding eligible activities node by node, one at a time; each AON node represents a partial weekly plan, and each expansion step appends another activity. The search is bounded to a one-week planning horizon. Risk indicators were included in candidate-plan selection because workable weekly commitments should reflect not only whether activities are nominally available, but also the consequences of committing them under unresolved exposure. The ranking formulation therefore discouraged alternatives that concentrated excessive schedule, cost, quality, or safety risk into the weekly work set. Time risk was linked to remaining duration relative to available float, cost risk reflected unresolved commercial exposure in relation to remaining work, quality risk captured interface and inspection-related uncertainty, and safety risk represented reported hazards and congestion pressure.

Three planning styles are used to generate alternatives: conservative, balanced, and aggressive. They differ in the minimum readiness level accepted and the intended weekly plan size. Conservative favors higher-confidence activities and smaller work packages, balanced allows moderate exposure and size, and aggressive admits lower readiness levels and larger work packages. MCTS generates alternatives consistent with these styles by adjusting selection preference and stopping criteria

After MCTS generates a pool of candidate weekly plans under each planning style, the alternatives are ranked using the risk indicators computed for eligible activities. Risk is not used to exclude activities from consideration; it is used only to distinguish between feasible plans. For each candidate plan, activity-level risk values are aggregated to a plan-level score, and plans with lower overall risk exposure under each planning style are retained as the representative (conservative, balanced, and aggressive) execution plans for subsequent evaluation.

EXECUTION AND FEEDBACK LOOP

Each alternative plan generated in a given week is evaluated using the next weekly execution report, which records what was actually completed and what issues occurred. The evaluation considers realized progress on planned activities, the persistence or resolution of constraints, safety incidents, and cost/scope signals. Progress credit is hybrid, meaning that executed activities receive observed progress, while planned-but-not-executed activities contribute a counterfactual progress term. Evaluation results are combined into a scalar reward reflecting managerial priorities. Selection among planning styles (i.e., conservative, balanced, aggressive) is modeled as a contextual multi-armed bandit (Lattimore and Szepesvári 2020), learning only at the policy level without changing activity eligibility, schedule logic, or execution rules. Each cycle outputs the updated project state, readiness and risk indicators, ranked plan alternatives, and execution-based evaluation signals, which are carried forward to the next cycle.

The same weekly execution report used for plan evaluation also provides feedback for calibrating the CRI scoring rules. Specifically, the framework compares the predicted readiness of planned activities with their observed execution outcomes in the subsequent report and uses systematic discrepancies to adjust the CRI penalty weights and thresholds over successive cycles. Updates are applied conservatively (with bounded step sizes) to improve consistency between readiness scores and observed executability, while preserving interpretability.

IMPLEMENTATION

The proposed framework was implemented on a building fit-out project at the University campus. The project involved multiple trades working concurrently across shared interior spaces, requiring frequent coordination of access, approvals, and sequencing. The study period covers seventeen consecutive weekly reporting cycles, corresponding to the project's routine planning pattern. The framework used routinely produced project documentation, including a baseline CPM schedule and a sequence of weekly project reports. Supplementary records were used where available, including approval and submittal logs, material lead-time information, and cost/scope tracking notes (e.g., change orders and variations).

Data extracted from the raw document inputs shown in Table 1 were used to compute the CRI and to derive activity readiness. The initial CRI rule set applied in this study is summarized in Table 2. Reported constraints identified from the weekly reports and supporting records are localized to the relevant schedule activities and maintained as per-activity constraint sets. CRI is initialized at 1.0 for each activity and reduced by the applicable penalty terms associated with the constraints linked to that activity, as specified in Table 2. The penalty scheme is intentionally simple and transparent and is treated as an initial configuration for illustration purposes. As execution evidence accumulates across reporting cycles, the penalty magnitudes

can be periodically recalibrated based on observed mismatches between readiness classification and subsequent execution outcomes.

Table 1: Summary of data extraction from all of the data used in this experiment

Activities with complete information, including durations, dependencies, floats, and planned dates	208 unique activities
Task update information extracted from weekly reports	All the 17-week reports, 209 items (196 numeric)
Information on Manhours on-site and Manpower deployed	All the 17-week reports contain combined per-week values
Specific mentions of safety violations	All the 17-week reports mention nine unique instances.
Authority approval submittal counts	All the 17-week reports mention 17 approval requests.
Scope/variations/VO/CR items (pending count/value/age)	72 pending approvals are mentioned in the 13 reports
Extracted constraints list	57 constraints are extracted from all 17-week reports
PO value, invoices (amounts)	16 reports mention valued invoices

Table 2: Initial CRI calculation rules

Rule group	Rule definition
Hard blocks (CRI = 0; label “Blocking”)	CRI is set to 0 if any of the following hold: - Any predecessor with FS dependency is not “complete” (progress < 90%); - Any linked constraint is marked blocking in the report. - Critical activity is stalled (float ≤ 0 and stagnation ≥ 2 weeks); - Critical approval is overdue (approval-type, float ≤ 0, and overdue ≥ 2 weeks); - Linked cost/scope request is on hold.
Non-approval activities	- Capacity pressure (site overcrowding): compute manhours per active task = weekly manhours/ (active tasks + 1). If <20 subtract 0.15; if <40 subtract 0.10; if <80 subtract 0.05; else 0. - Lateness penalty (float-sensitive): if delayed ≥ 1 week, subtract more when float is smaller (critical tasks penalized more). - Stagnation penalty: subtract 0.10 (1 week), 0.20 (2 weeks), 0.25 (≥ 3 weeks). - Material delay penalty: subtract 0.20 (1 week), 0.30 (2 weeks), 0.35 (≥ 3 weeks).
Approval/ milestone activities	Approval overdue penalty: subtract 0.20 (1 week), 0.30 (2 weeks), 0.35 (≥ 3 weeks).
Commercial/scope (pending only)	If a linked cost/scope request is pending, apply a value-band penalty with an age cap (age in weeks capped at 4): <20k: min (0.12, 0.03×age); 20k–100k: min (0.20, 0.05×age); >100k: min (0.20, 0.08×age).

The execution plan was generated using MCTS with a one-week planning window. Candidate plans were constructed from activities that satisfied the CRI readiness threshold and had all required predecessors completed in the AON network. The three planning styles, conservative, balanced, and aggressive, were implemented using the maximum number of activities per weekly plan, which was set to 6, 8, and 10 for the conservative, balanced, and aggressive styles, respectively. Moreover, only tasks labeled as Ready are considered in the conservative planning; balanced planning considers both Ready and near-ready; and aggressive planning

uses all activities with CRI values greater than or equal to 0.4, excluding only tasks labeled as not-ready. Each plan generated during the MCTS iterations is evaluated against the risk indicators and ranked, with the lowest-risk indicator as the best plan per style.

After executing the selected weekly plan, the subsequent weekly report is used to compute the realized reward for the chosen planning style. The contextual multi-armed bandit then updates its belief about each style's expected performance and selects the next style by trading off exploitation of high-performing styles against controlled exploration to avoid premature lock-in. Because learning operates only at the style level, the underlying activity eligibility criteria, precedence/schedule logic, and task-selection rules remain unchanged and auditable. This yields a lightweight adaptation mechanism that tracks changing project conditions while preserving governance and interpretability.

RESULTS AND DISCUSSIONS

Beyond demonstrating that the workflow can be executed on routine project records, the sample implementation highlights how readiness reasoning changes the interpretation of weekly work selection. Activities with reported progress or apparent field activity were not automatically treated as executable; rather, predecessor status, unresolved constraints, and downstream exposure materially affected their final readiness classification. Similarly, the comparison of generated planning policies shows that different commitment styles produce distinct reward trajectories, suggesting that short-horizon planning performance depends not only on what work is available, but also on how risk and readiness are balanced during selection. These results, therefore, indicate that the proposed framework contributes explanatory value in addition to automation.

Implementing the framework on the case project produced weekly CRI values for each task that did not achieve 100% completion. Figure 2 presents weekly boxplots of CRI values for constraint-linked, dispatch-ready activities across weeks 2024-W27 to 2024-W42. Because CRI is bounded between 0 and 1, the weekly distributions show the co-existence of ready, near-ready, and at-risk activities in the candidate pool. Across nearly all weeks, upper quartiles extend into the ready range, while medians typically remain in the near-ready band and lower tails fall into at-risk or not-ready ranges. This indicates that weekly planning consistently involves screening a heterogeneous set of activities with different readiness levels, reinforcing CRI's role as an eligibility signal for make-ready prioritization rather than a performance metric.




Figure 2: CRI values of tasks that were evaluated from project reports across the implementation timeline.

Prior work has advanced digital support for specific functions such as BIM-enabled execution support and automated look-ahead generation, but these capabilities are typically developed as partial interventions rather than as an integrated weekly reasoning loop for commitment

formation ((Agrawal et al. 2024; Schimanski et al. 2020; Soman and Molina-Solana 2022) and ((Soman and Molina-Solana 2022)). The present study contributes by combining project-state construction, readiness scoring, and alternative plan generation within a recurrent weekly cycle. The case results should therefore be interpreted not as proof of universal performance superiority, but as evidence that Lean make-ready logic can be operationalized in a traceable computational form suitable for refinement and future benchmarking.

Accordingly, the present case study should be interpreted as a feasibility and operationalization study rather than a definitive performance benchmark. Although the generated plans were compared with observed project outcomes and baseline planning practice, broader validation against additional projects and simpler benchmark strategies remains necessary and a part of the future work.

The activity-level traceability of the project state was assessed by checking whether readiness labels and constraint signals could be linked back to specific weekly report evidence. Table 3 summarizes this traceability for a representative reporting week. A key observation from Table 3 is that activities can be blocked due to incomplete predecessors and/or active localized constraints, even when progress is reported. For example, the gypsum partition submission is not-ready (CRI=0.30) due to an active long-lead material constraint despite 98% reported progress.

Table 3: CRI calculation back trackability sample activity set for week 38

Activity name	Predecessors complete?	Reported progress	Active?	Localized constraints	CRI (category)
Gypsum Partition – Material Submission	Y	98%	Y	Long-lead material (active)	0.30 (Not ready)
De-Mobilization	Y	93%	Y	–	0.85 (Ready)
Proposed Floor Plan – Shop Drawing Submission	N	0%	N	Approval (active)	0.00 (Blocking)
Access Control System – Material Submissions	N	0%	N	Long-lead material (active); Interface/coordination (active)	0.00 (Blocking)

Table 4 presents representative weekly execution plans for the conservative, balanced, and aggressive styles. Because each style applies a different readiness threshold, the resulting plan sizes differ across alternatives.

Table 4: Sample sets of execution plans for week 38

Plan	Wooden Wall Cladding works	Current completion	Remaining (days)	Float (days)	Predecessors done?	CRI	CRI class	Planned completion increase (%)
C	Demobilization of existing office furniture and demolition of materials	0%	7	7	Y	0.95	Ready	21.5
C	Mobilization Works Temporary working stations	93%	0.1	2	Y	0.85	Ready	7
C	Mobilization Works of Management and Staff (Organization Chart)	62%	5.3	14	Y	0.85	Ready	10.7
B	Demobilization of existing office furniture and demolition of materials	93%	0.1	2	Y	0.85	Ready	7
B	O & M Manual Submission	0%	5	5	Y	0.75	Near-ready	15
B	Warranty Submission	0%	5	5	Y	0.75	Near-ready	15
B	Supply and Installation of Fridge	0%	1	1	Y	0.7	Near-ready	75.2
B	Bottle filling station - Furniture Installations	0%	7	2	Y	0.7	Near-ready	10.7
B	Loose Furniture (CH-03, PD-01, PD-02 & Lockers) - Furniture Installations	0%	5	5	Y	0.7	Near-ready	15
B	Snagging and De-snagging	0%	5	5	Y	0.7	Near-ready	15
A	O & M Manual Submission	0%	5	5	Y	0.75	Near-ready	12
A	Warranty Submission	0%	5	5	Y	0.75	Near-ready	12
A	Supply and Installation of Fridge	0%	1	1	Y	0.7	Near-ready	60.1
A	MEP Work - Completion	0%	175	175	Y	0.7	Near-ready	0.3
A	Mobilization of Management and Staff (Organization Chart)	0%	7	7	Y	0.6	Near-ready	8.6
A	Site Taken Over Period	0%	23	23	Y	0.6	Near-ready	2.6
A	Contract Signed (Sub Contracts)	0%	1	1	Y	0.5	At-risk	60.1
A	Production site Taken over by Subcontractors	0%	1	1	Y	0.5	At-risk	60.1
A	Dilapidation Report	0%	2	2	Y	0.5	At-risk	30.1

Error! Reference source not found. compares the week-by-week reward outcomes produced by the multi-armed bandit evaluation of four planning policies, that is, three generated alternatives (conservative, balanced, and aggressive) and the baseline corresponding to the actually executed plan (PM-plan). The weekly reward is computed as the weighted sum of

actual task completions for activities committed in the execution plan, minus a penalty proportional to the remaining (unachieved) portion of those planned activities. This rewards delivery while discouraging systematic over-commitment. The x-axis spans the reporting weeks (W27-W42), while the y-axis reports the weekly reward signal assigned to each policy based on how well its proposed task set aligns with subsequent execution outcomes (higher reward for strong realized performance and a penalty for weak performance). Across the timeline, the aggressive policy (orange) attains the highest rewards in most weeks, indicating that committing to a larger volume of tasks can yield superior realized payoff when execution conditions are favourable; however, it also exhibits pronounced volatility, with sharp drops and rebounds that reflect greater sensitivity to constraint-driven disruption when committed work becomes blocked or delayed.



Figure 3: Weekly bandit reward trajectories for generated planning policies (conservative, balanced, aggressive) and the observed baseline (PM-plan)

In contrast, the balanced policy (green) follows a comparatively smoother and flatter trajectory, suggesting more consistent performance and a moderated exposure to near-ready work, trading peak reward potential for stability. The conservative policy (blue) generally yields lower rewards with occasional spikes, consistent with a cautious commitment strategy that limits exposure but can underperform when more work could have been reliably executed. Finally, the PM-plan (purple) closely follows the conservative trajectory across most weeks, indicating that the prevailing management practice resembles a conservative make-ready posture and that the conservative-generated alternatives are broadly aligned with observed planning behavior.

CONCLUSION

This paper presented an operational framework that translates conventional weekly project artefacts into a structured, auditable project state that supports core Lean lookahead functions: make-ready screening, explicit constraint reasoning, and evidence-based commitment formation. The case evaluation over 16 weekly cycles shows that readiness can be computed consistently from routine reporting, that blocking conditions can be traced to documented constraint and precedence evidence, and that weekly execution plans can be generated and compared using the same evidence-linked project state, which turns what is typically an informal reasoning process into a repeatable and inspectable workflow.

A key limitation is that the case project follows conventional reporting practices, and the weekly report framing is not optimized for machine interpretation; improving the reporting structure and standardizing constraint statements would likely strengthen mapping and readiness estimation. In addition, reliance on heuristic and fuzzy matching for constraint localization may miss or weakly link signals that are expressed implicitly or outside the weekly reports, which may partly explain why the dominant detected constraints concentrate on explicitly recorded issues such as approvals, predecessor completion, and material lead times.

Future work will integrate AI-enabled text understanding and additional modalities (e.g., RFIs, emails, meeting notes, images) to capture constraint information beyond weekly reports and strengthen the framework's robustness and generalizability.

ACKNOWLEDGEMENT

This research was partially supported by the SHORES Center funded by Tamkeen under NYUAD Research Institute Awards CG013. The authors also thank William Fulton and Mohamed Rasheed (NYUAD Campus Planning and Projects Office) for facilitating access to campus projects

REFERENCES

- Agrawal, A. K., Y. Zou, L. Chen, M. A. Abdelmegid, and V. A. González. 2024. "Moving toward lean construction through automation of planning and control in last planner system: A systematic literature review." *Developments in the Built Environment*, 18: 100419. <https://doi.org/10.1016/j.dibe.2024.100419>.
- Ballard, G. 2000. "The Last Planner System of Production Control."
- Ballard, G., and G. Howell. 1998. "Shielding Production: Essential Step in Production Control." *Journal of Construction Engineering and Management*, 124 (1): 11–17. American Society of Civil Engineers. [https://doi.org/10.1061/\(ASCE\)0733-9364\(1998\)124:1\(11\)](https://doi.org/10.1061/(ASCE)0733-9364(1998)124:1(11)).
- Bayhan, H. G., S. Demirkesen, C. Zhang, and A. Tezel. 2023. "A lean construction and BIM interaction model for the construction industry." *Production Planning & Control*, 34 (15): 1447–1474. Taylor & Francis. <https://doi.org/10.1080/09537287.2021.2019342>.
- Brady, D. A., P. Tzortzopoulos, J. Rooke, C. T. Formoso, and A. Tezel. 2018. "Improving transparency in construction management: a visual planning and control model." *Engineering, Construction and Architectural Management*, 25 (10): 1277–1297. <https://doi.org/10.1108/ECAM-07-2017-0122>.
- Dallasega, P., E. Rauch, and M. Frosolini. 2018. "A Lean Approach for Real-Time Planning and Monitoring in Engineer-to-Order Construction Projects." *Buildings*, 8 (3): 38. Multidisciplinary Digital Publishing Institute. <https://doi.org/10.3390/buildings8030038>.
- Dumrak, J., and S. A. Zarghami. 2023. "The role of artificial intelligence in lean construction management." *Engineering, Construction and Architectural Management*, 32 (1): 131–155. <https://doi.org/10.1108/ECAM-02-2022-0153>.
- Gingras, Y., and S. S. Schweber. 1986. "Constraints on Construction." *Social Studies of Science*. Sage Publications London. <https://doi.org/10.1177/0306312786016002011>.
- Heigermoser, D., B. García de Soto, E. L. S. Abbott, and D. K. H. Chua. 2019. "BIM-based Last Planner System tool for improving construction project management." *Automation in Construction*, 104: 246–254. <https://doi.org/10.1016/j.autcon.2019.03.019>.
- Jang, J. W., and Y.-W. Kim. 2007. "Use of Percent of Constraint Removal to Measure the Make Ready Process." In: *15th Annu. Conf. Int. Group Lean Constr.*, edited by C. L. Pasquire C. L. and P. Tzortzopoulos, 529–538. East Lansing, Michigan, USA.
- Kolberg, D., J. Knobloch, and D. Zühlke. 2017. "Towards a lean automation interface for workstations." *International Journal of Production Research*, 55 (10): 2845–2856. Taylor & Francis. <https://doi.org/10.1080/00207543.2016.1223384>.
- Koskela, L. 1992. "Application of the New Production Philosophy to Construction."
- Lattimore, T., and C. Szepesvári. 2020. *Bandit Algorithms*. Cambridge: Cambridge University Press.
- Masoetsa, T. G., B. F. Ogunbayo, C. O. Aigbavboa, and B. O. Awuzie. 2022. "Assessing Construction Constraint Factors on Project Performance in the Construction Industry." *Buildings*, 12 (8): 1183. Multidisciplinary Digital Publishing Institute. <https://doi.org/10.3390/buildings12081183>.

- Mengiste, E., M. S. Sonkor, Z. Zheng, S. A. P. Ayllón, and B. G. de Soto. 2025. “Automating Weekly Construction Activity Progress Reporting: Leveraging AI-Driven Workflows.” *International Symposium on Automation and Robotics in Construction (ISARC) Proceedings*, 2025 Proceedings of the 42nd ISARC, Montreal, Canada: 641–648. IAARC. <https://doi.org/10.22260/ISARC2025/0084>.
- Mengiste, E., B. G. de Soto, and T. Hartmann. 2023. “Automating lookahead planning using site appearance and space utilization.” arXiv.
- Mengiste, E. T., and B. Garcia de Soto. 2025. “Harmonizing Construction Flow: Integrating Takt Planning, Workspace Management, and Lookahead Strategies.” 856–868.
- Mouhoub, N. E., and A. Benhocine. 2012. “An efficient algorithm for generating AoA networks.” arXiv.
- Neumann, K., and C. Schwindt. 1997. “Activity-on-node networks with minimal and maximal time lags and their application to make-to-order production.” *Operations-Research-Spektrum*, 19 (3): 205–217. <https://doi.org/10.1007/BF01545589>.
- Scherer, R. J., and S.-E. Schapke. 2011. “A distributed multi-model-based Management Information System for simulation and decision-making on construction projects.” *Advanced Engineering Informatics*, Special Section: Advances and Challenges in Computing in Civil and Building Engineering, 25 (4): 582–599. <https://doi.org/10.1016/j.aei.2011.08.007>.
- Schimanski, C. P., C. Marcher, G. P. Monizza, and D. T. Matt. 2020. “The Last Planner® System and Building Information Modeling in Construction Execution: From an Integrative Review to a Conceptual Model for Integration.” *Applied Sciences*, 10 (3): 821. Multidisciplinary Digital Publishing Institute. <https://doi.org/10.3390/app10030821>.
- Singh, S., and K. Kumar. 2020. “Review of literature of lean construction and lean tools using systematic literature review technique (2008–2018).” *Ain Shams Engineering Journal*, 11 (2): 465–471. <https://doi.org/10.1016/j.asej.2019.08.012>.
- Sobieraj, J., and D. Metelski. 2022. “Project Risk in the Context of Construction Schedules—Combined Monte Carlo Simulation and Time at Risk (TaR) Approach: Insights from the Fort Bema Housing Estate Complex.” *Applied Sciences*, 12 (3): 1044. Multidisciplinary Digital Publishing Institute. <https://doi.org/10.3390/app12031044>.
- Soman, R. K., and M. Molina-Solana. 2022. “Automating look-ahead schedule generation for construction using linked-data based constraint checking and reinforcement learning.” *Automation in Construction*, 134: 104069. <https://doi.org/10.1016/j.autcon.2021.104069>.
- Wu, C., X. Li, R. Jiang, Y. Guo, J. Wang, and Z. Yang. 2023. “Graph-based deep learning model for knowledge base completion in constraint management of construction projects.” *Computer-Aided Civil and Infrastructure Engineering*, 38 (6): 702–719. <https://doi.org/10.1111/mice.12904>.
- Zhao, B., M. Cai, M. Zhou, and Y. Wang. 2025. “A Monte Carlo Tree Search-Based Algorithm for Workflow Scheduling.” In: *2025 IEEE Int. Symp. Parallel Distrib. Process. Appl. ISPA*, 557–566.