Zeng, X., Jiang, L., Cheng, Z., & Chen, Y. (2026). Lean post-typhoon recovery: multi-agent look-ahead and constraints for the Last Planner System. In F. Hamzeh, M. Poshdar, N. P. Garcia-Lopez, & V. Gan (Eds.), *Proceedings of the 34th Annual Conference of the International Group for Lean Construction* (IGLC34) (pp. 250–259). https://doi.org/10.24928/2026/0273

LEAN POST-TYPHOON RECOVERY: MULTI-AGENT LOOK-AHEAD AND CONSTRAINTS FOR THE LAST PLANNER SYSTEM

Xianghui Zeng[1], Liu Jiang[2], Zhiyuan Cheng[3], & Yu Chen[4]

ABSTRACT

Post-typhoon recovery involves high-variability work in which plan reliability is frequently undermined by late discovery of constraints, fragmented information, and rework. Rapid post-typhoon building recovery requires integrating heterogeneous evidence and converting it into actionable and auditable plans. This paper proposes a Lean Construction informed, knowledge-driven framework that combines a multi-modal knowledge graph (MMKG) with multi-agent orchestration for end-to-end building damage assessment and recovery planning, aligned with the Last Planner System (LPS) through look-ahead planning, make-ready checks, and explicit constraint management. The MMKG captures emergency management criteria, typhoon evolution and exposure status, building-level damage evidence, and reusable post-disaster recovery templates. The approach uses the MMKG as a shared and traceable evidence base, enabling multi-agent orchestration to generate a constraint log, assign make-ready status, and support weekly work plan release under explicit rule-based conditions. Typhoon Yagi is used as an illustrative proof-of-concept case to demonstrate how heterogeneous post-typhoon evidence can be translated into LPS-oriented planning artifacts. The results indicate that the framework is feasible for generating traceable planning outputs under multi-source uncertainty; however, quantitative field-scale validation and execution stage performance assessment remain future work.

KEYWORDS

Last Planner System, multi-agent, knowledge-driven, multi-modal knowledge graph, Typhoon disaster.

INTRODUCTION

Typhoons cause destructive winds, intense rainfall, flooding, and storm surges, which can compromise building structural safety, envelope integrity, and functionality, especially in coastal and low-lying regions (Shan et al., 2023). After landfall, emergency response and recovery planning must be carried out under severe time pressure, while post-disaster evidence remains heterogeneous, continuously updated, and difficult to align across departments, which weakens traceability and decision efficiency (Koldasbayeva et al., 2024). Beyond information

[1] PhD Student, School of Civil Engineering and Architecture, Hainan University, Haikou, China, 25110814000020@hainanu.edu.cn, orcid.org/0009-0008-5084-8575

[2] Associate Professor, School of Civil Engineering and Architecture, Hainan University, Haikou, China, oxazajl@hainanu.edu.cn, orcid.org/0000-0002-3355-2157

[3] Undergraduate student, School of Civil Engineering and Architecture, Hainan University, Haikou, China, 20223005212@hainanu.edu.cn

[4] Master Student, School of Civil Engineering and Architecture, Hainan University, Haikou, China, 19989716938@163.com, orcid.org/0009-0008-1120-6393

integration, a more critical operational bottleneck is that recovery plans become unreliable when prerequisites, resource constraints, access restrictions, and safety conditions are discovered late, leading to rework, delayed task release, and weak commitments (Agrawal et al., 2024).

From a lean perspective, post-typhoon recovery is a high-variability production system where plan reliability depends on explicit constraint management and reliable task release (Garcés et al., 2025). LPS provides look-ahead, make-ready, and weekly release logic (Ballard, 2000), yet existing post-disaster tools lack a shared, traceable evidence base for LPS-oriented planning.

This study addresses the lack of a shared and traceable evidence base for LPS-oriented post-typhoon recovery planning by integrating a multi-modal knowledge graph with multi-agent orchestration to generate evidence-linked planning artifacts. Typhoon Yagi is used as an illustrative proof-of-concept case at the building-recovery level. The study is limited to a proof-of-concept setting rather than full operational deployment, and focuses on two questions: how heterogeneous post-typhoon evidence can be organized into a shared and traceable knowledge base, and how it can support auditable look-ahead, constraint, and task-release decisions under uncertainty.

LITERATURE REVIEW

MULTI-MODAL KNOWLEDGE GRAPH FOR DISASTER AND DECISION-MAKING

The core challenge of post-disaster management is aligning visual evidence, statistics, time series, and reports within a unified semantic and spatiotemporal framework. A multi-modal knowledge graph (MMKG) addresses this by integrating heterogeneous data into entities, relationships, events, and constraints, enabling structured retrieval and cross-departmental collaboration (Song et al., 2024). In disaster scenarios, combining KG with large language models supports controlled, evidence-based reasoning. Chen et al. (2024) built an emergency KG and used chained prompts to guide LLM reasoning, reducing hallucinations via external structured constraints. For typhoons, Huang et al. (2026) proposed TyphoonKG, an LLM-agent framework for end-to-end KG construction and dynamic updates to support trajectory tracking and impact analysis.

MULTI-AGENT ORCHESTRATION FOR LEAN POST-TYPHOON RECOVERY PLANNING

Multi-agent systems (MAS) with large language models as their reasoning core support planning, execution, and reflection via message passing (Li et al., 2024). For post-typhoon recovery, the key challenge is production control under high variability: recovery activities must become executable work with explicit prerequisites, responsibilities, and verifiable evidence. Recent LLM-agent studies improve typhoon knowledge foundations (Huang et al., 2026) but emphasize knowledge organization and query services, offering limited support for lean-oriented recovery orchestration-specifically look-ahead planning, make-ready checks, and constraint management required by LPS (Soman & Molina-Solana, 2022). Existing disaster informatics lack a single source of truth for evidence-linked planning, explicit constraint-verification logic, and strong integration with LPS routines. These gaps define the specific target of the present study.

RESEARCH GAP

Based on the above, the gaps in publishable research on post-typhoon building damage and recovery, from a Lean Construction and LPS-oriented perspective, can be summarized as follows:

1. Lack of a single source of truth for LPS-oriented planning. Existing studies improve damage mapping but poorly align imagery, statistics, reports, and constraints within one

traceable structure (Oktari et al., 2020; Song et al., 2024; Xu et al., 2025), hindering constraint identification and auditable task release. This study addresses this via an MMKG linking entities, spatiotemporal anchors, provenance, confidence, and versioning.

2. Insufficient planning reliability evaluation for the recovery phase under high variability. Geographic data-driven models suffer bias, imbalance, and uncertainty, causing unstable cross-event transfer; MAS without verification may produce untraceable error chains, undermining lean planning reliability (Koldasbayeva et al., 2024; Wang et al., 2025).
3. Limited methodologies for multi-objective trade-offs and multi-risk coupling integrated into LPS routines. Recovery involves competing objectives and coupled risks, yet existing work lacks actionable mechanisms linking evidence governance to executable planning (e.g., look-ahead updates, make-ready checks, rule-based release) (Mohammadi et al., 2024).

METHODOLOGY

Based on the above gaps, this study frames post-typhoon building recovery as an evidence-governed planning and control problem under high variability. The method comprises two connected layers. First, MMKG construction unifies multi-source information (post-event imagery, loss statistics, typhoon tracks, meteorological sequences, field reports) while recording provenance, spatiotemporal scope, confidence, and versioning to support auditable planning. Second, multi-agent orchestration performs evidence alignment, look-ahead planning, constraint logging, make-ready verification, and weekly plan release under graph-grounded constraints. Figure 1 summarizes this end-to-end workflow.

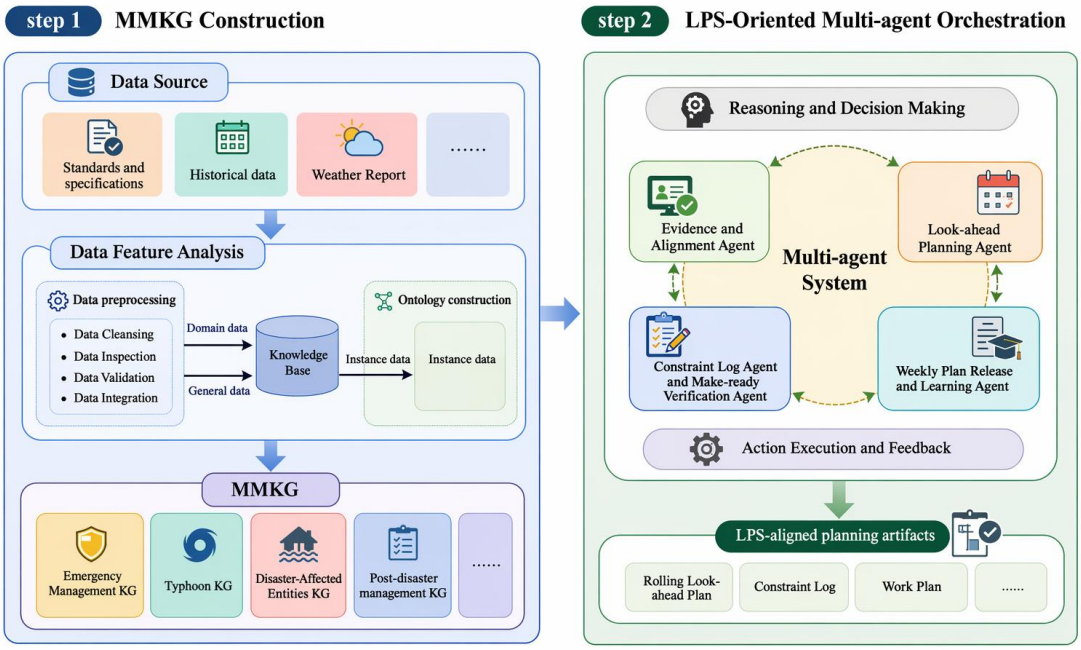

Figure 1: Framework of the methodology

MMKG CONSTRUCTION

Knowledge representation is the foundation and prerequisite for knowledge application. Understanding and reconstructing the typhoon-related knowledge framework is crucial for transforming widespread unstructured knowledge into practical structured resources (Chen et al., 2024; Oktari et al., 2020). This paper establishes a method for representing typhoon-related

knowledge from the perspectives of emergency management regulations, typhoon information, building information, and post-disaster recovery.

Emergency Management Standards KG

The role of the emergency management normative framework is to transform institutionalized norms into searchable and executable structured knowledge, thereby enabling multi-agent task decomposition and decision generation to have boundaries and auditability. The emergency management standards KG transforms institutional rules into executable planning constraints for LPS. Normative clauses are abstracted into a minimal structure that supports make-ready verification, including stage, responsible party, permitted action, and acceptance criteria. These clauses define feasibility boundaries and compliance constraints for task release, and enable auditable commitments by linking each planned action to explicit prerequisites and evidence.

Typhoon KG

The Typhoon KG represents event evolution and hazard exposure within a unified spatiotemporal structure. Following the general construction logic of Typhoon KG, event fragments extracted from meteorological bulletins, disaster reports, and best-track products are instantiated as structured graph entries, with each entry linked to its source, time, location, and confidence level. In the present framework, the Typhoon KG serves two functions. First, it provides spatiotemporal anchors for exposure reasoning and cross-source alignment. Second, it supports downstream planning queries by linking event evolution to affected regions, time windows, and impact thresholds that may influence task urgency, accessibility, and recovery sequencing.

Disaster-affected Entities KG

The Disaster-affected Entities KG is intended to address the difficulty of translating assessment results into actionable restoration targets. The key is to treat buildings as disaster-bearing entities, aligning imagery, statistics, and textual evidence at the building scale, and linking damage typification and restoration action templates to the same semantic chain. During construction, the Building entity, which includes location, footprint, usage/structure type, and criticality level, serves as the core entity, connecting ImageEvidence (such as post-disaster imagery and visual model output), LossRecord (such as county or street-level statistical losses), and post-disaster building damage assessments.

Post-disaster Management KG

The post-disaster management KG provides reusable recovery routines and control indicators consistent with Lean planning and improvement. It encodes recovery phases, task templates, handoff conditions, risk coupling, and uncertainty disclosures as queryable structures. The graph also stores planning-control indicators and diagnostic categories such as reasons for non-completion, enabling continuous improvement loops that refine constraint recognition and make-ready rules across events.

LPS-ORIENTED MULTI-AGENT ORCHESTRATION

Given the MMKG as a shared memory and constraint foundation, the MAS is designed to operationalize LPS routines rather than generate one-shot recommendations (Chen et al., 2024). Using knowledge in the MMKG as constraints, the MAS decomposes the objective into four types of tasks: information aggregation, conflict diagnosis, solution generation, and solution verification. The workflow is organized around four LPS-centered functions.

1. The Evidence and Alignment Agent consolidates cross-modal inputs and maintains the MMKG as a single source of truth, producing evidence-linked building states and flagging conflicts with provenance and confidence annotations.

2. Look-ahead Planning Agent generates a rolling look-ahead plan by decomposing recovery objectives into candidate work packages and sequencing them by risk, accessibility, and dependencies retrieved from the MMKG.
3. Constraint Log Agent and Make-ready Verification Agent explicitly identify prerequisites for each candidate task, record them in a living constraint log with responsible parties and due conditions, and perform make-ready checks using standards KG rules and evidence links. Tasks that fail readiness checks are not released and are returned to evidence alignment or constraint resolution steps.
4. Weekly Plan Release and Learning Agent releases only executable tasks to a weekly work plan based on readiness status and resource constraints. During execution, it records completion-to-commitment outcomes and reasons for non-completion, and writes back validated changes and diagnostic signals to the MMKG, enabling iterative improvement of constraint identification and planning rules.

The final outputs are LPS-aligned planning artifacts with traceable evidence, including a rolling look-ahead plan, a constraint log, make-ready status for each task, and a weekly work plan with evidence-linked commitments. Damage lists and recovery priorities are treated as intermediate products that support planning, and all released tasks retain auditable evidence paths and explicit constraint states.

CASE STUDY

EVENT AND DATASET

Typhoon Yagi (No. 11 of 2024) is used as an illustrative case to instantiate the proposed framework. A subset of damaged buildings in the affected region is selected as the analysis scope. The typhoon event information is summarized in Tables 1 and 2, and representative building records are presented in Table 3. After landfall, Typhoon Yagi produced strong winds, heavy rainfall, and compound hazard effects. These characteristics make the event suitable for examining whether the proposed framework can align heterogeneous evidence and generate executable, evidence-linked planning artifacts for post-typhoon building recovery.

Constructed around the end-to-end chain of post-disaster building damage assessment and recovery orchestration, the dataset contains five modalities uniformly mapped to MMKG entities: post-disaster image evidence generates building-level DamageObservations (imaging time, coverage area, model confidence); a loss statistics table with indicators (number of damaged items, damage type, damage assessment) serves as a macro-level constraint and consistency reference beyond building-level evidence; typhoon path and intensity sequences (from meteorological reports or best-track products) use path points and time windows as spatiotemporal anchors for exposure inference; disaster reports, emergency management regulations and related standards are decomposed into locatable ReportClaims (location, object, damage, demand, source paragraphs) with retained metadata tags for conflict interpretation and verification; general post-disaster recovery reference cases include historical data on personnel and resource allocation for post-disaster management and building reconstruction.

Table 1: Typhoon Yagi (No. 11 of 2024) summary information

Storm ID	Start	End	Max Intensity
2024244N09137	Aug 30 12Z	Sep 9 12Z	145 kt (Sep 5 18Z), 907 mb (Sep 5 18Z)

Table 2: Part of Typhoon Yagi (No. 11 of 2024) position and intensity

Time	Wind Speed	Center Position	Intensity
2024/9/4 10:00	35m/s	19.0N/117.7E	Typhoon
2024/9/4 13:00	38m/s	19.1N/117.5E	Typhoon
2024/9/4 17:00	48m/s	19.1N/117.2E	Strong typhoon
2024/9/4 21:00	50m/s	19.1N/116.6E	Strong typhoon
2024/9/4 23:00	52m/s	19.1N/116.3E	Super typhoon
2024/9/5 3:00	55m/s	19.2N/116.0E	Super typhoon

Table 3: Part of buildings information

Structure Type	Picture	Time	Position	Damage Situation
Steel structure		2024/9/10 16:00	110.039996N/ 19.937942E	Severe damage
Membrane structure		2024/9/10 16:00	110.287286N/ 20.041453E	Local damage

IMPLEMENTATION DETAILS

MMKG Schema for LPS-oriented Planning

The MMKG adopts a shared core-entity and sub-graph extension design, and the schema centers on eight core entity types: Event (typhoon event), Region (administrative region or grid), Building, TimeWindow, Evidence, Action (recovery action), Rule (standard rule), and Resource. Building serves as the focal object for post-disaster management, strongly linked to Evidence, DamageAssessment, and RecoveryPlan, while Event anchors the spatiotemporal framework through strong associations with Region and TimeWindow. Each sub-graph emphasizes distinct aspects: the normative KG defines stage-based responsibility chains, triggering conditions, and mandatory constraints; the typhoon KG provides sequences of path points, intensity evolution, and impact scopes; the building KG supplies post-disaster image evidence, damage levels, statistical records.

For multi-source data ingestion and information extraction, structured image evidence objects are first generated from post-disaster imagery via visual models or rule-based methods, containing building candidate boxes, damage segmentation results, damage-level probabilities, and quality labels, all stored as ImageEvidence nodes. Loss statistics tables are directly mapped to LossRecord nodes with field-level alignment to region via region_id. Disaster reports are processed through a chained pipeline of stepwise extraction and alignment: initially extracting event fragments (location, time, object, damage description, and handling actions), then linking entities to Building, Region, and finally binding sentence-level evidence to structured entries to form queryable evidence paths. A fundamental engineering requirement for the MMKG is to bind every conclusion to a traceable evidence chain, as shown in Figure 2.

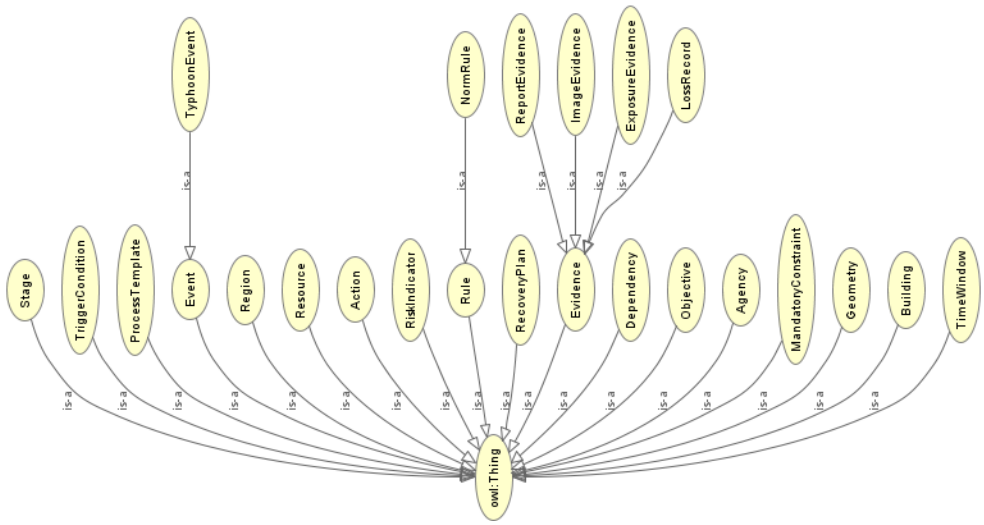

Figure 2: Ontology schema for the Typhoon Yagi.

LPS-Oriented Multi-Agent Implementation and Proof-of-concept Results

In the current proof-of-concept, the multi-agent workflow is implemented as a graph-grounded orchestration pipeline rather than as a set of independently trained specialist models. The MMKG serves as the shared memory and constraint base. LLM-assisted modules are used only where semantic interpretation is needed, such as structured extraction from textual reports and explanation of evidence-linked planning outputs. This separation is intended to limit unsupported free generation and to preserve auditability.

The operational workflow proceeds in five steps. First, multi-source inputs are ingested and aligned into MMKG entities and evidence links. Second, the Evidence and Alignment Agent consolidates building-level and region-level evidence into an evidence pack and flags inconsistencies. Third, the Look-ahead Planning Agent generates candidate work packages and preliminary sequencing based on damage severity, accessibility, dependencies, and criticality. Fourth, the Constraint Log and Make-ready Verification Agent checks each task against rule-based release criteria and records unresolved prerequisites in a living constraint log. Fifth, the Weekly Plan Release and Learning Agent releases only executable tasks to the weekly work plan and records diagnostic information for later refinement. In this way, the framework transforms heterogeneous evidence into auditable LPS-oriented planning outputs through a traceable and reproducible sequence. Conflict handling and make-ready verification are implemented conservatively. A conflict is triggered when two or more evidence items linked to the same building or time window provide incompatible statements on damage status, resource readiness, or recovery priority. In the current proof-of-concept, conflicts are handled by prioritizing evidence with clearer provenance and higher confidence. If the conflict remains unresolved and affects task executability, the task is not released and is returned for evidence supplementation or manual review.

An example is illustrated in Figure 3. To complete the proof-of-concept cycle, the case study is reported not only as a query example but as an evidence-to-plan trace. For a selected building case, the system retrieves imagery-derived damage observations, linked textual claims, regional loss records, and applicable rule nodes from the MMKG. These inputs are then consolidated into an evidence pack, after which the orchestration generates a candidate work package, identifies unresolved constraints, assigns make-ready status, and determines whether the task can be released to the weekly work plan. This sequence demonstrates the intended value of the framework, namely the conversion of heterogeneous disaster evidence into an auditable planning decision rather than an isolated query response.

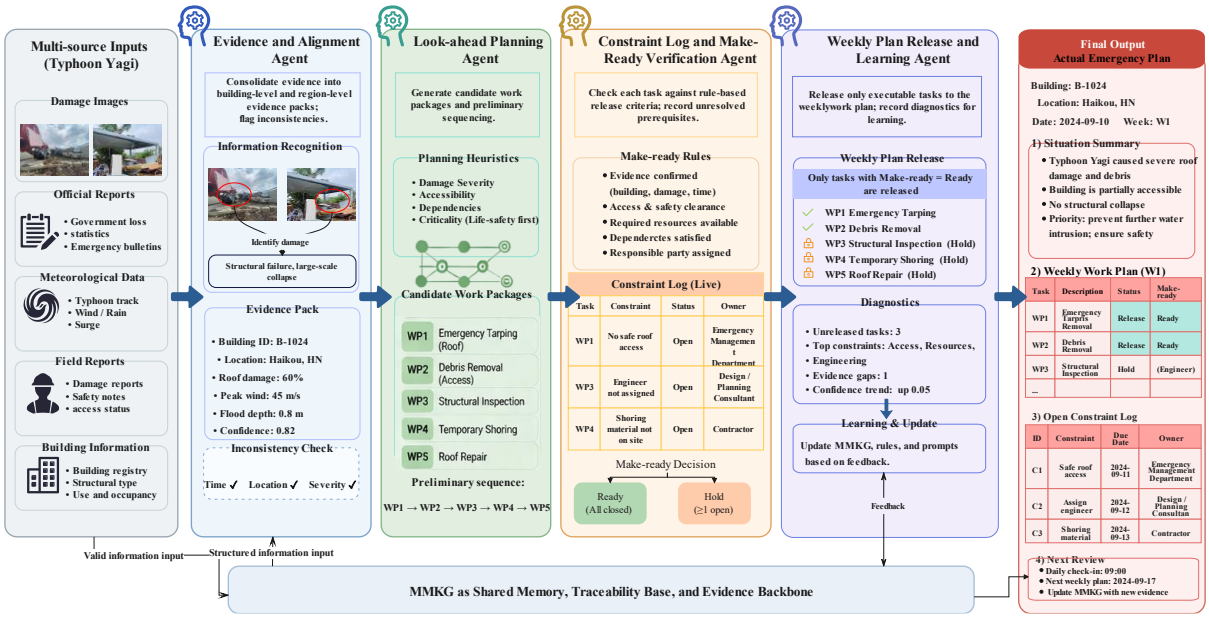

Figure 3: Workflow of the LPS-oriented multi-agent system in the Typhoon Yagi

To provide an initial validation beyond qualitative illustration, we compared the proposed workflow with a manual review baseline on 12 building cases from the Typhoon Yagi dataset. Three indicators were used. First, we measured the average time required to generate a constraint log and make-ready status per building. Second, we evaluated the quality of constraint identification by comparing system outputs against manual review annotations. Third, we assessed release-decision agreement, namely whether the system and manual review reached the same hold or release decision for each candidate task. The proposed workflow reduced the average preparation time from 11.3 min/building under manual review to 4.8 min/building. For constraint identification, it achieved a precision of 0.83, a recall of 0.77, and an F1 score of 0.80. For make-ready decisions, the agreement with manual review was 83.3%.

These quantitative results should be interpreted as preliminary proof-of-concept evidence rather than field execution performance. The current study does not yet include live recovery implementation or Percent Plan Complete (PPC) based production tracking. Accordingly, the evaluation is limited to pre-execution planning indicators, including time to generate planning artifacts, quality of constraint identification, and agreement on release decisions. Stronger claims regarding planning reliability and flow improvement will require field deployment and execution-stage observation in future work.

CONCLUSION

This study presented a lean-oriented and knowledge-driven framework for post-typhoon building recovery by integrating a multi-modal knowledge graph (MMKG) with multi-agent orchestration to support LPS-oriented planning. Its main contribution is to show how heterogeneous disaster evidence can be transformed into traceable planning artifacts, including a rolling look-ahead plan, a living constraint log, make-ready status, and an evidence-linked weekly work plan. Using Typhoon Yagi as a proof-of-concept case, the study demonstrated the feasibility of this evidence-to-plan workflow and provided an initial validation through pre-execution planning indicators.

The contribution of the present work is not a new damage dataset or an isolated LLM workflow. Rather, it is a planning-oriented integration that connects evidence governance with LPS routines. Compared with prior studies that mainly support damage interpretation, knowledge organization, or information retrieval, the proposed framework places stronger

emphasis on executable planning outputs, especially the generation of a living constraint log and auditable make-ready decisions. In this sense, the MMKG functions not only as a knowledge repository but also as a planning control basis for recovery coordination under high variability.

This work remains limited in several respects. The validation scope is currently restricted to building recovery. The quantitative evaluation is limited to planning-stage indicators and does not include live field execution, PPC based performance assessment, or avoided rework under deployment conditions. In addition, uncertainty is handled conservatively through rule-based release control rather than formal probabilistic modeling, and implementation settings for deployment-oriented agent orchestration have not yet been standardized.

Within these boundaries, the framework provides a feasible basis for extending knowledge-driven disaster recovery planning toward more reliable and auditable LPS-oriented workflows. Future work should focus on broader asset classes, formal uncertainty modeling, standardized implementation settings, and field-scale validation.

REFERENCES

Ballard, H. G. (2000). The last planner system of production control. https://api.semanticscholar.org/CorpusID:111047389

Chen, H., Song, J., Dietrich, O., Broni-Bediako, C., Xuan, W., Wang, J., Shao, X., Wei, Y., Xia, J., Lan, C., Schindler, K., & Yokoya, N. (2025). Bright: A globally distributed multimodal building damage assessment dataset with very-high-resolution for all-weather disaster response. Earth System Science Data , 17 (11), 6217–6253. https://doi.org/10.5194/essd-17-6217-2025

Chen, M., Tao, Z., Tang, W., Qin, T., Yang, R., & Zhu, C. (2024). Enhancing emergency decision-making with knowledge graphs and large language models. International Journal of Disaster Risk Reduction, 113, 104804. https://doi.org/10.1016/j.ijdrr.2024.104804

Garcés, G., Forcael, E., Osorio, C., Castañeda, K., & Sánchez, O. (2025). Systematic review of Lean Construction: An approach to sustainability and efficiency in construction management. Journal of Infrastructure Preservation and Resilience, 6(1), 6. https://doi.org/10.1186/s43065-025-00119-1

Huang, Y., Xia, Y., Tao, R., Jiao, D., Min, X., Zheng, J., Jiang, Y., Wu, W., & Du, P. (2026). A LLM-based agent for the construction of typhoon knowledge graphs. Environmental Modeling & Software , 197 , 106856. https://doi.org/10.1016/j.envsoft.2026.106856

Koldasbayeva, D., Tregubova, P., Gasanov, M., Zaytsev, A., Petrovskaia, A., & Burnaev, E. (2024). Challenges in data-driven geospatial modeling for environmental research and practice. Nature Communications , 15 (1), 10700. https://doi.org/10.1038/s41467-024-55240-8

Li, X., Wang, S., Zeng, S., Wu, Y., & Yang, Y. (2024). A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges. Vicinagearth, 1(1), 9. https://doi.org/10.1007/s44336-024-00009-2

Mohammadi, S., De Angeli, S., Boni, G., Pirlone, F., & Cattari, S. (2024). Review article: Current approaches and critical issues in multi-risk recovery planning of urban areas exposed to natural hazards. Natural Hazards and Earth System Sciences , 24 (1), 79–107. https://doi.org/10.5194/nhess-24-79-2024

Agrawal, A. K., Zou, Y., Chen, L., Abdelmegid, M. A., & González, V. A. (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

Oktari, R. S., Munadi, K., Idroes, R., & Sofyan, H. (2020). Knowledge management practices in disaster management: Systematic review. International Journal of Disaster Risk Reduction, 51, 101881. https://doi.org/10.1016/j.ijdrr.2020.101881

Shan, K., Lin, Y., Chu, P.-S., Yu, X., & Song, F. (2023). Seasonal advance of intense tropical cyclones in a warming climate. Nature , 623 (7985), 83–89. https://doi.org/10.1038/s41586-023-06544-0

Soman, R. K., & Molina-Solana, M. (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

Song, Y., Sun, P., Liu, H., Li, Z., Song, W., Xiao, Y., & Zhou, X. (2024). Scene-Driven Multimodal Knowledge Graph Construction for Embodied AI. IEEE Transactions on Knowledge and Data Engineering , 36 (11), 6962–6976. https://doi.org/10.1109/TKDE.2024.3399746

Wang, Z., Jin, Q., Wei, C.-H., Tian, S., Lai, P.-T., Zhu, Q., Day, C.-P., Ross, C., Leaman, R., & Lu, Z. (2025). GeneAgent: Self-verification language agent for gene-set analysis using domain databases. Nature Methods , 22 (8), 1677–1685. https://doi.org/10.1038/s41592-025-02748-6

Xiao, Y., & Mostafavi, A. (2025). DamageCAT: A deep learning transformer framework for typology-based post-disaster building damage categorization. International Journal of Disaster Risk Reduction , 128 , 105704. https://doi.org/10.1016/j.ijdrr.2025.105704

Xu, F., Ma, J., Li, N., & Cheng, JCP (2025). Large language model applications in disaster management: An interdisciplinary review. International Journal of Disaster Risk Reduction , 127 , 105642. https://doi.org/10.1016/j.ijdrr.2025.105642

Yuan, S., Yang, L., Mu, B., Qin, B., & Huang, Y. (2025). Assessment of Tropical Cyclone Disaster Damage Based on Learnable Inter-City Interaction GNN. Journal of Meteorological Research , 39 (5), 1146–1166. https://doi.org/10.1007/s13351-025-4239-6