Large Language Models for Intelligent Decision-Making in Robotics-Enabled Sustainable Construction

Muhammad Hamza
Department of Artificial Intelligence, Institute of Intelligent Technologies, Pakistan
https://doi.org/10.37547/ibast/Volume06Issue08-03

Abstract.
The integration of Large Language Models (LLMs) with robotics-enabled construction represents an emerging socio-technical approach to intelligent decision-making in complex and sustainability-oriented construction environments. This paper develops a conceptual research framework for understanding how language-based artificial intelligence can support robotic systems in interpreting operational information, coordinating activities, prioritizing construction decisions, and responding to uncertainty. Because the supplied literature primarily addresses data-driven prediction, customer-behavior modeling, class imbalance, feature selection, and social-network effects rather than construction robotics or LLMs directly, the study adopts a theory-transfer methodology. The analytical foundation is derived exclusively from the seven supplied references and extends their principles of predictive modeling, feature selection, imbalance management, survival-oriented reasoning, and relational analysis toward construction decision environments. The proposed framework conceptualizes LLM-enabled robotic decision-making as a layered process comprising data acquisition, contextual interpretation, predictive reasoning, decision generation, robotic execution, and feedback-based adaptation. The analysis indicates that intelligent construction systems require more than language generation: they require reliable data representation, uncertainty management, contextual prioritization, and mechanisms for translating analytical outputs into operational actions. The study identifies potential benefits in resource allocation, safety-oriented planning, equipment coordination, workflow adaptation, and sustainability management while emphasizing limitations involving domain transfer, data quality, interpretability, imbalance, and decision accountability. The paper contributes a theoretically grounded framework for future empirical investigation of LLM-driven decision intelligence in robotics-enabled sustainable construction.

Keywords: Large Language Models; Intelligent Decision-Making; Construction Robotics; Sustainable Construction; Predictive Analytics; Decision Intelligence; Human–Robot Collaboration; Resource Optimization; Artificial Intelligence; Socio-Technical Systems.

INTRODUCTION

1.1 Background

IBAST | Volume 6, Issue 08, August

Image Construction projects operate through interconnected decisions involving resources, equipment, labor, schedules, materials, environmental constraints, and changing site conditions. Robotics introduces an additional layer of operational capability because physical machines can increasingly perform inspection, transportation, assembly, monitoring, and repetitive activities. However, robotic capability alone does not establish intelligent decision-making. A construction robot must determine what action is appropriate under incomplete information, changing priorities, competing constraints, and potentially conflicting sustainability objectives. The central research problem is therefore not simply robotic automation but the development of decision mechanisms capable of connecting heterogeneous information with context-sensitive operational actions.

Large Language Models (LLMs) provide a potentially important interface for this challenge because they can process heterogeneous textual and structured information and convert complex inputs into interpretable decision representations. In a construction context, an LLM-based decision layer could theoretically combine project instructions, inspection observations, work schedules, resource information, equipment status, and sustainability requirements before recommending an action to a robotic or human-controlled system. Nevertheless, intelligent decision-making cannot be reduced to language processing. The predictive modeling literature supplied for this study demonstrates that useful decision systems depend on feature selection, class imbalance management, behavioral patterns, and appropriate analytical methods. Idris, Rizwan, and Khan, for example, demonstrate the importance of combining data balancing and feature-selection strategies in predictive classification (Idris, Rizwan and Khan, 2012). Such principles are relevant when construction decision systems must distinguish among numerous operational states.

The relevance of this research consequently lies in connecting two domains that are often discussed separately: language-based artificial intelligence and physical robotic execution. The proposed perspective treats an LLM not as an autonomous authority but as a decision-support component positioned within a broader analytical and robotic architecture. The objective is to develop a conceptual framework explaining how predictive reasoning principles can support LLM-enabled construction robotics while maintaining attention to sustainability, uncertainty, and operational accountability.

The paper has four principal objectives. First, it synthesizes the analytical principles available in the supplied literature. Second, it develops a conceptual architecture for LLM-assisted robotic decision-making. Third, it examines potential decision outcomes in sustainable construction operations. Fourth, it identifies theoretical, practical, and methodological limitations that should be addressed before such systems can be validated in real construction environments.

LITERATURE REVIEW

The supplied literature does not directly investigate LLMs, construction robotics, or sustainable construction. Therefore, the literature review does not claim that these studies empirically validate LLM-based construction systems. Instead, it extracts transferable methodological principles that can inform the conceptual design of such systems.

Hadden, Tiwari, and Roy examined computer-assisted customer churn management and emphasized the evolution of analytical approaches for identifying customers likely to discontinue a relationship (Hadden, Tiwari and Roy, 2007). Although the application domain is fundamentally different from construction, the study provides a useful theoretical foundation for decision-support architecture. Churn management transforms large amounts of heterogeneous information into a prioritized decision problem. A comparable construction system could transform project information into prioritized operational decisions, such as whether a robot should inspect an area, transport materials, modify a task sequence, or request human intervention. The transferable principle is therefore not the specific churn application but the conversion of complex data into actionable prioritization.

The studies by Idris, Rizwan, and Khan and by Burez and Van den Poel provide complementary perspectives on predictive data quality. Idris, Rizwan, and Khan investigate random forest, particle swarm optimization, data balancing, and feature-selection strategies for churn prediction (Idris, Rizwan and Khan, 2012). Burez and Van den Poel specifically address class imbalance in customer churn prediction (Burez and Van den Poel, 2009). Together, these studies suggest that decision performance depends strongly on how data characteristics influence model learning. In robotics-enabled construction, the equivalent problem may arise when normal operating conditions are substantially more frequent than critical events. For example, thousands of ordinary equipment observations may exist for every collision risk, structural anomaly, or severe resource conflict. A decision model trained without adequate treatment of such imbalance could underrepresent rare but operationally important conditions.

Prasad and Madhavi investigated data-mining approaches for predicting bank-customer churn behavior (Prasad and Madhavi, 2012). Their work reinforces the role of computational data analysis in transforming behavioral information into predictions. Applied conceptually to construction, this principle supports the development of predictive decision layers capable of identifying emerging workflow problems before they become operational failures.

Richter, Yom-Tov, and Slonim introduced social-group analysis into customer-churn prediction, demonstrating that relationships among individuals can provide predictive information (Richter, Yom-Tov and Slonim, 2010). This relational perspective is particularly relevant to construction because construction operations are inherently interconnected. A robot's decision may influence other robots, workers, equipment, material flows, and task dependencies. Thus, the state of one component cannot always be interpreted independently of the states of related components.

Wang, Yao, Cheng, and Liu applied survival-analysis reasoning to factors influencing customer churn in commercial banking (Wang, Yao, Cheng and Liu, 2014). The important transferable concept is temporal reasoning: a decision system should not only ask what state exists but also consider how conditions evolve over time. In sustainable construction, temporal reasoning can support questions such as whether material shortages are likely to persist, whether equipment degradation is increasing, or whether a workflow disruption is becoming increasingly probable.

Xiao, Liu, and He proposed GMDH-based integrated modeling for customer-churn prediction (Xiao, Liu and He, 2012). The relevance to the present framework is the concept of integrated predictive modeling rather than dependence on a single analytical representation. LLM-enabled construction intelligence may similarly require integration of language interpretation, numerical prediction, temporal information, relational information, and robotic state data.

Collectively, the literature identifies five transferable foundations: predictive decision support, feature relevance, imbalance management, relational reasoning, and temporal analysis. The major research gap is clear. These analytical principles have not, within the supplied references, been integrated with LLMs and construction robotics. The present study therefore positions LLMs as a potential coordination and reasoning layer rather than treating them as a replacement for predictive analytics or robotic control.

METHODOLOGY

3.1 Research Design

This study adopts a conceptual and theory-transfer research design. The methodology does not claim experimental validation of an LLM or robotic platform. Instead, it derives an analytical framework from the seven supplied studies and maps their methodological principles onto the decision requirements of robotics-enabled sustainable construction.

The framework follows a six-stage decision cycle:

Construction data → Contextual interpretation → Predictive reasoning → Decision generation → Robotic/human execution → Feedback and adaptation

The first stage represents data acquisition. Construction robots and associated systems may generate information concerning equipment states, task progress, environmental observations, material movement, inspection records, and operational instructions. The second stage involves contextual interpretation, where an LLM can theoretically organize heterogeneous information into a decision-relevant representation.

The third stage is predictive reasoning. Here, the principles identified in the literature become important. Feature selection can reduce irrelevant information, imbalance management can improve attention to rare operational events, temporal modeling can represent changing conditions, and relational reasoning can represent dependencies among workers, robots, equipment, and tasks.

The fourth stage generates candidate decisions. Importantly, the LLM should not be conceptualized as producing unrestricted physical commands. Instead, it should generate structured recommendations subject to operational constraints. A supervisory system can then verify those recommendations before robotic execution.

3.2 Data and Feature Representation

A construction decision system must distinguish between information that is merely available and information that is decision-relevant. The feature-selection perspective of Idris, Rizwan, and Khan is therefore transferable to construction analytics (Idris, Rizwan and Khan, 2012). For example, a robot deciding whether to transport material may require information about material location, destination, priority, route accessibility, equipment availability, and task deadlines. Excessive unrelated information could increase decision complexity without improving accuracy.

An LLM can provide contextual interpretation, but structured variables should remain available for quantitative reasoning. A hybrid representation could therefore contain textual context, numerical measurements, temporal indicators, relational dependencies, and operational constraints. Such an arrangement reduces the risk of treating language as a substitute for measurable project-state information.

3.3 Imbalance and Rare-Event Reasoning

Class imbalance is a particularly important theoretical concern. Burez and Van den Poel demonstrate that imbalanced classes can create difficulties in predictive churn modeling (Burez and Van den Poel, 2009). The corresponding construction problem occurs when safety-critical or high-impact events are rare compared with routine events.

Suppose a robotic inspection platform observes 100,000 normal surface conditions but only 500 potentially hazardous anomalies. A naive decision system could become disproportionately optimized for normal conditions. An LLM may describe the rare anomaly appropriately once detected, but it cannot compensate for an upstream system that systematically fails to identify rare events. Therefore, imbalance management should occur before or alongside language-based reasoning.

3.4 Temporal and Relational Decision Modeling

The temporal principle reflected in survival analysis can be transferred to construction decision-making. Wang et al. demonstrate that explanatory factors can be analyzed in relation to the timing of churn events (Wang, Yao, Cheng and Liu, 2014). In construction, the analogous question is whether an operational condition is stable, deteriorating, or approaching a threshold.

Relational reasoning is similarly important. Richter, Yom-Tov, and Slonim demonstrate that social-group information can contribute to prediction (Richter, Yom-Tov and Slonim, 2010). In construction robotics, a relational model could represent dependencies such as robot-to-robot coordination, worker-to-equipment interactions, and task-to-resource dependencies. This would allow an LLM to interpret a decision within a broader operational network rather than in isolation.

3.5 Decision Generation and Human Oversight

The final framework treats LLMs as decision-support agents operating under constraints. A robotic action should ideally pass through three logical checks: analytical validity, operational feasibility, and safety or sustainability compliance. For instance, if an LLM recommends accelerating material transportation because of a schedule delay, the system must verify whether the route is available, whether another robot occupies the path, whether the action violates safety constraints, and whether the proposed acceleration produces unacceptable resource or energy consumption.

This architecture addresses a fundamental limitation of purely generative decision-making: a linguistically plausible recommendation is not necessarily an operationally valid recommendation. The computer-assisted decision-support perspective identified by Hadden, Tiwari, and Roy therefore becomes relevant because decision systems should transform analytical outputs into controlled managerial actions rather than merely produce descriptive text (Hadden, Tiwari and Roy, 2007).

RESULTS

The conceptual analysis produces five principal findings. First, LLM-enabled construction robotics should be understood as a multi-layer decision architecture, not as a standalone language model. The supplied literature consistently indicates that predictive performance depends on the structure and quality of information used for decision-making. Consequently, the LLM layer should operate above data-processing and predictive components rather than replacing them.

Second, data quality and feature relevance are fundamental prerequisites. The findings derived from Idris, Rizwan, and Khan indicate that data balancing and feature-selection strategies can materially influence predictive modeling (Idris, Rizwan and Khan, 2012). In construction, this implies that LLM-based decision systems require mechanisms for filtering project information and identifying variables that genuinely affect a decision. Without such mechanisms, contextual richness may become informational noise.

Third, rare-event sensitivity represents a major design requirement. The imbalance problem identified by Burez and Van den Poel is transferable to construction scenarios where critical events occur infrequently (Burez and Van den Poel, 2009). A sustainable robotic construction system must therefore avoid optimizing solely for routine operations. The system should preserve sensitivity to unusual equipment behavior, exceptional resource conflicts, abnormal environmental conditions, and other low-frequency but high-consequence events.

Fourth, temporal and relational reasoning expand the usefulness of intelligent decision-making. The temporal perspective associated with Wang et al. indicates that decisions can benefit from understanding how conditions evolve rather than evaluating only a single observation (Wang, Yao, Cheng and Liu, 2014). Similarly, the relational approach of Richter, Yom-Tov, and Slonim demonstrates the value of interactions among connected entities (Richter, Yom-Tov and Slonim, 2010). For robotics-enabled construction, these principles imply that a decision should account for both the trajectory of a project condition and its effects on connected resources.

Fifth, the framework indicates that LLMs are most valuable as coordination and interpretation mechanisms. Their prospective contribution lies in translating heterogeneous project information into understandable decision contexts and supporting communication between human managers, analytical models, and robotic systems. However, the supplied literature does not empirically establish the performance of LLMs in this role. The finding is therefore conceptual rather than experimentally validated.

Overall, the proposed model suggests that intelligent construction decision-making should integrate language interpretation with predictive analytics, temporal reasoning, relational modeling, and controlled execution. The result is a socio-technical architecture in which computational intelligence improves decision preparation while humans and validated control mechanisms retain responsibility for high-consequence actions.

DISCUSSION

The findings have important theoretical implications for the emerging intersection of artificial intelligence, robotics, and sustainable construction. The central theoretical proposition is that LLM-based decision intelligence should not be conceptualized as a purely linguistic phenomenon. The supplied predictive-modeling literature indicates that effective decision systems depend on how information is selected, represented, balanced, and temporally interpreted. LLMs can potentially provide a flexible interface across these analytical layers, but their usefulness remains dependent on the quality of the underlying decision architecture.

One significant implication concerns the relationship between prediction and explanation. Traditional predictive systems may generate a probability or classification, whereas an LLM can potentially translate analytical outputs into human-readable reasoning. In a construction environment, this could improve communication between project managers and robotic systems. Nevertheless, explanatory fluency creates a potential risk: a convincing textual explanation may be mistaken for evidence of analytical correctness. The framework therefore requires separation between evidence generation and language-based explanation.

A second issue involves class imbalance. The findings from Burez and Van den Poel and Idris et al. suggest that predictive systems must explicitly address uneven data distributions (Burez and Van den Poel, 2009; Idris, Rizwan and Khan, 2012). In construction, rare safety or equipment events may be more consequential than common operational states. This creates a trade-off between optimizing average performance and preserving sensitivity to exceptional conditions. An LLM cannot independently resolve this trade-off because it depends on the upstream data and decision policies.

A third implication concerns organizational coordination. Richter et al.'s relational perspective suggests that predictive information may emerge from relationships among entities rather than isolated observations (Richter, Yom-Tov and Slonim, 2010). This is particularly significant for multi-robot construction environments. A decision that improves one robot's efficiency may create congestion elsewhere. Therefore, sustainable optimization should be system-oriented rather than individually optimized.

Temporal reasoning also introduces an important distinction between current-state decisions and trajectory-aware decisions. Wang et al.'s survival-analysis perspective supports the broader idea that timing matters in predictive decision-making (Wang, Yao, Cheng and Liu, 2014). A construction system should consequently consider whether a delay, equipment anomaly, or material shortage is transient or likely to persist. This can improve resource planning and reduce reactive decision-making.

The principal limitation is the domain mismatch of the supplied literature. None of the seven references directly evaluates LLMs, robotics, construction, or sustainability. Therefore, the framework represents a theoretically motivated transfer rather than empirical evidence. Claims concerning actual LLM accuracy, robotic performance, energy savings, safety improvements, or construction productivity cannot be validated from these references alone.

A second limitation is that language models may generate contextually plausible but operationally inappropriate recommendations. This is especially important when recommendations are translated into physical robotic actions. Human oversight, rule-based constraints, validated sensors, and independent control mechanisms are therefore necessary.

A third limitation concerns sustainability itself. Sustainability involves environmental, economic, and social objectives that may conflict. A decision that minimizes travel distance might increase waiting time; a decision that maximizes productivity might increase energy consumption. Consequently, sustainability should be represented as a multi-objective decision problem rather than a single optimization target.

Despite these limitations, the framework establishes a useful theoretical bridge between predictive decision support and LLM-enabled robotics. It suggests that future research should experimentally evaluate whether combining language interpretation with feature selection, imbalance handling, temporal modeling, and relational reasoning produces more robust construction decisions than either conventional analytics or LLM-only approaches.

CONCLUSION

This study developed a conceptual framework for Large Language Models in intelligent decision-making for robotics-enabled sustainable construction using only the supplied literature as its analytical foundation. Because the references primarily address customer-churn prediction and associated data-mining problems, their contribution was interpreted through transferable principles rather than as direct evidence about construction robotics.

The analysis identified five core requirements: relevant feature representation, imbalance-aware prediction, temporal reasoning, relational modeling, and controlled decision execution. These requirements support a layered architecture in which construction data are processed through analytical models, interpreted by an LLM-based reasoning interface, evaluated against operational constraints, and ultimately converted into human-approved or machine-controlled actions.

The principal contribution is therefore conceptual. The paper positions the LLM as a coordination and interpretation layer within a broader decision-intelligence system rather than as an autonomous replacement for predictive analytics or robotic control. This distinction is essential because linguistic plausibility does not guarantee operational validity.

Future research should empirically test the proposed architecture using real construction datasets and robotic platforms. Particular attention should be given to rare-event detection, temporal prediction, multi-robot coordination, sustainability trade-offs, explainability, human oversight, and safety validation. Comparative experiments could evaluate LLM-assisted systems against conventional predictive models and LLM-only approaches. Such studies would determine whether the theoretical integration proposed here can produce measurable improvements in construction efficiency, resilience, safety, and sustainability.

REFERENCES

1. A. Idris, M. Rizwan and A. Khan, "Churn prediction in telecom using random forest and PSO based data balancing in combination with various feature selection strategies," Computers & Electrical Engineering, vol. 38, no. 6, pp. 1808–1819, 2012.
2. J. Burez and D. Van den Poel, "Handling class imbalance in customer churn prediction," Expert Systems with Applications, vol. 36, no. 3, pp. 4626–4636, 2009.
3. J. Hadden, A. Tiwari and R. Roy, "Computer assisted customer churn management: State-of-the-art and future trends," Computers & Operations Research, vol. 34, no. 10, pp. 2902–2917, 2007.
4. D. Prasad and S. Madhavi, "Prediction of churn behavior of bank customer customers using data mining tools," Business Intelligence Journal, vol. 5, no. 1, pp. 96–101, 2012.
5. Y. Richter, E. Yom-Tov and N. Slonim, "Predicting customer churn in mobile networks through analysis of social groups," in The SIAM Int. Conf. on Data Mining, Columbus, USA, pp. 732–741, 2010.
6. W. Q. Wang, R. Yao, C. Cheng and C. Liu, "Influencing factors of customer churn in commercial banks: A study based on survival analysis method," Financial Forum, vol. 19, no. 1, pp. 73–79, 2014.
7. J. Xiao, D. H. Liu and H. Z. He, "GMDH-based one-step integrated modeling for customer churn prediction," Systems Engineering Theory and Practice, vol. 32, no. 4, pp. 808–813, 2012.
8. Ramamurthy, K. (2023). AI-Driven Test Automation Frameworks for the Modern Software Quality Engineering. International Journal of Emerging Trends in Computer Science and Information Technology, 4(4), 257-269.
9. Geo Philip, Paulson, Robotics-Enabled Sustainable Construction Management: A Socio-Technical Framework for Operational Efficiency, Digital Integration, and Sustainability Performance. Available at SSRN: https://ssrn.com/abstract=6845167 or http://dx.doi.org/10.2139/ssrn.6845167