Frontiers in Emerging Multidisciplinary Sciences (Open Access)

Volume 03, Issue 08, August 2026
Publish Date: 15-08-2026
doi: https://doi.org/10.64917/fems/Volume03Issue08-05 PageNo.62-68

Integrating Large Language Models and Robotics for Smart and Sustainable Construction Management

Chinedu Okafor
Department of Artificial Intelligence, Institute of Computing and Robotics, Nigeria

Amina Yusuf
Department of Intelligent Systems Engineering, Centre for AI Innovation, Nigeria

ABSTRACT

The integration of large language models (LLMs) with robotics represents an emerging approach for improving intelligence, automation, coordination, and sustainability in construction management. This research and review paper develops a conceptual framework for integrating language-based intelligence with robotic execution in construction environments. The analysis is grounded exclusively in the seven references supplied by the study, which primarily investigate Tibetan-language sorting, sequencing, recognition, Unicode-based processing, algorithmic ordering, and rule-based computational implementation. Although these references do not directly address construction, robotics, sustainability, or contemporary LLMs, their computational principles provide a useful theoretical basis for understanding how structured information can be transformed into machine-interpretable sequences and operational rules. The proposed framework therefore conceptualizes an LLM as an interaction and reasoning layer, a structured information layer as the translation mechanism, and robotics as the physical execution layer. Particular attention is given to information normalization, task sequencing, rule representation, human–robot coordination, resource optimization, and sustainability-oriented decision support. The findings indicate that the fundamental challenge is not merely connecting an LLM to a robot but establishing reliable transformations between ambiguous natural-language instructions and deterministic physical actions. The paper argues that rule-based sequencing and structured information processing remain important foundations for reliable intelligent construction automation. Limitations arise from the indirect relevance of the supplied literature and the absence of empirical construction datasets. Future research should therefore validate the framework through construction-specific experiments, digital twins, robotic field trials, and sustainability performance measurements.

KEYWORDS: Language Models; Robotics; Smart Construction; Sustainable Construction; Construction Management; Intelligent Automation; Task Sequencing; Human–Robot Collaboration; Information Processing.

INTRODUCTION

Construction management is increasingly influenced by digitalization, automation, intelligent decision support, and data-driven coordination. A particularly significant development is the potential integration of large language models with robotic systems. LLMs can interpret natural-language instructions, organize contextual information, generate structured plans, and support interaction between human operators and digital systems, while robots can execute physical activities with repeatability and measurable operational precision. The combination consequently creates the possibility of connecting managerial intent with physical construction activity through an intelligent computational interface.

The central problem, however, is not simply the availability of language models or robotic hardware. Construction operations involve heterogeneous information, incomplete instructions, changing environmental conditions, sequencing constraints, safety requirements, material dependencies, and coordination between multiple participants. An instruction such as "prepare the work area and install the next structural component" contains semantic information that must eventually be transformed into explicit, ordered, and executable actions. The computational problem therefore involves classification, normalization, sequencing, rule interpretation, and execution.

The supplied literature offers an indirect but useful foundation for this problem. Cereng-Zhaxi (1999) examines automatic sorting through computational rules, while Jiang and Zhou (2000) investigate the sequence of Tibetan words and methods for establishing ordering. Jiang and Kang (2004) subsequently formulate mathematical models and algorithms for written-language sorting. These studies demonstrate that apparently linguistic operations can be represented through explicit computational structures. Such a principle is relevant to intelligent construction because natural-language instructions similarly require transformation into structured sequences before they can reliably control automated processes.

This paper therefore investigates how the conceptual principles of computational ordering, recognition, standardization, and rule-based processing can inform an LLM–robotics architecture for smart and sustainable construction management. The objective is not to claim that the supplied references constitute direct evidence for LLMs or construction robotics. Rather, the objective is to extract transferable computational principles and construct a theoretically grounded framework from them.

The scope encompasses information processing, task sequencing, robotic coordination, resource management, sustainability-oriented decision support, and human–machine interaction. The significance of the research lies in identifying the information-processing layer required between language-based intelligence and physical robotic execution.

LITERATURE REVIEW

2.1 Computational Ordering as a Foundation for Intelligent Automation

The supplied literature consistently treats ordering as a computational problem rather than merely a linguistic convention. Cereng-Zhaxi (1999) describes rules for sorting Tibetan language and discusses automatic implementation through computers. The importance of this work for the present research is methodological: a system can transform complex symbolic information into an ordered computational representation when appropriate rules are defined.

For construction robotics, this principle can be interpreted as the need to transform high-level instructions into ordered operational units. An LLM may identify several possible activities from a project instruction, but a robot cannot safely execute those activities without explicit precedence relationships. Thus, computational ordering becomes an intermediary between semantic interpretation and physical execution.

Jiang and Zhou (2000) focus specifically on the sequence of Tibetan words and methods for establishing ordering. Their work reinforces the theoretical significance of sequence construction. In construction management, sequence is equally consequential because activities such as material delivery, positioning, assembly, inspection, and finishing are interdependent. A language model that identifies tasks without establishing their dependencies would provide incomplete operational intelligence.

2.2 Mathematical and Algorithmic Representation

Jiang and Kang (2004) examine a mathematical model and algorithm for sorting written Tibetan language. This contribution is important because it demonstrates the transition from qualitative ordering principles toward formal algorithmic processing. In an intelligent construction framework, the analogous transition is from natural-language instructions to machine-readable task representations.

A hypothetical instruction such as “move the prefabricated panel to the installation zone after the supporting frame has been inspected” contains at least three computational components: an object, an action, and a dependency. An LLM can potentially extract these elements, while a deterministic sequencing mechanism can verify whether the proposed order satisfies predefined constraints. The resulting architecture would prevent natural-language flexibility from directly becoming uncontrolled physical action.

Huang and Zhao (2008) present a DUCET-based Tibetan sorting algorithm, illustrating the importance of standardized ordering mechanisms. Standardization is particularly relevant to construction automation because different contractors, equipment manufacturers, and software systems may describe equivalent activities using different terminology. Without normalization, interoperability becomes difficult.

2.3 Recognition, Standardization, and Implementation

Tshedpa (2014) investigates automatic recognition and sorting of Tibetan word components using Unicode. This work extends the computational perspective from ordering toward recognition and standardized representation. The implication for LLM–robotic construction systems is that semantic recognition should be separated from physical execution through an intermediate representation.

Jiang (2006) discusses the status and standardization of sorting order in Tibetan dictionaries. Standardization is conceptually important for construction environments because intelligent automation requires shared representations of activities, objects, locations, and constraints. Without standardized representations, an LLM may produce semantically plausible but operationally incompatible outputs.

Zhong and Huang (2011) describe the design and implementation of a Tibetan syllable sorting rule using VBA. Their work demonstrates that abstract sorting rules can be translated into an implementable computational mechanism. This principle supports the proposed architecture in which language-derived instructions are converted into explicit rules before being transmitted to robotic systems.

2.4 Research Gap and Theoretical Positioning

The major gap in the supplied literature is also its principal limitation for this study: none of the seven references directly investigates LLMs, robotics, construction management, or sustainability. Consequently, they cannot empirically validate construction-specific claims. Nevertheless, collectively they establish a transferable computational foundation involving recognition, standardization, ordering, algorithmic representation, and implementation.

The theoretical position of this paper is therefore that reliable intelligent automation requires a semantic-to-structured-to-physical transformation chain. The LLM provides semantic interpretation; structured representations provide normalization and sequencing; and robotic systems provide physical execution. The supplied literature primarily informs the middle component of this chain. The resulting framework should consequently be understood as a conceptual synthesis rather than an experimentally validated construction system.

METHODOLOGY

3.1 Research Design

This study adopts a conceptual research-and-review methodology. Because the supplied references are concentrated on computational language sorting, the analysis identifies transferable computational principles rather than treating the sources as direct construction evidence. The methodological process consists of four stages: thematic extraction, abstraction of computational principles, construction of an integrated architecture, and analytical evaluation.

The first stage identifies recurring concepts across the references: sorting rules, ordering, sequence construction, mathematical modeling, automatic recognition, standardization, and implementation. The second stage abstracts these concepts into technology-independent functions. The third stage maps those functions onto an LLM–robotics construction-management architecture. The fourth stage evaluates the resulting framework in terms of operational intelligence, sustainability, reliability, interoperability, and limitations.

3.2 Proposed LLM–Robotics Architecture

The proposed framework consists of five interconnected layers.

The first layer is the human and project-information layer. It contains managerial instructions, work schedules, specifications, site observations, maintenance requirements, and resource information. Information may be expressed through natural language, structured records, or mixed formats.

The second layer is the LLM semantic layer. The LLM interprets natural-language information and identifies entities, actions, conditions, priorities, and relationships. For example, a project manager could state that a robotic system should transport a particular component after inspection and place it at a specified location. The LLM converts this instruction into a structured task description.

The third layer is the normalization and sequencing layer. This is the most important theoretical contribution derived from the supplied literature. Sorting and sequencing principles demonstrated in computational language processing suggest that symbolic information can be normalized according to explicit rules (Cereng-Zhaxi, 1999; Jiang and Zhou, 2000). In the construction framework, tasks are assigned identifiers, dependencies, priorities, constraints, and execution states.

A simplified representation can be expressed as:

Task = {Object, Action, Location, Preconditions, Priority, Constraint, Output}

The fourth layer is the robotic planning and execution layer. Structured tasks are converted into robotic commands, motion plans, inspection operations, or material-handling actions. Importantly, the LLM should not directly control safety-critical physical movements without an intermediate verification mechanism. Deterministic rules should validate whether an action is permitted before execution.

The fifth layer is the monitoring and feedback layer. Sensors and robotic systems generate information concerning task completion, environmental conditions, material status, and deviations. This information can be returned to the language layer, enabling the system to revise plans when conditions change.

3.3 Task Sequencing and Constraint Management

Construction activities are highly dependent on precedence relationships. Therefore, the framework represents project activities as a directed task graph. If task A must occur before task B, the system represents the relationship as:

$A \rightarrow B$

Multiple dependencies can create a graph such as:

Inspection → Preparation → Material Delivery → Robotic Installation → Verification

The purpose of the sequencing layer is to prevent the LLM from generating an apparently logical but physically infeasible sequence. Jiang and Kang (2004) demonstrate the broader principle that ordering can be represented through mathematical and algorithmic mechanisms. Applying this principle conceptually to construction suggests that task ordering should be governed by explicit constraints rather than generated solely through probabilistic language prediction.

For example, a robotic system should not receive an instruction to install a structural element when the supporting surface has not passed inspection. The language model can interpret the instruction, but a deterministic constraint engine should verify the prerequisite.

3.4 Standardization and Interoperability

Standardization is another essential component. Jiang (2006) emphasizes standardization in ordering systems, while Tshedpa (2014) addresses recognition and sorting through Unicode-based representations. The construction equivalent is a standardized vocabulary for activities, materials, equipment, locations, and states.

Consider the expressions “lift panel,” “raise prefabricated panel,” and “move panel upward.” An LLM may recognize their semantic similarity, but robotic systems require a normalized operational representation. The framework therefore converts equivalent expressions into standardized task codes. This reduces ambiguity and facilitates interoperability among planning software, robotic controllers, and monitoring systems.

3.5 Sustainability-Oriented Decision Support

Sustainability is integrated as a constraint rather than treated as a separate reporting activity. A task sequence can incorporate material waste, equipment utilization, energy consumption, transportation distance, and rework risk as optimization variables.

A conceptual objective function can be represented as:

Minimize: $C = \alpha E + \beta W + \gamma R + \delta T$

where E represents energy consumption, W material waste, R rework risk, and T operational time. The coefficients α, β, γ, and δ represent project-specific priorities.

The LLM can assist in interpreting sustainability requirements and generating candidate plans, whereas the structured optimization layer evaluates alternatives against measurable constraints. This separation is essential because persuasive natural-language output should not be confused with verified environmental performance.

3.6 Human–Robot Collaboration

Human oversight remains central. Zhong and Huang (2011) demonstrate how computational sorting rules can be implemented explicitly, reinforcing the value of transparent operational rules. In construction, similar transparency is necessary when robots execute tasks in dynamic environments.

A practical scenario illustrates the model. A site manager instructs the system to prepare a robotic installation sequence for prefabricated components. The LLM extracts the required activities. The sequencing layer identifies dependencies. The robot planning system checks workspace availability. If the sensor system identifies an obstruction, execution is suspended and the event is returned to the decision layer. The LLM can then communicate the reason for interruption and propose alternative scheduling, while final execution remains subject to verified constraints.

3.7 Reliability and Safety Mechanisms

Reliability requires separation between language generation and physical authorization. The framework therefore proposes three verification gates: semantic validation, constraint validation, and execution validation.

Semantic validation determines whether the instruction has been interpreted correctly. Constraint validation determines whether the requested operation is compatible with project rules. Execution validation determines whether current site conditions permit the operation.

This layered approach is consistent with the broader computational lesson of the supplied references: linguistic or symbolic information must be transformed through explicit rules before automated processing can be reliably implemented (Huang and Zhao, 2008).

RESULTS

The conceptual analysis produces five principal findings. First, structured sequencing is the central bridge between LLM intelligence and robotic action. The supplied references consistently demonstrate that language-related information can be computationally organized through sorting rules, algorithms, mathematical models, and implementation mechanisms. Transferred to construction, this suggests that an LLM should not directly translate unrestricted natural-language output into physical robotic commands. Instead, its outputs should pass through a structured sequencing layer.

Second, standardization is essential for interoperability. The literature's attention to sorting standards, recognition, and Unicode representation indicates that computational systems require consistent symbolic representations. For construction robotics, standardized representations can reduce ambiguity among project managers, digital platforms, and robotic systems. The practical implication is that different natural-language expressions should map to consistent operational concepts.

Third, hybrid intelligence is more appropriate than fully autonomous language-driven control. LLMs are suited to interpretation, explanation, contextual reasoning, and interaction, whereas deterministic algorithms are better suited to enforcing constraints. Robots are then responsible for physical execution. The proposed division of responsibilities reduces the risk that linguistic uncertainty will propagate directly into physical operations.

Fourth, sustainability can be embedded into task planning. Instead of evaluating sustainability after work has been completed, resource use and waste-related constraints can influence task sequencing before execution. For instance, a plan that minimizes unnecessary material movement may reduce equipment utilization and operational inefficiency. However, such benefits require measurable project data and cannot be inferred from language generation alone.

Fifth, the evidence base is insufficient for empirical claims about construction performance. The seven supplied references establish computational concepts but do not provide experimental evidence regarding construction robots, LLM accuracy, energy savings, productivity, or carbon reduction. Therefore, the proposed framework should be regarded as a theoretically derived architecture rather than a validated engineering solution.

Overall, the findings indicate that the most important research opportunity lies in developing reliable interfaces between semantic reasoning and deterministic execution. The relevance of computational sorting and sequencing is therefore architectural rather than domain-specific: information must be recognized, normalized, ordered, validated, and implemented before it becomes dependable machine action.

DISCUSSION

The findings position LLM–robotics integration as an information-transformation problem as much as an artificial-intelligence problem. The major theoretical contribution is the distinction between interpretation, structuring, and execution. Contemporary language-oriented systems are capable of interpreting complex instructions, but physical construction automation demands a stronger level of determinism. The supplied literature supports this distinction indirectly by demonstrating the importance of explicit sorting rules, sequence formation, mathematical algorithms, standardization, and implementation (Jiang and Kang, 2004; Jiang and Zhou, 2000).

The proposed framework also challenges an overly simplistic conception of autonomous construction. If an LLM is treated as a direct controller, ambiguity in language can become ambiguity in physical action. By contrast, the proposed architecture gives the LLM a decision-support and semantic role while assigning constraint enforcement to a structured computational layer. This division creates a more defensible architecture for safety-critical environments.

The framework has practical implications for construction management. Project managers could interact with robotic systems through natural language without manually programming every low-level operation. At the same time, the underlying system could convert managerial language into standardized task structures. Such an approach may reduce the communication gap between construction professionals and technical automation systems. The recognition and representation principles discussed by Tshedpa (2014) are particularly relevant to this transition because different symbolic forms must ultimately be mapped into computationally recognizable units.

A second implication concerns sustainability. Intelligent planning could potentially incorporate resource efficiency into the sequencing process. For example, if two task sequences achieve comparable completion times, the system could prefer the alternative requiring fewer equipment movements or less material handling. Nevertheless, this possibility depends on reliable measurements. The language model itself cannot establish that a plan is sustainable; sustainability must be evaluated using explicit operational indicators.

There are also important trade-offs. Greater automation can improve consistency but may reduce flexibility when unexpected site conditions occur. More elaborate rule systems can increase reliability but may make implementation more complex. Conversely, excessive dependence on probabilistic language generation can improve adaptability while introducing uncertainty. The appropriate solution is therefore not maximal autonomy but controlled integration.

The principal limitation is the indirect nature of the evidence base. The supplied references concern Tibetan language sorting and computational implementation rather than construction automation. Consequently, conclusions concerning productivity, safety, sustainability, or robotic performance remain conceptual. The paper should not be interpreted as empirical proof that LLMs or robots improve construction outcomes.

Future validation should use construction-specific datasets and measurable experimental variables. Research could compare conventional planning with LLM-assisted planning, measure task-sequencing accuracy, evaluate robotic execution errors, quantify material and energy consumption, and investigate human acceptance. Digital-twin environments could provide an intermediate experimental setting before deployment on active construction sites. Such studies would determine whether the theoretical principles extracted from computational language processing translate into measurable construction benefits.

CONCLUSION

This paper developed a conceptual framework for integrating large language models and robotics into smart and sustainable construction management using only the supplied literature as its theoretical foundation. The references, although focused on Tibetan-language sorting and computational processing, collectively demonstrate principles of recognition, ordering, mathematical modeling, standardization, algorithmic representation, and implementation. These principles were interpreted as transferable foundations for constructing a semantic-to-structured-to-physical architecture.

The proposed framework assigns distinct functions to five layers: project information, LLM-based semantic interpretation, structured normalization and sequencing, robotic execution, and monitoring feedback. This architecture addresses a central challenge in intelligent construction: converting flexible human language into reliable and verifiable physical operations.

The analysis further indicates that sustainability should be integrated into task planning through explicit constraints rather than treated solely as an after-the-fact evaluation. Resource movement, waste, energy consumption, rework, and operational time can conceptually become optimization variables within an intelligent planning system.

The principal contribution is therefore architectural and theoretical. The research identifies structured information processing as the critical intermediary between language intelligence and robotic execution. It also emphasizes that LLMs should not independently authorize safety-critical physical operations; deterministic validation and human oversight remain necessary.

However, the supplied literature does not empirically establish the effectiveness of LLMs, robotics, or sustainable construction technologies. Accordingly, the proposed framework requires substantial empirical validation. Future studies should develop construction-specific datasets, formal task ontologies, robotic test environments, digital twins, sustainability metrics, and human–robot evaluation protocols. Such research could transform the conceptual integration proposed here into a measurable and deployable framework for intelligent construction management.

REFERENCES

1. Cereng-Zhaxi, “The sorting rules of Tibetan language and the realization of automatic sorting by computer,” China Tibetology, vol. 12, no. 4, pp. 128–135, 1999.

2. D. Jiang, "The current status of sorting order of tibetan dictionaries and standardization," in The 20th Pacific Asia Conf. on Language, Information and Computation, Wuhan, China, 2006.
3. D. Jiang and C. J. Kang, "The sorting mathematical model and algorithm of written Tibetan language," Chinese Journal of Computer, vol. 27, no. 4, pp. 524–529, 2004.
4. D. Jiang and J. W. Zhou., "On the sequence of Tibetan words and the method of making sequence," Journal of Chinese Information Processing, vol. 14, no. 1, pp. 56–64, 2000.
5. H. M. Huang and C. X. Zhao., "A ducet-based Tibetan sorting algorithm," Journal of Chinese Information Processing, vol. 22, no. 4, pp. 109–113, 2008, 2008.
6. Tshedpa, "Research on the automatic recognition and sorting of Tibetan word components on the unicode," Journal of Tibet University, vol. 29, no. 2, pp. 81–86, 2014.
7. Z. X. Zhong and H. M. Huang, "The design and implementation of a Tibetan syllables sorting rule based on VBA," Journal of Qinghai Normal University (Natural Science), vol. 33, no. 3, pp. 46–49, 2011.
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