

Driven Digital Decision Support for Operational Efficiency in Sustainable Construction

Chinedu Okafor

Department of Artificial Intelligence and Machine
Learning, Institute of Computing Technologies, Nigeria

Fatima Ibrahim

Department of Intelligent Systems and
Robotics, Centre for AI Innovation, Nigeria

VOLUME 03 ISSUE 08 (2026)

Published Date: 15 August 2026 // Page no.: - 47-53

ABSTRACT

The increasing complexity of construction operations requires decision-support mechanisms capable of integrating heterogeneous information, interpreting operational conditions, and supporting timely managerial decisions. Large Language Models (LLMs) provide a potentially important interface for such decision support because they can transform natural-language information into structured recommendations, explanations, and decision alternatives. However, the effectiveness of an LLM-driven system depends not only on language-generation capability but also on the underlying computational architecture, learning strategy, data interpretation, and reliability of automated decisions. This research and review paper develops a conceptual framework for LLM-driven digital decision support in sustainable construction by synthesizing the computational foundations represented in the provided literature. The reviewed studies demonstrate the evolution of neural-network architectures, evolutionary computation, genetic programming, deep learning, and automated architecture optimization. These foundations are interpreted as enabling technologies for constructing adaptive decision-support pipelines, while the application of LLMs to sustainable construction is positioned as a forward-looking extension rather than a capability directly demonstrated by the cited studies. The proposed framework integrates data acquisition, LLM-based interpretation, computational optimization, decision evaluation, human validation, and sustainability-oriented operational feedback. The analysis indicates that LLM-driven decision support can potentially improve information accessibility, scenario comparison, coordination, and operational responsiveness, but its value depends on model reliability, explainability, data quality, computational efficiency, and human oversight. The study contributes a theoretically grounded framework for connecting advanced learning systems with operational decision-making in sustainable construction and identifies research directions for validating such systems in real construction environments.

Keywords: Large Language Models, Digital Decision Support, Sustainable Construction, Operational Efficiency, Artificial Intelligence, Deep Learning, Evolutionary Computing, Construction Management, Decision Optimization.

INTRODUCTION

1.1 Problem Statement

Construction projects generate large volumes of operational information concerning schedules, resources, equipment, workforce, material movement, site conditions, safety observations, and environmental performance. The difficulty is not merely the availability of information but the ability to transform fragmented information into timely and actionable decisions. Conventional decision-support approaches frequently depend on predefined rules, manually interpreted reports, or isolated analytical systems. Such approaches may become less effective when operational conditions change rapidly and when decisions require simultaneous consideration of multiple constraints.

The development of deep-learning systems has

demonstrated that computational models can learn complex representations from data. Goodfellow, Bengio, Courville and Bengio established the theoretical foundations of deep learning as a family of representation-learning methods capable of extracting increasingly complex patterns through multilayer neural architectures (Goodfellow et al., 2016). Subsequent research has demonstrated that neural architectures can also be automatically explored and evolved rather than designed exclusively through manual engineering. Bakhshi, Chalup and Noman investigated fast evolution of CNN architectures for image classification, illustrating the importance of automated architectural search (Bakhshi et al., 2020).

Within construction, the conceptual emergence of LLM-driven decision support extends this computational trajectory toward natural-language interaction. An LLM-based system can potentially act as an interpretive layer between construction

personnel and complex digital information systems. Instead of requiring managers to interact directly with multiple analytical interfaces, the system could receive questions expressed in natural language, retrieve relevant operational information, compare alternatives, and generate decision-oriented explanations.

However, LLM-based decision support should not be treated as an autonomous substitute for engineering or managerial judgment. The provided literature does not directly validate LLMs in construction environments. Therefore, this paper does not claim empirical performance of LLMs based on the cited studies. Instead, it develops a theoretically grounded framework by connecting established computational principles from deep learning and evolutionary optimization with the emerging concept of language-based digital decision support.

1.2 Research Relevance

The relevance of this research arises from the need to connect intelligent computational systems with operational efficiency and sustainability objectives. Operational efficiency requires decisions that minimize avoidable delays, inefficient resource utilization, unnecessary movement, and coordination failures. Sustainability introduces additional constraints because decisions may need to account for resource consumption and environmental consequences simultaneously.

Evolutionary computation provides an important theoretical foundation for this multi-constraint perspective. Darwish, Hassanien and Das reviewed swarm and evolutionary computing approaches for deep learning, demonstrating how evolutionary strategies can contribute to the optimization of complex learning systems (Darwish et al., 2019). David and Greental similarly demonstrated the use of genetic algorithms for evolving deep neural networks, highlighting the possibility of using evolutionary mechanisms to search complex model configurations (David and Greental, 2014).

Consequently, the central research problem is formulated as follows: How can an LLM-driven digital decision-support architecture be conceptually designed to improve operational decision-making while incorporating computational optimization and sustainability considerations?

1.3 Objectives

The primary objective is to develop a conceptual framework for LLM-driven digital decision support in sustainable construction. The study specifically aims to examine the computational foundations relevant to intelligent decision support, integrate LLM-based

interaction with optimization-oriented analytical components, identify potential operational and sustainability applications, and critically assess limitations associated with automated decision support.

1.4 Scope and Significance

The study is conceptual and review-oriented. It does not claim experimental validation of an LLM in a construction project. Its technical contribution is the integration of established concepts from deep learning, neural architecture evolution, genetic programming, and evolutionary optimization into an LLM-oriented construction decision-support framework.

The significance of the framework lies in its attempt to distinguish between language-based decision interaction and computational decision optimization. An LLM can potentially interpret information and communicate alternatives, whereas specialized analytical or optimization components may be required to evaluate technically constrained alternatives. This separation is important for designing trustworthy digital decision-support systems.

LITERATURE REVIEW

2.1 Computational Foundations of Intelligent Decision Support

Deep learning provides the fundamental computational background for modern intelligent systems. Goodfellow et al. (2016) describe deep neural networks as hierarchical representation-learning systems capable of transforming input data through multiple layers of learned representations. This principle is relevant to digital decision support because construction data can contain complex relationships that may not be captured effectively through simple rule-based mechanisms.

The evolution of neural architectures is particularly relevant to scalable intelligent systems. Bakhshi et al. (2020) examined the fast evolution of CNN architectures, demonstrating that architecture design itself can become an optimization problem. Suganuma, Shirakawa and Nagao similarly proposed a genetic-programming approach for designing CNN architectures, showing how evolutionary procedures can be used to explore alternative neural structures (Suganuma et al., 2017).

Sun, Xue, Zhang and Yen further investigated the evolution of deep convolutional neural networks for image classification and demonstrated the relevance of evolutionary search to deep-network design (Sun et al., 2020). Although these studies concern image-processing systems rather than construction management or LLMs, their methodological implication is important: intelligent computational systems can be designed as adaptive structures whose configurations are optimized according to defined objectives.

2.2 Evolutionary Optimization and Adaptive Architectures

Evolutionary computing introduces a second important foundation. Darwish et al. (2019) synthesized swarm and evolutionary approaches for deep learning and showed the broader relevance of population-based optimization strategies to computational learning. David and Greental (2014) examined genetic algorithms for evolving deep neural networks, emphasizing automated exploration of network configurations.

Such approaches can conceptually complement LLM-driven decision support. An LLM may generate or interpret candidate operational strategies, while an optimization layer evaluates those alternatives against constraints such as time, resources, equipment availability, or sustainability objectives. The evolutionary component could therefore function as a search mechanism rather than as the language interface itself.

The work of Such et al. on deep neuro-evolution is particularly relevant to this distinction. Their research presents genetic algorithms as an alternative approach for training deep neural networks for reinforcement learning (Such et al., 2017). The underlying principle suggests that adaptive computational systems may benefit from optimization mechanisms capable of searching large and complex solution spaces.

2.3 Perception and Operational Information

Construction decision support is inherently multimodal. Operational information can originate from written reports, visual observations, sensor systems, schedules, and monitoring platforms. Matsugu et al. (2003) demonstrated the application of convolutional neural networks to facial-expression recognition with robust face detection, while Selvi et al. (2022) investigated an enhanced CNN-based system for suspicious-action detection in surveillance video. Although these applications are outside construction, they demonstrate how deep-learning systems can transform visual information into higher-level interpretations.

This capability can conceptually support construction decision systems in which visual observations become inputs to operational reasoning. For example, a site-monitoring component could detect an operational condition and transmit a structured event to a decision-support layer. The LLM could then explain the event, relate it to current project information, and present potential actions. Importantly, this remains a proposed architecture rather than an empirically validated capability derived from the cited studies.

2.4 Research Gap and Theoretical Positioning

The reviewed literature reveals a significant gap between advanced computational learning and construction-oriented decision interaction. The cited research

establishes foundations for deep learning, neural architecture optimization, evolutionary computation, and visual interpretation, but it does not directly establish an LLM-based operational decision-support system for sustainable construction.

The research gap therefore concerns integration rather than the isolated development of another neural architecture. A construction-oriented framework requires a mechanism that can connect data interpretation, natural-language interaction, optimization, operational constraints, human validation, and sustainability objectives. The proposed research addresses this conceptual gap by positioning the LLM as an interaction and reasoning-support layer while retaining specialized computational modules for data processing and optimization.

METHODOLOGY

3.1 Research Design

The study adopts a conceptual research-and-review methodology. The methodology consists of systematic synthesis of the provided computational literature followed by architectural abstraction and framework development. Because the references do not provide empirical evidence concerning LLM deployment in construction, the study avoids presenting hypothetical outcomes as experimentally verified findings.

The analytical process has three stages. First, the computational capabilities represented by the references are categorized into learning, perception, architecture optimization, and evolutionary search. Second, these capabilities are mapped onto construction decision-support functions. Third, the resulting components are integrated into a conceptual LLM-driven framework.

3.2 Proposed Digital Decision-Support Architecture

The proposed architecture contains six interconnected functional layers: data acquisition, information normalization, LLM interaction, analytical optimization, decision validation, and operational feedback.

The data-acquisition layer collects project information from operational systems. Potential inputs include progress records, resource information, equipment status, schedules, material requirements, environmental observations, and site-monitoring outputs. Visual information can conceptually be processed through deep-learning models, following the broader principle demonstrated by CNN-based recognition systems (Matsugu et al., 2003; Selvi et al., 2022).

The information-normalization layer converts heterogeneous information into structured representations. This stage is essential because an LLM should not be expected to independently resolve every inconsistency in raw project data. Structured records, timestamps, identifiers, operational

constraints, and confidence indicators should accompany the information passed to the decision-support layer.

The LLM interaction layer functions as the natural-language interface. A construction manager could ask, for example, which available scheduling alternative is likely to reduce equipment idle time while maintaining a specified sustainability constraint. The LLM would interpret the request, identify relevant variables, retrieve structured project information, and formulate candidate decision scenarios.

The analytical optimization layer evaluates those scenarios. Evolutionary approaches are relevant here because construction decisions may involve competing objectives and large search spaces. The evolutionary optimization perspective developed by Darwish et al. (2019), David and Greental (2014), and Such et al. (2017) provides a theoretical basis for treating alternative configurations as candidate solutions within a search process.

3.3 Decision Evaluation Mechanism

A decision-support system should distinguish between generation and validation. The LLM may generate candidate explanations or alternatives, but technically constrained decisions should be evaluated through explicit criteria.

A conceptual decision score can be expressed as:

$$D = w_1E + w_2S + w_3R + w_4Q$$

where (E) represents operational efficiency, (S) sustainability performance, (R) resource feasibility, and (Q) decision-quality confidence. The weights (w_1, w_2, w_3, w_4) represent project-specific priorities.

This formulation does not prescribe a universal optimization equation. Rather, it illustrates how an LLM-driven system can translate qualitative managerial priorities into an explicit evaluation structure. Evolutionary optimization can then search candidate alternatives according to these criteria.

3.4 Human-in-the-Loop Control

Human oversight is a central architectural requirement. Construction decisions can involve contractual obligations, engineering judgments, safety considerations, and contextual knowledge that may not be represented in digital data. Therefore, the system should present recommendations with supporting information,

assumptions, constraints, and confidence indicators.

A practical example is equipment allocation. Suppose several machines are available but their use affects schedule performance and resource consumption differently. The LLM could summarize the available alternatives, while the optimization layer evaluates combinations against time and resource constraints. The construction manager would retain authority to approve, modify, or reject the proposed action.

3.5 Sustainability Integration

Sustainability should not be treated as an independent reporting module added after operational optimization. Instead, it should become part of the decision objective. A schedule that reduces project duration but creates excessive equipment movement may be operationally efficient under one metric while performing poorly under a broader sustainability assessment.

The proposed framework therefore uses multi-objective evaluation. The system can compare alternatives according to operational efficiency, resource utilization, and sustainability-related indicators. This approach is consistent with the broader optimization perspective represented by evolutionary and genetic approaches in the reviewed literature (Darwish et al., 2019; Suganuma et al., 2017).

3.6 Feedback and Continuous Adaptation

The final layer is operational feedback. Decisions produce outcomes, and those outcomes should be returned to the digital system. Over time, this creates an iterative loop:

Operational Data → Interpretation → Candidate Decisions → Optimization → Human Validation → Implementation → Outcome Monitoring → Feedback.

The architecture is conceptually consistent with the broader idea of adaptive computational systems. However, implementation would require domain-specific datasets, evaluation protocols, data-governance mechanisms, and empirical validation that are outside the scope of the provided literature.

RESULTS

The conceptual synthesis produces five principal findings. First, the reviewed literature indicates that intelligent computational performance depends on more than a single learning architecture. Deep learning provides representational capacity, while architecture-evolution approaches demonstrate that computational structures themselves can be optimized (Goodfellow et al., 2016; Bakhshi et al., 2020). For construction decision support, this implies that an LLM should be regarded as one component within a broader computational ecosystem rather than as a complete decision-making solution.

Second, evolutionary computation provides a useful conceptual mechanism for evaluating competing construction alternatives. The studies by Darwish et al. (2019), David and Greental (2014), and Such et al. (2017) collectively indicate that evolutionary methods can explore complex solution spaces. This is particularly relevant when operational decisions involve multiple interacting constraints. An LLM can formulate alternatives in human-readable form, while optimization mechanisms can assess whether those alternatives are computationally feasible.

Third, the literature supports the separation of perception from decision interpretation. CNN-based research demonstrates that machine-learning systems can extract meaningful information from visual inputs (Matsugu et al., 2003; Selvi et al., 2022). Within the proposed construction framework, this principle supports a pipeline in which visual or sensor-derived observations are transformed into structured operational events before reaching the language-based decision layer.

Fourth, automated architecture optimization indicates the importance of adaptability. Suganuma et al. (2017) and Sun et al. (2020) demonstrate that evolutionary methods can explore neural architectures rather than relying exclusively on fixed manually designed structures. The corresponding implication for construction decision support is that system components should be configurable according to project characteristics instead of assuming that one computational configuration is universally optimal.

Fifth, the principal potential benefit of an LLM-driven system lies in integration and accessibility rather than simply in prediction. A manager could interact with complex project information through natural language while specialized computational modules perform analytical evaluation. This can reduce the cognitive and technical barrier between operational personnel and digital analytical systems. Nevertheless, this finding is conceptual: the provided references do not empirically establish productivity improvements from LLM deployment in construction.

The resulting framework therefore positions LLMs as an interpretive and communicative layer within a larger decision-support architecture. Its effectiveness would depend on data quality, computational validation, domain adaptation, human oversight, and the explicit representation of sustainability objectives. The findings also indicate that replacing established analytical methods with an LLM alone would be theoretically weaker than integrating language-based interaction with specialized learning and optimization components.

DISCUSSION

The findings suggest that the most appropriate conceptual role of LLMs in sustainable construction is not unrestricted autonomous decision-making but decision augmentation. This distinction is important because construction environments contain constraints that cannot reliably be inferred from language alone. A language model may communicate a plausible recommendation, but plausibility does not establish engineering validity. Consequently, the proposed architecture places analytical optimization and human validation between recommendation generation and operational implementation.

This position is compatible with the computational literature. Goodfellow et al. (2016) establish the importance of learned representations, while Bakhshi et al. (2020), Suganuma et al. (2017), and Sun et al. (2020) demonstrate that model structures can be systematically explored or evolved. The implication is that intelligent construction systems should be architected as configurable computational pipelines. An LLM may provide a flexible interaction mechanism, but the underlying analytical system should retain explicit computational controls.

The evolutionary-computing literature provides another important theoretical contribution. Darwish et al. (2019) demonstrate the relevance of evolutionary and swarm approaches to deep learning, while David and Greental (2014) and Such et al. (2017) demonstrate evolutionary strategies for neural-network development and learning. For construction, these principles can be translated into multi-objective decision optimization. Instead of asking only whether a decision is feasible, the system can evaluate whether alternative decisions provide better combinations of efficiency, resource utilization, and sustainability.

The framework also exposes several trade-offs. Greater automation can increase responsiveness but may reduce transparency if recommendations cannot be traced to underlying data and constraints. More sophisticated optimization can improve the search for alternatives but can also increase computational complexity. Natural-language interaction can make analytical systems easier to use, but linguistic fluency can create an unwarranted perception of reliability. These tensions make human-in-the-loop control essential.

The visual-learning studies further strengthen the argument for modularity. Matsugu et al. (2003) and Selvi et al. (2022) demonstrate that specialized CNN systems can transform visual inputs into meaningful classifications. Rather than forcing an LLM to perform every computational task, a construction platform can use specialized perception models for site observations and allow the LLM to interpret their structured outputs. Such modularity can improve technical accountability because each subsystem has a defined responsibility.

Several limitations remain. Most importantly, the provided

references do not directly study LLMs, sustainable construction, or construction-site decision support. Therefore, the proposed framework should be regarded as a theoretically informed research architecture rather than a validated deployment model. Additional limitations include uncertainty in operational data, potential model errors, difficulty defining sustainability weights, interoperability between digital platforms, and the need for domain-specific evaluation datasets.

From a practical perspective, implementation should proceed incrementally. Organizations could initially deploy LLMs for information retrieval and report interpretation, then introduce scenario generation and optimization, followed by controlled decision-support applications. Each stage should be evaluated against measurable operational and sustainability indicators. Such staged implementation would reduce the risks associated with treating emerging language technologies as fully autonomous decision systems.

Overall, the central theoretical contribution is the integration of language interaction with established computational optimization principles. The framework extends the logic of adaptive learning and evolutionary search into a construction-management context while maintaining a clear distinction between conceptual applicability and empirical evidence.

CONCLUSION

This study developed a conceptual framework for LLM-Driven Digital Decision Support for Operational Efficiency in Sustainable Construction by synthesizing the computational principles represented in the provided literature. The review demonstrates that deep learning, evolutionary computation, genetic algorithms, and neural architecture optimization provide relevant theoretical foundations for constructing adaptive digital decision-support environments.

The proposed framework positions the LLM as an interaction, interpretation, and communication layer rather than as an independent autonomous decision-maker. Data-acquisition and perception components provide structured operational information; the LLM interprets managerial requirements and generates candidate alternatives; optimization mechanisms evaluate those alternatives; human decision-makers validate the recommendations; and operational feedback enables continuous improvement.

The principal contribution is therefore architectural rather than empirical. The framework provides a logical pathway for connecting language-based interaction with computational optimization and sustainability-oriented decision evaluation. It also emphasizes that operational

efficiency should not be separated from resource and sustainability considerations. A decision-support system should evaluate alternatives according to multiple project objectives rather than optimize a single operational metric.

Future research should empirically test the framework using real construction datasets and controlled project scenarios. Particular attention should be given to LLM reliability, retrieval accuracy, explainability, optimization performance, interoperability, sustainability measurement, and human acceptance. Comparative experiments could evaluate LLM-supported decision-making against conventional digital decision-support systems. Further work could also investigate how evolutionary optimization can be coupled with LLM-generated candidate solutions without compromising transparency or engineering control.

In conclusion, LLM-driven decision support represents a promising direction for sustainable construction management, but its value will depend on the integration of language intelligence with validated computational models, explicit optimization criteria, high-quality operational data, and responsible human oversight. The reviewed computational literature provides a foundation for this integration, while empirical construction-specific research remains necessary to establish its actual operational and sustainability benefits.

REFERENCES

1. A. Bakhshi, S. Chalup and N. Noman, "Fast evolution of CNN architecture for image classification," in *Natural Computing Series*, pp. 209–229, 2020.
2. A. Darwish, A. E. Hassanien and S. Das, "A survey of the swarm and evolutionary computing approaches for deep learning," *Artificial Intelligence Review*, vol. 53, pp. 1767–1812, 2019.
3. O. E. David and I. Greental, "Genetic algorithms for evolving deep neural networks," in *Proc. of the Companion Publication of the 2014 Annual Conf. on Genetic and Evolutionary Computation*, New York, USA, pp. 1451–1452, 2014.
4. I. Goodfellow, Y. Bengio, A. Courville and A. Bengio, *Deep Learning*, vol. 1, Cambridge: MIT Press, 2016.
5. M. Matsugu, K. Mori, Y. Mitari and Y. Kaneda, "Subject-independent facial expression recognition with robust face detection using a convolutional neural network," *Neural Network*, vol. 16, no. 5, pp. 555–559, 2003.
6. E. Selvi, M. Adimoolam, G. Karthi, K. Thinakaran, N. M. Balamurugan et al., "Suspicious actions detection system using enhanced CNN and surveillance video," *Electronics*, vol. 11, no. 24, pp. 1–20, 2022.
7. M. Suganuma, S. Shirakawa and T. Nagao, "A genetic

programming approach to designing convolutional neural network architectures,” in Proc. of the Genetic and Evolutionary Computation Conf., New York, USA, pp. 497–504, 2017.

8. Y. Sun, B. Xue, M. Zhang and G. G. Yen, “Evolving deep convolutional neural networks for image classification,” *IEEE Transactions on Evolutionary Computation*, vol. 24, no. 2, pp. 394–407, 2020.
9. F. P. Such, V. Madhavan, E. Conti, J. Lehman, K. O. Stanley et al., “Deep neuro-evolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning,” *arXiv:1712.06567*, 2017.
10. M. Azab, “Design approach and performance analysis of trap filter for three-phase PV grid integration systems using evolutionary search algorithms,” *Journal of King Saud University—Engineering Sciences*, vol. 33, no. 7, pp. 491–506, 2021.
11. 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.
12. Philip, P. G. (2026). Artificial Intelligence–Driven Intelligent Project Management: an integrated framework for planning, scheduling and control in engineering and construction projects. *Journal of Engineering and Artificial Intelligence*, 02(02), 01–09. <https://doi.org/10.64142/jeai.2.2.50>