COLOMBO SCIENTIFIC PUBLISHING

INTERNATIONAL JOURNAL OF COMPUTER SCIENCE & INFORMATION SYSTEM
E-ISSN: 2536-7919
P-ISSN: 2536-7900

https://doi.org/10.55640/ijcsis/Volume11Issue08-07 PAGE NO: 68-76

Robotics and Large Language Models for Digital Transformation in Sustainable Construction

Nethmi Fernando
Department of Robotics and
Automation, Centre for Intelligent Systems, Sri Lanka

ARTICLE INFO

Article history:
Submission: June,28 2026
Accepted: July 23, 2026
Published: August 15, 2026
VOLUME: Vol.11 Issue 08 2026

Keywords:
Robotics; Large Language Models; Digital Transformation; Sustainable Construction; Computer Vision; Intelligent Inspection; Machine Learning; Human-Robot Collaboration; Construction Automation; Cyber-Physical Systems.

ABSTRACT

The digital transformation of construction is increasingly shaped by the convergence of intelligent perception, automated physical systems, machine learning, and language-based artificial intelligence. This research and review paper examines how robotics and large language models (LLMs) can be conceptually integrated into sustainable construction operations through a layered digital transformation framework. Because the supplied literature primarily addresses computer vision, automated defect detection, machine learning, image classification, and human fatigue rather than construction-specific LLM applications, the paper develops a theoretically grounded synthesis that transfers these capabilities into construction contexts without attributing unsupported findings to the cited studies. The proposed framework connects four functional layers: visual perception and condition assessment, intelligent interpretation and reasoning, robotic execution, and sustainability-oriented operational feedback. Computer-vision research demonstrates the feasibility of automated visual recognition and defect classification, while machine-learning studies establish the importance of automated inspection and pattern recognition (Ghorai et al., 2012; Krummenacher et al., 2017). Large-scale image-recognition research further provides a foundation for scalable perception systems (Deng et al., 2009; Russakovsky et al., 2015). The paper argues that LLMs can function as an orchestration and knowledge-interface layer between human operators, perception systems, and robotic platforms, while robotics provides the physical execution capability required to transform digital decisions into construction actions. The resulting framework emphasizes interoperability, human oversight, safety, explainability, data quality, and sustainability metrics. Findings indicate that the strongest transformation potential arises not from replacing individual construction functions, but from integrating perception, reasoning, and physical action into a closed-loop cyber-physical workflow. The study contributes a research framework for future empirical validation of robotics-LLM systems in sustainable construction.

INTRODUCTION

1.1 Problem Statement

Construction remains a complex operational environment in which physical work, inspection, planning, resource allocation, safety management, and quality control are strongly interdependent. Unlike purely digital industries, construction requires intelligent systems to interpret heterogeneous information and subsequently interact with physical environments. Consequently, digital transformation cannot be understood only as the digitization of documents or the introduction of isolated software tools. It requires a coordinated transition from data collection to machine-supported interpretation and, ultimately, intelligent physical action.

Recent developments in machine learning and automated visual inspection demonstrate the technical feasibility of extracting meaningful information from complex physical environments. Automated defect-detection research has shown how image-based methods can identify anomalies in industrial materials and products (Ghorai et al., 2012). Similar work has explored machine-learning approaches for detecting wheel defects, illustrating the potential for algorithmic inspection in physically consequential environments (Krummenacher et al., 2017). These developments are relevant to construction because buildings, infrastructure components, equipment, and construction materials similarly generate visual and condition-related data that can be assessed automatically.

Large-scale visual-recognition research provides another important foundation. ImageNet established a benchmark-oriented paradigm for training and evaluating image-recognition systems at substantial scale (Deng et al., 2009), while the ImageNet Large Scale Visual Recognition Challenge subsequently demonstrated the importance of standardized datasets and systematic evaluation in computer vision (Russakovsky et al., 2015). Deep convolutional architectures further demonstrated how learned representations could substantially improve visual classification performance (Krizhevsky et al., 2017). Within construction, such capabilities can support automated site inspection, material recognition, progress monitoring, and condition assessment.

However, perception alone does not constitute digital transformation. A construction system must also interpret observations in relation to project objectives, communicate recommendations, coordinate tasks, and potentially initiate physical operations. This creates a conceptual role for LLMs. Rather than treating an LLM as an autonomous decision-maker, this study positions it as a language-based reasoning and orchestration layer that can translate human requirements and machine-generated observations into structured operational instructions. Robotics can then provide the physical interface through which validated instructions are executed.

1.2 Research Relevance

The integration of robotics and LLMs is particularly relevant to sustainable construction because sustainability requires simultaneous consideration of resource consumption, productivity, quality, waste, worker well-being, and lifecycle performance. An automated inspection system that detects a material defect may reduce rework, but its broader sustainability value depends on whether the information is connected to an appropriate intervention. Similarly, a robotic system may improve productivity but fail to produce sustainable outcomes if it increases material consumption or energy demand.

Research concerning human fatigue reinforces the importance of considering human factors within technological transformation. Mascord and Heath examined behavioral and physiological indices of fatigue during visual tracking, while Seghers, Jochem, and Spaepen investigated posture, muscle activity, and fatigue in prolonged visual-display work (Mascord and Heath, 1992; Seghers et al., 2003). Although these studies are not construction studies, they provide a useful theoretical reminder that technological systems should be designed around human performance and not solely around machine efficiency.

1.3 Objectives

This paper has four principal objectives. First, it synthesizes the supplied literature to identify technical foundations relevant to intelligent construction automation. Second, it develops an integrated conceptual framework connecting computer vision, machine learning, LLM-based reasoning, and robotics. Third, it evaluates how such integration could contribute to sustainable construction operations. Fourth, it identifies limitations and research requirements for empirical implementation.

1.4 Scope and Significance

The study is conceptual and review-oriented. It does not claim that the supplied references experimentally validate LLMs or robotics in construction. Instead, it uses the demonstrated capabilities in automated inspection, image recognition, machine learning, and human-performance research as technical foundations for constructing a research framework. This distinction is important because a rigorous literature synthesis should not extend empirical findings beyond their original domains without qualification.

The central proposition is that robotics and LLMs can create greater value when deployed as complementary components. Computer vision supplies environmental awareness, LLMs provide a flexible interface for contextual interpretation and coordination, and robotics supplies physical execution. The proposed integration is therefore represented as perception → interpretation → decision support → validated action → feedback.

LITERATURE REVIEW

2.1 Automated Visual Perception as a Digital Foundation

The development of automated visual inspection represents one of the most relevant foundations for intelligent construction systems. Ghorai et al. investigated automatic defect detection on hot-rolled flat steel products, demonstrating the application of computational methods to industrial visual-quality assessment (Ghorai et al., 2012). Feng, Gao, and Luo introduced the X-SDD benchmark for hot-rolled steel strip surface-defect detection, emphasizing the importance of benchmark datasets for evaluating defect-recognition systems (Feng et al., 2021). Kuo et al. examined the integration of image processing and classification technology for automated polarizing-film defect inspection (Kuo et al., 2018).

These studies collectively suggest that automated inspection requires more than image acquisition. Effective systems require image representation, classification mechanisms, suitable datasets, and evaluation procedures. In construction, this principle can be transferred to applications such as surface-defect identification, component inspection, material classification, and progress verification. The transfer, however, should be treated as a methodological proposition rather than as direct evidence of construction performance.

Kapsalas et al. examined optical inspection for quantifying decay on stone surfaces, providing another relevant example of visual assessment applied to physical materials (Kapsalas et al., 2007). Such work is conceptually important for sustainable construction because condition assessment can influence maintenance and intervention decisions. If deterioration is detected earlier and characterized more consistently, maintenance strategies may potentially become more targeted, although such sustainability benefits require direct empirical validation.

2.2 Machine Learning and Scalable Recognition

Deng et al. established ImageNet as a large-scale hierarchical image database, providing a major foundation for data-driven visual recognition (Deng et al., 2009). Russakovsky et al. subsequently documented the ImageNet Large Scale Visual Recognition Challenge and its role in systematic evaluation of visual-recognition methods (Russakovsky et al., 2015). Krizhevsky, Sutskever, and Hinton demonstrated the effectiveness of deep convolutional neural networks for large-scale image classification (Krizhevsky et al., 2017).

From a construction perspective, these contributions establish three theoretical requirements for scalable intelligent perception: sufficiently representative data, standardized evaluation, and effective learned representations. Construction environments are highly variable, however. Lighting, weather, occlusion, geometry, material diversity, equipment movement, and temporary site conditions can reduce the direct transferability of models developed in controlled industrial environments.

The major research gap is therefore not simply whether computer vision can recognize objects or defects. The more consequential question is whether recognition systems can remain reliable under dynamic construction conditions and whether their outputs can be converted into actionable operational knowledge.

2.3 Intelligent Inspection and Physical Systems

Krummenacher et al. investigated machine-learning-based wheel-defect detection, demonstrating how machine learning can be applied to physically relevant defect identification (Krummenacher et al., 2017). The significance for construction lies in the general architecture of automated condition monitoring:

physical assets generate observable signals, computational models identify abnormal conditions, and inspection becomes partially automated.

Such an architecture can be extended conceptually to robotic construction inspection. A mobile robot or robotic platform could collect images or sensor information, while computer-vision models identify anomalies. An LLM could then organize the findings into a human-readable inspection report, connect observations with predefined operational rules, and support task prioritization. Importantly, the LLM should not be assumed to replace deterministic safety or engineering controls.

2.4 Human Factors and Human-Robot Collaboration

The literature on fatigue provides a complementary theoretical dimension. Mascord and Heath investigated behavioral and physiological indices of fatigue, whereas Seghers et al. examined posture, muscle activity, and fatigue under prolonged visual-display conditions (Mascord and Heath, 1992; Seghers et al., 2003). Their relevance to digital construction is indirect but important: technological transformation can alter human cognitive and physical workloads.

An intelligent construction system should therefore optimize the human-machine system rather than maximize automation independently. Robots can potentially perform repetitive or hazardous inspection tasks, while workers retain responsibility for contextual judgment and exceptional situations. LLMs may facilitate this interaction by providing natural-language interfaces through which workers can query system status, request explanations, or receive task summaries.

2.5 Research Gap and Theoretical Positioning

The supplied references collectively provide evidence for automated visual inspection, machine learning, image classification, scalable recognition, and human-performance assessment. They do not, however, establish an empirical framework integrating LLMs and robotics for sustainable construction. This gap motivates the present conceptual framework.

Theoretically, the proposed model combines three technological functions. Perception is represented by computer vision and machine learning. Cognitive orchestration is represented by LLM-based language processing and task coordination. Physical execution is represented by robotics. Sustainability operates as a cross-cutting evaluation dimension rather than as an isolated technological component.

METHODOLOGY

3.1 Research Design

The study uses a structured conceptual synthesis methodology. The provided references were examined according to their demonstrated technical contribution rather than their disciplinary labels. Each reference was mapped to one or more functional capabilities: visual inspection, image classification, machine learning, benchmark development, condition assessment, or human-performance evaluation.

The methodological objective is not to statistically aggregate findings because the references are heterogeneous in purpose, domain, and methodology. Instead, the study develops a transferable systems framework. This approach is appropriate when the research problem concerns technological integration across previously separated technical functions.

3.2 Proposed Digital Transformation Architecture

The proposed framework consists of five interconnected layers.

Layer 1: Physical Construction Environment.

The construction site contains workers, materials, equipment, structures, robots, and changing environmental conditions. The system begins by treating the physical site as a dynamic source of data rather than a static workspace.

Layer 2: Perception and Data Acquisition.

Robotic platforms and sensing systems collect images and other operational observations. Computer-vision techniques then transform raw observations into structured information. The literature on industrial defect detection and image classification provides the methodological basis for this layer (Ghorai et al., 2012; Kuo et al., 2018).

Layer 3: Intelligent Interpretation.

The resulting observations are passed to analytical models and an LLM-based interface. The LLM can contextualize recognized events, summarize inspection results, compare observations against project-defined requirements, and convert technical outputs into understandable instructions. Its function should be constrained by validated data sources and predefined operational policies.

Layer 4: Robotic Execution.

Where an action is authorized, the robotic system performs the corresponding physical operation. Potential functions include inspection, material movement, repetitive measurement, surface assessment, or data collection. Human approval should remain mandatory for high-risk or safety-critical actions.

Layer 5: Feedback and Sustainability Monitoring.

The outcome of each action is returned to the digital system. This feedback can be used to determine whether the intervention achieved its intended result. Sustainability indicators can include material waste, unnecessary rework, energy consumption, equipment utilization, and worker exposure to repetitive or hazardous activities.

3.3 LLM as an Orchestration Layer

The LLM is positioned neither as a conventional visual classifier nor as an unrestricted autonomous controller. Instead, it acts as an orchestration layer. For example, a robotic inspection unit may identify a surface anomaly. A vision model provides the detected location and classification. The LLM can transform this structured information into an inspection summary, identify the appropriate workflow, and request additional inspection data if required.

This architecture separates perception from language-based reasoning. The distinction is essential because an LLM-generated interpretation may be linguistically convincing while still being technically incorrect. Therefore, numerical measurements, safety thresholds, engineering specifications, and robot-control commands should originate from validated systems rather than from unconstrained language generation.

3.4 Robotics as the Physical Execution Layer

Robotics transforms digital recommendations into physical operations. Consider a hypothetical construction-site inspection workflow. A mobile robot traverses a designated zone and captures images. A computer-vision system identifies potential surface anomalies. The LLM receives structured outputs, organizes them according to project priorities, and generates an inspection summary. A human supervisor validates the proposed intervention. The robot subsequently revisits the location to collect additional evidence or perform an authorized inspection task.

This workflow illustrates why robotics and LLMs should not be viewed as substitutes. The robot has physical mobility and repeatability but limited contextual language capability. The LLM has flexible language-based reasoning and communication capabilities but lacks reliable physical embodiment. Their integration therefore creates a complementary system.

3.5 Sustainability Decision Model

Sustainability is incorporated through a multi-dimensional decision function:

Sustainable Action Value = Quality Benefit + Productivity Benefit + Safety Benefit − Resource Cost − Energy Cost − Waste Cost.

This equation is conceptual rather than a validated quantitative metric. Its purpose is to establish that an automated action should not be evaluated exclusively according to speed or accuracy. For example, a robotic inspection that reduces inspection time but consumes excessive energy may not provide a net sustainability benefit.

Similarly, automated defect recognition can potentially reduce unnecessary replacement by distinguishing localized defects from broader deterioration. Research on optical inspection and surface-condition quantification provides methodological support for this logic (Kapsalas et al., 2007). Nevertheless, construction-specific experiments are required to determine actual environmental benefits.

3.6 Human Oversight and Safety

Human oversight is a central methodological requirement. The fatigue literature indicates that technological environments affect human behavioral and physiological conditions (Mascord and Heath, 1992; Seghers et al., 2003). In the proposed system, automation should therefore reduce unnecessary repetitive workload without eliminating human authority.

Three levels of autonomy are proposed conceptually. At the first level, the system recommends actions while humans execute them. At the second level, robots execute low-risk, pre-authorized tasks after automated validation. At the third level, autonomous operation is permitted only within narrowly defined operational boundaries. Safety-critical decisions should remain subject to human authorization and deterministic control rules.

3.7 Evaluation Framework

Evaluation should occur across five dimensions: perception accuracy, reasoning reliability, robotic execution performance, human interaction quality, and sustainability impact. Perception can be evaluated through precision, recall, classification accuracy, and false-positive rates. Robotic performance can be assessed through task completion rate, localization accuracy, downtime, and intervention frequency. LLM performance requires evaluation of factual consistency, instruction adherence, explainability, and resistance to unsupported conclusions.

Sustainability should be measured using project-specific indicators rather than generic claims. Relevant metrics could include material waste avoided, rework hours reduced, energy consumed per inspection, equipment utilization, and worker exposure to repetitive tasks.

RESULTS

The conceptual synthesis produces five principal findings.

First, automated perception is the most mature technological foundation within the supplied literature. The studies on industrial defect detection, optical inspection, image processing, and machine learning demonstrate that physical conditions can be transformed into machine-interpretable information (Ghorai et al., 2012; Kapsalas et al., 2007; Kuo et al., 2018). The ImageNet-related studies further show the importance of large-scale datasets, benchmark evaluation, and learned visual representations (Deng et al., 2009; Russakovsky et al., 2015). For sustainable construction, this suggests that reliable digital transformation begins with dependable environmental data.

Second, the principal integration gap exists between perception and action. Computer vision can detect an anomaly, but detection alone does not determine what should happen next. A construction manager must interpret the significance of the anomaly, determine its priority, coordinate resources, and decide whether intervention is necessary. The proposed LLM layer addresses this coordination problem by translating structured analytical outputs into operationally meaningful language. The LLM therefore functions primarily as an interface and orchestration mechanism rather than as a replacement for engineering models.

Third, robotics provides the missing physical-action capability. A digital recommendation has limited operational value unless it can be converted into an inspection, measurement, material-handling, or maintenance action. Robotics can provide repeatable physical execution while simultaneously acting as a mobile data-collection platform. The machine-learning inspection literature supports the broader proposition that physical asset assessment can be partially automated (Krummenacher et al., 2017).

Fourth, sustainability benefits depend on system-level integration rather than individual technological performance. High classification accuracy does not automatically produce lower waste, and high robotic productivity does not necessarily imply lower environmental impact. Sustainable value emerges when improved inspection, reduced rework, optimized resource use, and safer work allocation occur together. The proposed framework consequently treats sustainability as an evaluation layer spanning the entire digital workflow.

Fifth, human oversight remains a structural requirement. The human-factor references demonstrate why technological systems should account for human performance and workload (Mascord and Heath, 1992; Seghers et al., 2003). The framework therefore favors controlled autonomy in which robots perform repeatable tasks, LLMs facilitate communication and coordination, and humans retain authority over ambiguous, safety-critical, or high-consequence decisions.

Overall, the findings support the proposition that robotics and LLMs have greater transformation potential as an integrated cyber-physical system than as isolated technologies. The framework extends the automated-inspection paradigm toward a closed-loop architecture in which physical observations are converted into information, information is interpreted through contextual reasoning, authorized actions are executed physically, and outcomes are fed back into the digital system. This represents the central contribution of the study.

DISCUSSION

The findings indicate that digital transformation in sustainable construction should be conceptualized as an integration problem rather than a technology-adoption problem. Existing inspection research demonstrates that automated perception can identify physical abnormalities, while large-scale computer-vision research demonstrates how data and learned representations support scalable recognition (Deng et al., 2009; Krizhevsky et al., 2017; Russakovsky et al., 2015). The unresolved challenge is how such recognition becomes operationally meaningful within a complex construction workflow.

The proposed LLM layer addresses this challenge by providing contextual translation between machine outputs and human-oriented project decisions. However, this capability introduces a significant trade-off. Natural-language flexibility increases accessibility but may also increase ambiguity. An LLM can summarize and coordinate information effectively, yet language fluency cannot be equated with engineering correctness. Consequently, the framework requires strict separation between generated language and authoritative technical data. The LLM should interpret validated information rather than independently establish safety limits, structural requirements, or robotic control parameters.

The second major implication concerns robotics. Automated inspection studies show the value of machine-supported identification of defects and conditions (Ghorai et al., 2012; Kapsalas et al., 2007). Robotics extends this logic by enabling the sensing platform to move through the environment. This creates a potentially powerful feedback loop: robots collect data, perception systems identify conditions, LLMs coordinate interpretation, and robots return to the relevant location for verification. Nevertheless, dynamic construction sites introduce occlusion, temporary structures, changing terrain, dust, weather, and human movement. These factors create conditions substantially different from controlled inspection environments. Therefore, transfer from industrial inspection to construction should be empirically tested rather than assumed.

A third implication concerns sustainability. The literature supplied does not directly measure the environmental performance of robotics-LLM systems. Therefore, claims of sustainability must remain conditional. The strongest theoretical pathway is indirect: better inspection can reduce avoidable rework; improved coordination can reduce unnecessary movement and waiting; robotic inspection can potentially reduce exposure to repetitive tasks; and more accurate condition assessment can support targeted intervention. Yet these benefits may be offset by energy consumption, equipment production, maintenance, computational requirements, or inefficient automation. Sustainability assessment must therefore use lifecycle and operational metrics.

Human factors constitute another critical consideration. The fatigue studies indicate that technological work environments can influence physical and behavioral performance (Mascord and Heath, 1992; Seghers et al., 2003). In construction, robotics and LLMs should consequently be evaluated according to whether they improve human work quality rather than simply reduce human involvement. A poorly designed interface could increase cognitive burden by presenting excessive alerts or unclear recommendations. An effective interface should instead summarize relevant information, explain the basis of recommendations, and allow rapid human intervention.

The framework also reveals an important organizational limitation: digital integration requires standardized data structures and interoperable systems. A robot, vision model, project-management system, and LLM cannot operate as an effective digital ecosystem if their data representations are incompatible. The benchmark orientation evident in large-scale visual recognition research reinforces the importance of standardized evaluation and data organization (Deng et al., 2009; Russakovsky et al., 2015).

Finally, the study's conceptual nature limits empirical generalization. The proposed framework has not been validated through controlled construction-site experiments, longitudinal projects, or lifecycle sustainability assessments. Future research should therefore develop construction-specific datasets, test multimodal LLM interfaces, evaluate robotic navigation and inspection under realistic site conditions, and quantify sustainability outcomes. Particular attention should be given to failure cases because the reliability of an autonomous system is determined not only by average performance but also by its behavior under uncertainty.

CONCLUSION

This research and review paper developed a conceptual framework for integrating robotics and large language models into sustainable construction digital transformation. The analysis of the supplied literature demonstrates that automated visual inspection, machine learning, image classification, scalable recognition, and human-performance assessment provide important technological and theoretical foundations. Industrial defect-detection research establishes the feasibility of automated physical-condition assessment, while ImageNet-related research demonstrates the importance of large datasets, benchmark evaluation, and learned visual representations (Deng et al., 2009; Russakovsky et al., 2015). Machine-learning-based inspection further supports the proposition that physical assets can be monitored through computational methods (Krummenacher et al., 2017).

The central contribution is the proposed five-layer architecture comprising the physical construction environment, perception and data acquisition, intelligent interpretation, robotic execution, and sustainability-oriented feedback. Within this architecture, robotics provides physical embodiment, computer vision provides environmental perception, and LLMs provide a flexible language-based coordination and interpretation layer. The framework therefore avoids positioning any single technology as a complete solution.

The study also emphasizes that sustainability cannot be inferred merely from automation. A sustainable digital construction system must demonstrate measurable reductions in waste, rework, unnecessary resource consumption, or harmful human exposure while maintaining quality and safety. Human oversight remains essential, particularly where uncertainty or high-consequence decisions are involved. The human-factor literature reinforces the need to design technological systems around human performance rather than treating workers as peripheral components (Mascord and Heath, 1992; Seghers et al., 2003).

Ultimately, the most significant opportunity lies in closing the loop between observation, interpretation, action, and feedback. Robotics and LLMs should therefore be investigated not as independent digital technologies but as complementary components of an intelligent cyber-physical construction ecosystem. Such integration could provide a pathway toward construction operations that are more data-driven, responsive, resource-conscious, and human-centered, provided that technical reliability, safety, interoperability, and sustainability are demonstrated through rigorous empirical validation.

REFERENCES

1. JAI, 2022, vol. 4, no. 4, 257.
2. J. Deng, W. Dong, R. Socher, L. Li, K. Li et al., "Imagenet: A large-scale hierarchical image database," in 2009 IEEE Conf. on Computer Vision and Pattern Recognition, Miami, FL, USA, pp. 248–255, 2009.
3. S. Ghorai, A. Mukherjee, M. Gangadaran and P. K. Dutta, "Automatic defect detection on hot-rolled flat steel products," IEEE Transactions, vol. 62, no. 3, pp. 612–621, 2012.
4. X. Feng, X. Gao and L. Luo, "X-SDD: A new benchmark for hot rolled steel strip surface defects detection," Symmetry, vol. 13, no. 4, pp. 706, 2021.
5. P. Kapsalas, P. Maravelaki, M. Zervakis, E. Delegou and A. Moropoulou, "Optical inspection for quantification of decay on stone surfaces," NDT & E International, vol. 40, no. 1, pp. 2–11, 2007.
6. G. Krummenacher, C. Ong, S. Koller, S. Kobayashi and J. Buhmann, "Wheel defect detection with machine learning," IEEE Transactions, vol. 19, no. 4, pp. 1176–1187, 2017.
7. Krizhevsky, I. Sutskever and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," Communication ACM, vol. 60, no. 6, pp. 84–90, 2017.
8. Kuo, C. Lai, C. Kao and C. Chiu, "Integrating image processing and classification technology into automated polarizing film defect inspection," Optical Lasers Engineering, vol. 104, pp. 204–219, 2018.
9. J. Mascord and R. A. Heath, "Behavioral and physiological indices of fatigue in a visual tracking task," Journal of Safety Research, vol. 23, no. 1, pp. 19–25, 1992.
10. O. Russakovsky et al., "Imagenet large scale visual recognition challenge," International Journal of Computer Vision, Santiago, Chile, vol. 115, pp. 211–252, 2015.
11. J. Seghers, A. Jochem and A. Spaepen, "Posture, muscle activity and muscle fatigue in prolonged VDT work at different screen height settings," Ergonomics, vol. 46, no. 7, pp. 714–730, 2003.
12. Y. Xie, Y. Ye, J. Zhang, L. Liu and L. Liu, "A Physics-based defects model and inspection algorithm for automatic visual inspection," Optical Lasers Engineering, vol. 52, pp. 218–223, 2014.
13. 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.
14. 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