COLOMBO SCIENTIFIC PUBLISHING

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

https://doi.org/10.55640/ijcsis/Volume11Issue08-03 PAGE NO: 111-118

Scale Solve: A Generative LLM Framework for Large-Scale Constraint Reasoning and Decision Optimization

Faisal Al-Harbi
Department of Artificial Intelligence and Data
Analytics Institute of Advanced Technology, Riyadh, Saudi Arabia

ARTICLE INFO

Article history:
Submission: June,27 2026
Accepted: July 18, 2026
Published: August 19, 2026
VOLUME: Vol.11 Issue 08 2026

Keywords:
Supervised Machine Learning, Railway Bridge Safety, Collision Detection, Structural Health Monitoring, Predictive Analytics, Intelligent Transportation Systems, Wireless Sensor Networks, Infrastructure Monitoring, Bridge Damage Detection, Safety Optimization.

ABSTRACT

Large-scale decision problems increasingly involve heterogeneous data, interacting constraints, uncertain operating conditions, and multiple competing objectives. Conventional optimization systems can provide mathematically rigorous solutions, but their effectiveness may depend on structured problem representations that are difficult to construct when operational knowledge is expressed in natural language, visual observations, or domain-specific descriptions. This paper proposes ScaleSolve, a generative large language model (LLM) framework designed to connect natural-language problem specifications with constraint reasoning, optimization formulation, solution evaluation, and decision interpretation. The framework is conceptually positioned around five functional layers: problem interpretation, constraint extraction, optimization modeling, candidate-solution evaluation, and decision explanation. The literature supplied for this study primarily concerns food processing, drying kinetics, computer vision, image analysis, preprocessing, and quality optimization. These studies provide an important application foundation because drying and postharvest processing involve interacting physical, quality, process, and resource constraints. The proposed framework therefore demonstrates how heterogeneous observations can be translated into structured decision variables and constraints. The conceptual analysis indicates that LLM-based reasoning can improve accessibility and adaptability of large-scale optimization workflows, while numerical validation, constraint verification, and domain-specific objective functions remain necessary to prevent plausible but infeasible decisions. The framework contributes a research-oriented architecture for integrating generative reasoning with optimization rather than treating an LLM as an autonomous optimization engine.

INTRODUCTION

1.1 Background

Large-scale decision optimization concerns the selection of actions from potentially enormous combinations of alternatives while satisfying operational, physical, economic, and quality constraints. In practical environments, the difficulty is not limited to solving an optimization equation. Decision-makers must first identify relevant variables, interpret observations, define constraints, establish objective functions, and determine which information is reliable enough to enter a formal model. This becomes particularly challenging when information originates from heterogeneous sources such as sensor measurements, images, process descriptions, experimental observations, and expert knowledge.

Food-processing applications provide a useful example of this complexity. Research on carrots and other agricultural products has examined phytochemical characteristics, drying behavior, pretreatment effects, image-based quality assessment, and process optimization (Ahmad et al., 2019; Mudahar et al., 1989; Wang et al., 2021). These factors do not operate independently. A change in drying conditions can influence moisture removal, color, structural characteristics, phytochemical retention, and antioxidant capacity. Consequently, selecting a process configuration becomes a constrained multi-objective decision problem rather than a single-variable prediction task.

Computer vision further increases the dimensionality of such decision environments. Jahanbakhshi and Kheiralipour (2020) examined image-processing and discriminant-analysis approaches for postharvest carrot processing, while Sharma et al. (2018) investigated computer vision for monitoring physical quality during solar drying. Kakani et al. (2020) identified the broader importance of computer vision and artificial intelligence in food-industry applications. Such studies indicate that visual observations can become decision-relevant information rather than merely descriptive measurements.

Recent work on combinatorial LLM frameworks similarly motivates the integration of generative reasoning with scalable constraint-solving workflows. Ramamurthy, Bellamkonda, and Amanmadov (2026) provide a relevant conceptual basis for examining LLM-supported combinatorial constraint processing. The central premise of ScaleSolve is therefore not that an LLM replaces mathematical optimization, but that it acts as an intelligent interface between unstructured information and formally verifiable optimization procedures.

1.2 Problem Statement

Traditional optimization workflows generally assume that the problem has already been formalized. Variables, constraints, objectives, and feasible domains must be specified before a solver can operate. However, real-world decision environments frequently communicate requirements through natural language, experimental reports, images, or partially structured records. The gap between these representations can create substantial modeling effort and introduce interpretation errors.

The research problem addressed by this paper is therefore the design of a generative LLM architecture capable of translating heterogeneous operational information into structured constraint-reasoning and optimization workflows while maintaining feasibility and interpretability.

1.3 Research Relevance

The relevance of this problem lies in the increasing convergence of AI-based perception, natural-language reasoning, and optimization. The provided literature demonstrates that computer vision can monitor physical quality, pretreatment can alter rehydration behavior, and drying conditions can affect multiple quality attributes (Dey et al., 2023; Sharma et al., 2018; Wang et al., 2021). ScaleSolve extends this application logic by considering how such information could be transformed into machine-interpretable decision constraints.

1.4 Objectives

The primary objective is to propose a scalable architecture for LLM-assisted constraint reasoning and decision optimization. Secondary objectives are to establish a systematic transformation from natural-language requirements to formal constraints, integrate heterogeneous observations into optimization models, define mechanisms for constraint validation, and establish a human-interpretable decision layer.

1.5 Scope and Significance

The framework is conceptual and research-oriented rather than an experimentally validated production system. Food processing and drying are used as an illustrative application domain because the supplied references contain diverse evidence concerning process parameters, quality measurements, visual information, and optimization. The architecture can potentially be generalized to other domains involving large-scale constrained decision-making.

LITERATURE REVIEW

2.1 Process Optimization and Constraint Complexity

Early optimization research in food dehydration illustrates the fundamental relationship between process variables and measurable outcomes. Mudahar et al. (1989) applied response surface methodology to carrot dehydration, demonstrating the relevance of systematic parameter optimization in drying operations. Such approaches establish an important theoretical foundation for ScaleSolve because they represent process optimization as a structured relationship between controllable variables and desired outcomes.

Later research expands this perspective by examining the effects of preprocessing and drying configurations. Dhalsamant et al. (2017) investigated pretreatment effects on rehydration, color, and nanoindentation properties of potato cylinders dried using a mixed-mode solar dryer. Dey et al. (2023) similarly examined different pretreatments and their effects on rehydration kinetics of solar- and hot-air-dried Fuji apple slices. These findings demonstrate that preprocessing decisions can propagate through subsequent process outcomes.

2.2 Quality, Nutritional Characteristics, and Multi-Objective Reasoning

Ahmad et al. (2019) reviewed phytochemicals in Daucus carota and their associated health benefits. Wang et al. (2021) investigated vacuum-steam pulsed blanching and its effects on drying kinetics, color, phytochemical contents, antioxidant capacity, and structural characteristics of carrots. Together, these studies illustrate why a decision framework must account for multiple objectives simultaneously.

A process configuration that maximizes drying speed may not maximize nutritional or sensory quality. Therefore, an optimization framework must distinguish between hard constraints and soft objectives. For example, a minimum acceptable quality level can be modeled as a constraint, while drying time can be treated as an objective to minimize. This distinction is central to constraint reasoning.

2.3 Computer Vision and AI-Based Monitoring

Jahanbakhshi and Kheiralipour (2020) demonstrated the applicability of image processing and discriminant analysis to postharvest carrot processing. Sharma et al. (2018) investigated computer vision for monitoring turmeric quality during direct solar drying, while Sharma et al. (2021) analyzed drying characteristics and quality under hot-air and direct-solar drying conditions. Kakani et al. (2020) positioned computer vision and artificial intelligence as important technologies in the food industry.

These studies establish the feasibility of converting visual characteristics into decision-relevant information. Within ScaleSolve, such information can be interpreted as observations that contribute to constraint evaluation or objective estimation rather than being treated as isolated classification outputs.

2.4 Research Gap and Theoretical Positioning

The principal gap identified from the supplied literature is the absence of an integrated reasoning layer connecting heterogeneous observations, natural-language requirements, formal constraints, optimization, and human-readable decisions. Existing studies emphasize individual components such as process optimization, drying behavior, computer vision, or quality assessment. ScaleSolve proposes an architecture that conceptually connects these components.

The framework is consistent with the direction represented by Ramamurthy et al. (2026), where LLM-based combinatorial reasoning is positioned as a mechanism for addressing scalability constraints. The proposed contribution is to extend that reasoning perspective toward heterogeneous, domain-oriented decision workflows rather than treating combinatorial reasoning as an isolated language-model task.

METHODOLOGY

3.1 Research Design

This study adopts a conceptual framework-development methodology. The methodology synthesizes the functional characteristics represented across the supplied literature and maps them into an LLM-assisted decision architecture. The resulting framework is not presented as empirical proof of optimization performance; instead, it defines a testable architecture and establishes the mechanisms required for future implementation.

3.2 ScaleSolve Architecture

ScaleSolve consists of five interconnected layers.

Layer 1: Problem Interpretation.

The LLM receives natural-language requirements, operational descriptions, process specifications, and contextual information. Its task is to identify decision entities, objectives, limitations, and dependencies. For example, a requirement such as reducing drying time while preserving acceptable color and nutritional characteristics can be decomposed into optimization objectives and quality constraints.

Layer 2: Constraint Extraction.

The extracted information is converted into explicit constraint candidates. Constraints may represent minimum quality, maximum processing time, allowable temperature ranges, equipment capacity, resource availability, or process dependencies. A key design principle is that generated constraints must remain distinguishable from assumptions. The system should therefore label each constraint according to its source and confidence.

Layer 3: Formal Optimization Modeling.

The structured representation is transformed into a mathematical or algorithmic optimization model. A generalized formulation can be expressed as:

$$
\min_{x} F(x)=\sum_{i=1}^{m} w_i f_i(x)
$$

subject to

$$
g_j(x)\leq b_j,\quad j=1,\ldots,p
$$

and

$$
x_i\in X_i.
$$

Here, (x) represents decision variables, ($f_i(x)$) represents objective functions, (w_i) represents objective weights, and ($g_j(x)\leq b_j$) represents operational constraints.

The LLM should not directly determine whether a generated solution is mathematically feasible. Instead, it should produce the structured model and transfer it to a deterministic validation or optimization mechanism. This separation is fundamental to ScaleSolve.

Layer 4: Candidate Generation and Verification.

Candidate solutions can be produced through optimization algorithms, combinatorial search, or hybrid procedures. The LLM can assist with strategy selection, constraint interpretation, decomposition, and explanation. Each candidate must subsequently pass deterministic constraint verification.

Layer 5: Decision Explanation.

The final layer translates solver outputs into human-readable decisions. The explanation should identify the selected configuration, objective trade-offs, active constraints, violated alternatives, and uncertainty. This enables decision-makers to understand not merely what solution was selected, but why it was selected.

3.3 Constraint Ontology

A central methodological component is the development of a constraint ontology. Constraints can be categorized into four groups: hard constraints, soft constraints, relational constraints, and contextual constraints.

Hard constraints define non-negotiable feasibility conditions. Soft constraints represent preferences that can be relaxed through penalties. Relational constraints define dependencies between variables, while contextual constraints represent conditions derived from observations or operating environments.

This classification is particularly useful for food-processing scenarios. Drying equipment capacity can be represented as a hard constraint, while minimizing drying duration can be an objective. Acceptable color quality can be modeled as a threshold constraint, while nutritional preservation may function as a competing objective.

3.4 Multimodal Information Integration

The framework is designed to accept structured and semi-structured observations from computer vision, process records, and textual descriptions. The literature on image processing and computer vision demonstrates that physical quality can be assessed computationally (Jahanbakhshi & Kheiralipour, 2020; Sharma et al., 2018).

ScaleSolve converts these observations into structured attributes. For example, an image-analysis component could estimate a quality indicator, after which the reasoning layer associates that indicator with a predefined constraint. This prevents the LLM from directly inventing visual measurements.

3.5 Scalability Mechanism

Scalability is addressed through hierarchical decomposition. A large optimization problem can be divided into subproblems according to process stages, resource groups, geographical units, or decision categories. The LLM can determine decomposition strategies while numerical components solve the resulting subproblems.

This design follows the principle that generative reasoning should reduce the complexity of model construction rather than assume responsibility for all combinatorial computation. Ramamurthy et al. (2026) similarly provide a useful basis for considering LLM-supported combinatorial constraint scalability.

3.6 Validation and Reliability

ScaleSolve requires three validation mechanisms: syntactic validation, semantic validation, and numerical validation. Syntactic validation checks whether generated constraints are formally well-structured. Semantic validation verifies whether the constraints correctly represent the original requirements. Numerical validation determines whether candidate solutions satisfy all mathematical constraints.

A further safeguard is counterfactual testing. The framework can ask whether a proposed decision remains feasible if a parameter changes. This is particularly relevant in drying applications where environmental or process conditions can alter outcomes (Wang et al., 2021).

RESULTS

The conceptual analysis produces five principal findings. First, the supplied literature demonstrates that operational decisions in food processing are inherently multi-dimensional. Drying kinetics, rehydration, color, structural properties, phytochemical content, antioxidant capacity, and processing conditions may interact rather than vary independently (Ahmad et al., 2019; Dhalsamant et al., 2017; Wang et al., 2021). Consequently, a scalable decision architecture must support multiple objectives and constraints simultaneously.

Second, the literature indicates that process optimization and intelligent monitoring represent complementary rather than competing capabilities. Response surface methodology provides a structured optimization perspective, while computer vision provides observational information that can inform quality assessment (Mudahar et al., 1989; Jahanbakhshi & Kheiralipour, 2020). ScaleSolve integrates these functions by placing an LLM between heterogeneous information sources and formal decision mechanisms.

Third, the proposed architecture identifies constraint extraction as a critical bottleneck. Natural-language requirements may contain implicit conditions that conventional optimization systems cannot directly interpret. The LLM can assist in identifying these conditions, but the framework requires subsequent deterministic verification. This finding supports the conceptual position that generative reasoning should function as a modeling and coordination layer rather than an uncontrolled numerical solver.

Fourth, hierarchical decomposition emerges as the principal scalability mechanism. Large-scale constraint problems can be separated into manageable subproblems, allowing optimization components to operate on smaller decision spaces while the LLM maintains semantic coordination between them. This approach is conceptually aligned with the scalability concerns associated with combinatorial LLM reasoning discussed by Ramamurthy et al. (2026).

Fifth, the framework suggests that explainability should be integrated into the optimization pipeline rather than added after solution generation. A useful decision explanation must connect the selected solution to its objectives, active constraints, and relevant observations. In practical applications, this can improve human oversight, especially when competing objectives produce non-obvious trade-offs.

The findings therefore position ScaleSolve as a hybrid architecture in which generative intelligence supports interpretation and coordination, while formal optimization and validation maintain computational reliability.

DISCUSSION

The findings have important theoretical implications for the relationship between generative AI and optimization. Conventional optimization begins with a formal mathematical representation, whereas LLM-based systems begin naturally with linguistic and heterogeneous information. ScaleSolve addresses this representational mismatch by assigning different responsibilities to different computational layers. The LLM interprets requirements and coordinates reasoning, while deterministic optimization mechanisms establish feasibility and numerical validity. This separation reduces the risk of treating fluent language generation as equivalent to mathematical correctness.

The literature on food drying illustrates the value of this separation. Mudahar et al. (1989) demonstrate the usefulness of systematic process optimization, whereas later studies reveal that pretreatment and drying configurations can affect multiple quality dimensions (Dey et al., 2023; Dhalsamant et al., 2017). Wang et al. (2021) further demonstrates the interconnected nature of drying kinetics, color, phytochemical content, antioxidant capacity, and structural characteristics. A generative framework can represent these relationships in natural language, but objective evaluation still requires formal definitions and measurable variables.

Computer vision provides another important implication. Studies by Jahanbakhshi and Kheiralipour (2020) and Sharma et al. (2018) indicate that visual information can support quality monitoring. However, the output of a vision model should not automatically become a hard constraint. Measurement uncertainty, classification error, and environmental variation may cause an apparently precise observation to be unreliable. ScaleSolve should consequently maintain provenance and confidence information for every observation.

A major trade-off concerns flexibility versus reliability. Greater LLM autonomy can reduce modeling effort but may increase the possibility of incorrect constraints or unsupported assumptions. Greater deterministic control improves reliability but reduces the flexibility that makes generative systems attractive. The proposed architecture resolves this tension through staged verification: the LLM proposes representations, while formal mechanisms validate them.

The approach also has limitations. First, the framework has not been empirically benchmarked against established optimization solvers. Second, the supplied literature does not directly provide experimental evidence about LLM-based optimization performance. Third, real-world deployment would require domain-specific ontologies, numerical datasets, validation protocols, and carefully defined objective functions. Fourth, computational scalability cannot be inferred solely from linguistic reasoning capability. Large combinatorial spaces may still require specialized algorithms, decomposition, heuristics, or mathematical programming.

Nevertheless, the framework contributes a coherent conceptual bridge between AI-assisted interpretation and formal optimization. The core contribution is therefore architectural rather than a claim of experimentally demonstrated superiority. Its research value lies in defining a verifiable pathway through which generative models can participate in large-scale decision optimization without being treated as unrestricted autonomous solvers.

CONCLUSION

This paper proposed ScaleSolve, a generative LLM framework for large-scale constraint reasoning and decision optimization. The framework integrates problem interpretation, constraint extraction, formal optimization, candidate verification, and decision explanation into a unified architecture. Analysis of the supplied literature demonstrates that complex food-processing decisions involve interacting process, quality, visual, and nutritional considerations, providing a representative environment for examining multi-objective constraint reasoning.

The principal contribution is the separation of generative reasoning from deterministic validation. LLMs can assist with translating heterogeneous requirements into structured optimization representations, identifying dependencies, decomposing large problems, and explaining solutions. However, mathematical feasibility and numerical validity should remain subject to formal verification. This hybrid design addresses a fundamental limitation of purely generative decision systems: linguistic plausibility does not guarantee constraint satisfaction.

Future research should implement ScaleSolve with benchmark optimization problems, multimodal quality datasets, and controlled comparisons against conventional optimization pipelines. Particular attention should be given to constraint-extraction accuracy, solution feasibility, computational scalability, explanation quality, and robustness under changing constraints. Such evaluation would establish whether LLM-assisted reasoning can provide measurable advantages while retaining the reliability required for large-scale decision optimization.

REFERENCES

1. Ahmad T, Cawood M, Iqbal Q, et al. Phytochemicals in Daucus carota and Their Health Benefits. Foods. 2019; 8(9): 424.

2. Dey D, Dhalsamant K, Tripathy PP. Effect of different pre-treatments on rehydration kinetics of solar and hot-air dried Fuji apple slices. Trends in Horticulture. 2023; 6(2): 3185.

3. Dhalsamant K, Tripathy PP, Shrivastava SL. Effect of pretreatment on rehydration, colour and

4. Computing and Artificial Intelligence 2023; 1(1): 383.17 nanoindentation properties of potato cylinders dried using a mixed-mode solar dryer. Journal of the Science of Food and Agriculture. 2017; 97(10): 3312–3322.

5. Food and Agriculture Organization of the United Nations Carrots and Turnips, 2020. Available online: http: //www.fao.org/faostat/en/#data/QC/visualize. (Accessed on 10 October 2022).

6. Jahanbakhshi A, Kheiralipour K. Evaluation of image processing technique and discriminant analysis methods in postharvest processing of carrot fruit. Food Science & Nutrition. 2020; 8(7): 3346-3352.

7. Kakani V, Nguyen VH, Kumar BP, et al. A critical review on computer vision and artificial intelligence in food industry. Journal of Agriculture and Food Research. 2020; 2: 100033.

8. Mudahar GS, Toledo RT, Floros JD, et al. Optimization of Carrot Dehydration Process using Response Surface Methodology. Journal of Food Science. 1989; 54(3): 714-719.

9. Sharma S, Dhalsamant K, Tripathy PP. Application of computer vision technique for physical quality monitoring of turmeric slices during direct solar drying. Journal of Food Measurement and Characterization. 2018; 13(1): 545-558.

10. Sharma S, Dhalsamant K, Tripathy PP, et al. Quality analysis and drying characteristics of turmeric (Curcuma longa L.) dried by hot air and direct solar dryers. LWT. 2021; 138: 110687.

11. Wang H, Fang XM, Sutar PP, et al. Effects of vacuum-steam pulsed blanching on drying kinetics, colour, phytochemical contents, antioxidant capacity of carrot and the mechanism of carrot quality changes revealed by texture, microstructure and ultrastructure. Food Chemistry. 2021; 338: 127799

12. K. Ramamurthy, N. Bellamkonda and N. Amanmadov, "ScalePulse: Combinatorial Llm Framework for Scalability Constraint," SoutheastCon 2026, Huntsville, AL, USA, 2026, pp. 1-6, doi: 10.1109/SoutheastCon63549.2026.11476075

13. Geo Philip, Paulson, Integrated Intelligent Building Energy Management: A Multi-Layer Framework for Renewable Energy, HVAC Optimization, and Smart Electrical Network Coordination. Available at SSRN: https://ssrn.com/abstract=6993209 or http://dx.doi.org/10.2139/ssrn.6993209