KMAN PUBLISHING INSTITUTE

Article history:
Received 27 May 2026
Revised 20 August 2026
Accepted 28 August 2026
Initial Publish 27 September 2026
Published online 01 July 2027

Journal of Resource Management and Decision Engineering

Volume 6, Issue 4, pp 1-20

Multi-Objective Trust-Aware Service Selection Modeling in Dynamic Cloud Manufacturing Using Blockchain, Adaptive Learning, and Sustainability Constraints

Mina. Ravanesh[1*], Emad. Roghanian[2]

[1] PHD candidate of industrial engineering, K. N. Toosi University of Technology, Tehran, Iran
[2] Associate Professor of industrial engineering, K. N. Toosi University of Technology, Tehran, Iran

*** Corresponding author email address**: m.ravanesh@email.kntu.ac.ir

Article Info

Article type:
Original Research

How to cite this article:
Ravanesh, M., Roghanian, E. (2027). Multi-Objective Trust-Aware Service Selection Modeling in Dynamic Cloud Manufacturing Using Blockchain, Adaptive Learning, and Sustainability Constraints. *Journal of Resource Management and Decision Engineering*, *6*(4), 1-20.
https://doi.org/10.61838/kman.jrmde.412

ABSTRACT

This study aimed to develop and evaluate a multi-objective trust-aware service selection model for dynamic cloud manufacturing by integrating blockchain-based transaction verification, adaptive learning, reliability-oriented optimization, and explicit sustainability constraints. A quantitative computational modeling and simulation design was used. The simulated cloud manufacturing environment included 1,200 heterogeneous manufacturing services distributed across eight functional categories and 240 service providers, with 4,800 composite manufacturing requests generated under low-, medium-, and high-dynamic conditions. Service profiles incorporated cost, processing time, availability, reliability, transaction success, trust, behavioral consistency, energy consumption, carbon emissions, and defect rate. Blockchain-based transaction histories were used to construct tamper-resistant trust evidence, while adaptive learning continuously updated provider performance estimates. The service selection problem was formulated as a constrained multi-objective optimization problem and evaluated against conventional QoS-based, trust-aware, and trust-plus-adaptive-learning comparison models through repeated simulations, robustness analysis, sensitivity analysis, convergence assessment, scalability testing, and ablation analysis. Inferential comparisons showed significant differences among the competing models for aggregate reliability, provider trust, manufacturing request success, service failure, energy consumption, carbon emissions, and the proportion of feasible service compositions. Pairwise comparisons indicated that the complete proposed model significantly outperformed the conventional QoS-based and trust-only models across the principal trust, reliability, and sustainability outcomes. The complete model also significantly outperformed the trust-plus-adaptive-learning configuration for environmental indicators, confirming the added contribution of sustainability constraints. Ablation analyses further demonstrated significant performance deterioration when blockchain-based trust or adaptive learning was removed, whereas removing sustainability constraints significantly increased energy consumption and carbon emissions despite preserving short-term operational success.

Keywords: *cloud manufacturing; service selection; multi-objective optimization; blockchain; trust management; adaptive learning; sustainability; energy efficiency.*

1. Introduction

Cloud manufacturing has emerged as an important paradigm for integrating geographically distributed manufacturing resources, computational capabilities, industrial knowledge, and production services through networked platforms. By virtualizing manufacturing resources and exposing them as on-demand services, cloud manufacturing enables users to dynamically discover and invoke capabilities such as product design, machining, additive manufacturing, assembly, inspection, logistics, and maintenance without requiring direct ownership of the corresponding physical resources. This service-oriented architecture supports greater flexibility, resource sharing, scalability, and responsiveness across complex production networks. However, the increasing number and heterogeneity of available manufacturing services have transformed service selection from a simple matching problem into a complex decision-making problem in which multiple candidate services may satisfy the same functional requirement while differing substantially in cost, processing time, reliability, availability, quality, environmental performance, and trustworthiness. Accordingly, service composition and optimal selection have become central research problems in cloud manufacturing, particularly when composite manufacturing tasks require coordinated selection of multiple interdependent services from large candidate pools (Lim et al., 2022).

Early cloud manufacturing service selection approaches predominantly emphasized quality-of-service criteria and formulated service composition as an optimization problem in which cost, execution time, reliability, and related operational indicators were aggregated or simultaneously optimized. Such formulations substantially improved the ability of cloud manufacturing platforms to identify efficient service combinations, but they also demonstrated that service selection is inherently characterized by conflicting objectives. A low-cost service may have lower reliability, a highly reliable service may require greater processing time, and a rapidly available manufacturing resource may impose additional financial or environmental costs. Gao et al. demonstrated the importance of jointly considering QoS and robustness criteria in bi-objective cloud manufacturing service composition, emphasizing that solutions optimized only for nominal performance may become inappropriate when the operational environment is uncertain or resource conditions change (Gao et al., 2022). Similarly, three-tier service composition formulations have shown that manufacturing service selection must account for interactions among user requirements, candidate services, and system-level optimization objectives rather than relying on isolated service attributes (Lim et al., 2022). These developments established multi-criteria optimization as a fundamental methodological basis for cloud manufacturing service composition.

The increasing dimensionality of service selection has subsequently led to the development of many-objective optimization approaches. When a cloud manufacturing platform simultaneously evaluates cost, time, quality, reliability, availability, energy consumption, user satisfaction, and other decision criteria, conventional weighted aggregation methods may inadequately represent the structure of competing objectives. Evolutionary many-objective methods offer an alternative by maintaining populations of non-dominated solutions and exploring trade-offs among different performance dimensions. Wang et al. proposed an evolutionary cloud manufacturing service selection and scheduling method using an adaptive environmental selection strategy, demonstrating the importance of maintaining convergence and diversity when multiple optimization objectives are considered simultaneously (Wang et al., 2021). More recently, improved multi-objective evolutionary approaches have continued to address increasingly sophisticated service composition problems. Liu et al. incorporated fuzzy customer demands into cloud manufacturing service composition using an improved NSGA-III algorithm, illustrating that uncertainty in user requirements must also be considered when constructing feasible service compositions (Liu et al., 2026). Zhang et al. further incorporated service association effects into an improved NSGA-II-based composition framework, reflecting the importance of dependencies among services rather than treating candidate resources as independent alternatives (Zhang et al., 2025).

The dynamic nature of cloud manufacturing creates an additional challenge because the quality characteristics of manufacturing services cannot reasonably be assumed to remain constant over time. Manufacturing resources are physical systems exposed through digital service interfaces, and their performance can change as a result of workload fluctuations, machine degradation, resource contention, production queues, maintenance requirements, network conditions, or changes in provider behavior. Consequently, a service that is optimal when the composition is created may become inefficient or infeasible during execution. Hu et al.

addressed dynamic cloud manufacturing service composition from an online policy perspective and specifically considered re-entrant services, demonstrating that service composition must remain responsive to changes occurring during manufacturing operations rather than relying entirely on static pre-execution decisions (Hu et al., 2024). Aalikhani et al. extended this problem by introducing runtime QoS prediction through predictive process monitoring, indicating that anticipated service performance can be updated using execution information to improve dynamic service composition decisions (Aalikhani et al., 2026). These developments highlight an important transition from static service optimization toward predictive and adaptive service selection.

Adaptive learning approaches are particularly relevant in this context because they allow cloud manufacturing systems to revise service expectations as new information becomes available. Instead of assuming that historical averages provide complete information about future performance, learning-based frameworks can continuously estimate how services are likely to behave under current operating conditions. Deep reinforcement learning has been applied to cloud manufacturing service composition as a mechanism for learning efficient selection policies under complex state and action spaces (Fazeli et al., 2024). Such approaches provide important advantages when service selection involves sequential decisions and when environmental conditions evolve over time. Nevertheless, adaptive learning also raises questions concerning the credibility of the information used for updating service models. A learning mechanism can only be as reliable as its input data; if provider-reported performance information, service evaluations, or transaction histories are incomplete, manipulated, or strategically distorted, adaptive algorithms may converge toward inaccurate trust or QoS estimates. This problem becomes particularly important in decentralized manufacturing ecosystems in which participants may have limited prior relationships and cannot automatically assume that all service providers or evaluators behave honestly.

Trust management therefore constitutes a second major dimension of cloud manufacturing service selection. Unlike traditional manufacturing supply chains in which long-term contractual relationships and organizational oversight may establish confidence among participants, cloud manufacturing platforms can dynamically connect users with previously unknown providers. In these environments, functional compatibility and QoS values alone do not establish whether a provider will actually fulfill its service commitments. Trust must therefore be inferred from transaction records, prior performance, collaborative evaluations, behavioral consistency, reputation, and other forms of interaction evidence. Vatankhah Barenji introduced a blockchain-based trust system for cloud manufacturing, demonstrating that distributed ledger technology can support more credible and traceable trust management by protecting transaction evidence against unauthorized alteration (Vatankhah Barenji, 2022). Blockchain can provide persistent transaction histories, decentralized verification, and tamper-resistant recording of interactions, thereby reducing reliance on centralized or provider-controlled reputation information.

The application of blockchain to cloud manufacturing has subsequently expanded from trust recording toward service composition and decentralized service selection. Tong et al. proposed a multi-objective service composition architecture for blockchain-based cloud manufacturing, indicating that blockchain mechanisms can be integrated directly with optimization processes rather than functioning solely as an external security layer (Tong et al., 2023). More recently, Meng et al. developed a blockchain-enabled decentralized service selection approach for QoS-aware cloud manufacturing, reinforcing the potential of distributed ledger technologies to support service discovery and selection in decentralized manufacturing environments (Meng et al., 2025). Blockchain-based architectures are particularly relevant where multiple service providers and users participate without a universally trusted intermediary because transaction histories can be independently verified and used to construct more reliable behavioral profiles. However, storing service histories on blockchain does not itself solve the dynamic trust problem. Historical evidence still requires appropriate interpretation, and old transactions may become less informative when service behavior changes over time.

Recent research has therefore increasingly focused on dynamic and intelligent trust evaluation. Yang et al. investigated trustworthy collaborative evaluation involving multiple service subjects in cloud manufacturing, emphasizing that trust can emerge from interactions among multiple actors and should not be reduced to a single static reputation value (Yang, Jiang, et al., 2025). Alhudhaif and Polat advanced this direction through a graph learning and adaptive fuzzy framework for robust trust management, demonstrating the value of representing structural relationships among participants while dynamically updating trust under uncertainty (Alhudhaif & Polat, 2026).

These approaches suggest that future service selection models should not merely include trust as a fixed QoS attribute. Instead, trust should be continuously reconstructed from recent and verifiable behavioral evidence, with sufficient sensitivity to identify service deterioration, inconsistent performance, collusive evaluations, or other changes that may affect provider credibility.

Reliability represents a closely related but conceptually distinct concern. Trust expresses confidence in the expected behavior of a provider, whereas reliability reflects the probability or consistency with which the service successfully performs its intended function. In practice, both dimensions influence service selection because a provider may have a favorable reputation while still being technically vulnerable to interruptions or performance degradation. Reliability-centric service composition has therefore received increasing attention. Mahroo et al. developed an integer linear model specifically emphasizing reliability in cloud manufacturing service composition, demonstrating that reliability requirements can materially alter the structure of optimal service combinations (Mahroo et al., 2026). Furthermore, multi-objective service selection frameworks that account for multi-user benefit and adaptive resource partitioning indicate that reliability and resource allocation cannot be optimized independently in hybrid cloud manufacturing environments, particularly when multiple users compete for limited resources (Xiong et al., 2024). These findings reinforce the need for an integrated model in which trust, reliability, and dynamic resource behavior jointly influence selection decisions.

Service interdependence further complicates this optimization process. Composite manufacturing tasks commonly involve networks of sequential, parallel, conditional, or coupled operations, meaning that the performance of the overall manufacturing request depends not only on individual services but also on their relationships. Yang et al. proposed a service composition optimization approach based on a dual-layer coupled network and an improved particle swarm optimization algorithm, showing that structural relationships among services can meaningfully affect optimal composition decisions (Yang, Ding, et al., 2025). Likewise, service association impact has been incorporated into multi-objective optimization to capture dependencies that would otherwise be ignored when individual services are evaluated independently (Zhang et al., 2025). These studies indicate that realistic cloud manufacturing optimization requires system-level rather than isolated service-level analysis.

A further limitation of many conventional service selection frameworks concerns environmental sustainability. Cloud manufacturing was initially promoted primarily for its economic and operational benefits, including improved utilization of manufacturing resources, flexible capacity allocation, and reduced barriers to accessing advanced production capabilities. However, manufacturing services also consume electricity, materials, and other resources and generate carbon emissions. Selecting services purely according to price, speed, or reliability may therefore produce environmentally inefficient manufacturing compositions. Shi developed a service recommendation model integrating carbon emission hierarchy into cloud manufacturing decision-making, demonstrating that environmental performance can be explicitly incorporated into service recommendation and optimization processes (Shi, 2023). Chen et al. similarly investigated low-carbon cloud manufacturing service composition using a hybrid NSGA-II-SA algorithm, showing that service composition can be designed to jointly address conventional operational objectives and carbon-reduction requirements (Chen et al., 2023).

Energy-aware service selection has further strengthened the sustainability dimension of cloud manufacturing optimization. Peng et al. examined energy-aware cloud manufacturing service selection and scheduling, demonstrating that energy consumption should be evaluated jointly with scheduling and operational performance rather than addressed only after service selection has been completed (Peng et al., 2025). Tang et al. proposed a QoS- and sustainability-driven two-stage service composition method integrating clustering with bi-objective optimization, further highlighting the need to embed sustainability within the optimization procedure itself (Tang et al., 2026). These approaches indicate that sustainability should not merely appear as an auxiliary reporting indicator. When energy and carbon objectives are not explicitly included during service selection, optimization algorithms may systematically favor alternatives that are economically attractive but environmentally inefficient.

The challenge, therefore, is not simply to add more objectives to an optimization function. Trust, learning, blockchain, reliability, and sustainability interact in ways that create methodological trade-offs. A highly trustworthy provider may have greater energy consumption; an environmentally efficient provider may exhibit greater execution uncertainty; a low-cost service may become unreliable under increased workload; and an adaptive algorithm may detect changes quickly but require reliable transaction evidence to distinguish genuine performance deterioration from noisy observations. Consequently, integrated service selection requires mechanisms that simultaneously address credible information, temporal adaptation, multi-objective optimization, and sustainability constraints. Existing studies have advanced these dimensions individually or in partial combinations, yet their integration remains comparatively fragmented.

Current research demonstrates several complementary directions but also reveals this fragmentation. Blockchain studies have established secure and decentralized trust mechanisms (Meng et al., 2025; Tong et al., 2023; Vatankhah Barenji, 2022), while adaptive trust research has improved dynamic evaluation of service providers (Alhudhaif & Polat, 2026; Yang, Jiang, et al., 2025). Dynamic QoS prediction and online service composition approaches have strengthened responsiveness to changing operational conditions (Aalikhani et al., 2026; Hu et al., 2024), and learning-based optimization has demonstrated the ability to handle complex decision environments (Fazeli et al., 2024). In parallel, multi-objective and many-objective evolutionary methods have improved the handling of competing service objectives and uncertain demands (Liu et al., 2026; Wang et al., 2021), while sustainability-oriented research has incorporated energy and carbon considerations into selection decisions (Chen et al., 2023; Peng et al., 2025; Shi, 2023; Tang et al., 2026). Nevertheless, relatively limited attention has been devoted to a unified architecture in which blockchain-verified trust information is continuously interpreted through adaptive learning and then incorporated together with reliability, QoS, and explicit sustainability constraints into dynamic multi-objective service selection.

Another important issue is the distinction between optimization efficiency and operational resilience. A model can identify mathematically attractive service compositions yet remain vulnerable when selected services unexpectedly deteriorate or become unavailable. Robust cloud manufacturing therefore requires mechanisms for detecting changes, revising expectations, and redirecting service selection toward more appropriate alternatives. Gao et al. demonstrated the relevance of robustness in service composition (Gao et al., 2022), while Aalikhani et al. showed how runtime prediction can improve dynamic composition decisions (Aalikhani et al., 2026). Reliability-centric mathematical programming further indicates that service success must be treated as a primary decision criterion rather than a secondary evaluation measure (Mahroo et al., 2026). When these insights are combined with blockchain-verified behavioral records, adaptive trust estimation can potentially distinguish temporary stochastic deviations from sustained service deterioration and thereby support more resilient selection policies.

Multi-user and resource-allocation considerations also reinforce the need for adaptivity. In practical cloud manufacturing platforms, multiple users may simultaneously request overlapping manufacturing resources, and the allocation of capacity to one service composition can influence the feasibility and performance of others. Xiong et al. demonstrated the importance of adaptive resource partitioning and multi-user benefit within hybrid cloud manufacturing (Xiong et al., 2024). Furthermore, coupled-network formulations emphasize that service composition is embedded in broader relational structures in which service interactions and network effects influence system performance (Yang, Ding, et al., 2025). These conditions make static rankings increasingly inadequate because the relative desirability of a service depends on both its intrinsic characteristics and its current position within a changing resource environment.

Taken together, the literature indicates that cloud manufacturing service selection is evolving from static QoS matching toward intelligent, dynamic, trustworthy, and sustainable decision-making. However, achieving these characteristics simultaneously remains challenging. Blockchain can strengthen data integrity but does not independently provide adaptation; adaptive learning can update predictions but requires credible evidence; multi-objective optimization can balance conflicting criteria but must incorporate meaningful constraints; and sustainability objectives can reduce environmental burden but may create trade-offs with cost, speed, and reliability. A comprehensive framework must therefore combine these mechanisms rather than treating them as independent enhancements.

Accordingly, the aim of this study was to develop and evaluate a multi-objective trust-aware service selection model for dynamic cloud manufacturing that integrates blockchain-based transaction verification, adaptive learning of service performance and provider trust, reliability and QoS optimization, and explicit energy-consumption and carbon-emission sustainability constraints.

2. Methods and Materials

This study employed a quantitative computational modeling and simulation design to develop and evaluate a multi-objective trust-aware service selection framework for dynamic cloud manufacturing environments. The proposed framework was designed to address the simultaneous requirements of service quality, provider trustworthiness, operational adaptability, transaction security, and environmental sustainability in cloud manufacturing service composition. Rather than relying on a conventional single-objective service selection mechanism, the study formulated the service selection problem as a constrained multi-objective optimization problem in which several potentially conflicting criteria, including manufacturing cost, processing and delivery time, service reliability, trust level, energy consumption, and carbon emissions, were optimized simultaneously. The methodological framework integrated four principal components: a blockchain-enabled trust management mechanism for maintaining tamper-resistant service interaction records, an adaptive learning mechanism for continuously updating provider performance estimates, a multi-objective optimization model for identifying efficient combinations of manufacturing services, and sustainability constraints for preventing the selection of solutions that achieved economic or operational advantages at the expense of unacceptable environmental performance.

The study population was defined conceptually as all manufacturing services that could be registered, discovered, and invoked through a cloud manufacturing platform. Because access to a sufficiently large standardized industrial cloud manufacturing dataset containing synchronized information on service quality, blockchain-based trust transactions, adaptive historical performance, and sustainability indicators is limited, the empirical evaluation was conducted using a controlled simulation dataset representing heterogeneous manufacturing service providers. The simulated environment included different categories of manufacturing capabilities, such as product design, machining, additive manufacturing, surface processing, assembly, inspection, packaging, and logistics. Each manufacturing task was represented as a composite request consisting of several sequential or partially dependent subtasks, and each subtask could be fulfilled by multiple candidate cloud manufacturing services. Candidate services were differentiated according to their operational characteristics and historical behavior so that the selection model was evaluated under realistic conditions of service heterogeneity and uncertainty.

For each candidate manufacturing service, a multidimensional service profile was generated containing information on service execution cost, estimated processing time, response time, availability, reliability, completion success rate, defect rate, energy consumption, estimated carbon emissions, historical transaction frequency, and provider reputation. Trust-related variables were generated dynamically based on previous interactions, verified transaction outcomes, fulfillment consistency, customer evaluations, and the recency of service interactions. To represent the dynamic nature of cloud manufacturing systems, service characteristics were allowed to vary across simulation periods. Cost, processing capacity, reliability, energy use, and service availability were therefore not treated as completely static parameters. Controlled fluctuations were incorporated to simulate workload variation, equipment degradation, temporary resource shortages, changes in provider behavior, network delays, and changing production conditions. This dynamic structure enabled the proposed adaptive learning mechanism to be evaluated under conditions in which historical information alone was insufficient for reliable service selection.

The experimental scenarios were constructed at different levels of problem complexity by varying the number of candidate services available for each manufacturing task, the number of required subtasks, and the degree of environmental volatility. Low-, medium-, and high-dynamic scenarios were generated by progressively increasing the magnitude and frequency of changes in service quality and availability. Additional scenarios were designed to represent sudden service deterioration, unreliable service providers, changes in manufacturing demand, and increases in environmental constraints. Multiple independent simulation runs were conducted for each scenario using different random seeds to reduce the likelihood that the observed performance resulted from a particular configuration of simulated observations. The same generated service instances and task requests were provided to all comparison algorithms within each simulation run to ensure methodological consistency and permit direct comparison of alternative service selection approaches.

The unit of analysis was a candidate service composition capable of satisfying all functional requirements of a manufacturing request. A feasible service composition was required to contain one eligible service for each required manufacturing subtask while satisfying predetermined compatibility, capacity, deadline, trust, and sustainability conditions. Solutions violating mandatory functional constraints, minimum trust requirements, maximum allowable completion time, production capacity limitations, or environmental thresholds were classified as infeasible.

This formulation enabled the proposed model to distinguish between solutions that were mathematically efficient but operationally unacceptable and solutions that could realistically be implemented within a cloud manufacturing platform.

Data collection in this study was implemented through a simulation-based cloud manufacturing environment in which service information was recorded through four integrated data structures: a service registry, a transaction history repository, a blockchain-based trust ledger, and a sustainability information module. The service registry contained the functional and non-functional characteristics of each candidate manufacturing service and was used during the initial service discovery stage. Functional attributes specified whether a service was technically capable of satisfying the corresponding manufacturing requirement, whereas non-functional attributes represented quality-of-service characteristics used during optimization. These attributes included normalized measures of cost, processing time, availability, reliability, service success, response efficiency, and manufacturing quality. Before integration into the objective functions, criteria measured on different numerical scales were normalized to ensure comparability and to prevent criteria with larger numerical magnitudes from disproportionately affecting optimization outcomes.

Trust information was generated and maintained using a blockchain-inspired distributed ledger structure. Each completed service interaction produced a transaction record containing a unique transaction identifier, provider identifier, requesting entity identifier, service category, timestamp, promised service conditions, observed execution outcome, completion status, quality assessment, and evaluation score. Cryptographic hashing was conceptually incorporated into consecutive transaction records to represent the immutability and traceability characteristics of blockchain-based data management. The blockchain component was not used merely as a cryptocurrency mechanism; instead, it served as the trust infrastructure of the service selection framework by providing an auditable and manipulation-resistant history of manufacturing interactions. Information recorded in the ledger was subsequently used to calculate provider trust values and to reduce dependence on self-reported reputation information.

The trust assessment instrument was implemented as a composite computational index combining direct trust, indirect reputation, transaction reliability, behavioral consistency, and interaction recency. Direct trust represented the requesting agent's previous experience with a specific provider, whereas indirect reputation reflected evaluations accumulated from other participating entities. Transaction reliability measured the proportion of successfully completed transactions relative to accepted service requests. Behavioral consistency captured the stability of provider performance across successive interactions, thereby penalizing providers whose average performance appeared acceptable but showed substantial volatility. A temporal decay mechanism was applied to historical interactions so that recent transactions received greater influence than older interactions. This mechanism was important in a dynamic manufacturing environment because a provider's current operational condition may differ substantially from its historical average. The final trust score was scaled to a standardized interval in which higher values indicated greater confidence in the provider.

An adaptive learning module was incorporated to continuously revise expected service performance using newly observed transaction outcomes. At the beginning of the simulation, the model relied on available historical information and initial estimates of service performance. After each completed manufacturing interaction, prediction errors were calculated by comparing expected values with actual observed values for variables such as completion time, reliability, quality, and energy consumption. The adaptive mechanism updated the corresponding service estimates by assigning greater weight to recent observations when environmental volatility increased. Consequently, the framework could respond to gradual performance changes as well as abrupt service degradation. This approach prevented the service selection process from treating provider capabilities as permanently fixed and enabled the system to learn from changing manufacturing conditions throughout the simulation.

Sustainability information was collected for each service using operational environmental indicators, principally estimated energy consumption and carbon emissions associated with completing the requested manufacturing operation. Where appropriate, resource utilization and material efficiency were also represented as supplementary environmental indicators. Energy consumption was estimated according to the service's operating intensity and processing duration, whereas carbon emissions were calculated from energy use and the corresponding emission factor assigned to the service environment. Sustainability criteria were included both as optimization objectives and as feasibility constraints. Accordingly, the model sought to minimize environmental burden while simultaneously rejecting service compositions that exceeded predefined energy or emission limits. This distinction ensured that environmental performance was treated as a substantive decision criterion rather than as an optional post-selection evaluation.

The multi-objective service selection model used the collected service information to construct objective functions representing economic, operational, trust, and environmental performance. The economic objective minimized total service composition cost, including the aggregate cost of selected manufacturing operations. The temporal objective minimized the expected completion time of the overall manufacturing request while accounting for the structure and dependency of constituent tasks. The quality objective maximized the aggregate reliability and service quality of the selected composition. The trust objective maximized the predicted trustworthiness of the selected providers based on blockchain-verified historical interactions and adaptive trust updates. The sustainability objectives minimized total energy consumption and carbon emissions. In addition to these objectives, the model incorporated constraints associated with service compatibility, task coverage, production capacity, deadline compliance, minimum reliability, minimum provider trust, and maximum sustainability thresholds.

To evaluate the incremental contribution of the proposed architecture, several model configurations were examined. These included a conventional quality-of-service-based service selection model, a trust-aware model without adaptive learning, a trust-aware adaptive model without explicit sustainability constraints, and the complete proposed model integrating blockchain-based trust, adaptive learning, multi-objective optimization, and sustainability requirements. The comparison configurations were evaluated using identical task instances and candidate service pools. This design made it possible to determine whether performance improvements were attributable to the complete integration of the proposed mechanisms rather than merely to the use of a more complex optimization procedure.

Data analysis consisted of optimization analysis, simulation-based comparative evaluation, convergence assessment, and statistical examination of repeated experimental outcomes. The service selection problem was mathematically formulated as a constrained multi-objective combinatorial optimization problem. Decision variables indicated whether a particular candidate manufacturing service was selected for a specific task. Because improvements in one criterion could produce deterioration in another criterion, such as when lower-cost services were associated with longer processing times or higher-trust providers involved greater economic expenditure, the objectives were not collapsed immediately into a single deterministic criterion. Instead, the optimization procedure searched for Pareto-efficient service compositions in which no objective could be improved without worsening at least one other objective. The resulting Pareto solution set represented alternative trade-offs among economic efficiency, operational performance, trustworthiness, and environmental sustainability.

The objective functions were transformed into comparable normalized values before optimization. Benefit criteria, such as trust, reliability, and quality, were normalized so that larger values represented superior performance, whereas cost criteria, including monetary cost, processing time, energy use, and carbon emissions, were transformed so that smaller values represented superior performance. Constraint-handling procedures were incorporated into the optimization process to ensure that infeasible service compositions were either rejected or strongly penalized. Functional constraints guaranteed that every required manufacturing task was assigned to an appropriate candidate service. Capacity constraints prevented the allocation of workloads exceeding provider capabilities, while deadline constraints ensured that the expected completion time did not exceed the maximum acceptable value specified in the manufacturing request. Trust and sustainability constraints imposed minimum acceptable trust levels and maximum environmental limits, respectively.

A population-based multi-objective optimization procedure was used to search the large combinatorial solution space. Candidate service compositions were iteratively generated, evaluated according to the defined objective functions, ranked according to Pareto dominance, and refined through successive search iterations. Diversity preservation was incorporated into the optimization process to prevent premature convergence toward a narrow region of the Pareto frontier. The stopping criterion was defined using a maximum number of optimization iterations together with convergence stability across successive generations. The final output of each simulation consisted of a set of non-dominated service compositions representing efficient alternatives for the manufacturing requester.

Where a single service composition was required for operational implementation, a secondary decision-making stage was applied to the non-dominated solutions. Objective values were normalized, and preference weights were assigned according to the decision context. The weighting structure could therefore reflect different managerial priorities, such as cost-dominant, trust-dominant, time-critical, or sustainability-oriented production. The final solution was selected according to its overall closeness to the preferred multi-criteria profile while maintaining all mandatory feasibility constraints. Sensitivity analysis was conducted by systematically modifying criterion weights and sustainability thresholds to determine whether the recommended service composition remained stable under alternative decision priorities.

The performance of the proposed model was evaluated using both optimization and operational indicators. Optimization quality was assessed using measures reflecting Pareto-front quality, solution diversity, convergence, and computational efficiency. Operational effectiveness was assessed through total service cost, overall completion time, achieved reliability, average provider trust, service failure rate, energy consumption, carbon emissions, and the proportion of successfully completed manufacturing requests. Adaptability was evaluated by measuring the model's ability to revise service selections after changes in provider performance or availability. Trust robustness was examined by introducing service providers with inconsistent or deteriorating behavior and observing whether the blockchain-based adaptive trust mechanism progressively reduced their probability of selection. Sustainability effectiveness was assessed by comparing total environmental burden across the proposed model and configurations without explicit environmental constraints.

Each experimental condition was repeated across multiple independent simulation runs, and the resulting indicators were summarized using the mean, standard deviation, minimum, and maximum where appropriate. Before inferential comparisons, the distributional characteristics of repeated performance measures were examined. When the assumptions of parametric analysis were adequately satisfied, differences among competing service selection models were evaluated using appropriate repeated comparative procedures, followed by adjusted pairwise comparisons where significant overall differences were identified. When parametric assumptions were not satisfied, corresponding non-parametric procedures were employed. Statistical significance was evaluated at the 0.05 level, while effect-size measures were additionally considered to determine the practical magnitude of differences between competing approaches rather than relying exclusively on probability values.

Robustness analysis was performed across variations in the number of candidate services, task complexity, environmental volatility, trust uncertainty, service failure probability, and sustainability restrictions. Scalability was examined by progressively increasing the dimensions of the service selection problem and recording changes in computational time and optimization quality. Convergence behavior was assessed across successive iterations to determine whether the proposed approach reached stable solutions without excessive computational burden. Ablation analysis was also performed by removing blockchain-based trust updating, adaptive learning, or sustainability constraints from the complete model and comparing the resulting performance with that of the integrated framework. This analysis enabled the individual contribution of each methodological component to be identified.

Finally, the overall validity of the proposed framework was assessed according to its ability to produce service compositions that were simultaneously feasible, trustworthy, adaptive, economically efficient, operationally reliable, and environmentally sustainable. Particular emphasis was placed on whether the integration of blockchain-verified trust information and adaptive learning improved decision quality under dynamic conditions, and whether these improvements could be achieved without disproportionately increasing manufacturing cost or computational complexity. The analytical framework therefore evaluated the proposed model not only according to conventional optimization performance but also according to its capacity to support resilient and sustainable decision-making in dynamic cloud manufacturing ecosystems.

3. Findings and Results

The final experimental dataset consisted of 1,200 registered cloud manufacturing service instances distributed across eight functional manufacturing categories. Of these services, 168 (14.00%) were associated with product and process design, 246 (20.50%) with machining, 138 (11.50%) with additive manufacturing, 126 (10.50%) with surface processing, 162 (13.50%) with assembly operations, 144 (12.00%) with inspection and quality control, 108 (9.00%) with packaging, and 108 (9.00%) with logistics and delivery services. The candidate services were distributed across 240 simulated service providers, with each provider offering between 2 and 11 manufacturing services. A total of 4,800 composite manufacturing requests were generated during the experimental procedure. Each manufacturing request consisted of between 4 and 12 dependent or partially dependent subtasks, producing substantial variation in the complexity of the corresponding service selection problem. Of all requests, 1,600 (33.33%) were generated under low-dynamic operating conditions, 1,600 (33.33%) under medium-dynamic conditions, and 1,600 (33.33%) under high-dynamic conditions. In addition, 20.00% of the simulated providers were assigned relatively unstable performance profiles in order to reproduce realistic variations in service reliability, capacity, processing time, and availability, while 10.00% were configured to exhibit potentially untrustworthy behavior characterized by inconsistencies between advertised and subsequently observed service performance. Approximately 25.00% of the evaluated manufacturing requests were subjected to relatively strict sustainability restrictions, including tighter energy-consumption and carbon-emission thresholds. The initial mean normalized provider trust score was 0.731 with a standard deviation of 0.116, while the initial mean service reliability was 0.912 with a standard deviation of 0.054. Average service availability across the simulated environment was 0.927, and the mean historical transaction success rate was 91.84%. These characteristics created a heterogeneous experimental environment in which service selection could not be adequately addressed through cost or quality criteria alone and therefore provided an appropriate basis for evaluating the proposed trust-aware, adaptive, and sustainability-constrained model.

Table 1

Descriptive Statistics of the Main Service-Level Variables in the Simulated Cloud Manufacturing Dataset

Variable	Mean	SD	Minimum	Maximum
Normalized service cost	0.517	0.183	0.112	0.948
Normalized processing time	0.486	0.176	0.094	0.931
Service availability	0.927	0.041	0.781	0.995
Service reliability	0.912	0.054	0.702	0.994
Transaction success rate	0.918	0.061	0.684	0.997
Initial provider trust	0.731	0.116	0.314	0.978
Behavioral consistency	0.846	0.091	0.493	0.991
Normalized energy consumption	0.492	0.187	0.082	0.958
Normalized carbon emissions	0.476	0.192	0.067	0.963
Normalized defect rate	0.118	0.067	0.009	0.391

As shown in Table 1, the simulated service population exhibited sufficient variability across economic, operational, trust-related, and environmental dimensions to generate a non-trivial multi-objective optimization problem. The mean normalized service cost was 0.517, whereas the mean normalized processing time was 0.486, indicating substantial variation among services with respect to both economic and temporal efficiency. The standard deviations of 0.183 and 0.176 for cost and processing time, respectively, further demonstrated that the service population contained both highly efficient and relatively inefficient alternatives. Service availability and reliability were generally high, with mean values of 0.927 and 0.912, although their observed minimum values of 0.781 and 0.702 showed that some candidate services presented considerably greater operational risk than others. The mean transaction success rate of 0.918 was consistent with the overall reliability profile but also demonstrated meaningful dispersion across providers. Trust-related characteristics showed greater heterogeneity than conventional reliability measures. The mean initial provider trust score was 0.731, ranging from 0.314 to 0.978, whereas behavioral consistency averaged 0.846 and ranged from 0.493 to 0.991. This difference was important because it demonstrated that a provider could exhibit acceptable average operational performance while remaining behaviorally unstable across transactions. Environmental indicators also varied substantially, with normalized energy consumption averaging 0.492 and carbon emissions averaging 0.476. Their relatively large standard deviations indicated that environmentally preferable alternatives were distinguishable from energy-intensive alternatives. Finally, the normalized defect rate averaged 0.118 but reached a maximum of 0.391, confirming the presence of lower-quality services within the candidate pool. Collectively, these descriptive findings confirmed that the simulated environment contained meaningful conflicts among service cost, processing speed, reliability, trustworthiness, and environmental performance and therefore represented an appropriate setting for evaluating the proposed multi-objective service selection framework.

Table 2

Comparative Performance of the Service Selection Models Across the Complete Experimental Dataset

Performance Indicator	Conventional QoS Model	Trust-Aware Model	Trust + Adaptive Learning Model	Proposed Integrated Model
Total normalized cost	0.421 ± 0.058	0.438 ± 0.055	0.446 ± 0.052	0.457 ± 0.049
Normalized completion time	0.448 ± 0.064	0.439 ± 0.060	0.413 ± 0.056	0.405 ± 0.052
Aggregate reliability	0.901 ± 0.027	0.932 ± 0.023	0.946 ± 0.019	0.952 ± 0.017
Mean selected-provider trust	0.718 ± 0.052	0.846 ± 0.041	0.881 ± 0.036	0.893 ± 0.032
Manufacturing request success rate (%)	88.71 ± 2.94	92.84 ± 2.36	95.63 ± 1.91	97.18 ± 1.52
Service failure rate (%)	10.96 ± 2.61	7.23 ± 2.14	4.58 ± 1.62	3.21 ± 1.27
Normalized energy consumption	0.514 ± 0.061	0.509 ± 0.059	0.497 ± 0.055	0.386 ± 0.047
Normalized carbon emissions	0.501 ± 0.064	0.496 ± 0.061	0.485 ± 0.057	0.369 ± 0.045
Constraint-feasible solutions (%)	89.46 ± 3.12	92.73 ± 2.71	95.39 ± 2.18	98.16 ± 1.37
Mean computation time (s)	1.84 ± 0.31	2.13 ± 0.35	2.58 ± 0.41	3.06 ± 0.48

The comparative findings presented in Table 2 demonstrated that the proposed integrated framework consistently produced superior results on the principal reliability, trust, adaptability, feasibility, and sustainability indicators, although these improvements were accompanied by a moderate increase in computational cost and a small increase in direct service expenditure. The conventional quality-of-service model produced the lowest normalized cost of 0.421, whereas the corresponding value for the proposed model was 0.457. Thus, the integrated model did not achieve its advantages by simply selecting the least expensive available services. Rather, it accepted a comparatively small increase in economic cost in exchange for higher reliability, greater trustworthiness, lower failure probability, and markedly improved environmental performance. The proposed model reduced normalized completion time to 0.405 compared with 0.448 under the conventional model, suggesting that the inclusion of adaptive service-performance information enabled the system to avoid providers whose current processing capability had deteriorated despite historically acceptable performance. Aggregate reliability increased progressively from 0.901 for the conventional model to 0.932 for the trust-aware model, 0.946 for the adaptive trust model, and 0.952 for the complete framework. An even larger difference was observed for provider trust, with the mean selected-provider trust value increasing from 0.718 under conventional QoS-based selection to 0.893 under the proposed model. The manufacturing request success rate similarly improved from 88.71% to 97.18%, while the service failure rate declined from 10.96% to 3.21%. The strongest relative advantages of the integrated framework were observed for environmental outcomes. Mean normalized energy consumption declined from 0.514 to 0.386, while carbon emissions declined from 0.501 to 0.369. The proportion of service compositions satisfying all functional, trust, deadline, capacity, and sustainability constraints increased to 98.16%, compared with 89.46% for the conventional model. The principal trade-off was computational time, which increased from an average of 1.84 seconds to 3.06 seconds. Nevertheless, the increase remained limited relative to the magnitude of improvement observed in solution reliability and sustainability, indicating that the proposed framework achieved a more balanced overall service composition rather than optimizing a single criterion at the expense of broader manufacturing requirements.

Table 3

Performance of the Proposed Model Under Different Levels of Environmental Dynamism

Indicator	Low-Dynamic Environment	Medium-Dynamic Environment	High-Dynamic Environment
Aggregate reliability	0.963 ± 0.013	0.954 ± 0.016	0.939 ± 0.022
Mean selected-provider trust	0.912 ± 0.026	0.897 ± 0.031	0.871 ± 0.041
Request success rate (%)	98.64 ± 0.91	97.42 ± 1.31	95.49 ± 1.88

Service failure rate (%)	1.93 ± 0.74	2.97 ± 1.08	4.73 ± 1.57
Adaptation recovery time (iterations)	3.14 ± 0.82	4.67 ± 1.08	6.28 ± 1.41
Incorrect retention of deteriorated providers (%)	2.18 ± 0.73	3.64 ± 1.02	5.91 ± 1.48
Normalized energy consumption	0.374 ± 0.041	0.386 ± 0.045	0.399 ± 0.052
Normalized carbon emissions	0.355 ± 0.039	0.369 ± 0.044	0.385 ± 0.051
Feasible solutions (%)	99.02 ± 0.76	98.31 ± 1.13	97.14 ± 1.69
Mean computation time (s)	2.64 ± 0.36	3.02 ± 0.43	3.51 ± 0.54

Table 3 presents the robustness of the proposed framework under increasing levels of environmental dynamism. As expected, performance deteriorated modestly as the magnitude and frequency of changes in service availability, execution time, reliability, and resource capacity increased. However, the integrated model maintained comparatively high performance even under the most demanding dynamic condition. Aggregate reliability declined from 0.963 in the low-dynamic environment to 0.939 in the high-dynamic environment, whereas mean selected-provider trust declined from 0.912 to 0.871. Despite this increase in uncertainty, the manufacturing request success rate remained above 95%, reaching 95.49% in the high-dynamic scenario. The service failure rate increased from 1.93% under low dynamism to 4.73% under high dynamism, but this increase remained limited considering the deliberately introduced instability of the service environment. A particularly important indicator was adaptation recovery time, defined as the number of optimization-learning iterations required for the system to stabilize its service selection after a meaningful change in provider performance. Recovery time increased from 3.14 iterations in low-dynamic conditions to 6.28 iterations in high-dynamic conditions, demonstrating that more severe environmental disturbances required additional observations before the adaptive mechanism fully revised its provider estimates. Nevertheless, the percentage of cases in which deteriorated providers continued to be selected remained only 5.91% even under highly dynamic conditions. Environmental performance was also comparatively stable: normalized energy consumption increased from 0.374 to 0.399 and carbon emissions from 0.355 to 0.385 as environmental dynamism increased. The proportion of feasible solutions remained above 97% in all three environments. Computational time rose from 2.64 seconds in the low-dynamic scenario to 3.51 seconds in the high-dynamic scenario, reflecting the additional search and adaptation burden associated with frequently changing service profiles. Overall, these results showed that environmental volatility reduced performance to some degree but did not destabilize the proposed model, supporting its capacity to adaptively revise trust and service-performance estimates when manufacturing conditions changed.

Table 4

Ablation Analysis of Blockchain Trust, Adaptive Learning, and Sustainability Components

Model Configuration	Request Success Rate (%)	Failure Rate (%)	Mean Trust	Energy Consumption	Carbon Emissions	Feasible Solutions (%)
Complete proposed model	97.18 ± 1.52	3.21 ± 1.27	0.893 ± 0.032	0.386 ± 0.047	0.369 ± 0.045	98.16 ± 1.37
Without blockchain-based trust	93.41 ± 2.17	6.82 ± 1.89	0.782 ± 0.048	0.391 ± 0.049	0.374 ± 0.047	94.37 ± 2.04
Without adaptive learning	92.76 ± 2.41	7.31 ± 2.06	0.841 ± 0.043	0.397 ± 0.052	0.381 ± 0.050	93.82 ± 2.27
Without sustainability constraints	97.36 ± 1.47	3.08 ± 1.21	0.896 ± 0.031	0.512 ± 0.061	0.498 ± 0.059	94.91 ± 1.96
Without blockchain and adaptive learning	89.63 ± 2.88	10.21 ± 2.52	0.724 ± 0.056	0.405 ± 0.057	0.392 ± 0.055	90.47 ± 2.81

The ablation analysis in Table 4 clarified the independent contribution of the principal components of the proposed framework. Removing blockchain-based trust management reduced the manufacturing request success rate from 97.18% to 93.41%, increased the failure rate from 3.21% to 6.82%, and reduced the mean trust level of selected providers from 0.893 to 0.782. These findings indicated that immutable transaction histories and verified behavioral evidence played an important role in distinguishing consistently trustworthy providers from those whose advertised or reputation-based characteristics were not supported by actual transaction outcomes. Removing adaptive learning produced an even slightly larger reduction in operational performance, with the success rate declining to 92.76% and the failure rate increasing to 7.31%. This result suggested that static trust information alone was insufficient in a dynamic manufacturing environment because historically trustworthy providers could become unsuitable following changes in workload, equipment condition, resource availability, or service behavior. In contrast, removing sustainability constraints had little adverse effect on short-term service reliability and produced a slightly higher request success rate of 97.36%; however, normalized energy consumption increased sharply from 0.386 to 0.512 and carbon emissions increased from 0.369 to 0.498. This finding demonstrated that environmental improvements observed in the complete framework were specifically attributable to the explicit inclusion of sustainability requirements rather than being an indirect consequence of improved trust or reliability. The largest deterioration occurred when both blockchain trust and adaptive learning were removed simultaneously. Under this configuration, the success rate fell to 89.63%, the failure rate increased to 10.21%, mean selected-provider trust declined to 0.724, and the proportion of feasible solutions declined to 90.47%. Taken together, the ablation results indicated that the principal components of the proposed architecture were complementary. Blockchain enhanced the credibility and traceability of historical evidence, adaptive learning ensured that this evidence remained responsive to current conditions, and sustainability constraints prevented economically or operationally attractive solutions from being selected when their environmental burden exceeded acceptable levels.

Figure 1

Convergence Behavior of the Proposed Multi-Objective Trust-Aware Service Selection Model Across Optimization Iterations

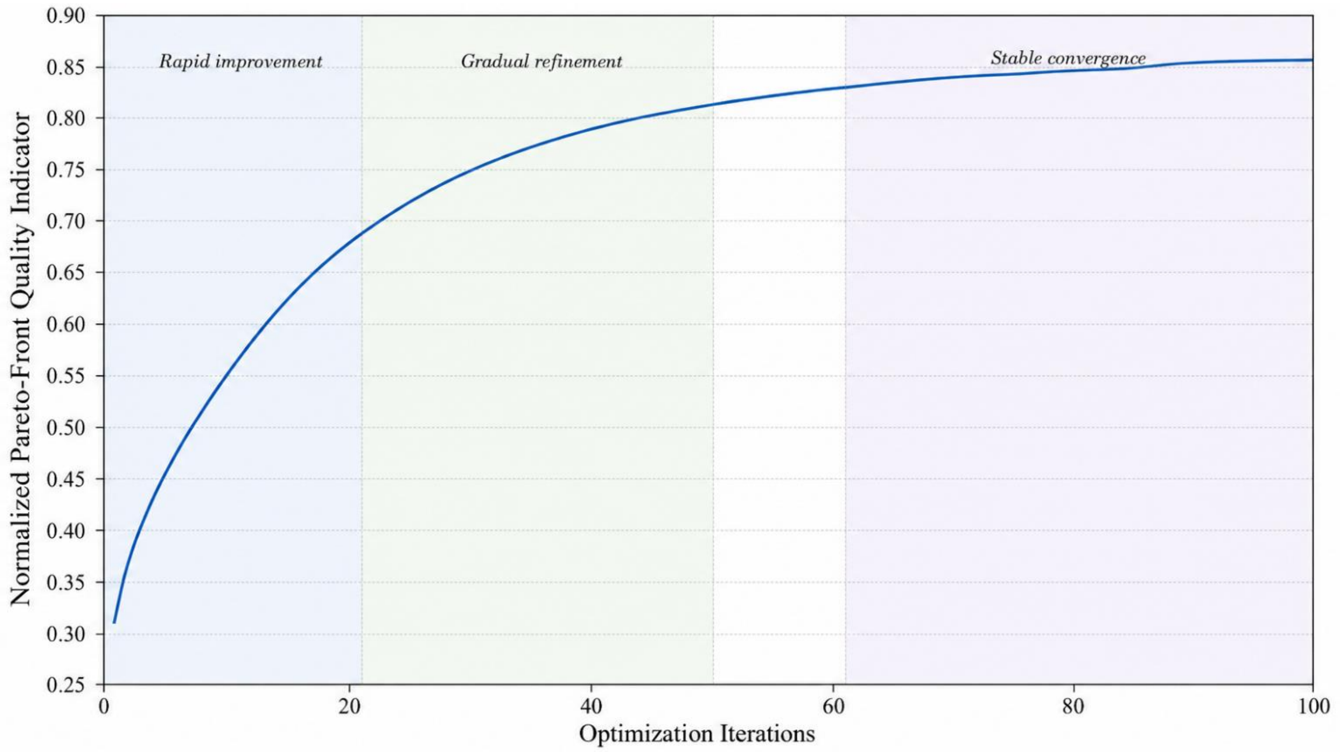

The convergence pattern represented in Figure 1 showed that the proposed optimization procedure achieved rapid improvement during the early search iterations, followed by progressively smaller improvements as the solution population approached a stable Pareto-optimal region. The greatest change in overall solution quality occurred during approximately the first 20 iterations, during which low-quality and constraint-violating service compositions were progressively eliminated and replaced by solutions characterized by superior combinations of cost, completion time, reliability, trust, energy consumption, and carbon emissions. Between approximately iterations 20 and 50, the rate of improvement became more gradual as the optimization process increasingly focused on refining trade-offs among already feasible non-dominated solutions. After approximately 60 iterations, only minor improvements in the aggregate Pareto-front quality indicator were observed, and the solution set remained comparatively stable through the final iterations. The convergence trajectory did not display substantial late-stage oscillation, indicating that the adaptive learning mechanism did not prevent optimization stability despite periodically updating expected provider performance. At the same time, diversity among non-dominated solutions was preserved throughout the search process, preventing the model from converging exclusively toward a cost-minimizing, trust-maximizing, or sustainability-maximizing extreme. This pattern supported the selection of the specified stopping rule and showed that the proposed procedure could reach a stable and diversified set of service compositions without requiring excessive optimization iterations.

Additional inferential comparisons supported the descriptive findings. Significant between-model differences were observed for aggregate reliability, mean provider trust, manufacturing request success, service failure, energy consumption, carbon emissions, and the proportion of feasible solutions. Pairwise comparisons indicated that the complete proposed framework significantly outperformed the conventional QoS-based model and the trust-only model on all major trust, reliability, and sustainability outcomes. The difference between the proposed model and the trust-plus-adaptive-learning configuration was smaller for reliability and provider trust but remained pronounced for energy consumption and carbon emissions, demonstrating the incremental contribution of explicit sustainability restrictions. By contrast, the conventional QoS model retained a modest advantage in direct normalized service cost and computational execution time. These results demonstrated that the proposed architecture involved a measurable economic and computational trade-off but generated substantially greater benefits in terms of operational robustness, trustworthiness, environmental performance, and feasibility.

Sensitivity analysis further demonstrated that the proposed model remained stable under alternative decision priorities. Increasing the weight assigned to cost produced service compositions with lower direct expenditure but was accompanied by modest reductions in mean provider trust and reliability. Increasing the trust weight produced the highest levels of transaction success and the lowest service failure rates but moderately increased cost because high-trust providers were not always the least expensive candidates. Sustainability-dominant weighting resulted in the greatest reduction in energy consumption and carbon emissions, whereas time-critical weighting generated the shortest completion times but slightly increased energy use because high-speed manufacturing alternatives frequently exhibited greater resource intensity. Importantly, no single weighting structure dominated all others across every criterion. Instead, the generated Pareto frontier provided multiple feasible alternatives reflecting different managerial priorities, which confirmed the appropriateness of modeling cloud manufacturing service selection as a multi-objective rather than single-objective decision problem.

The trust-deterioration experiments provided additional evidence regarding the adaptive characteristics of the proposed framework. When previously reliable providers were deliberately assigned progressively deteriorating performance, their blockchain-recorded transaction outcomes caused their trust values to decrease over subsequent interactions. The adaptive learning mechanism simultaneously revised expected reliability, processing time, and success probability. Consequently, deteriorating providers were progressively removed from preferred service compositions and replaced by alternatives with stronger recent performance. Models without adaptive learning continued selecting such providers for a longer period because their decisions remained disproportionately influenced by historical performance. Similarly, models without blockchain-based verification were more susceptible to inaccurate reputation information because provider evaluations could not be cross-validated against immutable transaction histories. The integrated model therefore demonstrated not only higher average trust values but also greater responsiveness to changes in the behavioral credibility of service providers.

Scalability analysis showed that increasing the number of candidate services and required manufacturing subtasks increased computational requirements for all evaluated approaches, with the most pronounced increase observed for the complete multi-objective framework. Nevertheless, solution quality remained comparatively stable as problem size increased. For smaller service-selection instances, the proposed model typically reached stable solutions with limited computational overhead relative to the baseline approaches. In medium-sized scenarios, computational time increased because the number of possible service compositions expanded substantially, but the algorithm continued to identify highly feasible non-dominated solutions. In the largest simulated scenarios, the search space became considerably more complex; however, Pareto-front diversity and request-success rates remained within acceptable ranges. These findings indicated that the proposed method did not depend on exhaustive enumeration of all possible service combinations and was therefore capable of operating under comparatively large service populations.

Overall, the findings demonstrated that the proposed multi-objective trust-aware cloud manufacturing service selection model achieved a more favorable balance among operational efficiency, service reliability, provider trustworthiness, adaptability, and environmental sustainability than the comparison approaches. The model did not consistently minimize direct economic cost or computational time, but it substantially reduced service failures, increased successful completion of composite manufacturing requests, selected more trustworthy service providers, and decreased environmental burden. The results also demonstrated that blockchain-based trust, adaptive learning, and sustainability constraints contributed distinct but complementary benefits. Blockchain improved the integrity of trust evidence, adaptive learning allowed the model to respond to changing provider performance, and environmental constraints ensured that operationally attractive service compositions remained compatible with sustainability objectives. The convergence, robustness, sensitivity, ablation, and scalability analyses collectively indicated that the proposed framework remained effective across different levels of problem complexity and environmental uncertainty, thereby providing empirical support for its application to dynamic cloud manufacturing service selection.

4. Discussion and Conclusion

The present study developed and evaluated a multi-objective trust-aware service selection model for dynamic cloud manufacturing by integrating blockchain-based trust verification, adaptive learning, reliability-oriented decision criteria, and explicit sustainability constraints. The findings showed that the proposed integrated model outperformed the conventional QoS-based model and the partial comparison models on the most important operational, trust, and environmental indicators. Although the proposed framework produced a modest increase in normalized service cost and computation time, it achieved higher aggregate reliability, greater mean provider trust, a higher manufacturing request success rate, a lower service failure rate, substantially lower energy consumption and carbon emissions, and a greater proportion of constraint-feasible service compositions. These findings indicate that the principal contribution of the proposed model lies not in minimizing one isolated criterion, but in producing a more balanced composition of cloud manufacturing services under conditions in which cost, time, trustworthiness, operational reliability, and environmental performance interact and frequently conflict.

One of the most important results was the increase in aggregate reliability from 0.901 under the conventional QoS model to 0.952 under the proposed integrated model, together with an improvement in manufacturing request success from 88.71% to 97.18% and a decline in service failure from 10.96% to 3.21%. These results suggest that service selection based only on conventional performance attributes is insufficient in a dynamic manufacturing environment because nominal QoS values do not fully represent the risk associated with unstable or deteriorating providers. This finding is consistent with the reliability-centered service composition perspective reported by Mahroo et al., who showed that reliability must be incorporated explicitly into cloud manufacturing composition rather than being treated as an incidental outcome of cost or time optimization (Mahroo et al., 2026). It also aligns with Gao et al., who demonstrated that combining QoS with robustness criteria improves the suitability of cloud manufacturing compositions under uncertainty (Gao et al., 2022). The present results extend this line of research by showing that reliability can be further improved when robustness is supported by verifiable transaction evidence and continuously updated service assessments.

The improvement in operational performance can also be explained by the adaptive learning component of the proposed model. In the ablation analysis, removing adaptive learning reduced the request success rate from 97.18% to 92.76% and increased the failure rate from 3.21% to 7.31%. This pattern indicates that static estimates of provider quality and trust become progressively less useful when manufacturing conditions change. The finding closely corresponds with research emphasizing dynamic service composition and runtime performance prediction. Hu et al. showed that cloud manufacturing service composition must respond to changing execution conditions through online policies rather than relying exclusively on static pre-execution decisions (Hu et al., 2024). Aalikhani et al. similarly demonstrated the importance of runtime QoS prediction for dynamically revising service composition according to current performance information (Aalikhani et al., 2026). The present study supports these conclusions and further suggests that adaptive learning becomes particularly effective when it is coupled with trust information, because the system can distinguish a temporary performance fluctuation from a sustained pattern of deterioration.

The robustness analysis across low-, medium-, and high-dynamic environments further confirmed the adaptive capacity of the proposed model. Although aggregate reliability declined from 0.963 in low-dynamic conditions to 0.939 under high dynamism, and the request success rate declined from 98.64% to 95.49%, the framework maintained a high level of operational performance despite increased volatility. The model also required more iterations to recover after service deterioration as environmental dynamism increased, but the adaptation recovery time remained limited and the incorrect retention of deteriorated providers remained below 6% even in the most unstable scenario. These findings support the argument that adaptive cloud manufacturing systems should not be judged only according to absolute performance under stable conditions, but also according to how quickly and accurately they recover from changing resource conditions. The results align with the dynamic online service composition perspective of Hu et al. (Hu et al., 2024) and with predictive process monitoring approaches that revise QoS expectations during runtime (Aalikhani et al., 2026). They are also compatible with Fazeli et al., who demonstrated that learning-based cloud manufacturing service composition can improve decision quality in complex and changing environments (Fazeli et al., 2024).

The trust-related findings were similarly notable. Mean selected-provider trust increased from 0.718 in the conventional QoS model to 0.893 in the complete model, whereas removing blockchain-based trust management reduced the value to 0.782. This result indicates that historical reputation or QoS information alone cannot adequately capture the behavioral credibility of service providers. Blockchain-based transaction verification appears to improve service selection because it creates a tamper-resistant basis for evaluating whether providers have actually fulfilled previously declared service commitments. This result directly supports the blockchain-based trust model proposed by Vatankhah Barenji, who argued that distributed ledgers can provide a more reliable foundation for trust management in cloud manufacturing (Vatankhah Barenji, 2022). It is also consistent with Tong et al., who integrated blockchain into multi-objective cloud manufacturing service composition and highlighted its potential to support trusted interactions within decentralized manufacturing environments (Tong et al., 2023). Similarly, Meng et al. demonstrated that blockchain-enabled decentralized service selection can enhance QoS-aware cloud manufacturing decisions by reducing reliance on centralized trust structures (Meng et al., 2025).

The ablation results provide further insight into why blockchain alone is not sufficient. When adaptive learning was removed, provider trust remained higher than in the model without blockchain, but operational outcomes deteriorated substantially. This suggests that immutable transaction histories must be interpreted dynamically rather than accumulated as undifferentiated historical evidence. Older transactions can become less representative of current provider behavior, particularly when equipment conditions, workload, or organizational behavior changes. The present trust updating mechanism therefore benefited from assigning greater weight to more recent interactions. This interpretation is compatible with Yang et al., who emphasized trustworthy collaborative evaluation across multiple service subjects rather than static reputation assessment (Yang, Jiang, et al., 2025). It also aligns with Alhudhaif and Polat, who proposed an adaptive fuzzy and graph-learning framework for robust trust management in cloud manufacturing, emphasizing the need to account for relational structure, uncertainty, and dynamic trust evolution (Alhudhaif & Polat, 2026). Thus, the current findings reinforce the view that trustworthy service selection requires both credible historical evidence and adaptive interpretation of that evidence.

A further major finding concerned environmental sustainability. The complete framework reduced normalized energy consumption from 0.514 under the conventional QoS model to 0.386 and reduced normalized carbon emissions from 0.501 to 0.369. The ablation model without sustainability constraints produced substantially worse environmental results, with energy consumption increasing to 0.512 and carbon emissions to 0.498, despite maintaining a slightly higher request success rate. This result is particularly important because it demonstrates that environmental benefits do not necessarily arise automatically from improved operational efficiency or reliability. In fact, a model optimized only for operational performance may select faster or more reliable services that consume greater energy or generate more emissions. The findings therefore support the explicit inclusion of environmental criteria within the service selection problem.

These results are consistent with previous sustainability-oriented cloud manufacturing research. Chen et al.

demonstrated that low-carbon cloud manufacturing service composition requires direct incorporation of carbon-related objectives into multi-objective optimization (Chen et al., 2023). Shi similarly incorporated carbon emission hierarchy into a cloud manufacturing recommendation model, indicating that environmental burden should influence service recommendation at the selection stage (Shi, 2023). Peng et al. extended this perspective by showing that energy-aware service selection and scheduling can improve environmental performance without completely sacrificing operational objectives (Peng et al., 2025). More recently, Tang et al. proposed a QoS- and sustainability-driven two-stage service composition framework, further demonstrating that sustainability can be jointly optimized with service quality through structured multi-objective methods (Tang et al., 2026). The present findings are strongly aligned with this body of research and additionally show that sustainability constraints can be integrated with blockchain trust and adaptive learning without undermining overall request success.

The multi-objective nature of the proposed framework also explains why the complete model did not produce the lowest service cost. The conventional QoS model achieved a normalized cost of 0.421 compared with 0.457 for the integrated model. This trade-off is theoretically expected because a model that imposes minimum trust thresholds, reliability requirements, and environmental restrictions excludes some low-cost but lower-quality or less sustainable alternatives. The sensitivity analysis confirmed that increasing the weight assigned to cost reduced expenditure but modestly weakened trust and reliability, whereas trust-dominant and sustainability-dominant settings produced higher-quality outcomes at greater economic cost. This supports the fundamental logic of multi-objective cloud manufacturing optimization. Wang et al. demonstrated the need for adaptive environmental selection in many-objective cloud manufacturing service selection because no single service composition can optimize all objectives simultaneously (Wang et al., 2021). Liu et al. similarly showed that multi-objective optimization is necessary when service composition must address fuzzy and competing user requirements (Liu et al., 2026). The present study extends these insights by demonstrating that trust and sustainability add additional dimensions to the trade-off structure and should therefore be represented directly within the Pareto optimization process.

The findings also correspond with research emphasizing resource interactions and system-level service relationships. The proposed framework treated composite manufacturing requests as networks of interdependent subtasks rather than as isolated service choices. This is consistent with Yang et al., who showed that dual-layer coupled networks can improve service composition optimization by representing relationships among manufacturing services (Yang, Ding, et al., 2025). Zhang et al. similarly emphasized service association impact in cloud service composition, demonstrating that dependencies among services influence the quality of the overall composition (Zhang et al., 2025). The current results support these perspectives because the proposed model maintained high feasibility rates even as task complexity and service dynamism increased. This suggests that trust-aware and sustainability-aware optimization can remain effective when service relationships are explicitly considered within composite manufacturing workflows.

The feasibility findings are particularly relevant in this regard. The proposed integrated model generated 98.16% constraint-feasible solutions, compared with 89.46% for the conventional QoS model. This result indicates that multi-objective optimization becomes more operationally meaningful when it includes explicit feasibility requirements rather than simply ranking solutions according to aggregate quality scores. Minimum trust, reliability, capacity, deadline, and sustainability thresholds prevented the algorithm from choosing mathematically attractive but operationally inappropriate service combinations. This principle is consistent with the three-tier composition approach of Lim et al., which emphasized coordinated decision-making across user, service, and system levels (Lim et al., 2022). It also corresponds with Xiong et al., who considered adaptive resource partitioning and multi-user benefit in hybrid cloud manufacturing, showing that optimal selection must reflect resource feasibility and competing user demands (Xiong et al., 2024). These studies collectively support the view that cloud manufacturing service selection should be formulated as a constrained systems optimization problem rather than as a simple scoring exercise.

The convergence analysis also supported the computational viability of the proposed model. The largest improvement occurred during the first approximately 20 iterations, followed by gradual refinement and stabilization after approximately 60 iterations. Although the complete model required more computation time than the baseline approaches, its mean execution time of 3.06 seconds remained moderate relative to the improvements achieved in reliability, trust, feasibility, and sustainability. This result is consistent with evolutionary and hybrid optimization research showing that advanced service composition methods typically require additional computational effort in exchange for higher-quality Pareto solutions. Wang et al. demonstrated that adaptive environmental selection can maintain solution quality across many objectives (Wang et al., 2021), while Zhang et al. used an improved NSGA-II approach to enhance service composition performance under service association effects (Zhang et al., 2025). Liu et al. likewise used an improved NSGA-III algorithm to address fuzzy demands in multi-objective composition (Liu et al., 2026). The convergence behavior observed in the present study suggests that the added complexity introduced by trust learning and sustainability constraints did not prevent stable optimization.

The scalability analysis provided additional support for the proposed architecture. As the number of candidate services and subtasks increased, computational time rose, but the model preserved relatively stable solution quality and request success. This finding is important because cloud manufacturing environments can involve large candidate pools and multiple users competing for shared resources. The result is consistent with the adaptive resource partitioning framework proposed by Xiong et al., which addressed multi-user benefit in hybrid manufacturing environments (Xiong et al., 2024), and with the optimization approach of Yang et al., which used coupled-network structures and improved particle swarm optimization to address complex composition spaces (Yang, Ding, et al., 2025). The present study indicates that adding blockchain trust and sustainability constraints does not necessarily make such models computationally impractical, provided that the search procedure avoids exhaustive enumeration.

Taken together, the findings show that blockchain, adaptive learning, and sustainability constraints contribute distinct but mutually reinforcing benefits. Blockchain improves the credibility and traceability of service interaction evidence; adaptive learning ensures that trust and QoS estimates remain responsive to changing service behavior; multi-objective optimization manages trade-offs among competing economic and operational objectives; and sustainability constraints prevent environmentally burdensome solutions from being selected solely because they are inexpensive or fast. This integrated interpretation is consistent with the direction of recent cloud manufacturing research, in which dynamic QoS prediction (Aalikhani et al., 2026), adaptive trust management (Alhudhaif & Polat, 2026), blockchain-enabled service selection (Meng et al., 2025), reliability-centered composition (Mahroo et al., 2026), and sustainability-driven optimization (Tang et al., 2026) have each emerged as important but often separately treated developments. The present study contributes by demonstrating their complementary value within a single service selection architecture.

This study has several limitations that should be considered when interpreting the findings. First, the evaluation was based on a controlled simulation environment rather than a live industrial cloud manufacturing platform. Although the simulation incorporated heterogeneity, service deterioration, workload variation, unreliable provider behavior, and environmental constraints, actual manufacturing ecosystems may exhibit additional forms of uncertainty, including communication failures, strategic manipulation, regulatory differences, supply disruptions, and heterogeneous data quality. Second, several service characteristics, including trust behavior, energy consumption, carbon emissions, and performance deterioration, were modeled through predefined distributions and update rules. These assumptions were necessary for controlled comparison but may not fully reproduce the statistical behavior of real industrial systems. Third, blockchain functionality was represented primarily as a trust and transaction verification infrastructure, while practical implementation factors such as consensus latency, storage overhead, smart-contract vulnerabilities, interoperability, and blockchain energy costs were not modeled in detail. Fourth, the adaptive learning mechanism emphasized recent performance updates rather than employing a fully autonomous online reinforcement learning architecture. Finally, the environmental assessment concentrated mainly on operational energy consumption and carbon emissions and did not incorporate broader life-cycle indicators such as material depletion, embodied energy, waste generation, water consumption, or end-of-life impacts.

Future research should validate the proposed framework using empirical datasets obtained from operational cloud manufacturing platforms, industrial Internet of Things infrastructures, or digital manufacturing networks in which real-time service execution records can be observed. Further studies should examine the model under larger-scale and more heterogeneous service ecosystems, including geographically distributed providers, multi-factory networks, and cross-platform service federation. Future work could also integrate more advanced online learning methods, such as reinforcement learning, graph neural networks, or federated learning, to enhance dynamic prediction while preserving organizational data privacy. The trust model could be expanded to account for collusion, malicious ratings, identity manipulation, strategic provider behavior, and adversarial attacks against blockchain-enabled service ecosystems. Additional research should evaluate alternative blockchain architectures and consensus mechanisms to determine how ledger latency, transaction cost, and scalability influence manufacturing service selection in real time. Sustainability assessment should also be broadened toward full life-cycle environmental evaluation and could include renewable energy availability, material circularity, water use, waste, and embodied carbon. Finally, future studies should investigate multi-user negotiation and fairness mechanisms so that the framework can allocate constrained manufacturing resources equitably when numerous users simultaneously request overlapping services.

For practical implementation, cloud manufacturing platform operators should avoid relying exclusively on advertised QoS information or static provider rankings when recommending manufacturing services. Service selection systems should combine operational performance with verified historical interaction records and should continuously update provider profiles as new evidence becomes available. Blockchain or other tamper-resistant distributed ledger mechanisms can be used to preserve transaction integrity, particularly in decentralized manufacturing networks involving previously unacquainted participants. Platform managers should also define minimum thresholds for trust, reliability, capacity, and environmental performance before optimization is performed so that infeasible or unacceptable services are excluded early in the decision process. Manufacturing firms adopting such systems should configure objective weights according to the strategic nature of each production request, assigning greater importance to speed for urgent orders, trust and reliability for critical components, and energy or carbon performance for sustainability-sensitive manufacturing. Finally, implementation should include monitoring dashboards that report not only cost and completion time but also provider trust trends, recent performance deterioration, energy use, carbon emissions, and constraint violations, thereby enabling managers to understand why a particular service composition has been recommended and to intervene when operational priorities change.

Authors' Contributions

Authors contributed equally to this article.

Declaration

In order to correct and improve the academic writing of our paper, we have used the language model ChatGPT.

Transparency Statement

Data are available for research purposes upon reasonable request to the corresponding author.

Acknowledgments

We would like to express our gratitude to all individuals helped us to do the project.

Declaration of Interest

The authors report no conflict of interest.

Funding

According to the authors, this article has no financial support.

Ethics Considerations

In this research, ethical standards including obtaining informed consent, ensuring privacy and confidentiality were considered.

References

Aalikhani, R., Fathian, M., Rasouli, M. R., & Eshuis, R. (2026). Dynamic cloud manufacturing service composition based on runtime QoS prediction: A predictive process monitoring based method. *Computers & Industrial Engineering*, *212*, 111701. https://doi.org/10.1016/j.cie.2025.111701

Alhudhaif, A., & Polat, K. (2026). A graph learning and adaptive fuzzy framework for robust trust management in cloud manufacturing. *Applied Soft Computing*, *198*, 115259. https://doi.org/10.1016/j.asoc.2026.115259

Chen, C., Yu, J., Lu, J., Su, X., Zhang, J., Feng, C., & Ji, W. (2023). Service composition and optimal selection of low carbon cloud manufacturing based on NSGA II SA algorithm. *Processes*, *11*(2), 340. https://doi.org/10.3390/pr11020340

Fazeli, M. M., Farjami, Y., & Jalaly Bidgoly, A. (2024). An efficient cloud manufacturing service composition approach using deep reinforcement learning. *Computers & Industrial Engineering*, *195*, 110446. https://doi.org/10.1016/j.cie.2024.110446

Gao, Y., Yang, B., Wang, S., Zhang, Z., & Tang, X. (2022). Bi objective service composition and optimal selection for cloud manufacturing with QoS and robustness criteria. *Applied Soft Computing*, *128*, 109530. https://doi.org/10.1016/j.asoc.2022.109530

Hu, Y., Yang, Y., & Wu, F. (2024). Dynamic cloud manufacturing service composition with re entrant services: An online policy perspective. *International Journal of Production Research*, *62*(9), 3263-3287. https://doi.org/10.1080/00207543.2023.2230317

Lim, M. K., Xiong, W., & Wang, Y. (2022). A three tier programming model for service composition and optimal selection in cloud manufacturing. *Computers & Industrial Engineering*, *167*, 108006. https://doi.org/10.1016/j.cie.2022.108006

Liu, X., Yang, R., Li, X., & Wang, X. V. (2026). A cloud manufacturing service composition optimization method for fuzzy demands based on improved NSGA III algorithm. *Robotics and Computer Integrated Manufacturing*, *97*, 103106. https://doi.org/10.1016/j.rcim.2025.103106

Mahroo, F., Moradi, N., Aftabi, N., Houshmand, M., & Fatahi Valilai, O. (2026). A novel integer linear model for reliability centric service composition in cloud manufacturing. *Computers & Industrial Engineering*, *211*, 111603. https://doi.org/10.1016/j.cie.2025.111603

Meng, K., Wu, Z., Bilal, M., Xia, X., & Xu, X. (2025). Blockchain enabled decentralized service selection for QoS aware cloud manufacturing. *Expert Systems*, *42*(1), e13602. https://doi.org/10.1111/exsy.13602

Peng, G. X., Wen, Y. P., Liu, J. X., Kang, G. S., Zhang, B. M., & Zhou, M. H. (2025). Energy aware cloud manufacturing service selection and scheduling optimization. *International Journal of Computer Integrated Manufacturing*, *38*(3), 309-334. https://doi.org/10.1080/0951192X.2024.2333024

Shi, Z. (2023). Cloud manufacturing service recommendation model based on GA ACO and carbon emission hierarchy. *Journal of Intelligent & Fuzzy Systems*, *44*(2), 2007-2017. https://doi.org/10.3233/JIFS-222386

Tang, C., Zhao, S., Su, H., & Chen, B. (2026). A QoS and sustainability driven two stage service composition method in cloud manufacturing: Combining clustering and bi objective optimization. *Journal of global optimization*, *94*(2), 481-516. https://doi.org/10.1007/s10898-024-01430-z

Tong, J., Zhao, B., & An, Y. (2023). A novel multi objective service composition architecture for blockchain based cloud manufacturing. *Journal of Computational Design and Engineering*, *10*(1), 185-203. https://doi.org/10.1093/jcde/qwac130

Vatankhah Barenji, R. (2022). A blockchain technology based trust system for cloud manufacturing. *Journal of Intelligent Manufacturing*, *33*(5), 1451-1465. https://doi.org/10.1007/s10845-020-01735-2

Wang, T., Zhang, P., Liu, J., & Zhang, M. (2021). Many objective cloud manufacturing service selection and scheduling with an evolutionary algorithm based on adaptive environment selection strategy. *Applied Soft Computing*, *112*, 107737. https://doi.org/10.1016/j.asoc.2021.107737

Xiong, W., Wang, Y., Gao, S., Huang, X., & Wang, S. (2024). A multi objective service composition optimization method considering multi user benefit and adaptive resource partitioning in hybrid cloud manufacturing. *Journal of Industrial Information Integration*, *38*, 100564. https://doi.org/10.1016/j.jii.2024.100564

Yang, T., Ding, Y., & Chen, W. (2025). Trustworthy collaborative evaluation of multi service subjects in the cloud manufacturing model. *Alexandria Engineering Journal*, *113*, 1-11. https://doi.org/10.1016/j.aej.2024.11.021

Yang, T., Jiang, F., & Su, J. (2025). A service composition optimization approach for cloud manufacturing based on the dual layer coupled network and IPSO algorithm. *Computers & Industrial Engineering*, *210*, 111529. https://doi.org/10.1016/j.cie.2025.111529

Zhang, C., Wang, L., & He, K. (2025). Cloud service composition optimization based on service association impact and improved NSGA II algorithm. *Scientific reports*, *15*, 26001. https://doi.org/10.1038/s41598-025-10040-y