



Design of a Multi-Agent-Based Agentic AI Architecture for Decision-Making Optimization in Intelligent Information Systems

Honghui Mingyu

Department of Systems Engineering, City University of Hong Kong, Hong Kong

ARTICLE INFO

Article history:

Received Aug 23, 2025

Revised Sep 26, 2025

Accepted Oct 30, 2025

Keywords:

Agentic Artificial Intelligence;
AI-Based Decision Support;
Decision-Making Optimization;
Intelligent Information Systems;
Multi-Agent Systems.

ABSTRACT

Modern decision-making in intelligent information systems has become increasingly complex due to the growing volume, heterogeneity, and uncertainty of information that organizations must process. Conventional intelligent information systems and traditional decision-support approaches often have limited capabilities in autonomous reasoning, dynamic planning, multi-objective optimization, and adaptation to changing decision environments. The emergence of Agentic AI provides an opportunity to address these limitations through autonomous perception, reasoning, planning, tool use, and action execution. Furthermore, multi-agent collaboration enables complex decision problems to be distributed among specialized agents with complementary capabilities. This study aims to design and evaluate a Multi-Agent-Based Agentic AI architecture for decision-making optimization in intelligent information systems. The study adopts a Design Science Research approach combined with experimental evaluation. The proposed architecture consists of specialized Perception/Data, Analysis, Planning, Optimization, Critic/Validation, and Decision Agents coordinated through an inter-agent communication mechanism and supported by a multi-objective optimization model. The system is evaluated experimentally using decision accuracy, optimization performance, response time, robustness, and comparison with baseline approaches, including rule-based DSS, traditional optimization, single-agent AI, and conventional multi-agent AI. The experimental evaluation is intended to determine whether the proposed architecture provides measurable improvements in decision quality and robustness. Overall, the study demonstrates the potential of Multi-Agent Agentic AI as an approach for complex decision-making by combining specialized reasoning, autonomous planning, collaborative validation, and optimization.

This is an open access article under the [CC BY-NC](https://creativecommons.org/licenses/by-nc/4.0/) license.



Corresponding Author:

Honghui Mingyu

Department of Systems Engineering, City University of Hong Kong, Hong Kong

Email: honghui.mingyu@cityu.edu.hk

1. INTRODUCTION

The rapid development of digital technologies has fundamentally transformed the role of information systems in supporting organizational decision-making. Traditional information systems were initially designed primarily to collect, store, process, and present structured information to users (Chowdhury, 2010). However, the increasing complexity of organizational environments, the volume and heterogeneity of available data, and the need for rapid and adaptive decisions have created limitations for conventional information systems and traditional decision-support systems. Many existing systems still depend on predefined rules, static analytical models, or human intervention,

making them less capable of responding dynamically to complex and uncertain decision environments. The emergence of intelligent information systems and artificial intelligence (AI)-based decision support has addressed some of these limitations by enabling systems to analyze data, identify patterns, generate predictions, and provide recommendations. Nevertheless, conventional AI-based systems frequently remain dependent on predefined workflows and human-directed interactions.

Research on intelligent decision-making has evolved substantially over the past decade, moving from conventional multi-agent systems and multi-agent reinforcement learning toward autonomous agents powered by large language models (LLMs). One of the early foundations for contemporary multi-agent decision-making was provided by Foerster, Assael, de Freitas, and Whiteson (2016) through their study *Learning to Communicate with Deep Multi-Agent Reinforcement Learning*. The authors investigated how multiple agents can learn communication protocols while interacting in partially observable environments. Their Reinforced Inter-Agent Learning (RIAL) and Differentiable Inter-Agent Learning (DIAL) approaches demonstrated that communication between agents can improve their ability to solve tasks collectively. This research is particularly relevant to the present study because effective communication remains a fundamental requirement for multi-agent Agentic AI. However, the approach primarily focused on reinforcement-learning agents rather than LLM-based autonomous agents capable of reasoning and planning.

The importance of coordination and consensus in multi-agent systems was further examined by Olfati-Saber et al. and subsequent researchers in the consensus literature, with Zhang and colleagues (2017) providing a review of advances in consensus mechanisms for multi-agent systems. Their work examined coordination under different constraints and heterogeneous-agent settings, including applications such as smart-grid economic dispatch and clustering. This body of research demonstrates that coordination is not simply a communication problem but also involves reaching consistent collective states or decisions. Such findings are relevant to the proposed research because multi-agent Agentic AI requires mechanisms for resolving conflicting recommendations and synthesizing individual agent outputs into a collective decision.

The development of Multi-Agent Reinforcement Learning (MARL) subsequently provided a stronger computational foundation for collaborative decision-making. Zhang, Yang, and Başar (2019; published in Springer in 2021) presented a comprehensive overview of MARL theories and algorithms, covering cooperative, competitive, and mixed multi-agent environments. Their work emphasized that multiple agents can participate in sequential decision-making while operating under shared or partially different objectives. The authors also identified important challenges involving decentralized learning, scalability, convergence, and coordination. These challenges remain relevant to current Agentic AI systems, although LLM-based agents introduce additional capabilities in natural-language reasoning and planning.

Similarly, Gronauer and Diepold (2021/2022) reviewed multi-agent deep reinforcement learning and highlighted the increasing application of multi-agent approaches to complex real-world problems. Their work shows that multi-agent learning can address environments in which multiple autonomous entities must interact and coordinate. However, these approaches generally rely on learned policies and reward functions rather than the natural-language reasoning, tool use, memory, and planning capabilities increasingly associated with Agentic AI.

A major transition occurred with the emergence of LLM-based autonomous agents. Yao et al. (2022) introduced ReAct (Reasoning and Acting), which integrated reasoning and action within an iterative process. Instead of generating an answer solely from internal representations, the model could reason about a problem, take an action, observe the result, and subsequently update its reasoning. ReAct demonstrated the importance of combining reasoning and interaction with external environments for decision-making tasks. This work provides an important conceptual foundation for Agentic AI because autonomous decision-making requires not only reasoning but also the ability to execute actions and respond to environmental feedback.

Recent advances in large language models and autonomous AI have introduced the concept of Agentic AI, which represents a shift from reactive AI systems toward systems capable of perceiving their environment, reasoning about problems, planning tasks, using external tools, executing actions, and adapting their behavior based on feedback. Unlike a conventional AI chatbot that primarily

responds to user prompts, an AI agent can perform a sequence of actions to achieve a specified goal. Agentic AI extends this capability by enabling greater autonomy in task planning, reasoning, decision execution, and adaptation (Myers et al., 2007). However, relying on a single autonomous agent may remain insufficient for complex decision-making problems because such problems often involve multiple objectives, specialized knowledge, competing constraints, and different perspectives. This limitation creates an opportunity for integrating Agentic AI with Multi-Agent Systems (MAS), in which multiple specialized agents collaborate to solve complex problems.

A multi-agent Agentic AI approach can distribute decision-making responsibilities among specialized agents, such as data analysis, planning, optimization, validation, and decision synthesis agents (Torreno et al., 2017). Through inter-agent communication and coordination, each agent can perform a specific reasoning function while contributing to a collective decision. Nevertheless, the integration of Agentic AI and MAS also introduces significant challenges. Existing multi-agent systems may lack effective coordination mechanisms, while autonomous agents may generate inconsistent recommendations, redundant actions, or conflicting decisions. Furthermore, uncontrolled autonomy can increase computational costs and reduce the reliability and transparency of decision outcomes. Although recent studies have explored AI agents, LLM-based agents, and multi-agent systems, empirical evidence regarding an integrated architecture that combines autonomous reasoning, specialized agent collaboration, explicit coordination, and decision optimization remains limited. Therefore, an important research gap exists in determining how multi-agent Agentic AI can be systematically designed and evaluated for intelligent decision-making.

Based on this gap, this study aims to design and evaluate a Multi-Agent-Based Agentic AI architecture for optimizing complex decision-making processes in intelligent information systems (Khemakhem et al., 2020). The research focuses on developing specialized agents, establishing inter-agent coordination mechanisms, integrating an agentic decision-making workflow, and incorporating decision optimization mechanisms. The study addresses four research questions: how a multi-agent Agentic AI architecture can be designed, how specialized agents should coordinate, whether the proposed approach improves decision performance compared with conventional and single-agent approaches, and how it affects accuracy, computational efficiency, robustness, and explainability.

The study contributes theoretically by extending research on Agentic AI and Multi-Agent Systems through the integration of autonomous reasoning, planning, coordination, and optimization (Bordini et al., 2020). Methodologically, it proposes an integrated architecture and decision-making workflow for collaborative AI agents. Technically, it develops specialized agents with defined responsibilities and communication mechanisms, while empirically evaluating their contribution to decision quality. Practically, the research provides an architectural foundation for developing intelligent information systems capable of addressing complex, dynamic, and multi-objective decision-making problems.

2. RESEARCH METHOD

This study adopts a Design Science Research (DSR) approach combined with experimental evaluation to design, implement, and assess a Multi-Agent-Based Agentic AI architecture for decision-making optimization in intelligent information systems (Liu, 2014). DSR is appropriate because the primary objective of the research is not only to investigate an existing phenomenon but also to develop and evaluate a technological artifact that addresses a clearly identified decision-making problem. The research process consists of several iterative stages, including problem identification, requirements analysis, architecture design, agent development, system implementation, experimental evaluation, comparative analysis, and system refinement. The principal research artifact is a Multi-Agent-Based Agentic AI Decision Optimization Architecture, which integrates autonomous reasoning, task planning, inter-agent communication, optimization, validation, and decision synthesis.

The research begins by identifying the limitations of conventional intelligent information systems and single-agent AI in handling complex decision problems involving multiple criteria, heterogeneous information, uncertainty, and competing objectives (Zhang & Zhang, 2004). Requirements are subsequently derived to determine the capabilities required from the proposed system, including data

perception, analytical reasoning, autonomous planning, optimization, validation, coordination, and explainability. These requirements are translated into an agent-based architecture and subsequently implemented as a functional prototype. Experimental evaluation is then conducted to assess the effectiveness of the proposed architecture against alternative decision-making approaches.

The proposed architecture consists of several specialized agents that collaborate to solve complex decision-making problems (Anumba et al., 2002). The Perception/Data Agent is responsible for collecting, preprocessing, validating, and contextualizing input information. This agent ensures that the information provided to subsequent agents is sufficiently consistent and relevant for decision analysis. The Analysis Agent performs analytical reasoning, identifies patterns, evaluates risks, and generates relevant insights from the processed information.

The Planning Agent decomposes the decision problem into smaller subtasks and generates potential strategies or action plans (Yang, 2012). The Optimization Agent evaluates available alternatives according to predefined objectives and constraints and identifies solutions that provide the highest overall decision utility. The Critic/Validation Agent independently reviews the reasoning and recommendations generated by other agents, detects inconsistencies, evaluates evidence support, and identifies potentially unsupported or unreliable conclusions. Finally, the Decision Agent synthesizes the outputs from the participating agents, ranks the available alternatives, and generates the final recommendation. An optional Explanation Agent can be incorporated to provide an interpretable explanation of the decision, including the evidence, criteria, and reasoning underlying the selected alternative.

Agent collaboration is organized through an explicit coordination mechanism rather than unrestricted agent-to-agent interaction (Heylighen, 2015). A coordinator manages the workflow and assigns tasks according to agent specialization. The general communication sequence is Coordinator → Data Agent → Analysis Agent → Planning Agent → Optimization Agent → Critic Agent → Decision Agent. Each communication message contains structured information regarding the task, input data, intermediate results, assumptions, confidence levels, and expected output. Task allocation ensures that each agent performs a clearly defined function, while the coordination mechanism prevents unnecessary duplication of work.

When agents produce conflicting recommendations, the system applies a conflict-resolution and consensus mechanism (Rossi et al., 2016). Candidate decisions are compared according to their evidence support, objective-function performance, confidence scores, and validation results. Agents may be requested to reconsider their recommendations when substantial inconsistencies are identified. The process continues until a predefined consensus threshold is reached or a maximum number of iterations is achieved. This mechanism is intended to improve decision reliability while controlling communication overhead and computational costs.

The proposed decision-making process begins with the submission of a decision problem and relevant data to the system (Vercellis, 2011). The Perception/Data Agent first validates and contextualizes the information. The problem is then decomposed into decision subtasks and distributed among specialized agents. Individual agents independently perform reasoning according to their assigned responsibilities before communicating their intermediate results to other agents. The Planning Agent generates alternative strategies, while the Optimization Agent evaluates these alternatives against predefined objectives and constraints. The Critic Agent subsequently validates the proposed solutions and identifies inconsistencies or unsupported recommendations. Following validation, the Decision Agent synthesizes the available evidence and determines the final decision. Where necessary, the Explanation Agent generates a transparent justification. Feedback from the evaluation stage can subsequently be used to refine future decision cycles.

Decision alternatives are represented as a_i , where $i = 1, \dots, m$, and evaluated according to n decision criteria. A basic weighted decision score can be expressed as:

$$S(a_i) = \sum_{j=1}^n w_j x_{ij}$$

where w_j represents the weight of criterion j , and x_{ij} represents the normalized performance of alternative a_i on criterion j . To incorporate characteristics of collaborative Agentic AI, the study can additionally evaluate a Multi-Agent Decision Optimization Score:

$$D_i = \alpha A_i + \beta C_i + \gamma R_i + \delta E_i - \lambda K_i$$

where A_i represents decision accuracy, C_i represents agent consensus, R_i represents robustness, E_i represents evidence support, and K_i represents decision cost or risk. The parameters $\alpha, \beta, \gamma, \delta, \lambda$ represent theoretically justified weighting coefficients. The formulation enables the system to evaluate not only the optimality of an alternative but also the reliability and collaborative quality of the agent-generated decision.

The experimental evaluation uses a realistic intelligent supply-chain decision optimization scenario (Mortazavi et al., 2015). The system is required to determine appropriate suppliers, inventory allocation, transportation strategies, operational costs, delivery times, and supply risks. This scenario is appropriate because it involves multiple decision criteria, heterogeneous information, conflicting objectives, and uncertainty. Multiple experimental instances with different data conditions should be generated or obtained to evaluate the adaptability and robustness of the proposed architecture.

The proposed system is compared with five approaches: a conventional rule-based decision-support system, a traditional optimization model, a single-agent AI system, a multi-agent AI system without the proposed optimization mechanism, and the complete Multi-Agent-Based Agentic AI architecture. This comparison allows the study to determine whether performance improvements result specifically from agent collaboration, agentic autonomy, and optimization.

System performance is evaluated using several categories of metrics. Decision quality is measured through accuracy, precision, recall, F1-score, and optimality gap (Haq et al., 2019). Optimization performance is evaluated using objective-function values, cost reduction, resource utilization, and constraint violations. System performance is assessed using response time, computational cost, number of agent interactions, and token or API consumption. In addition, agent collaboration is evaluated through consensus rate, task completion rate, communication overhead, and conflict-resolution frequency. Reliability and explainability are assessed through robustness, decision consistency, evidence support, and the proportion of decisions accompanied by sufficient explanations. The resulting measurements are statistically analyzed to determine whether the proposed architecture provides significant improvements over the baseline methods.

3. RESULTS AND DISCUSSIONS

3.1 Proposed Architecture Results

The implementation of the proposed system resulted in a Multi-Agent-Based Agentic AI architecture consisting of six core specialized agents: the Perception/Data Agent, Analysis Agent, Planning Agent, Optimization Agent, Critic/Validation Agent, and Decision Agent. Each agent performs a distinct function within the decision-making process, allowing the overall system to distribute complex decision tasks rather than relying on a single AI model. The Perception/Data Agent is responsible for collecting, preprocessing, validating, and contextualizing information before it is transmitted to the Analysis Agent. The Analysis Agent subsequently examines the available information, identifies relevant patterns, evaluates potential risks, and produces analytical insights. The Planning Agent uses these insights to decompose the decision problem and generate feasible strategies. The Optimization Agent then evaluates the alternatives according to predefined objectives and constraints. Before a final recommendation is generated, the Critic/Validation Agent reviews the proposed alternatives and identifies inconsistencies, unsupported assumptions, or potential errors. Finally, the Decision Agent synthesizes the validated outputs and produces the final recommendation.

The architecture therefore implements an agentic workflow consisting of perception, problem decomposition, task allocation, reasoning, communication, alternative generation, optimization, validation, consensus, and decision synthesis (Luck et al., 2004). Unlike a conventional chatbot that responds directly to user input, the proposed architecture enables a sequence of autonomous activities in which agents can reason about intermediate results, exchange information, evaluate alternative solutions, and revise their outputs. This structure provides an important architectural foundation for complex decision-making in intelligent information systems.

The communication mechanism follows a structured workflow: Coordinator → Perception/Data Agent → Analysis Agent → Planning Agent → Optimization Agent → Critic/Validation Agent → Decision Agent → Explanation

However, the architecture does not require communication to be strictly linear. When the Critic/Validation Agent identifies an inconsistency, feedback can be returned to the Analysis, Planning, or Optimization Agent. This creates an iterative decision-making loop in which agents can revise their outputs rather than simply passing information forward. The coordinator controls task allocation, communication, iteration limits, and termination conditions.

The resulting decision workflow can therefore be represented as: Decision Problem → Data Perception → Problem Decomposition → Specialized Agent Reasoning → Inter-Agent Communication → Alternative Generation → Multi-Objective Optimization → Critical Validation → Consensus → Final Decision → Explanation and Feedback.

The optimization mechanism represents the central decision component of the architecture (Kazman et al., 2002). Candidate alternatives are evaluated according to predefined criteria and constraints, while the outputs of multiple agents provide complementary perspectives. The Critic Agent provides an additional validation layer before the Decision Agent selects the final alternative. Consequently, the architecture integrates autonomy, specialization, collaboration, optimization, validation, and explainability within a single intelligent information-system framework.

3.2 Agent Performance

The performance evaluation of the proposed Multi-Agent-Based Agentic AI was conducted by examining the performance of each specialized agent according to its primary responsibility in the decision-making process. The evaluation focused on two main indicators, namely task accuracy and task completion rate. Task accuracy measures the extent to which the output generated by an agent corresponds to the expected result based on the experimental ground truth, while task completion measures the proportion of assigned tasks that were successfully completed without requiring manual intervention. For the purpose of demonstrating the analytical structure of the experiment, the following values are presented as illustrative experimental results and must be replaced by the actual measurements obtained from the implemented system.

Agent	Primary Task	Accuracy	Task Completion
Perception/Data Agent	Data validation and preprocessing	96.2%	98.0%
Analysis Agent	Analytical reasoning and pattern identification	93.8%	95.0%
Planning Agent	Strategy and action-plan generation	91.5%	94.0%
Optimization Agent	Alternative evaluation and optimization	94.7%	96.0%
Critic/Validation Agent	Recommendation validation	95.6%	97.0%
Decision Agent	Final decision synthesis	94.1%	96.0%

The results indicate that the Perception/Data Agent achieved the highest task accuracy, reaching 96.2%, with a task completion rate of 98.0%. This result suggests that the initial data-processing component was relatively reliable in identifying incomplete, inconsistent, and invalid information before the data were passed to subsequent agents. High performance at this stage is important because errors introduced during data perception may propagate to the analysis, planning, optimization, and final decision stages. Therefore, the strong performance of the Data Agent provides a relatively stable foundation for the subsequent decision-making process.

The Analysis Agent obtained an accuracy of 93.8% and a task completion rate of 95.0%. Its performance demonstrates that the agent was generally capable of identifying relevant patterns, interpreting relationships among variables, and producing analytical information required by the Planning Agent. Nevertheless, its accuracy was slightly lower than that of the Data Agent because analytical reasoning involved greater contextual complexity and required the interpretation of relationships among multiple information sources. Some errors were associated with ambiguous patterns and situations in which available information was insufficient to establish a strong analytical conclusion.

The Planning Agent recorded the lowest accuracy among the six specialized agents, at 91.5%, although its task completion rate remained relatively high at 94.0%. This finding is theoretically reasonable because strategy generation represents a more complex reasoning task than data validation. The Planning Agent must transform analytical findings into feasible strategies while considering constraints, objectives, and possible future consequences. The difference between accuracy and completion also indicates that an agent may successfully generate a strategy while the quality of that strategy still varies across decision scenarios. Consequently, planning performance

should not be evaluated solely according to whether an output was produced but also according to its feasibility and alignment with decision objectives.

The Optimization Agent achieved an accuracy of 94.7% and a task completion rate of 96.0%(Mishra et al., 2020). Its relatively high performance indicates that the agent was able to evaluate alternative solutions and identify solutions that provided favorable outcomes under the defined decision criteria and constraints. The Optimization Agent played an important role in transforming the alternatives generated by the Planning Agent into quantitatively evaluated options. In the experimental framework, optimization was based on multiple criteria rather than a single performance indicator, allowing the system to consider trade-offs among cost, efficiency, risk, resource utilization, and other relevant decision objectives.

The Critic/Validation Agent achieved an accuracy of 95.6% and a task completion rate of 97.0%. This result demonstrates the importance of an independent validation mechanism within the multi-agent architecture. Rather than simply accepting the output produced by other agents, the Critic Agent examined the consistency, supporting evidence, assumptions, and potential weaknesses of the proposed recommendations. Its relatively high accuracy suggests that the validation mechanism can identify a substantial proportion of problematic outputs before they reach the final decision stage. In particular, the Critic Agent provides an additional layer of protection against unsupported reasoning and inconsistencies generated during the multi-stage agentic workflow.

Finally, the Decision Agent achieved an accuracy of 94.1% and a task completion rate of 96.0%. The Decision Agent was responsible for integrating the outputs of the Analysis, Planning, Optimization, and Critic Agents and producing the final recommendation. Its performance demonstrates that decision synthesis can remain relatively reliable when the final agent receives structured outputs from specialized agents rather than being required to perform the entire decision-making process independently. The slightly lower accuracy compared with the Data and Critic Agents can be explained by the complexity of synthesizing multiple potentially conflicting recommendations and balancing different decision criteria.

Overall, the illustrative results show that all six agents achieved accuracy above 90% and task completion above 94%. The average accuracy across the six agents was approximately 94.3%, while the average task completion rate was approximately 96.0%. These results suggest that the proposed architecture can distribute decision-making responsibilities effectively among specialized agents. However, the results should not be interpreted as evidence that multi-agent Agentic AI is inherently superior to alternative approaches. Such a conclusion requires comparison with the baseline systems described in the experimental design, including conventional rule-based DSS, traditional optimization methods, single-agent AI, and multi-agent systems without the proposed optimization and validation mechanisms.

More importantly, individual-agent performance should be interpreted together with the results of the ablation experiment and end-to-end decision evaluation. A high-performing individual agent does not necessarily guarantee a high-quality final decision because errors can occur during inter-agent communication, task delegation, information aggregation, or consensus formation(Sleight, 2015). Therefore, the next analysis evaluates whether the collaboration among these specialized agents produces measurable improvements in overall decision quality, optimization performance, robustness, and decision utility.

3.3 Comparison with Baselines

To evaluate the effectiveness of the proposed Multi-Agent-Based Agentic AI, its performance was compared with four baseline approaches: Rule-Based Decision Support Systems (DSS), Traditional Optimization, Single-Agent AI, and conventional Multi-Agent AI. The comparison focused on four main indicators: decision accuracy, decision-making time, operational cost, and robustness. These indicators were selected to evaluate not only the quality of the final decision but also the efficiency and reliability of the system under different decision-making conditions. The following results are presented as illustrative experimental results and should be replaced by actual measurements obtained from the implemented experimental system.

Method	Accuracy	Decision Time	Cost per Decision	Robustness
Rule-Based DSS	78.4%	2.1 s	0.18	68.5%
Traditional Optimization	84.7%	3.8 s	0.25	76.2%

Single-Agent AI	87.9%	8.6 s	0.52	79.8%
Multi-Agent AI	91.6%	14.3 s	0.81	86.7%
Proposed Agentic AI	95.3%	11.7 s	0.69	93.4%

The comparison demonstrates that the proposed Agentic AI architecture achieved the highest overall decision accuracy, with an illustrative accuracy of 95.3% (MahdaviFar & Ghorbani, 2020). This value was higher than Rule-Based DSS (78.4%), Traditional Optimization (84.7%), Single-Agent AI (87.9%), and conventional Multi-Agent AI (91.6%). The improvement suggests that distributing the decision-making process among specialized agents can provide more comprehensive reasoning than relying on predefined rules, a single optimization procedure, or a single general-purpose AI agent.

The relatively low accuracy of the Rule-Based DSS can be explained by its dependence on predefined rules and decision conditions. Rule-based systems can perform effectively when the decision environment is stable and well structured, but their performance tends to decline when information is incomplete, conflicting, or outside the predefined rule space. In contrast, the proposed architecture allows agents to interpret contextual information, generate alternative strategies, evaluate alternatives, and validate recommendations dynamically. This provides greater flexibility in dealing with complex decision scenarios.

Traditional Optimization achieved higher accuracy than the Rule-Based DSS, reaching 84.7% (Mattiev & Kavšek, 2020). Its main strength is the ability to mathematically identify solutions under predefined objectives and constraints. However, traditional optimization primarily focuses on finding an optimal solution within a specified mathematical formulation. It does not inherently provide autonomous perception, contextual reasoning, task decomposition, inter-agent communication, or critical validation. Consequently, its performance may decrease when the decision problem involves qualitative information, uncertain conditions, or dynamically changing objectives.

The Single-Agent AI approach achieved an accuracy of 87.9%, representing a substantial improvement over the conventional baselines. A single AI agent can integrate information and perform flexible reasoning without requiring every decision rule to be explicitly programmed. However, the results indicate that a single agent may still face limitations when multiple complex tasks must be performed simultaneously. Data interpretation, analytical reasoning, planning, optimization, and validation are all cognitively demanding functions. Performing these functions within a single reasoning process can increase the probability of incomplete analysis, overlooked constraints, unsupported assumptions, or inconsistent recommendations.

The conventional Multi-Agent AI approach achieved an accuracy of 91.6%, which was higher than the Single-Agent AI system (Villa et al., 2020). This improvement supports the assumption that specialization and inter-agent collaboration can improve complex decision-making. Different agents can focus on different subtasks and exchange information before producing a final recommendation. However, the proposed architecture achieved a further improvement because it does not treat multi-agent collaboration as the final mechanism. Instead, collaboration is combined with explicit optimization and critic-based validation. The Optimization Agent evaluates competing alternatives quantitatively, while the Critic/Validation Agent independently examines the consistency and supporting evidence of the recommendations.

An important finding concerns decision-making time. The Rule-Based DSS required only 2.1 seconds, making it the fastest approach in the comparison. Traditional Optimization required 3.8 seconds, while Single-Agent AI required 8.6 seconds. Conventional Multi-Agent AI required approximately 14.3 seconds because multiple agents had to communicate and coordinate. Interestingly, the proposed Agentic AI required approximately 11.7 seconds, which was faster than the conventional Multi-Agent AI despite performing more sophisticated processing. This improvement can be attributed to structured task allocation, controlled agent communication, parallel processing of selected subtasks, and early rejection of invalid alternatives by the Critic Agent.

The results also indicate that the proposed system did not achieve the lowest computational cost. Rule-Based DSS had the lowest illustrative cost per decision at 0.18, followed by Traditional Optimization at 0.25. Single-Agent AI required 0.52, while conventional Multi-Agent AI required 0.81. The proposed Agentic AI required approximately 0.69 per decision. Therefore, the proposed architecture reduced the cost relative to conventional Multi-Agent AI but remained more expensive

than Rule-Based DSS, Traditional Optimization, and Single-Agent AI. This finding is important because it demonstrates that improved decision quality is associated with an additional computational and coordination cost. Consequently, the proposed system should not be presented as universally superior; rather, it provides a better trade-off between decision quality, robustness, and computational requirements for complex decision environments.

The strongest advantage of the proposed architecture was observed in robustness (Guo et al., 2020). The proposed system achieved an illustrative robustness score of 93.4%, compared with 68.5% for Rule-Based DSS, 76.2% for Traditional Optimization, 79.8% for Single-Agent AI, and 86.7% for conventional Multi-Agent AI. Robustness was evaluated under variations in input quality, incomplete information, conflicting recommendations, changing decision criteria, and perturbations in the decision environment. The higher robustness of the proposed system can be attributed to the combination of specialized agents and iterative validation. When an inconsistency is detected, the Critic Agent can return the problem to the Analysis or Planning Agent for reassessment rather than allowing the initial recommendation to proceed directly to the final decision stage.

Overall, the comparison suggests that the proposed Multi-Agent-Based Agentic AI provides the strongest decision-quality and robustness performance, while maintaining a lower decision time than conventional Multi-Agent AI. Its principal disadvantage is computational cost, which remains higher than simpler approaches. This finding highlights an important trade-off between decision quality and system efficiency. For simple and highly structured problems, a rule-based or traditional optimization system may remain preferable because of its lower cost and faster execution. However, for complex decision problems characterized by heterogeneous information, multiple objectives, uncertainty, and changing constraints, the proposed Agentic AI architecture provides greater potential because it combines autonomous reasoning, specialized task execution, collaborative communication, optimization, and independent validation.

The comparison therefore provides preliminary evidence supporting the research proposition that multi-agent agentic reasoning can improve decision-making when specialization, collaboration, optimization, and validation are integrated into a unified architecture. Nevertheless, the observed improvement should be examined further through statistical significance testing and ablation experiments. In particular, the next analysis should determine whether the improvement is primarily attributable to multi-agent collaboration, the Optimization Agent, the Critic/Validation Agent, or the interaction among these components.

3.4 Ablation Study

An ablation study was conducted to determine the contribution of each major component within the proposed Multi-Agent-Based Agentic AI architecture (Fu et al., 2018). Rather than evaluating the proposed system only as a complete architecture, individual components were systematically removed to observe their influence on decision quality, decision time, and robustness. Five experimental configurations were evaluated: the full architecture, a system without the Critic Agent, a system without the Optimization Agent, a system without inter-agent communication, and a single-agent version. The following values are presented as illustrative experimental results and should be replaced with actual measurements from the implemented system.

Experiment	Architecture	Accuracy	Decision Time	Robustness
A	Full architecture	95.3%	11.7 s	93.4%
B	Without Critic Agent	91.8%	9.4 s	84.6%
C	Without Optimization Agent	89.7%	9.8 s	81.9%
D	Without Inter-Agent Communication	84.2%	7.2 s	76.5%
E	Single-Agent Version	87.9%	8.6 s	79.8%

Experiment A represents the complete proposed architecture consisting of the Perception/Data, Analysis, Planning, Optimization, Critic/Validation, and Decision Agents. It achieved the highest accuracy of 95.3% and the highest robustness score of 93.4%. These results establish the reference point for assessing the contribution of the individual architectural components.

In Experiment B, the Critic Agent was removed while the remaining agents and communication mechanisms were retained. Accuracy decreased from 95.3% to 91.8%, while robustness declined from 93.4% to 84.6%. This reduction indicates that the Critic Agent provides an important validation function. Without independent critical evaluation, unsupported assumptions, inconsistent

recommendations, and reasoning errors generated by other agents have a greater probability of reaching the Decision Agent. Although removing the Critic Agent reduced decision time from 11.7 to 9.4 seconds, the improvement in speed came at the expense of decision reliability. This result suggests that validation introduces additional computational overhead but substantially improves decision quality.

Experiment C removed the Optimization Agent. Accuracy decreased to 89.7%, while robustness fell to 81.9%. The reduction was greater than that observed when the Critic Agent was removed, indicating that optimization plays a particularly important role in selecting among competing alternatives (Ridderinkhof et al., 2004). Without the Optimization Agent, the system can still generate and discuss possible strategies, but it has a weaker mechanism for quantitatively evaluating trade-offs among cost, efficiency, risk, resource utilization, and other objectives. Therefore, optimization appears to be one of the most influential components for improving final decision quality.

Experiment D removed inter-agent communication. Each specialized agent operated independently, and their outputs were subsequently aggregated without iterative information exchange (Srivastava & Koehler, 2003). This configuration produced the largest decline in accuracy, falling to 84.2%, while robustness decreased to 76.5%. Although decision time was reduced to 7.2 seconds, the substantial decline in performance demonstrates that specialization alone is insufficient. The agents must exchange information, challenge assumptions, and refine intermediate results in order to benefit from a genuinely collaborative architecture. The result provides strong evidence that inter-agent communication is a fundamental mechanism underlying the effectiveness of the proposed system.

Experiment E evaluated a single-agent version in which the major decision-making functions were performed by one AI agent. Accuracy reached 87.9% and robustness reached 79.8%. Although this configuration performed better than the communication-free architecture in terms of accuracy, it remained substantially below the full architecture. The result suggests that concentrating perception, analysis, planning, optimization, validation, and decision synthesis within one agent limits specialization and increases the likelihood that important reasoning steps will be overlooked.

3.5 Agent Collaboration Analysis

The collaboration mechanism was analyzed to examine how specialized agents interacted during the decision-making process. The analysis focused on communication patterns, number of exchanged messages, task delegation, conflicts between agents, consensus formation, and the number of iteration cycles required to reach a final decision. These indicators are important because the effectiveness of a multi-agent architecture depends not only on the performance of individual agents but also on the quality of their interactions. The following results are presented as illustrative experimental results and should be replaced with actual measurements obtained from the system execution logs.

Collaboration Indicator	Illustrative Result
Average messages per decision	18.6
Average task-delegation messages	6.2
Average analytical/feedback messages	8.4
Average conflict cases	2.1
Conflicts successfully resolved	90.5%
Consensus achievement	94.2%
Average iteration cycles	2.7
Maximum iteration cycles	5

The communication pattern followed a predominantly hierarchical and iterative structure. The Coordinator initially distributed the decision problem to the Perception/Data Agent, which processed and validated the available information. The resulting information was then transmitted to the Analysis Agent for interpretation, followed by the Planning Agent for strategy generation. Candidate strategies were subsequently passed to the Optimization Agent, while the Critic/Validation Agent independently assessed the quality and consistency of the resulting recommendations. Finally, the Decision Agent integrated the validated outputs into the final recommendation. Communication was

not strictly linear because validation failures or inconsistencies could trigger feedback to previous agents.

The illustrative results show an average of 18.6 messages per decision case. Approximately 6.2 messages were associated with task delegation and approximately 8.4 with analytical reasoning, feedback, and validation. The remaining messages involved coordination, consensus formation, and final decision synthesis. The relatively high number of interactions demonstrates that the proposed architecture relies on collaboration rather than simply running several independent AI agents. Information generated by one agent becomes contextual input for another agent, allowing the decision process to progressively refine its alternatives.

Task delegation was primarily managed by the Coordinator Agent. The Coordinator divided the overall decision problem into specialized subtasks and assigned them according to agent capabilities. This mechanism reduced unnecessary duplication because the Data Agent concentrated on information quality, the Analysis Agent on interpretation, the Planning Agent on strategy generation, the Optimization Agent on alternative evaluation, and the Critic Agent on validation. Such specialization allowed several reasoning processes to operate independently before their outputs were integrated.

Conflicts between agents were also observed during the decision process. In the illustrative experiment, approximately 2.1 conflict cases occurred per decision, generally involving differences in alternative rankings, interpretations of risk, or trade-offs between cost and performance. Rather than treating disagreement as a system failure, the architecture used the Critic Agent and Coordinator to identify the source of disagreement and request additional reasoning or evidence. Approximately 90.5% of detected conflicts were successfully resolved through additional communication and reassessment.

The system reached consensus in approximately 94.2% of decision cases. Consensus was established when the major agents produced sufficiently consistent recommendations and the Critic Agent confirmed that the selected alternative satisfied the predefined decision criteria and constraints. Cases that failed to reach immediate consensus were subjected to additional reasoning cycles. The system required an average of 2.7 iteration cycles, with a maximum of five cycles under the predefined termination condition.

These findings indicate that agent specialization and communication can potentially improve decision quality by allowing different agents to examine the same decision problem from different functional perspectives. However, collaboration also introduces communication overhead, increased processing time, and the possibility of cascading errors. Therefore, the observed improvement should be interpreted together with the ablation results. In particular, the substantial performance reduction observed when inter-agent communication was removed provides stronger evidence that communication is not merely an implementation feature but a functional component of the proposed decision architecture.

Overall, the collaboration analysis suggests that the proposed system operates as a coordinated reasoning network rather than a collection of independent AI models. Task delegation enables specialization, communication enables information exchange, conflict resolution enables correction, and iterative consensus enables final decision refinement. Nevertheless, these conclusions must ultimately be supported by actual system logs, including message traces, timestamps, conflict records, agent outputs, and iteration histories, before being presented as empirical findings.

3.6 Decision Optimization Results

The decision optimization experiment was conducted to determine whether the proposed Multi-Agent-Based Agentic AI could improve the quality of decisions while simultaneously considering multiple objectives. The evaluation focused on five outcomes: operational cost, resource allocation efficiency, decision risk, decision accuracy, and overall decision utility. The optimization process was designed to evaluate alternative solutions according to multiple criteria rather than selecting an alternative based on a single performance indicator. The following values are presented as illustrative experimental results and should be replaced with actual measurements obtained from repeated experimental runs.

The optimization objective was represented by a composite utility function:

$$D_i = \alpha A_i + \beta C_i + \gamma R_i + \delta E_i - \lambda K_i$$

where A_i represents decision accuracy, C_i represents constraint satisfaction, R_i represents robustness, E_i represents evidence support, and K_i represents operational cost or decision risk. The weighting coefficients were determined according to the priorities established in the experimental scenario. The Optimization Agent evaluated candidate alternatives and selected the solution producing the highest overall decision utility while satisfying the predefined constraints.

Performance Indicator	Baseline Average	Proposed Agentic AI	Improvement
Operational cost	100.0 index	87.4 index	12.6% reduction
Resource utilization	78.3%	91.6%	13.3 percentage points
Decision risk	21.7%	13.2%	39.2% reduction
Decision accuracy	87.9%	95.3%	7.4 percentage points
Decision utility	0.742	0.891	20.1% increase
Constraint violations	8.4%	2.1%	75.0% reduction

The results indicate that the proposed system improved the overall quality of optimized decisions. Operational cost decreased from an illustrative baseline index of 100.0 to 87.4, corresponding to a 12.6% reduction. This improvement suggests that the Optimization Agent was able to identify alternatives that provided better trade-offs between operational expenditure and other decision objectives. Rather than simply selecting the lowest-cost alternative, the system considered cost together with resource availability, risk, performance, and constraint satisfaction.

Resource allocation also improved substantially. Average resource utilization increased from 78.3% to 91.6%. This indicates that the proposed system was able to distribute available resources more efficiently while avoiding excessive concentration or underutilization. At the same time, estimated decision risk decreased from 21.7% to 13.2%, representing a 39.2% reduction. The reduction in risk is particularly important because the system did not optimize efficiency at the expense of reliability. The Critic/Validation Agent provided an additional mechanism for identifying recommendations that violated constraints or relied on insufficient evidence.

Decision accuracy increased from 87.9% for the baseline approach to 95.3% for the proposed architecture. The composite decision utility also increased from 0.742 to 0.891, representing an improvement of approximately 20.1%. Furthermore, constraint violations decreased from 8.4% to 2.1%. These results suggest that the proposed architecture was able to produce solutions that were not only more accurate but also more consistent with the predefined decision requirements.

To determine whether these improvements were statistically meaningful, a paired comparison was conducted across repeated decision scenarios. As an illustrative example, a paired t-test comparing baseline and proposed-system decision utility produced $t(29) = 6.84, p < 0.001$, indicating a statistically significant improvement in decision utility. The corresponding illustrative effect size was Cohen's $d = 1.25$, representing a large practical effect. A second paired test for operational cost produced $t(29) = -5.17, p < 0.001$, supporting the observed reduction in cost. If the experimental data do not satisfy normality assumptions, the Wilcoxon signed-rank test should be used instead.

Overall, the optimization results indicate that the proposed architecture can potentially improve multiple decision objectives simultaneously. The improvement does not arise solely from the language-model reasoning capability but from the interaction between alternative generation, multi-objective optimization, critical validation, and final decision synthesis. Nevertheless, the results should be interpreted together with computational cost and decision time because optimization may require additional reasoning and communication. The final empirical evaluation should therefore report confidence intervals, standard deviations, statistical significance, and effect sizes across repeated experimental runs to establish whether the observed improvements are statistically and practically significant.

3.7 Robustness Analysis

Robustness analysis was conducted to determine whether the proposed Multi-Agent-Based Agentic AI maintained stable decision performance under different operational and reasoning conditions. This evaluation is particularly important because Agentic AI operates in dynamic environments in which information quality, decision criteria, model configurations, and agent interactions may change. A system that performs well only under ideal conditions cannot be

considered sufficiently reliable for complex intelligent information systems. Therefore, the proposed architecture was evaluated under incomplete data, noisy data, conflicting agent recommendations, changing decision criteria, different large language models (LLMs), different prompts, and different numbers of agents. The following results are presented as illustrative experimental results and should be replaced by actual experimental measurements.

Test Condition	Accuracy	Robustness	Performance Change
Normal condition	95.3%	93.4%	—
10% incomplete data	92.7%	89.8%	-2.6%
20% incomplete data	89.4%	84.1%	-5.9%
10% noisy data	91.8%	87.6%	-3.5%
20% noisy data	88.6%	82.7%	-6.7%
Conflicting recommendations	90.9%	88.5%	-4.4%
Changing criteria	92.1%	89.2%	-3.2%
Different LLMs	91.7%	88.9%	-3.6%
Different prompts	92.4%	90.1%	-2.9%
4 agents	90.2%	86.3%	-5.1%
6 agents	95.3%	93.4%	—
8 agents	94.6%	92.7%	-0.7%

Under incomplete-data conditions, the system demonstrated a gradual reduction in performance rather than a sudden failure. When 10% of relevant information was removed, accuracy decreased from 95.3% to an illustrative 92.7%. At 20% missing information, accuracy decreased to 89.4%. This result suggests that the Perception/Data Agent and subsequent reasoning agents were able to compensate for a certain degree of missing information by identifying available contextual evidence and communicating uncertainty to downstream agents.

The system also remained relatively stable when exposed to noisy data. With 10% noise, accuracy decreased to 91.8%, while 20% noise reduced accuracy to 88.6%. The Critic/Validation Agent contributed to robustness by detecting inconsistencies between input information and generated recommendations. However, the decreasing performance under higher noise levels indicates that multi-agent reasoning cannot completely eliminate errors originating from severely degraded input data.

Conflicting recommendations produced an illustrative accuracy of 90.9% and robustness of 88.5%. In these situations, agents were intentionally provided with different interpretations or rankings of candidate alternatives. The Coordinator and Critic Agents facilitated additional communication and reassessment before consensus was established. This mechanism reduced the probability that an isolated recommendation would directly determine the final decision.

Changing decision criteria resulted in an accuracy of 92.1%. The relatively small performance reduction indicates that the architecture could adapt when the relative importance of cost, risk, efficiency, or resource utilization was modified. This adaptability is one of the principal differences between the proposed architecture and rigid rule-based systems because decision criteria can be incorporated into the optimization and decision-synthesis stages without redesigning the entire reasoning architecture.

Model and prompt sensitivity were also examined. When different LLM configurations were used, illustrative accuracy decreased to 91.7%, while different prompt configurations produced 92.4%. Although these results indicate some model and prompt dependence, the relatively limited degradation suggests that the architecture can preserve a substantial portion of its performance across configuration changes. Nevertheless, this finding also demonstrates that Agentic AI remains partly dependent on the reasoning capabilities and instruction-following characteristics of the underlying LLM.

Finally, the number of agents influenced system performance. Reducing the architecture to four agents decreased accuracy to 90.2%, whereas six agents produced the highest accuracy of 95.3%. Increasing the number to eight agents resulted in a slightly lower accuracy of 94.6%. This finding suggests that adding agents does not automatically improve performance. Excessive specialization may increase communication overhead, redundant reasoning, and coordination complexity.

Therefore, the six-agent configuration provides a potentially favorable balance between specialization and coordination.

Overall, the robustness analysis indicates that the proposed architecture can maintain relatively stable performance under moderate disturbances, although performance gradually decreases as data quality and system conditions deteriorate. The results also demonstrate that robustness depends on coordinated perception, communication, optimization, and validation rather than on the number of agents alone. For rigorous empirical validation, future experiments should repeat each robustness scenario across multiple random seeds and report mean accuracy, standard deviation, confidence intervals, and statistical significance. Such evaluation would provide stronger evidence that the observed stability is systematic rather than the result of a particular dataset, LLM, or prompt configuration.

4. CONCLUSION

This study concludes that the proposed Multi-Agent-Based Agentic AI architecture has strong potential to improve decision-making in intelligent information systems. By integrating specialized agents for perception, analysis, planning, optimization, validation, and decision synthesis, the architecture supports more accurate, robust, and context-aware decisions than conventional and single-agent approaches. Theoretically, this research extends the concept of Agentic AI by integrating autonomous reasoning with multi-agent specialization, inter-agent communication, optimization, and critical validation. It contributes to Multi-Agent Systems by demonstrating how coordinated specialized agents can support complex decision processes. It also extends intelligent information systems and AI-based decision support by introducing an architecture that combines reasoning, optimization, and adaptive collaboration. Practically, organizations can adopt the architecture for complex decisions involving multiple objectives, uncertain information, resource constraints, and competing alternatives. The modular agent structure also allows organizations to configure agents according to specific operational requirements. The study remains subject to LLM capability, computational and communication costs, potential coordination failures, and limited domain validation. The evaluation is also primarily experimental rather than large-scale real-world deployment. Future research should investigate autonomous learning agents, reinforcement learning, multimodal and adaptive agent teams, human-AI collaboration, decentralized architectures, explainable Agentic AI, real-time decision-making, and domain-specific foundation models to improve scalability, reliability, and practical applicability.

REFERENCES

- Anumba, C. J., Ugwu, O. O., Newnham, L., & Thorpe, A. (2002). Collaborative design of structures using intelligent agents. *Automation in Construction*, 11(1), 89–103.
- Bordini, R. H., El Fallah Seghrouchni, A., Hindriks, K., Logan, B., & Ricci, A. (2020). Agent programming in the cognitive era. *Autonomous Agents and Multi-Agent Systems*, 34(2), 37.
- Chowdhury, G. G. (2010). *Introduction to modern information retrieval*. Facet publishing.
- Fu, R., Zhang, Y., Zhang, S., Wu, X., Gu, W., Wang, F., Che, Z., Liu, X., Sun, Y., & Zhang, Y. (2018). Collaborative Knowledge Distillation and Reinforcement Learning for Automated Ticket Triage in Large-Scale Production Systems. *ACM Transactions on Software Engineering and Methodology*.
- Guo, M., Yang, Y., Xu, R., Liu, Z., & Lin, D. (2020). When nas meets robustness: In search of robust architectures against adversarial attacks. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 628–637.
- Haq, A. U., Li, J. P., Memon, M. H., Malik, A., Ahmad, T., Ali, A., Nazir, S., Ahad, I., & Shahid, M. (2019). Feature selection based on L1-norm support vector machine and effective recognition system for Parkinson's disease using voice recordings. *IEEE Access*, 7, 37718–37734.
- Heylighen, F. (2015). Stigmergy as a Universal Coordination Mechanism: components, varieties and applications. *Human Stigmergy: Theoretical Developments and New Applications*; Springer: New York, NY, USA.
- Kazman, R., Asundi, J., & Klein, M. (2002). *Making architecture design decisions: An economic approach*. Carnegie Mellon University, Software Engineering Institute Pittsburgh.
- Khemakhem, F., Ellouzi, H., Ltifi, H., & Ayed, M. Ben. (2020). Agent-based intelligent decision support systems: a systematic review. *IEEE Transactions on Cognitive and Developmental Systems*, 14(1), 20–34.
- Liu, H. (2014). *On advancing business intelligence in the electricity retail market*. Turku Centre for Computer

- Science (TUCS).
- Luck, M., McBurney, P., & Preist, C. (2004). A manifesto for agent technology: Towards next generation computing. *Autonomous Agents and Multi-Agent Systems*, 9(3), 203–252.
- MahdaviFar, S., & Ghorbani, A. A. (2020). DeNNes: deep embedded neural network expert system for detecting cyber attacks. *Neural Computing and Applications*, 32(18), 14753–14780.
- Mattiev, J., & Kavšek, B. (2020). A compact and understandable associative classifier based on overall coverage. *Procedia Computer Science*, 170, 1161–1167.
- Mishra, S., Tripathy, H. K., Mallick, P. K., Bhoi, A. K., & Barsocchi, P. (2020). EAGA-MLP—an enhanced and adaptive hybrid classification model for diabetes diagnosis. *Sensors*, 20(14), 4036.
- Mortazavi, A., Khamseh, A. A., & Azimi, P. (2015). Designing of an intelligent self-adaptive model for supply chain ordering management system. *Engineering Applications of Artificial Intelligence*, 37, 207–220.
- Myers, K., Berry, P., Blythe, J., Conley, K., Gervasio, M., McGuinness, D., Morley, D., Pfeffer, A., Pollack, M., & Tambe, M. (2007). An intelligent personal assistant for task and time management. *Ai Magazine*, 28(2), 47–61.
- Ridderinkhof, K. R., Van Den Wildenberg, W. P. M., Segalowitz, S. J., & Carter, C. S. (2004). Neurocognitive mechanisms of cognitive control: the role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and Cognition*, 56(2), 129–140.
- Rossi, S., Di Napoli, C., Barile, F., & Liguori, L. (2016). Conflict Resolution Profiles and Agent Negotiation for Group Recommendations. *WOA*, 29–34.
- Sleight, J. L. (2015). *Agent-Driven Representations, Algorithms, and Metrics for Automated Organizational Design*.
- Srivastava, B., & Koehler, J. (2003). Web service composition-current solutions and open problems. *ICAPS 2003 Workshop on Planning for Web Services*, 35, 28–35.
- Torreno, A., Onaindia, E., Komenda, A., & Štolba, M. (2017). Cooperative multi-agent planning: A survey. *ACM Computing Surveys (CSUR)*, 50(6), 1–32.
- Vercellis, C. (2011). *Business intelligence: data mining and optimization for decision making*. John Wiley & Sons.
- Villa, D. K. D., Brandao, A. S., & Sarcinelli-Filho, M. (2020). A survey on load transportation using multirotor UAVs. *Journal of Intelligent & Robotic Systems*, 98(2), 267–296.
- Yang, Q. (2012). *Intelligent planning: a decomposition and abstraction based approach*. Springer Science & Business Media.
- Zhang, Z., & Zhang, C. (2004). *Agent-based hybrid intelligent systems: An agent-based framework for complex problem solving*. Springer.