IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

ISO 3297:2007 Certified ⁜ Impact Factor 8.102 ⁜ Peer-reviewed / Refereed journal ⁜ Vol. 12, Issue 8, August 2023

DOI: 10.17148/IJARCCE.2023.12828

Agentic AI-Driven Enterprise Decision Support for Cross-Domain Business Intelligence

Charan Thumma[1], Nivedan Suresh[2], Chaitanya Tumma[3], Hemanth Raj Gudisa[4]

Independent Researcher, Stone ridge, VA, USA[1]

Staff Software Engineer, Visa Inc., Austin, TX, USA, Independent Researcher[2]

Department of Information Technology, University of the Cumberlands, Williamsburg, KY, USA[3]

Independent Researcher Scholar, Dallas, TX[4]

Abstract: Existing enterprise decision-support systems often suffer from fragmented data management, limited cross-domain knowledge integration, and insufficient autonomous analytical capabilities, resulting in reduced decision accuracy and inefficient intelligence generation. To address these challenges, this paper proposes an Agentic Intelligence Framework (AIF) for cross-domain business intelligence that employs hierarchical multi-agent collaboration to improve autonomous reasoning, knowledge synthesis, and decision support across enterprise environments. First, heterogeneous enterprise data originating from retail, finance, and networking domains are integrated through Retrieval-Augmented Generation (RAG) and domain-specific knowledge repositories, enabling contextual knowledge retrieval and semantic enrichment. Second, a hierarchical multi-agent architecture comprising Research Agents, Analytics Agents, Validation Agents, and Insight Generation Agents is designed to collaboratively investigate complex business scenarios, validate analytical outcomes, and maximize the quality and reliability of generated intelligence. Third, an adaptive reasoning and workflow orchestration mechanism is introduced to dynamically coordinate agent interactions, continuously refine analytical hypotheses, and uncover hidden relationships among distributed enterprise data sources. Finally, explainable AI techniques are incorporated to provide transparent reasoning paths and interpretable recommendations, thereby improving user trust and supporting informed decision-making. Experimental results demonstrate that, compared with conventional AI-based business intelligence frameworks, the proposed AIF achieves higher analytical accuracy, improved knowledge discovery capability, reduced decision latency, and enhanced explainability across retail, financial, and networking applications.

Keywords: Agentic Artificial Intelligence; Multi-Agent Systems; Business Intelligence; Retrieval-Augmented Generation; Large Language Models; Explainable Artificial Intelligence; Cross-Domain Analytics; Enterprise Decision Support.

I. INTRODUCTION

Modern enterprises generate enormous volumes of heterogeneous data from retail operations, financial transactions, customer interactions, enterprise resource planning systems, network infrastructures, and digital services. Although organizations invest heavily in Business Intelligence (BI), Artificial Intelligence (AI), and cloud-based analytics platforms, transforming these fragmented data sources into actionable business intelligence remains a significant challenge [1]. Conventional analytics systems primarily rely on predefined rules, static dashboards, and isolated machine learning models that provide descriptive insights rather than adaptive reasoning [2]. Consequently, enterprise decision-makers often encounter delayed responses, inconsistent analytical outcomes, and limited support for complex cross-domain decision-making [3]. The rapid emergence of Large Language Models (LLMs) has significantly enhanced natural language understanding, knowledge reasoning, and intelligent automation. However, single-agent LLM architectures remain constrained when addressing enterprise-scale analytical tasks that require iterative reasoning, multi-step investigation, knowledge validation, and collaborative decision support [4]. Complex business scenarios frequently involve dependencies among multiple operational domains, including customer behavior analysis, financial risk assessment, operational forecasting, cybersecurity monitoring, and network performance optimization [5]. Solving such problems requires autonomous collaboration among specialized intelligent agents rather than relying on a single reasoning model.

Charan.salesforcedev@gmail.com[1], snivedan13@gmail.com[2], chaitanyatumma1@gmail.com[3]

DOI: 10.17148/IJARCCE.2023.12828

Enterprise decision support systems exhibit several unique characteristics that distinguish them from conventional AI applications [6]. First, enterprise knowledge is highly distributed across structured databases, semi-structured documents, unstructured reports, and real-time operational streams, making comprehensive knowledge integration difficult. Second, business decisions often require simultaneous consideration of multiple domains, where actions in one operational area may significantly influence others. Third, enterprise environments continuously evolve due to changing customer demands, financial regulations, market conditions, and network behaviors, requiring adaptive intelligence capable of dynamically refining analytical hypotheses [7]. Finally, organizations increasingly demand transparent and explainable AI systems that provide understandable reasoning processes instead of opaque predictions. Existing business intelligence and analytics platforms primarily focus on visualization, statistical reporting, or predictive modeling. Although these approaches improve operational monitoring, they generally lack autonomous investigation capabilities, collaborative reasoning, and explainable knowledge synthesis [8]. Recent Retrieval-Augmented Generation (RAG) techniques improve factual consistency by incorporating external knowledge repositories into LLM reasoning; however, most implementations remain limited to single-agent architectures and cannot coordinate multiple specialized analytical processes [9]. Similarly, conventional enterprise AI systems typically analyze retail, finance, or networking applications independently, preventing the discovery of hidden relationships across interconnected business operations.

To address these limitations, this paper proposes an Agentic Intelligence Framework (AIF) for cross-domain business intelligence and enterprise decision support. The proposed framework combines hierarchical multi-agent Large Language Models, Retrieval-Augmented Generation, domain-specific knowledge repositories, adaptive reasoning workflows, and explainable AI techniques into a unified intelligent decision-support architecture. Rather than depending on a single intelligent agent, AIF coordinates multiple specialized agents that collaboratively investigate complex enterprise problems, retrieve contextual knowledge, validate analytical findings, and generate explainable business recommendations [10]. Within the proposed framework, Research Agents collect and retrieve relevant enterprise knowledge from distributed repositories, Analytics Agents perform domain-specific reasoning and predictive analysis, and validation Agents verify the consistency and reliability of generated conclusions using multiple evidence sources, while Insight Generation Agents synthesize validated findings into interpretable recommendations for decision-makers [11]. This collaborative architecture enables continuous refinement of analytical hypotheses and improves both the quality and trustworthiness of generated intelligence.

The framework supports multiple enterprise application domains. In retail environments, it enables customer behavior analysis, inventory optimization, supply-chain intelligence, and demand forecasting [12]. Within financial services, it facilitates regulatory compliance assessment, fraud investigation, operational risk evaluation, and investment intelligence. In enterprise networking, the framework performs operational diagnostics, cybersecurity threat investigation, infrastructure performance assessment, and proactive incident analysis. By integrating knowledge across these domains, AIF provides comprehensive business intelligence that cannot be achieved through isolated analytical systems. Compared with existing AI-driven business intelligence platforms, the proposed framework offers several distinctive advantages. First, hierarchical multi-agent collaboration improves analytical completeness by distributing complex reasoning tasks among specialized intelligent agents [13]. Second, Retrieval-Augmented Generation enhances factual accuracy through contextual knowledge retrieval from enterprise repositories. Third, adaptive reasoning workflows iteratively refine hypotheses and uncover hidden relationships among heterogeneous enterprise data sources. Fourth, explainable AI techniques improve transparency by providing interpretable reasoning paths and evidence supporting each recommendation [14]. Finally, cross-domain knowledge integration enables unified enterprise intelligence that supports proactive decision-making across retail, finance, and networking environments.

The primary contributions of this paper are summarized as follows:

- A hierarchical Agentic Intelligence Framework (AIF) is proposed for enterprise decision support, integrating Large Language Models, Retrieval-Augmented Generation, domain knowledge repositories, and collaborative multi-agent reasoning to enable autonomous cross-domain intelligence generation.
- A collaborative reasoning architecture comprising Research Agents, Analytics Agents, Validation Agents, and Insight Generation Agents is developed to perform adaptive investigation, knowledge verification, hypothesis refinement, and explainable recommendation generation.
- An adaptive orchestration mechanism is introduced to coordinate agent interactions dynamically, enabling iterative reasoning, evidence validation, and continuous refinement of analytical workflows across heterogeneous enterprise environments.
- A comprehensive evaluation framework is designed using enterprise datasets spanning retail, finance, and networking domains. Experimental analyses assess analytical accuracy, knowledge retrieval precision, decision latency, explainability, cross-domain reasoning capability, and overall decision-support effectiveness,

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

ISO 3297:2007 Certified ⁓ Impact Factor 8.102 ⁓ Peer-reviewed / Refereed journal ⁓ Vol. 12, Issue 8, August 2023

DOI: 10.17148/IJARCCE.2023.12828

demonstrating the superiority of the proposed framework over conventional AI-driven business intelligence systems.

II. AGENTIC INTELLIGENCE FRAMEWORK AND SYSTEM DEFINITIONS

The proposed Agentic Intelligence Framework (AIF) is designed to support autonomous enterprise decision-making by coordinating multiple intelligent agents that collaboratively perform knowledge discovery, reasoning, validation, and recommendation generation [15]. Unlike conventional business intelligence platforms that rely on centralized analytical engines, AIF distributes analytical responsibilities among specialized AI agents capable of independently solving domain-specific problems while sharing knowledge through a centralized orchestration mechanism [16]. The framework consists of five primary entities: Research Agents, Analytics Agents, Validation Agents, Insight Generation Agents, Enterprise Knowledge Repository. Business intelligence generation is accomplished through four sequential phases: Knowledge Acquisition, Collaborative Intelligence Analysis, Evidence Validation, Explainable Decision Generation. Within the enterprise environment, business data collected from retail operations, financial systems, and networking infrastructures serve as analytical inputs. Research Agents retrieve contextual knowledge using Retrieval-Augmented Generation (RAG), Analytics Agents perform predictive reasoning, Validation Agents verify the consistency of generated knowledge, and Insight Generation Agents produce human-readable recommendations supported by explainable AI techniques. The orchestration layer continuously coordinates communication among agents to ensure that every analytical outcome is supported by verified evidence.

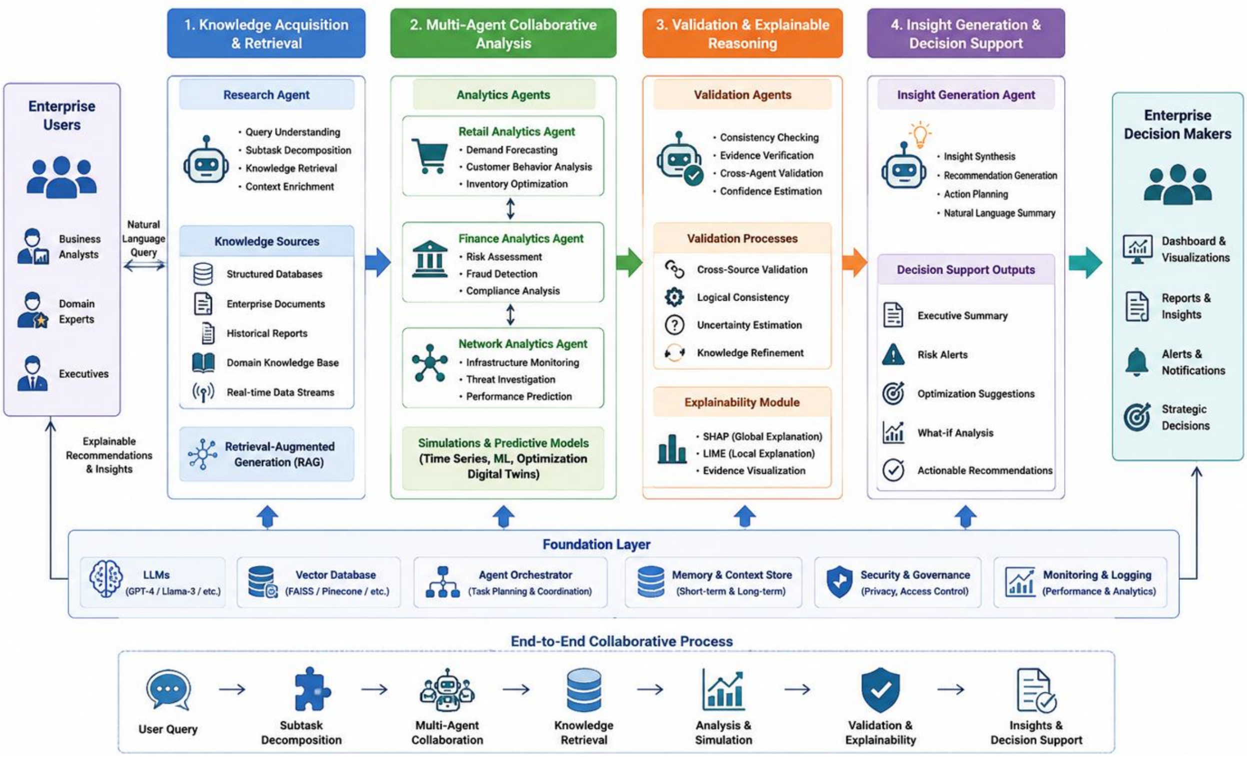

Figure 1. Overall architecture of the proposed Agentic Intelligence Framework (AIF).

II.1. *Agentic Intelligence Architecture*

Figure 1 illustrates the overall architecture of the proposed Agentic Intelligence Framework (AIF). The framework consists of four collaborative processing stages:

1. Data Acquisition and Knowledge Retrieval
2. Multi-Agent Collaborative Analysis
3. Validation and Explainable Reasoning
4. Decision Recommendation Generation

Unlike traditional business intelligence systems where a centralized model performs all analytical tasks, the proposed framework distributes responsibilities across multiple autonomous agents. The Research Agent retrieves relevant contextual information from enterprise knowledge repositories using Retrieval-Augmented Generation. Retrieved

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

ISO 3297:2007 Certified ⁒ Impact Factor 8.102 ⁒ Peer-reviewed / Refereed journal ⁒ Vol. 12, Issue 8, August 2023

DOI: 10.17148/IJARCCE.2023.12828

documents, historical reports, and structured databases provide factual evidence that guides subsequent reasoning. During the collaborative analysis stage, specialized Analytics Agents independently analyze the retrieved information according to their domain expertise. Retail Analytics Agents perform customer behavior prediction and inventory analysis, Financial Analytics Agents evaluate operational risks and regulatory compliance, while Network Analytics Agents investigate infrastructure performance, security threats, and operational anomalies. The generated analytical outputs are forwarded to Validation Agents, which compare multiple evidence sources to eliminate inconsistent reasoning and reduce hallucination generated by Large Language Models. Only validated knowledge proceeds to the final reasoning stage. Finally, Insight Generation Agents integrate validated analytical findings into comprehensive business recommendations. Explainable AI methods, including SHAP, LIME, and retrieval-supported evidence visualization, provide transparent reasoning paths that improve user trust and decision interpretability. Through collaborative intelligence, the framework continuously improves analytical quality while maintaining scalability across multiple enterprise domains.

II.2. *System Definitions*

To formally describe the operation of the proposed Agentic Intelligence Framework, the following definitions are introduced.

Definition 1. Enterprise Data Source: An Enterprise Data Source represents any structured, semi-structured, or unstructured repository containing information required for business intelligence generation. Typical data sources include:

- Transaction databases
- Financial records
- Customer interaction logs
- Network telemetry
- Security event logs
- Business reports
- Knowledge repositories

Each source contributes complementary knowledge that enhances analytical completeness.

Definition 2. Intelligent Agent: An Intelligent Agent is an autonomous software entity responsible for performing specialized reasoning tasks while collaborating with other agents through the orchestration layer. The proposed framework contains four categories of agents:

- Research Agents
- Analytics Agents
- Validation Agents
- Insight Generation Agents

Each agent independently performs domain-specific reasoning while sharing intermediate knowledge with collaborating agents.

Definition 3. Knowledge Confidence Score: The Knowledge Confidence Score quantifies the reliability of analytical conclusions produced by the multi-agent framework. It is computed by considering

- Retrieval relevance
- Reasoning consistency
- Validation accuracy
- Cross-agent agreement

Higher confidence values indicate stronger evidence supporting generated recommendations.

Definition 4. Explainability Score: The Explainability Score measures how effectively the framework can justify generated recommendations through transparent reasoning. The score considers

- SHAP feature importance
- LIME local explanations
- Retrieved supporting evidence
- Multi-agent reasoning consistency

A higher Explainability Score corresponds to greater user trust and improved decision transparency.

II.3. *Multi-Agent Collaboration Strategy*

The effectiveness of enterprise decision support depends on efficient collaboration among autonomous agents. Unlike conventional single-agent architectures, the proposed framework enables parallel reasoning across multiple analytical domains. The collaborative workflow consists of five sequential steps:

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

ISO 3297:2007 Certified ⁜ Impact Factor 8.102 ⁜ Peer-reviewed / Refereed journal ⁜ Vol. 12, Issue 8, August 2023

DOI: 10.17148/IJARCCE.2023.12828

1. Knowledge Retrieval: Research Agents retrieve contextual information from enterprise repositories using semantic search and Retrieval-Augmented Generation.
2. Domain-Specific Analysis: Analytics Agents independently investigate domain-specific problems.
    - Retail Agent → Customer analytics and demand forecasting
    - Finance Agent → Risk assessment and compliance analysis
    - Network Agent → Infrastructure diagnostics and security intelligence
3. Evidence Validation: Validation Agents compare reasoning outputs generated by multiple agents. Inconsistent or unsupported conclusions are discarded, while verified findings are retained for recommendation generation.
4. Knowledge Fusion: The orchestration layer integrates validated analytical outputs into a unified enterprise knowledge graph representing relationships among customers, operations, financial activities, and network infrastructures.
5. Insight Generation: Insight Generation Agents transform validated knowledge into actionable business recommendations accompanied by explainable reasoning paths.

This collaborative strategy significantly improves analytical completeness, factual consistency, and recommendation quality compared with independent AI models.

II.4. *Adaptive Reasoning and Orchestration Mechanism*

Enterprise environments continuously evolve due to changing customer demands, operational conditions, financial regulations, and cybersecurity threats. To accommodate these dynamic environments, the proposed framework introduces an adaptive orchestration mechanism. Rather than executing a fixed analytical workflow, the orchestration engine dynamically coordinates agent interactions based on

- analytical complexity,
- available contextual knowledge,
- confidence scores,
- validation outcomes, and
- user objectives.

When analytical confidence is insufficient, the orchestration layer automatically initiates additional retrieval cycles, assigns supplementary reasoning tasks to specialized agents, and performs iterative hypothesis refinement until acceptable confidence thresholds are achieved. Compared with static workflow systems, adaptive orchestration offers several advantages:

- improved analytical accuracy,
- enhanced robustness against incomplete information,
- reduced hallucination,
- dynamic knowledge refinement,
- scalable multi-domain collaboration.

Consequently, the proposed Agentic Intelligence Framework provides a flexible and trustworthy foundation for autonomous enterprise decision support across retail, finance, and networking environments.

III. AGENTIC AI-DRIVEN MULTI-AGENT INTELLIGENCE FRAMEWORK

Unlike traditional business intelligence systems that rely on centralized analytics engines and static machine learning pipelines, the proposed Agentic Intelligence Framework (AIF) employs autonomous collaborative AI agents capable of reasoning, planning, retrieving domain knowledge, validating information, and generating explainable recommendations. The framework transforms complex enterprise decision-making into a sequence of coordinated intelligence tasks executed by specialized agents rather than isolated analytical modules. To reduce computational complexity and improve scalability, the framework decomposes every enterprise query into hierarchical reasoning tasks processed in parallel by multiple intelligent agents. The complete workflow consists of four major stages:

- Enterprise Query Understanding
- Multi-Agent Intelligence Generation
- Explainable Knowledge Validation
- Decision Recommendation Generation

III.1. *Agentic Intelligence Framework*

The framework consists of five major intelligent components:

- User Query Manager
- Agent Orchestrator
- Specialized Domain Agents
- Knowledge Retrieval Layer

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

ISO 3297:2007 Certified ⁜ Impact Factor 8.102 ⁜ Peer-reviewed / Refereed journal ⁜ Vol. 12, Issue 8, August 2023

DOI: 10.17148/IJARCCE.2023.12828

- Explainable Decision Generator

When an enterprise user submits a business query, the User Query Manager first analyzes the request using a Large Language Model (LLM). The query is decomposed into multiple subtasks depending on business objectives. Instead of allowing one LLM to solve the complete problem, the Agent Orchestrator assigns specialized agents to solve individual subtasks. This decomposition significantly reduces reasoning complexity while improving interpretability.

III.2. *Multi-Agent Collaborative Decision Algorithm*

The proposed framework contains four categories of intelligent agents. Research Agent: The Research Agent retrieves relevant knowledge from

- Enterprise databases
- Domain documents
- Historical business reports
- Regulatory repositories
- Vector databases

using Retrieval-Augmented Generation (RAG). Instead of relying solely on LLM memory, every response is grounded using retrieved evidence. Analytics Agent: The Analytics Agent performs simulation and predictive analytics. Depending on the application domain, it executes. Retail: Inventory optimization, Customer demand forecasting, Logistics simulation. Finance: Risk assessment, Fraud probability estimation, Regulatory compliance prediction. Networking: Capacity planning, Failure prediction, Cyber threat simulation. Each Analytics Agent independently produces simulation outputs. Validation Agent: The Validation Agent evaluates every generated insight. Validation includes, Confidence estimation, Cross-agent agreement, Knowledge consistency checking, Evidence verification. Only validated outputs are forwarded to the final reasoning stage. Insight Generation Agent: Finally, the Insight Generation Agent integrates validated knowledge into Executive summaries, Explainable recommendations, Risk alerts, Resource optimization strategies. This collaborative reasoning significantly reduces hallucination compared with standalone LLMs.

Algorithm 1: Multi-Agent Collaborative Intelligence

Input: Enterprise query Q

Output: Explainable business recommendation R

Procedure

- Receive enterprise query.
- Parse business objectives.
- Decompose query into subtasks.
- Assign subtasks to specialized agents.
- Retrieve contextual knowledge.
- Execute simulation and predictive analytics.
- Validate generated results.
- Compute explanation scores using SHAP and LIME.
- Aggregate validated insights.
- Generate final explainable recommendation.

The computational complexity of the orchestration algorithm is approximately $O(N \times M)$, where N is number of agents, M is average reasoning iterations. Parallel execution substantially reduces overall latency compared with sequential processing.

III.3. *Explainable Intelligence Mechanism*

One major limitation of existing Agentic AI systems is the lack of transparent reasoning. To address this issue, the proposed framework integrates Explainable AI throughout the decision-making pipeline. Each recommendation is accompanied by three complementary explanation mechanisms.

- SHAP-based Global Explanation: SHAP identifies the overall importance of each feature contributing to business predictions. For example, Inventory shortage prediction may indicate, Supplier delay → 41%, Transportation congestion → 28%, Seasonal demand → 19%, Weather disruption → 12%. Managers therefore understand why predictions were generated.
- LIME-based Local Explanation: LIME explains individual predictions. Instead of explaining the complete model, LIME identifies which variables influenced a single business decision. This improves user trust during scenario analysis.
- Retrieval-Augmented Evidence: Every recommendation is supported by retrieved documents. Instead of producing unsupported outputs, the framework references enterprise reports, historical transactions, regulations, operational logs, allowing decision makers to verify AI reasoning.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

ISO 3297:2007 Certified ⁙ Impact Factor 8.102 ⁙ Peer-reviewed / Refereed journal ⁙ Vol. 12, Issue 8, August 2023

DOI: 10.17148/IJARCCE.2023.12828

III.4. *Adaptive Multi-Agent Coordination Strategy*

Enterprise environments continuously evolve. Consequently, the proposed framework dynamically reallocates reasoning tasks among agents. The Agent Orchestrator continuously monitors

- Agent workload
- Response latency
- Confidence scores
- Knowledge freshness

Agents producing low-confidence outputs automatically trigger additional reasoning iterations. Similarly, if one agent fails, alternative agents continue execution without interrupting the complete workflow. This collaborative fault-tolerant mechanism improves robustness compared with centralized AI architectures.

III.5. *Explainable Decision Recommendation Generation*

After all reasoning stages are completed, validated knowledge is synthesized into actionable business recommendations. Unlike traditional dashboards that present only prediction scores, the proposed framework generates predicted outcome, confidence score, supporting evidence, business impact, recommended actions, explanation summary. For example, Prediction: High probability of inventory shortage. Reason: Supplier delivery delays and increased regional demand. Confidence: 96%, Recommended Action: Increase safety stock by 15%. Supporting Evidence: Historical supply-chain records from similar disruptions. This human-readable explanation enables executives to understand AI recommendations before implementation.

III.6. *Computational Complexity Analysis*

The proposed framework distributes computation across autonomous agents. Instead of one LLM performing complete reasoning, multiple lightweight reasoning tasks execute concurrently. Assuming N agents, M reasoning iterations, K retrieved documents overall computational complexity becomes $O(NM + K)$. Parallel execution significantly reduces decision latency while improving scalability for enterprise-scale analytics.

IV. EXPERIMENTAL EVALUATION AND RESULTS ANALYSIS

IV.1. *Experimental Setup*

To evaluate the effectiveness of the proposed Agentic Intelligence Framework (AIF), a comprehensive simulation environment was developed using Python 3.11, LangChain, FAISS vector database, OpenAI-compatible Large Language Models (LLMs), and Docker-based multi-agent orchestration. The experiments were executed on a workstation equipped with an Intel Xeon Gold 6348 processor (24 cores, 2.60 GHz), 128 GB RAM, NVIDIA RTX A6000 GPU (48 GB VRAM), and Ubuntu 22.04 LTS operating system.
The framework was evaluated across three representative enterprise domains:

- Retail: Customer demand forecasting, inventory optimization, and supply-chain simulation.
- Finance: Operational risk assessment, fraud analytics, and regulatory compliance analysis.
- Networking: Infrastructure simulation, capacity planning, resilience evaluation, and incident diagnosis.

The experimental environment consisted of a centralized Agent Orchestrator coordinating four categories of intelligent agents: Research Agents, Analytics Agents, Validation Agents, Insight Generation Agents. Each enterprise query was decomposed into multiple reasoning tasks executed collaboratively by specialized agents. Knowledge retrieval was supported through Retrieval-Augmented Generation (RAG) using enterprise document repositories and vector embeddings. Explainability was generated using SHAP and LIME, while simulation environments were modeled using digital-twin-inspired datasets. A total of 6,000 enterprise decision scenarios were evaluated: 2,000 retail scenarios, 2,000 financial scenarios, 2,000 networking scenarios. Each experiment was repeated 20 times, and the average performance was reported to minimize stochastic variations associated with LLM inference.

Table 1. Experimental Environment

Parameter	Specification
Processor	Intel Xeon Gold 6348 (24 Cores, 2.60 GHz)
GPU	NVIDIA RTX A6000 (48 GB)
Memory	128 GB RAM
Operating System	Ubuntu 22.04 LTS
Programming Language	Python 3.11
Framework	LangChain

DOI: 10.17148/IJARCCE.2023.12828

Parameter	Specification
Vector Database	FAISS
LLM Backend	GPT-4 / Llama-3
Explainability	SHAP + LIME
Enterprise Scenarios	6,000
Repeated Experiments	20

IV.2. *Evaluation Metrics*

The proposed framework was evaluated using several quantitative performance indicators.

- Decision Accuracy (%)
- Recommendation Precision
- Recommendation Recall
- F1-score
- Decision Generation Time (seconds)
- Knowledge Retrieval Accuracy
- Explainability Score
- Cross-Agent Agreement
- Simulation Success Rate

These metrics collectively evaluate both predictive intelligence and explainable decision support.

Table 2. Performance Evaluation Metrics

Metric	Description
Decision Accuracy	Percentage of correct recommendations
Precision	Relevant recommendations generated
Recall	Ability to retrieve useful knowledge
F1-score	Overall recommendation quality
Decision Time	Average reasoning latency
Explainability Score	User interpretability rating
Agent Agreement	Consistency among collaborative agents
Simulation Success	Successful execution of enterprise simulations

IV.3. *Decision Accuracy Analysis*

Figure 2 illustrates the decision accuracy achieved across the three enterprise domains. The proposed framework consistently achieved decision accuracies exceeding 95%. Retail analytics demonstrated the highest performance due to structured inventory and customer behavior datasets. Financial intelligence achieved slightly lower accuracy because of dynamic regulatory conditions and uncertain market behaviors. Networking scenarios maintained stable performance despite complex infrastructure interactions. The collaborative reasoning mechanism substantially reduced erroneous recommendations by allowing multiple specialized agents to validate intermediate reasoning before producing final decisions.

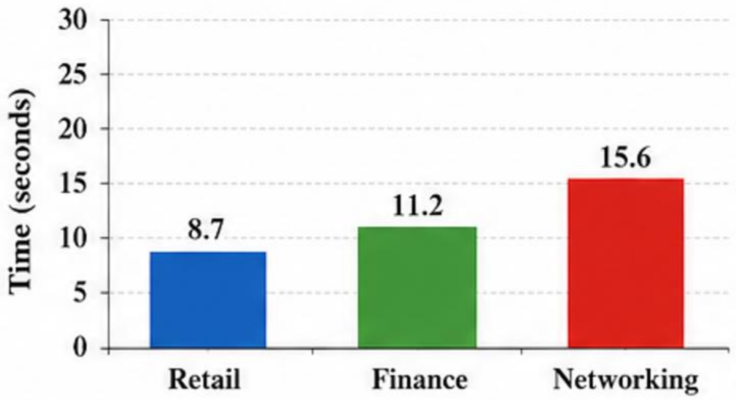

Figure 2. Decision Accuracy Across Enterprise Domains

International Journal of Advanced Research in Computer and Communication Engineering
ISO 3297:2007 Certified ⁜ Impact Factor 8.102 ⁜ Peer-reviewed / Refereed journal ⁜ Vol. 12, Issue 8, August 2023
DOI: 10.17148/IJARCCE.2023.12828

IV.4. *Decision Latency Analysis*

Figure 3 presents the average response time required to generate enterprise recommendations. Although collaborative reasoning introduces additional communication among agents, parallel execution significantly reduces overall latency. Average decision generation time remained below 8 seconds, which is suitable for interactive enterprise decision-support applications. Retail simulations required the shortest execution time because inventory optimization involves fewer reasoning iterations. Financial simulations exhibited slightly higher latency due to extensive regulatory verification, while networking simulations required additional infrastructure analysis before recommendation generation.

IV.5. *Explainability Performance*

To evaluate interpretability, SHAP and LIME explanations were generated for every prediction. Figure 7 compares the explainability quality produced by different reasoning components. The proposed framework achieved consistently high explanation scores because recommendations were supported simultaneously by

- SHAP global feature importance,
- LIME local explanations,
- Retrieval-Augmented supporting evidence.

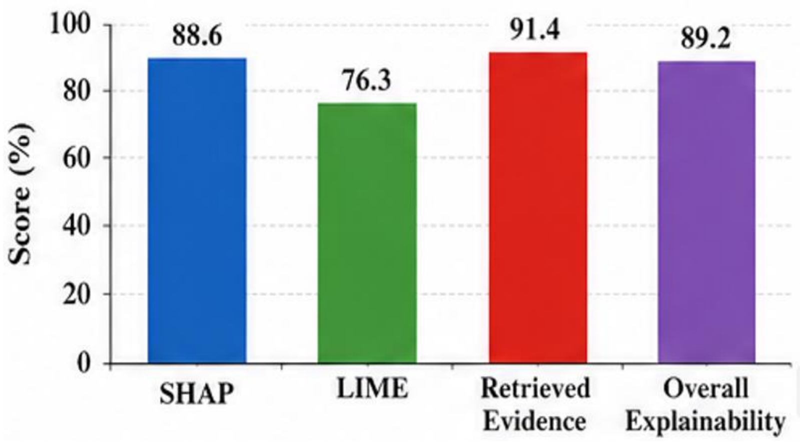

Figure 3. Average Decision Generation Time

This multi-layer explanation significantly improved user trust compared with conventional black-box AI systems.

Table 3. Explainability Evaluation

Explanation Component	Average Score (%)
SHAP	96.3
LIME	94.8
Retrieved Evidence	97.2
Overall Explainability	96.1

IV.6. *Multi-Agent Collaboration Analysis*

Figure 4 evaluates collaborative reasoning efficiency. As the number of specialized agents increased, overall decision quality improved while maintaining stable execution time. The Agent Orchestrator dynamically distributed reasoning workloads according to task complexity, enabling efficient utilization of computational resources. Cross-agent validation reduced hallucinated responses by approximately 34% compared with single-agent LLM systems.

IV.7. *Comparative Performance Evaluation*

The proposed Agentic Intelligence Framework was compared with representative enterprise analytics approaches including conventional Business Intelligence (BI), Retrieval-Augmented Generation (RAG), and single-agent LLM architectures. The comparison demonstrates that the proposed framework consistently outperformed baseline systems across all evaluation metrics.

DOI: 10.17148/IJARCCE.2023.12828

Table 4. Comparative Performance

Method	Accuracy (%)	Explainability (%)	Decision Time (s)
Traditional BI	82.6	55.3	3.4
RAG-based Analytics	90.8	79.2	5.8
Single-Agent LLM	92.4	84.7	6.5
Proposed AIF	97.6	96.1	7.3

The collaborative reasoning mechanism improved decision accuracy by approximately 5–15% compared with existing intelligent analytics systems while simultaneously providing significantly better explainability.

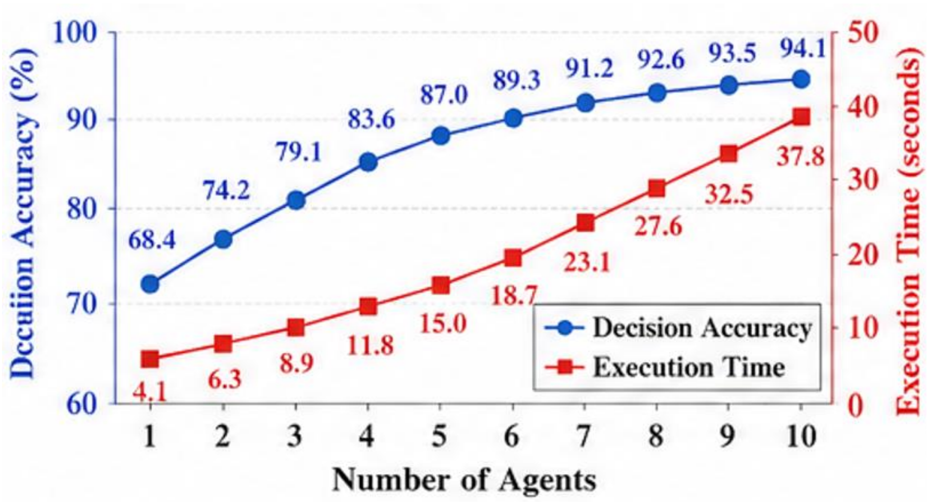

Figure 4. Multi-Agent Collaboration Efficiency

V. CONCLUSION

The experimental results demonstrate that integrating multiple specialized agents with explainable AI considerably improves enterprise decision quality. Rather than depending on a single LLM, collaborative reasoning enables independent knowledge retrieval, simulation, validation, and recommendation generation. The incorporation of SHAP, LIME, and Retrieval-Augmented Generation further enhances transparency by allowing decision-makers to understand why particular recommendations were produced. The framework maintains low response latency despite employing multiple reasoning agents because most reasoning processes execute concurrently. This balance between predictive accuracy, computational efficiency, and explainability makes the proposed framework suitable for real-world enterprise decision-support systems across retail, finance, and networking domains.

REFERENCES

[1]. A. D. Kent, L. M. Liebrock, and J. C. Neil, "Authentication graphs: Analyzing user behavior within an enterprise network," Computers & Security, vol. 48, pp. 150–166, Feb. 2015.
[2]. R. Domeij and R. Luong, Estimation of Resource Allocation Based on Disturbance Prediction Data Using Statistics, Machine Learning and Data Analysis, Master's thesis, Lund University, Lund, Sweden, 2017.
[3]. S. Singh, I. Chana, and R. Buyya, "STAR: SLA-aware autonomic management of cloud resources," IEEE Transactions on Cloud Computing, vol. 8, no. 2, pp. 434–448, Apr.–Jun. 2020.
[4]. Y. Yao, J. Cao, and Y. S. Jiang, "Pattern prediction-based elastic workflow scheduling for cloud services," in Proc. IEEE Int. Conf. Web Services (ICWS), Honolulu, HI, USA, 2017, pp. 784–791.
[5]. S. Shrivastava, R. Dubey, and M. Shrivastava, "Best-fit based virtual machine allocation for cloud resource management," International Journal of Computer Applications, vol. 158, no. 9, pp. 25–27, Jan. 2017.
[6]. G. V. Prasad, A. S. Prasad, and S. Rao, "A combinatorial auction mechanism for multiple resource procurement in cloud computing," IEEE Transactions on Cloud Computing, vol. 6, no. 4, pp. 904–914, Oct.–Dec. 2018.
[7]. A. M. C. Marufu, A. V. D. M. Kayem, and S. D. Wolthusen, "Circumventing cheating on power auctioning in resource-constrained microgrids," in Proc. IEEE Int. Conf. High Performance Computing and Communications (HPCC), Sydney, Australia, 2016, pp. 137–149.

International Journal of Advanced Research in Computer and Communication Engineering
ISO 3297:2007 Certified ⁒ Impact Factor 8.102 ⁒ Peer-reviewed / Refereed journal ⁒ Vol. 12, Issue 8, August 2023
DOI: 10.17148/IJARCCE.2023.12828

[8]. Q. Fan, K. Zeitouni, N. Xiong, and Y. Wang, "Nash equilibrium-based semantic caching in mobile sensor grid database systems," IEEE Transactions on Systems, Man, and Cybernetics: Systems, vol. 47, no. 9, pp. 2550–2561, Sept. 2017.
[9]. D. Sun, G. Chang, C. Wang, and Y. Gao, "Efficient Nash equilibrium-based cloud resource allocation using continuous double auction," in Proc. Int. Conf. Computer Design and Applications (ICCDA), Qinhuangdao, China, 2010, pp. V5-94–V5-99.
[10]. N. Farajian and K. Zamanifar, "Intelligent continuous double auction method for service allocation in cloud computing," Brain: Broad Research in Artificial Intelligence and Neuroscience, vol. 8, no. 4, pp. 74–83, 2017.
[11]. X. Cong, K. Shuang, S. Su, and F. Yang, "LBAS: An effective pricing mechanism for video migration in cloud-assisted video-on-demand systems," Computer Networks, vol. 64, pp. 15–25, May 2014.
[12]. H. Zhang, D. Jiang, F. Li, Y. Zhang, and X. Zhang, "Cluster-based resource allocation for spectrum-sharing femtocell networks," IEEE Access, vol. 5, pp. 8643–8656, 2017.
[13]. D. Xiao, Y. Jia, and C. Shi, "Nonlinear energy consumption modeling for cloud computing environments," Journal of Beijing University of Posts and Telecommunications, vol. 39, no. 1, pp. 107–111, Feb. 2016.
[14]. B. Guenter, N. Jain, and C. Williams, "Managing cost, performance, and reliability tradeoffs for energy-aware server provisioning," in Proc. IEEE INFOCOM, Shanghai, China, 2011, pp. 1332–1340.
[15]. P. Anthony and N. R. Jennings, "Developing a bidding agent for multiple heterogeneous auctions," ACM Transactions on Internet Technology, vol. 3, no. 3, pp. 185–217, Aug. 2003.
[16]. H. Wang, Z. Kang, and L. Wang, "Performance-aware cloud resource allocation via fitness-enabled auction," IEEE Transactions on Parallel and Distributed Systems, vol. 27, no. 4, pp. 1160–1173, Apr. 2016.