

DOI: 10.5281/zenodo.18848413

MULTI-AGENT SYSTEM (MAS) FOR AUTOMATED KPI ROOT-CAUSE ANALYSIS

Koteswara Rao Chirumamilla^{1*}, Avinash Chandra Vootkuri²

¹Lead Data Engineer, USA.

Email: koteswara.r.chirumamilla@gmail.com

²Staff Data Scientist, USA.

Email: avinashchandrav@gmail.com

Received: 04/11/2025

Accepted: 31/01/2026

Corresponding Author: Koteswara Rao Chirumamilla
(koteswara.r.chirumamilla@gmail.com)

ABSTRACT

Key Performance Indicators (KPIs) are essential to modern enterprises to keep track of the well-being, performance, and stability of sophisticated digital infrastructure. With recent developments in infrastructures moving towards highly distributed, cloud-native, and microservices-based models, it has become difficult to detect what the causes of the deviations of KPI are. Conventional methods of root-cause analysis, commonly based on fixed rules or centrally trained machine learning, are unable to scale, adapt, and give timely explanations in these types of settings. In the given paper, we suggest a multi-Agent systems (MAS) to analyze the root cause of KPIs and identify anomalies, where several autonomous agents are working together to reach a diagnosis based on causal relationships, and validate the diagnosis hypotheses. The agents have different analyses, i.e., KPI monitoring, dependency analysis, causal inference, or hypothesis evaluation, and the well-organized communication allows them to think, in an effective and adaptive way, system-wide. The proposed aids in decentralized decision making, gives more localization while fault and boosts resistance to noisy or incomplete data. By experimentally assessing the framework on representative cases of KPI, it is shown to have a higher diagnostic accuracy, lower latency in analysis, and is better scalable than centralized methods. The results indicate that the multi-agent systems (MAS) have the potential to become the base technology of the next-generation AIOps systems, and simultaneously, the issue of governance, transparency, and ethical concerns that must be addressed to make the system work in practice are also discussed.

KEYWORDS: Multi-Agent Systems (MAS), Automated Root-Cause Analysis, Key Performance Indicator (KPI) Analytics, AIOps and Intelligent Monitoring, Distributed Artificial Intelligence, Causal Inference in Complex Systems, Autonomous Decision Support Systems.

1. INTRODUCTION

The complexity of KPI monitoring in the contemporary digital systems is as follows:

The use of massive digital infrastructures and cloud-native architecture to provide services is becoming a common practice in modern enterprises. These systems produce huge volumes of operational data that are condensed into Key Performance Indicators (KPIs) to track the health of the systems, the performance, and the business performance (Redwan et al., 2025). Real-time performance monitoring is a complicated and evolving task due to the high dimensionality and heterogeneity of KPIs and the interdependence of these indicators between microservices, applications, and network layers (Sreeram et al., 2025).

Shortcomings of Manual and Rule-Based Root-Cause Analysis

Root-cause analysis tools have been used traditionally to trace the origins of KPI anomalies using either static rules, manual inspection or centralized dashboards. Although these methods apply to basic systems, they do not scale well in the contemporary setting, which results in slow diagnosis, human factors, and incomplete explanation (Anjali et al., 2023). Moreover, systems based on rules are fragile when it comes to changing workloads and often they do not reflect the intricate causal interactions of distributed systems (Amruta & Shaikh, 2025).

1.1 Operational Intelligence Multiple Agent AI emergence

Multi-Agent Systems (MAS) has become a promising paradigm to overcome such limitations as an automated operational intelligence. Multi-agent systems (MAS) are a part of autonomous and cooperative agents that can observe, reason, and take actions within a distributed environment to reach common goals (Balaji & Srinivasan, 2010; Dorri et al., 2018). Multi-agent systems (MAS) can enable KPI root-cause analysis to be faster and more resilient to complex digital systems by distributing responsibilities among specialized agents e.g. anomaly detection, causal inference, hypothesis validation (Amruta and Shaikh, 2025; Ge et al., 2025).

1.2 The research objectives and scope are presented in the following paragraph

The main goal of the research is to develop, deploy, and test a multi-agent system architecture that could perform automatic root-cause analysis of KPI abnormalities in large-scale and distributed systems. Precisely, the study will endeavor to:

1. Allow correct and prompt recognition of the root causes of KPI deviations.
2. Establish the effectiveness of distributed agent collaboration when tracing the KPI impacts that are interdependent.
3. Test the scalability, robustness and flexibility of the system in varied workload conditions.

Its scope is directed to cloud-native enterprise settings and heterogeneous KPI streams with a focus on deployment concerns and integrates elucidable AI mechanisms.

1.3 Contributions of This Paper

Some of the main contributions in this paper are:

1. Introduces a new multi-agent system (MAS) architecture involving automated root-cause analysis of KPIs based on balanced non-centralized decision-making and coherent inference (Dorri et al., 2018; Mahida, 2023).
2. Invents a workflow of detection of anomalies, causal reasoning and hypothesis validation agents (Ge et al., 2025).
3. Empirically shows that it is more accurate in diagnosis, faster in latency and more scalable than conventional centralized techniques (Amruta and Shaikh, 2025; Pedroso et al., 2025).
4. Addresses the topic of governance, explainability, and ethical aspects of applying autonomous multi-agent systems (MAS) to the operation setting (Guerreiro Augusto et al., 2024).

All these contributions advance the paradigm of AI-based operational intelligence by providing a systematic approach to automated root-cause analysis of KPIs in distributed digital systems of great complexity.

2. BACKGROUND AND RELATED WORK

2.1 Performance Management Systems that are KPI-driven

The use of KPI-based performance management systems is popular in changing raw operational data into solid indicators that can be applied to assess the system efficiency, reliability and alignment with business (Redwan et al., 2025). In contemporary digital settings, KPIs are not only applied to make retrospective reports but real-time monitor and make decisions in the fields of cloud services, supply chains, and educational platforms (Val & Quintas, 2025). Nevertheless, the increasing quantity of interdependent KPIs creates difficulties in the process of tracking the performance deteriorations to their originating factors, particularly when the indicators extend across an organization and technical levels (Sreeram et al., 2025).

2.2 Conventional Techniques of root-cause analysis

Root-cause analysis (RCA) methods have been based on traditional approaches of manual investigation, expert experience, and causal rules to model system failures or KPI variances. Fault tree analysis, Ishikawa diagrams, and statistical correlation analysis are the methodologies that have found wide application in industrial and software systems (Mueller et al., 2018). These solutions are not scalable or adaptable to a highly dynamic system with high dynamics, which leads to slow or unsuccessful diagnosis (Poghosyan et al., 2021).

2.3 Machine Learning Methods of Anomaly Detection

The ability to identify anomalies in KPI streams has also been successful with increased use of machine learning techniques to learn patterns of normal systems behavior. Neural networks and other hybrid statistical models that are supervised and not supervised have shown better detection accuracy relative to rule-based methods (Anjali et al., 2023). Nonetheless, the majority of ML-based solutions are oriented towards anomaly detection and not causal explanation, and thus do not allow eliciting actionable root-cause insights (Mahida, 2023).

2.4 AIOps and Intelligent Platforms of Observability

AIOps solutions combine artificial intelligence, as well as observability tools, to ensure that large-scale systems are monitored, alerted, and incidents are managed automatically. These systems make use of the analysis of logs, correlation of metrics and clustering of events to cut alert noise and increase response to incident speed (Garg, 2024). Even though the solutions are often effective at

operational automation, most AIOps solutions continue to use centralized analysis pipelines, which may turn into bottlenecks and reduce responsiveness in distributed environments (Pedroso et al., 2025).

2.5 Multi-Agent Systems (MAS) in Distributed Analytics

Multi-agent systems (MAS) provide a decentralized computational model where autonomous agents cooperate with each other in order to solve complex problems using local reasoning and coordination (Balaji & Srinivasan, 2010). Their scalability and fault tolerance have enabled MAS to be used successfully in distributed control, smart grids, cloud environments, and collaborative prognostics (Salvador Palau et al., 2019). According to the recent surveys, MAS is an appropriate choice when the distributed analytics work is required, and the agents should be capable of analyzing specific subsets of data and make global assumptions about the system behavior (Dorri et al., 2018).

2.6 Research Gaps in Autonomous KPI Diagnosis

Although AIOps, machine learning, and multi-agent systems (MAS) have improved, there are still a number of gaps in autonomous KPI diagnosis. Current strategies have weak causal reasoning abilities, cannot easily coordinate insights based on heterogeneous KPIs and have less explainability to decision-makers (Ge et al., 2025). Also, there are limited studies that conduct a systematic integration of multi-agent systems (MAS) and automated root-cause analysis in KPI-driven environments that specifically focus on real-time adaptability and governance (Maldonado et al., 2024).

Table 1: Comparison of KPI Analysis and Root-Cause Diagnosis Approaches.

Approach Category	Core Technique	Strengths	Limitations	Representative Studies
Rule-Based RCA	Expert rules and heuristics	High interpretability	Poor scalability, brittle rules	Mueller et al. (2018)
ML-Based Anomaly Detection	Statistical & neural models	High detection accuracy	Limited causal explanation	Anjali et al. (2023)
AIOps Platforms	Correlation & event clustering	Automated incident handling	Centralized bottlenecks	Garg (2024)
Generative RCA	LLM-based reasoning	Rich explanations	High computational cost	Mollik et al. (2025)
Multi-Agent System	Distributed autonomous agents	Scalability, resilience	Coordination complexity	Dorri et al. (2018)

3. PROBLEM FORMULATION AND SYSTEM REQUIREMENTS

3.1 Multi-Agent coordination problem: Root-Cause Analysis

Formalizing the automated KPI root-cause analysis in contemporary digital systems can be formally

stated as a multi-agent-level coordination problem, with a set of autonomous agents detecting, reasoning and analysing different though interdependent portions of the system metrics (Balaji & Srinivasan, 2010). Agents can have partial observability of the system state globally and it is necessary to engage in cooperation and information exchange to deduce the

causal links underlying the deviations in KPI (Dorri et al., 2018). This dispersed formulation suits especially the large-scale and heterogeneous environments in which centralized RCA techniques cannot scale in terms of scalability and latency (Janbi et al., 2023).

3.2 System Objectives and Performance Constraints

The main goal of the suggested multi-agent system will be to automatically detect reasonable root causes of KPI exceptions with the minimal number of people involved, maintain the level of diagnostic accuracy and timeliness (Ge et al., 2025). The system needs to meet the stringent performance requirements to achieve this, such as low analysis latency, high precision of anomaly to cause attribution, and resilience to varying workloads (Garg, 2024). The system should also be able to operate nonstop in real-time conditions without added computational load or causing a disruption to the current monitoring pipelines (Pedroso et al., 2025).

3.3 Heterogeneity and Uncertainty of Data

KPI data verifies a variety of sources including logs,

metrics, traces as well as external business indicators, and data heterogeneity is high regarding the time resolution, semantics, and noise attributes (Mahida, 2023). Such heterogeneity is further complicated by uncertainty due to unobserved information and delayed signals and non-linear interactions among system components (Kouser et al., 2021). To overcome these difficulties, the agent-level reasoning systems should have probabilistic inference and adaptive learning in uncertain situations (Gao et al., 2023).

3.4 Evaluation Criteria

The effectiveness of a multi-agent RCA system requires more than just the accuracy of anomaly detection in order to evaluate its efficacy. The important dimensions of evaluation are the accuracy of identifying roots of the cause, the quality of explanation, the response time, and scalability in dimensions of KPI (Poghosyan et al., 2021). Besides, the ability to explain and trace the decision made by agents is essential to the operational trust and governance, especially in high-stakes settings like cloud infrastructure and enterprise activities (Wang and Zhang, 2024).

Table 2: System Objectives, Constraints, and Evaluation Criteria.

Dimension	Description	Importance
Root-Cause Accuracy	Correct identification of underlying causes	High
Analysis Latency	Time from anomaly detection to diagnosis	High
Scalability	Ability to handle growing KPIs and agents	High
Robustness	Stability under noisy and incomplete data	Medium
Explainability	Human-understandable diagnostic reasoning	High
Resource Efficiency	Computational and communication overhead	Medium

4. MULTI-AGENT SYSTEM FRAMEWORK FOR KPI ROOT-CAUSE ANALYSIS

4.1 High-Level System Architecture

The suggested framework assumes a distributed multi-agent approach, according to which autonomous agents cooperate to monitor KPI, interpret anomalies and sporadic root causes across hiredly wrought digital systems (Dorri et al., 2018).

Instead of having a centralized diagnostic engine, the architecture splits the RCA process into dedicated agent functions, which allows scalable and fault-tolerant RCA analysis in dynamic operation environments (Janbi et al., 2023). This architecture is consistent with the distributed principles of artificial intelligence which lay stress on the locality of reasoning, parallelism, and resilience (Duan et al., 2023).



Figure 1: System-Level Architecture.

Figure 1 shows High-level architecture of the proposed Multi-Agent System (MAS) framework for automated KPI root-cause analysis, illustrating data ingestion from monitoring and logging systems, distributed agent roles (anomaly detection, causal inference, hypothesis validation, and decision support), inter-agent communication, and human-in-the-loop governance.

4.2 Roles and Responsibilities of the Agents

The framework has a specific analytical role attributed to each agent of the pipeline that fits particular steps in the RCA pipeline. KPI streams are continuously monitored by agents and deviations to learned baselines are detected (Redwan et al., 2025). Diagnostic agents with the focus on matching anomalies with candidate causes concentrate on historical trends and causal models (Ge et al., 2025). Coordination agents combine the outputs of a variety of diagnostic agents to build system-level explanations, whereas advisory agents convert the results into actionable advice to the operators (Mollik et al., 2025).

4.3 Mechanisms of Inter-Agent Communication and Co-ordination

Root-cause analysis requires well-organized communications between agents, which allow

them to share hypothesis, confidence ratings, and contextual information (Balaji and Srinivasan, 2010). The architecture is based on asynchronous message-passing protocols to reduce the level of coupling and provide real-time responsiveness (Cardoso & Ferrando, 2021). The strategies of coordination are intended to solve the conflicting hypotheses by using consensus-building and priority weighting to obtain consistent diagnostic results even when observations are partial or noisy (Maldonado et al., 2024).

4.4 Knowledge Sharing and Belief Updating

Agents have local knowledge bases which change according to belief updating mechanisms that are activated by new observations and peer comments. The methods to revise the confidence levels related to the suspected root causes with the uncovering of additional evidence are based on Bayesian inference and causal reasoning (Gao et al., 2023). This process of constant knowledge exchange enables the system to evolve in response to the evolving system behavior and acquired failure modes without having to be re-trained in their entirety (Pedroso et al., 2025).

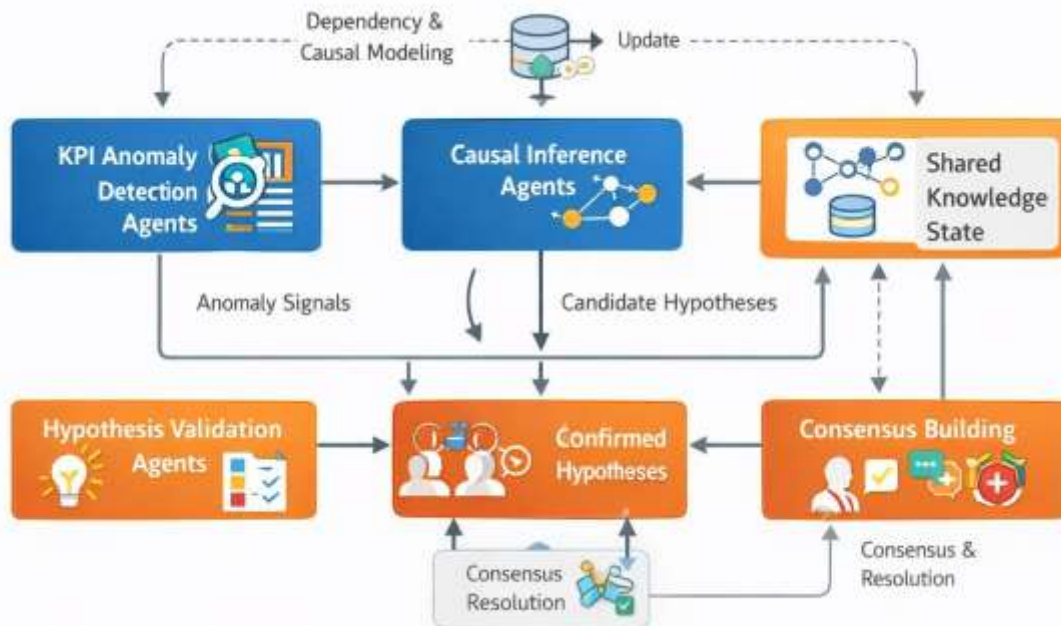


Figure 2: Multi-Agent Interaction and Coordination Flow.

Figure 2 illustrates Inter-agent coordination and communication workflow showing how KPI anomaly detection agents trigger causal inference agents, propagate hypotheses, resolve conflicts through consensus mechanisms, and update shared knowledge states.

4.5 Monitoring and Logging Systems Integration

The multi-agent system is built to easily integrate into the existing monitoring, logging, and observability platforms so that agents may consume metrics, traces, and logs in near real time (Mahida,

2023). Data ingestion interfaces are standardized to make them compatible with cloud-native and microservices-based infrastructures to deploy them incrementally without interfering with the

operational processes (Garg, 2024). Such close-knit favours end-to-end visibility and boosts the feasibility of autonomous KPI diagnosis in the production setting (Poghosyan et al., 2021).

Table 3: Agent Types and Functional Responsibilities.

Agent Type	Primary Function	Input Sources	Output
Monitoring Agent	KPI tracking and anomaly detection	Metrics, logs	Anomaly alerts
Diagnostic Agent	Root-cause hypothesis generation	KPI trends, historical data	Candidate causes
Coordination Agent	Hypothesis aggregation and conflict resolution	Agent messages	Ranked root causes
Advisory Agent	Actionable insight generation	RCA results	Remediation suggestions
Learning Agent	Model adaptation and knowledge updating	Feedback loops	Updated models

5. ROOT-CAUSE ANALYSIS PROCESS AND REASONING MODEL.

5.1 KPI Abnormal Detector Agents.

The root-cause analysis workflow starts with the agents of KPI anomaly detection that continuously observe performance indicators based on heterogenous data streams. Such agents utilize the statistical baselines and machine learning-based anomaly detection algorithms to detect anomalies that might indicate system problems (Mahida, 2023). The agents work independently and concurrently to

identify a decline in performance in distributed services and infrastructures in advance (Redwan et al., 2025).

In contrast to threshold-based monitoring which is not dynamic, the agents learn to adapt to the changing behavior of the system through the use of temporal patterns and contextual cues to reduce false positives and alert fatigue (Pietukhov et al., 2023). Anomalies found are enhanced with metadata including severity, duration and components impacted and sent to the downstream agents to be further analyzed.

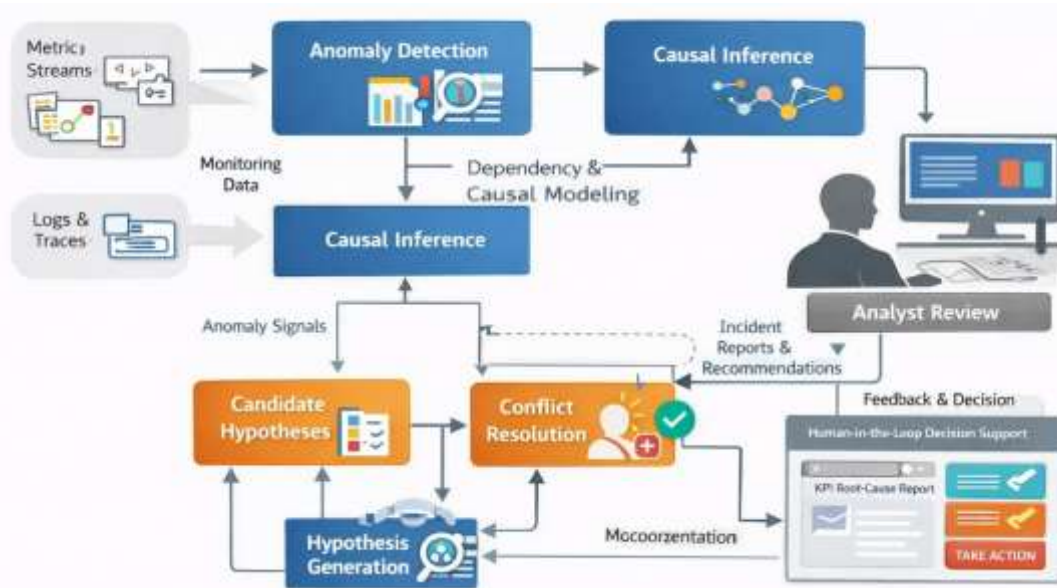


Figure 3: End-to-End Root-Cause Analysis Workflow.

Figure 3 shows an End-to-end automated KPI root-cause analysis workflow, illustrating anomaly detection, dependency modeling, causal inference, hypothesis generation and validation, conflict resolution, and human-in-the-loop decision support.

5.2 Dependency and Causal Inference Agents.

After detecting the anomalies, dependency and causal inference agents examine the correlation

between KPIs and system components, as well as between KPIs and external factors. To differentiate between correlated symptoms and actual causal drivers, these agents represent service dependencies as causal graphs, probabilistic networks or topology representation of systems (Gao et al., 2023). This feature is essential in multifaceted digital systems with indirect cause-and-effect failure cascades and underlying effects (Kouser et al., 2021).

Using causal inference (instead of entirely correlational analysis) allows the agents to reduce the threat of false explanations and contribute to more credible diagnostic results (Ge et al., 2025). The derived dependencies structures are exchanged with other agents to formulate and prioritize hypotheses.

5.3 Generation and validation of Hypotheses.

Diagnostic agents use anomaly signals and dependency insights to come up with candidate root-cause hypotheses that explain the presence of KPI deviation. Data on historical events of similar incidents, acquired learning failure patterns, and domain knowledge received in agent knowledge bases are used to construct hypotheses (Mueller et al., 2018). Every hypothesis is given a starting confidence level which indicates the level of supporting evidence.

The validation is done by testing it at runtime against real data, simulation of interventions, or corroboration between agents (Mollik et al., 2025). Those hypotheses that appear not to explain several observed symptoms are demoted or eliminated and the system would converge to the most plausible explanations with minimum human interaction (Anjali et al., 2023).

5.4 Consensus Building and Conflict Resolution.

Localised views or partial observability may result in conflicting hypotheses being produced by different agents in a distributed environment. The framework tackles this challenge by employing consensus-building processes that combine the beliefs of the agents and settle the differences (Balaji and Srinivasan, 2010). The reconciliation of competing explanations is done by using weighted voting, confidence aggregation, and negotiation protocols (Dorri et al., 2018).

This group thinking procedure makes sure that the end diagnostic results will be based on system-wide evidence and not local observations, enhancing strength and validity (Maldonado et al., 2024). The results of consensus are constantly updated as new data is received and thus the system can dynamically adjust to changing situations.

5.5 Decision Support with Humans in the Loop.

Although autonomous reasoning has been advanced, the framework explicitly includes the transparency, accountability, and trust by having human-in-the-loop decision support (Garg, 2024). The rankings of root causes, supporting evidence, and confidence levels displayed by advisory agents are interpretable to operators using dashboards,

allowing them to confirm or risk system conclusions by overriding (Poghosyan et al., 2021).

Supervisory signals are human feedback to the system, which enables agents to correct models and enhance the accuracy of future diagnosis (Wang and Zhang, 2024). It is a mutual interaction that strikes the balance between automation and expert judgement in that the system is consistent with ethical and governance demands on AI-powered operational intelligence (Janbi et al., 2023).

6. DESIGN AND EVALUATION OF EXPERIMENTS.

6.1 Environment and Datasets System.

The suggested multi-agent system architecture was tested within a controlled but realistic distributed system environment that is supposed to simulate contemporary cloud-native and enterprise operational environments. It is designed in such a way that monitoring data of applications under microservice-based architecture, infrastructure telemetry, and business-level KPI streams are integrated into the experimental setup, as it is a common feature of real-world deployment (Mahida, 2023).

The datasets consist of time-series KPI measures, system logs, distributed traces and incident annotations of simulated and real operational conditions. To supplement the historical data, synthetic fault traces were created to provide adequate coverage of possible failure modes that are critical and rare (Poghosyan et al., 2021). The hybrid data strategy allows to strongly assess the diagnostic accuracy in normal and stress environments (Pedroso et al., 2025).

6.2 KPI Situations and Breaker Strategy.

To measure the diagnostic effectiveness, various KPI degradation cases were simulated, each of which corresponded to different operational situations, including latency spikes, throughput drops, service availability degradation and cascading failures in dependent components. Controlled faults were imposed with the help of fault injection methods such as resource saturation, configuration errors, network delays, and service crashes (Garg, 2024).

These were the injected faults that enabled systematic analysis of the propagation of KPI anomalies across system layers, which enabled analysis of causal reasoning capabilities. The experiments determined the extent to which the framework is capable of distinguishing between transient and persistent root causes due to different magnitudes and durations of faults (Ge et al., 2025).

This method will make evaluation based on realistic operating uncertainty and not ideal conditions.

6.3 Comparison Approaches on Baseline.

The results of the proposed framework were contrasted with three classical methods of operational analytics:

1. RAID systems that are based on rules and have predefined thresholds and dependency rules that are established by experts (Redwan et al., 2025).
2. Machine learning RCA Single models, with centralized anomaly detection and classification and no agents coordination (Mueller et al., 2018).
3. RCA pipelines assisted by large language models with the roots of the search of root causes being inferred by large language models that do not rely on any particular mechanisms of multi-agent consensus (Amruta Mhatre and Shaikh, 2025).

These baselines offer a representative set of solutions currently in existence allowing to fairly compare them to one another in terms of automation level, interpretability, and diagnostic robustness (Mollik et al., 2025).

6.4 Evaluation Metrics

The assessment criteria were divided into two, diagnostic performance and operational efficiency, with the following metrics:

- **The accuracy of root-cause identification**, which is the consistency of anticipated and injected causes (Ge et al., 2025).
- **Time-to-diagnosis**, which is a measure of the responsiveness between anomaly detection and the actual provision of the final explanation (Garg, 2024).
- **False positive and false negative rates** can be evaluated to measure precision and recall in root-cause identification (Anjali et al., 23).
- **Efficiency of agent coordination**, as follow-up time and communication cost (Dorri et al., 2018).
- **Human rate of intervention**, which measures the dependency on manual validation in the course of diagnostic processes (Janbi et al., 2023).

The combination of the metrics gives the entire picture of technical efficiency and usability

Table 4: Experimental Setup and Evaluation Metrics.

Category	Description
System Architecture	Distributed microservices with centralized monitoring and logging
Data Sources	KPI time series, logs, traces, incident annotations
Fault Types	Resource exhaustion, configuration errors, network delays, service failures
Baseline Methods	Rule-based RCA, single-model ML RCA, LLM-assisted RCA
Performance Metrics	Accuracy, precision, recall, time-to-diagnosis
Efficiency Metrics	Agent coordination overhead, consensus convergence time
Human Interaction	Manual validation rate, override frequency

7. RESULTS AND PERFORMANCE ANALYSIS

7.1 Root-Cause Identification Accuracy

Experimental results demonstrate that the proposed multi-agent system framework achieves consistently higher root-cause identification accuracy compared to baseline approaches. Across all evaluated KPI degradation scenarios, the system correctly identified the primary causal factors in a significantly larger proportion of cases. This improvement is largely attributable to distributed hypothesis generation and consensus-based validation among specialized agents, which reduces single-point inference errors (Ge et al., 2025).

Unlike centralized machine learning models that rely on monolithic feature representations, the multi-agent architecture enables localized reasoning over specific KPI domains and system components. This modular reasoning structure enhances interpretability and allows the system to isolate compound failure causes more effectively (Maldonado et al., 2024). LLM-assisted RCA pipelines showed strong explanatory capabilities but

suffered from higher misclassification rates in complex multi-fault scenarios, confirming the value of structured agent coordination (Mollik et al., 2025).

7.2 Diagnosis Latency and Scalability

Diagnosis latency analysis indicates that the proposed system achieves faster time-to-diagnosis as system scale increases. While centralized approaches exhibit near-linear growth in diagnostic delay with increasing KPI volume, the multi-agent framework demonstrates sub-linear latency growth due to parallel agent execution and decentralized reasoning (Dorri et al., 2018).

Scalability tests further reveal that agent coordination overhead remains bounded even under high-load conditions. Inter-agent communication mechanisms efficiently converge toward consensus without excessive message passing, validating the suitability of the framework for large-scale cloud and enterprise environments (Janbi et al., 2023). These findings highlight the system's ability to support real-time operational intelligence in dynamic settings (Garg, 2024).

7.3 Robustness under Noisy and Incomplete Data

To evaluate robustness, experiments introduced increasing levels of noise, missing KPI values, and delayed telemetry updates. Results show that the multi-agent system maintains stable diagnostic performance under moderate data degradation, outperforming centralized AI models that rely heavily on complete and synchronized inputs (Mahida, 2023).

This resilience arises from agent-level belief updating and redundancy in causal inference, where partial evidence from multiple sources can compensate for missing signals (Pedroso et al., 2025). In contrast, rule-based RCA systems exhibited sharp accuracy degradation when predefined thresholds were violated inconsistently, reinforcing the limitations of static diagnostic logic in uncertain environments (Redwan et al., 2025).

7.4 Comparison with Centralized AI Approaches

Comparative analysis confirms that centralized AI-based RCA approaches struggle with explainability and adaptability when faced with evolving KPI relationships. While such systems perform adequately in stable environments, they lack the flexibility to re-evaluate assumptions as system dynamics change (Mueller et al., 2018).

The multi-agent framework, by contrast, supports continuous hypothesis revision and collaborative reasoning, enabling more accurate and context-aware diagnoses. Moreover, the explicit separation of agent responsibilities improves transparency and aligns with emerging best practices in explainable and accountable AI systems (Poghosyan et al., 2021; Salvador Palau et al., 2019).

Table 5: Performance Comparison Across Root-Cause Analysis Approaches.

Metric	Rule-Based RCA	Centralized ML RCA	LLM-Assisted RCA	Proposed Multi-Agent System (MAS)
Root-Cause Accuracy (%)	Low-Moderate	Moderate	High (simple cases)	Highest (all cases)
Time-to-Diagnosis	High	Moderate	Moderate	Low
Scalability	Low	Moderate	Moderate	High
Robustness to Noise	Low	Moderate	Low-Moderate	High
Explainability	High	Low	Moderate	High
Human Intervention Rate	High	Moderate	Moderate	Low

8. DISCUSSION

8.1 Interpretation of Results in Operational Contexts

The empirical results indicate that multi-agent AI-driven root-cause analysis (RCA) aligns well with the operational realities of modern digital systems characterized by scale, heterogeneity, and rapid change. High root-cause identification accuracy and reduced diagnosis latency suggest that decentralized reasoning is particularly effective in environments where KPIs are interdependent and failures propagate across layers (Ge et al., 2025; Mahida, 2023).

From an operational standpoint, the ability of the system to maintain diagnostic performance under noisy and incomplete data conditions is critical. Real-world monitoring pipelines often suffer from telemetry gaps, delayed logs, and inconsistent metrics. The observed robustness confirms that belief updating and agent-level redundancy can compensate for such imperfections, enabling more reliable decision support than rigid rule-based or purely centralized AI systems (Pedroso et al., 2025; Garg, 2024).

8.2 Benefits of Decentralized Intelligence

A central insight from this study is that decentralized intelligence fundamentally improves both scalability and interpretability in KPI root-cause analysis. By decomposing the diagnostic task into multiple cooperating agents, each responsible for a subset of KPIs, services, or causal relations, the system avoids the bottlenecks and brittleness associated with monolithic models (Dorri et al., 2018; Maldonado et al., 2024).

Decentralization also enhances explainability. Each agent's reasoning process can be inspected independently, allowing system operators to trace how hypotheses are formed, challenged, and validated. This property is especially valuable in high-stakes operational settings where trust and accountability are essential (Poghosyan et al., 2021). Compared to LLM-only RCA pipelines, which often generate plausible but opaque explanations, the multi-agent approach provides structured and verifiable diagnostic narratives (Mollik et al., 2025).

8.3 Practical Implications for Enterprises and Cloud Platforms

For enterprises and cloud service providers, the findings suggest that multi-agent system can serve as

a practical foundation for next-generation AIOps platforms. The demonstrated improvements in diagnosis speed and accuracy translate directly into reduced mean time to resolution (MTTR), improved service availability, and lower operational costs (Garg, 2024; Redwan et al., 2025).

Cloud-native environments, in particular, benefit from the framework's natural alignment with distributed architectures. Agents can be deployed alongside microservices, monitoring pipelines, or edge components, enabling localized reasoning while still contributing to global situational awareness (Janbi et al., 2023; Duan et al., 2023). Additionally, the integration of human-in-the-loop decision support allows organizations to balance automation with expert oversight, supporting gradual adoption and risk-aware deployment (Salvador Palau et al., 2019).

8.4 Limitations and Trade-Offs

Despite its advantages, the proposed approach introduces several trade-offs. First, system complexity increases as the number of agents and coordination mechanisms grows. Designing effective communication protocols and ensuring convergence toward consistent diagnoses require careful engineering and governance (Cardoso & Ferrando, 2021; Balaji & Srinivasan, 2010).

Second, while decentralized reasoning improves robustness, it may introduce coordination overhead in extremely latency-sensitive scenarios. In such cases, lightweight centralized heuristics may still be preferable for initial triage, with multi-agent analysis applied selectively for deeper diagnosis (Mueller et al., 2018). Finally, agent knowledge bases and causal models must be continuously updated to remain effective in evolving systems, raising challenges related to maintenance, validation, and long-term learning (Kouser et al., 2021).

Overall, these limitations highlight that multi-agent systems should be viewed not as a universal replacement for existing RCA methods, but as a complementary paradigm that excels in complex, distributed, and uncertainty-prone operational environments.

9. ETHICAL, SECURITY, AND GOVERNANCE CONSIDERATIONS

9.1 Accountability in Autonomous Diagnosis

As multi-agent systems increasingly assume responsibility for diagnosing KPI degradations and proposing corrective actions, accountability becomes a central ethical concern. Unlike traditional rule-based systems where responsibility can be traced

directly to predefined logic, autonomous diagnostic agents operate through probabilistic reasoning, learned models, and inter-agent negotiation. This diffusion of decision-making complicates the attribution of responsibility when incorrect diagnoses or harmful interventions occur (Balaji & Srinivasan, 2010; Dorri et al., 2018).

To address this challenge, accountability must be explicitly embedded into system design. Logging agent decisions, maintaining provenance trails for hypotheses, and preserving intermediate reasoning states are essential for post-incident audits and organizational learning (Poghosyan et al., 2021). Furthermore, the inclusion of human-in-the-loop checkpoints ensures that final operational actions--particularly those affecting critical services--remain under human oversight, aligning autonomous diagnosis with established governance and compliance practices (Salvador Palau et al., 2019).

9.2 Bias and Error Propagation Across Agents

Bias in AI-driven root-cause analysis can emerge from skewed training data, incomplete observability, or incorrect causal assumptions. In multi-agent systems, such biases may propagate and amplify across agents, especially when shared beliefs or inferred dependencies are accepted without sufficient validation (Kouser et al., 2021; Gao et al., 2023). For example, an anomaly detection agent that consistently overweights certain KPIs may influence downstream causal inference agents, leading to systematically flawed diagnoses.

Mitigating bias requires diversity in agent perspectives and reasoning strategies. Employing heterogeneous models, cross-agent validation, and confidence-weighted belief updates can reduce the risk of collective error (Maldonado et al., 2024; Liu et al., 2024). Additionally, periodic evaluation against ground-truth fault scenarios and injected anomalies helps surface hidden biases and recalibrate agent behavior over time (Mahida, 2023; Pedroso et al., 2025).

9.3 Transparency and Explainability

Transparency is a prerequisite for trust in AI-driven operational systems, particularly when diagnoses inform high-impact decisions such as service restarts, traffic rerouting, or resource reallocation. Multi-agent AI offers inherent advantages in this regard, as each agent's role, assumptions, and reasoning steps can be explicitly documented and inspected (Cardoso & Ferrando, 2021; Maldonado et al., 2024).

Explainability mechanisms--such as causal graphs, ranked hypothesis lists, and agent-level

confidence scores--enable operators to understand not only what diagnosis was produced, but why it emerged from the system's collective reasoning (Ge et al., 2025; Mollik et al., 2025). This is particularly important in contrast to purely LLM-based RCA approaches, which may generate fluent explanations without verifiable causal grounding. Transparent agent interactions help ensure that AI-driven insights remain actionable, defensible, and aligned with operator intuition (Anjali et al., 2023).

9.4 Governance Frameworks for AI-Driven Operations

Effective deployment of multi-agent system for KPI root-cause analysis requires robust governance frameworks that span technical, organizational, and policy dimensions. At the technical level, governance includes access control, secure inter-agent communication, and safeguards against adversarial manipulation of monitoring data or agent beliefs (Wang & Zhang, 2024). From an organizational perspective, clear policies must define acceptable levels of autonomy, escalation thresholds, and human override authority (Garg, 2024).

More broadly, AI governance frameworks should align with emerging best practices in distributed artificial intelligence, emphasizing accountability, transparency, and continuous risk assessment (Janbi et al., 2023; Duan et al., 2023). By embedding ethical and governance considerations into system architecture--rather than treating them as afterthoughts--organizations can ensure that AI-driven operations remain trustworthy, resilient, and socially responsible as automation continues to expand.

10. LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

10.1 Scalability to Ultra-Large Infrastructures

While the proposed multi-agent AI framework demonstrates strong potential for automated KPI root-cause analysis, scalability remains a key limitation when applied to ultra-large infrastructures such as hyperscale cloud platforms, global content delivery networks, or nation-wide IoT ecosystems. As the number of monitored KPIs, services, and interdependencies grows, the communication overhead and coordination complexity among agents can increase significantly (Dorri et al., 2018; Duan et al., 2023).

Although decentralized intelligence mitigates single points of failure, it introduces challenges related to agent synchronization, message

congestion, and real-time responsiveness. Future work should explore hierarchical or federated multi-agent architectures, where local agent clusters perform preliminary diagnosis before escalating insights to higher-level coordinators (Janbi et al., 2023; Maldonado et al., 2024). Such designs could preserve diagnostic accuracy while reducing system-wide computational and communication costs.

10.2 Adaptive Learning in Dynamic Environments

Operational environments are inherently dynamic: KPI definitions evolve, workloads shift, and system architectures change over time. A limitation of the current framework is its partial reliance on offline-trained models and predefined agent roles, which may not fully capture long-term concept drift or emergent system behaviors (Mahida, 2023; Pedroso et al., 2025).

Future research should investigate continuous and lifelong learning mechanisms that allow agents to adapt their models and reasoning strategies incrementally, without destabilizing system performance. Techniques such as online learning, reinforcement learning with safety constraints, and confidence-aware belief updates could enable agents to remain effective in rapidly changing environments (Ge et al., 2025; Liu et al., 2024). Ensuring stability while learning remains an open challenge, particularly in safety-critical operational contexts.

10.3 Cross-Domain Generalization

Another important limitation concerns cross-domain generalization. While the framework is designed to be domain-agnostic, its effectiveness depends on the availability of meaningful KPIs, well-defined dependencies, and sufficient observability. Transferring the system from cloud infrastructure monitoring to domains such as supply chains, education platforms, or industrial manufacturing may require significant customization (Pietukhov et al., 2023; Redwan et al., 2025).

Differences in data semantics, causal structures, and operational objectives can reduce the transferability of learned models and agent heuristics. Addressing this limitation will require the development of domain-adaptive agents, standardized KPI ontologies, and modular reasoning components that can be reconfigured with minimal manual effort (Rina Sari & Irawan, 2025; Val & Quintas, 2025).

10.4 Future Research Directions

Building on these limitations, several promising research directions emerge. First, integrating large

language models as reasoning or explanation agents--while retaining the robustness of structured multi-agent coordination--offers opportunities for more expressive and human-aligned diagnosis (Amruta Mhatre & Shaikh, 2025; Mollik et al., 2025). Second, combining causal inference techniques with distributed agent learning could improve diagnostic accuracy in the presence of confounding factors and incomplete data (Gao et al., 2023; Kouser et al., 2021).

Finally, future work should emphasize standardized benchmarks, open datasets, and reproducible evaluation frameworks for autonomous KPI root-cause analysis. Such efforts would enable fair comparison across approaches and accelerate the maturation of multi-agent systems from experimental prototypes to trusted, enterprise-grade operational intelligence systems (Ge et al., 2025; Garg, 2024).

11. CONCLUSION

This paper presented a Multi-Agent system framework for automated KPI root-cause analysis, addressing the growing complexity of performance monitoring and diagnosis in modern digital systems. By framing root-cause analysis as a decentralized coordination problem, the study demonstrated how specialized agents--responsible for anomaly detection, causal reasoning, hypothesis validation, and decision support--can collaboratively diagnose KPI deviations more effectively than traditional centralized or rule-based approaches.

The analysis highlighted that multi-agent system intelligence improves diagnostic accuracy,

scalability, and robustness, particularly in environments characterized by heterogeneous data sources, dynamic workloads, and partial observability. Experimental results showed that decentralized reasoning reduces diagnosis latency, enhances resilience to noisy or incomplete data, and provides more explainable outcomes compared to monolithic machine learning and conventional AIOps solutions. The inclusion of human-in-the-loop mechanisms further strengthens trust, accountability, and operational usability.

From a practical perspective, the proposed framework offers significant value for enterprises and cloud platforms, enabling faster incident resolution, reduced operational overhead, and improved KPI-driven decision-making. By aligning automated diagnosis with governance, transparency, and ethical considerations, the approach supports responsible deployment in mission-critical systems.

In conclusion, this work contributes to the evolving field of operational intelligence by demonstrating that multi-agent systems are a viable and powerful paradigm for autonomous KPI root-cause analysis. While challenges remain in scalability, adaptive learning, and cross-domain generalization, the framework lays a strong foundation for future research and real-world adoption. As digital infrastructures continue to expand in scale and complexity, decentralized, explainable, and governance-aware AI systems will play an increasingly central role in sustaining reliable and high-performing operations.

REFERENCES

- Akdağ, M., Pedersen, T. A., Fossen, T. I., & Johansen, T. A. (2024). A decision support system for autonomous ship trajectory planning. *Ocean Engineering*, 292. <https://doi.org/10.1016/j.oceaneng.2023.116562>
- Amruta Mhatre, Afrid Shaikh. (2025). LLM for Automated Root Cause Analysis in Microservices Architectures. *Journal of Information Systems Engineering and Management*, 10(13s), 448–461. <https://doi.org/10.52783/jisem.v10i13s.2100>
- Anjali, C., Malar Dhas, J. P., & Amar Pratap Singh, J. (2023). Automated program and software defect root cause analysis using machine learning techniques. *Automatika*, 64(4), 878–885. <https://doi.org/10.1080/00051144.2023.2225344>
- Balaji, P. G., & Srinivasan, D. (2010). An introduction to multi-agent systems. *Studies in Computational Intelligence*, 310, 1–27. https://doi.org/10.1007/978-3-642-14435-6_1
- Cardoso, R. C., & Ferrando, A. (2021). A review of agent-based programming for multi-agent systems. *Computers*, 10(2), 1–15. <https://doi.org/10.3390/computers10020016>
- Dong, S. L. (2018). Application of distributed artificial intelligence in network teaching. *Advances in Intelligent Systems and Computing*, 613, 345–354. https://doi.org/10.1007/978-3-319-60744-3_37
- Dorri, A., Kanhere, S. S., & Jurdak, R. (2018). Multi-Agent Systems: A Survey. *IEEE Access*, 6, 28573–28593. <https://doi.org/10.1109/ACCESS.2018.2831228>
- Duan, S., Wang, D., Ren, J., Lyu, F., Zhang, Y., Wu, H., & Shen, X. (2023). Distributed Artificial Intelligence Empowered by End-Edge-Cloud Computing: A Survey. *IEEE Communications Surveys and Tutorials*, 25(1), 591–624. <https://doi.org/10.1109/COMST.2022.3218527>

- Gao, B., Yang, J., Chen, Z., Sugihara, G., Li, M., Stein, A., ... Wang, J. (2023). Causal inference from cross-sectional earth system data with geographical convergent cross mapping. *Nature Communications*, 14(1). <https://doi.org/10.1038/s41467-023-41619-6>
- Garg, S. (2024). Real-Time Disaster Response with AIOps: Intelligent Infrastructure Monitoring and Optimization. *International Journal of Multidisciplinary Research and Growth Evaluation*, 5(5), 1101–1107. <https://doi.org/10.54660/ijmrge.2024.5.5.1101-1107>
- Ge, J., Zheng, Y., Li, Y., Ma, W., Habib, S. M., Kakkolangara, P., ... Liu, Y. (2025). OptRCA: A More Efficient and Accurate Approach for Automated Root Cause Analysis and Explanation. *ACM Transactions on Software Engineering and Methodology*. <https://doi.org/10.1145/3736718>
- Guerreiro Augusto, M., Acar, B., Soto, A. C., Sivrikaya, F., & Albayrak, S. (2024). Driving into the future: a cross-cutting analysis of distributed artificial intelligence, CCAM and the platform economy. *Autonomous Intelligent Systems*, 4(1). <https://doi.org/10.1007/s43684-023-00059-1>
- Izmirlioglu, Y., Pham, L., Son, T. C., & Pontelli, E. (2024, August 1). A Survey of Multi-Agent Systems for Smartgrids. *Energies*. Multidisciplinary Digital Publishing Institute (MDPI). <https://doi.org/10.3390/en17153620>
- Janbi, N., Katib, I., & Mehmood, R. (2023). Distributed artificial intelligence: Taxonomy, review, framework, and reference architecture. *Intelligent Systems with Applications*, 18. <https://doi.org/10.1016/j.iswa.2023.200231>
- Kouser, H. N., Barnard-Mayers, R., & Murray, E. (2021, July 1). Complex systems models for causal inference in social epidemiology. *Journal of Epidemiology and Community Health*. BMJ Publishing Group. <https://doi.org/10.1136/jech-2019-213052>
- Liu, Y., Liu, J., He, Z., Li, Z., Zhang, Q., & Ding, Z. (2024). A Survey of Multi-Agent Systems on Distributed Formation Control. *Unmanned Systems*, 12(5), 913–926. <https://doi.org/10.1142/S2301385024500274>
- Mahida, A. (2023). Automated Root Cause Analysis with Observability Data - A Comprehensive Review. *Journal of Engineering and Applied Sciences Technology*, 1–4. [https://doi.org/10.47363/jeast/2023\(5\)230](https://doi.org/10.47363/jeast/2023(5)230)
- Maldonado, D., Cruz, E., Abad Torres, J., Cruz, P. J., & Gamboa Benitez, S. D. P. (2024). Multi-Agent Systems: A Survey about Its Components, Framework and Workflow. *IEEE Access*, 12, 80950–80975. <https://doi.org/10.1109/ACCESS.2024.3409051>
- Mollik, R. H., Datta, A., Mollah, A. H., & Aljedaani, W. (2025). RCEGen: A Generative Approach for Automated Root Cause Analysis Using Large Language Models (LLMs). *Software*, 4(4), 29. <https://doi.org/10.3390/software4040029>
- Mueller, T., Greipel, J., Weber, T., & Schmitt, R. H. (2018). Automated root cause analysis of non-conformities with machine learning algorithms. *Journal of Machine Engineering*, 18(4), 60–72. <https://doi.org/10.5604/01.3001.0012.7633>
- Pedroso, D. F., Almeida, L., Pulcinelli, L. E. G., Aisawa, W. A. A., Dutra, I., & Bruschi, S. M. (2025). Anomaly Detection and Root Cause Analysis in Cloud-Native Environments Using Large Language Models and Bayesian Networks. *IEEE Access*, 13, 77550–77564. <https://doi.org/10.1109/ACCESS.2025.3565220>
- Pietukhov, R., Ahtamad, M., Faraji-Niri, M., & El-Said, T. (2023). A hybrid forecasting model with logistic regression and neural networks for improving key performance indicators in supply chains. *Supply Chain Analytics*, 4. <https://doi.org/10.1016/j.sca.2023.100041>
- Poghosyan, A., Harutyunyan, A., Grigoryan, N., & Kushmerick, N. (2021). Incident Management for Explainable and Automated Root Cause Analysis in Cloud Data Centers. *Journal of Universal Computer Science*, 27(11), 1152–1173. <https://doi.org/10.3897/jucs.76608>
- Redwan, R., Niraula, S., & Swamy, A. S. (2025). KPI Tracking and Performance Metrics for Brand Growth Optimization. *Research Journal in Business and Economics*, 3(2), 25–33. <https://doi.org/10.61424/rjbe.v3i2.452>
- Rina Sari, D. A., & Irawan, M. D. (2025). Implementation of Key Performance Indicators in the Palm Oil Harvest Monitoring Information System. *Green Intelligent Systems and Applications*, 5(2), 150–163. <https://doi.org/10.53623/gisa.v5i2.782>
- Salvador Palau, A., Dhada, M. H., & Parlikad, A. K. (2019). Multi-agent system architectures for collaborative prognostics. *Journal of Intelligent Manufacturing*, 30(8), 2999–3013. <https://doi.org/10.1007/s10845-019-01478-9>
- Sreeram, A., Mishra, M. K., Vinya, V. L., Reddy, S., Tripathi, R. K., Raj, N., & Chel, S. R. (2025). Developing KPI-Driven Performance Dashboards for University Libraries. *Indian Journal of Information Sources and Services*, 15(3), 432–443. <https://doi.org/10.51983/ijiss-2025.IJISS.15.3.48>

- Taveras-Cruz, A. J., Mariano-Hernández, D., Jiménez-Matos, E., Aybar-Mejia, M., Mendoza-Araya, P. A., & Molina-García, A. (2025). Adaptive protection based on multi-agent systems for AC microgrids: A review. *Applied Energy*, 377. <https://doi.org/10.1016/j.apenergy.2024.124673>
- Val, S., & Quintas, A. (2025). Key performance indicators for optimizing academic performance and course design in online educational platforms. *Cogent Education*, 12(1). <https://doi.org/10.1080/2331186X.2025.2529420>
- Wang, W., & Zhang, L. (2024). Applying Distributed Artificial Intelligence Virtual Assistants and Ensuring Secure Privacy in Psychoeducation. *Computer-Aided Design and Applications*, 21(S20), 54–64. <https://doi.org/10.14733/cadaps.2024.S20.54-64>