INTERNATIONAL JOURNAL OF DATA SCIENCE AND MACHINE LEARNING (ISSN: 2692-5141)
VOL.06 ISSUE 01 (2026)

Research Article

Root-Cause Knowledge Extraction Enabling Automated Resilience in Decentralized Organizational Platforms with Generative AI Techniques

Dr. Chanin Srisawat [1]

[1]Department of Enterprise Software Systems Bangkok Institute of Artificial Intelligence, Bangkok, Thailand

Received: 12 January 2026
Revised: 2 February 2026
Accepted: 20 March 2026
Published: 31 May 2026

Abstract

Modern decentralized organizational platforms, particularly those built on distributed cloud-native and microservices architectures, face increasing complexity in maintaining operational resilience. Failures in such systems are often non-linear, multi-factorial, and difficult to diagnose due to the absence of centralized control and the presence of heterogeneous telemetry sources. This research proposes a conceptual and analytical framework for root-cause knowledge extraction (RCKE) enhanced by generative AI techniques to enable automated resilience in decentralized organizational platforms.

The study synthesizes advancements in generative artificial intelligence, deep learning, and system-level fault analysis to construct an integrated model capable of identifying latent failure dependencies, extracting causal knowledge from system logs, and enabling autonomous recovery mechanisms. Prior work demonstrates that generative models, particularly GANs and diffusion-based architectures, can significantly enhance data synthesis and diagnostic precision in medical imaging domains (Akbar et al., 2024; Alalwan et al., 2024). These findings are extended into the domain of distributed system resilience, where synthetic data generation and pattern reconstruction are essential for diagnosing rare or unseen system failures.

Furthermore, recent research emphasizes the role of generative AI in complex adaptive systems, including brain-computer interfaces and medical decision systems, highlighting its capacity for contextual reasoning and pattern abstraction (Eldawlatly, 2024; Celard et al., 2023). Building upon these foundations, this paper introduces a structured framework that combines causal inference graphs, large language models (LLMs), and Kubernetes-based self-healing mechanisms to automate root-cause detection and remediation.

A critical dependency in this framework is the integration of post-mortem analytical intelligence for system recovery, as highlighted in recent studies on multi-cloud self-healing systems (Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes, 2026), which underscores the importance of retrospective learning for continuous resilience improvement.

The proposed approach contributes to the fields of distributed computing, artificial intelligence, and organizational systems engineering by offering a unified methodology for automated fault diagnosis, knowledge extraction, and resilience orchestration in decentralized environments.

Keywords: Root-cause analysis, generative AI, decentralized systems, automated resilience, Kubernetes, large language models, causal inference, self-healing systems, distributed computing, post-mortem intelligence.

INTRODUCTION

The increasing adoption of decentralized organizational platforms has fundamentally transformed how modern enterprises design, deploy, and manage computational systems. These platforms, typically built upon microservices architectures, multi-cloud deployments, and container orchestration frameworks such as Kubernetes, offer scalability, flexibility, and fault isolation. However, they also introduce significant challenges in system observability, failure detection, and root-cause analysis due to their inherently distributed nature.

In such environments, failures rarely originate from a single component. Instead, they emerge from complex interdependencies between services, network layers, data pipelines, and resource schedulers. This complexity makes traditional debugging and monitoring approaches insufficient for ensuring continuous system resilience. Consequently, there is a growing need for intelligent systems capable of autonomously extracting root-cause knowledge and enabling automated recovery mechanisms.

Recent advances in artificial intelligence, particularly generative AI, have demonstrated remarkable capabilities in pattern recognition, data synthesis, and contextual reasoning. Studies in medical imaging and biomedical engineering highlight how generative adversarial networks (GANs) and transformer-based architectures can reconstruct missing information, enhance diagnostic accuracy, and model complex distributions in high-dimensional data spaces (Akbar et al., 2024; Alalwan et al., 2024). These capabilities are directly transferable to distributed systems, where missing telemetry, noisy logs, and incomplete failure traces are common challenges.

Moreover, deep learning-based frameworks have shown significant promise in extracting structured knowledge from unstructured datasets. For instance, research on brain tumor detection and segmentation demonstrates how AI systems can integrate multi-source data to produce accurate classification and diagnostic outputs (Humayun et al., 2024). Similarly, in decentralized organizational platforms, multi-source telemetry data can be leveraged to infer system health and identify latent failure patterns.

A key limitation in current distributed system management is the lack of automated causal reasoning. While monitoring tools can detect anomalies, they often fail to explain the underlying cause or predict cascading failures. This gap motivates the integration of large language models (LLMs) and causal inference mechanisms into system observability pipelines. LLMs, with their ability to interpret structured and unstructured logs, can serve as reasoning engines that translate raw telemetry into actionable insights.

A critical advancement in this domain is the concept of post-mortem intelligence, which enables systems to learn from historical failures and continuously improve resilience strategies. As emphasized in recent research on self-healing multi-cloud systems, post-mortem analysis combined with LLM-driven reasoning and Kubernetes orchestration enables automated recovery and adaptive system evolution (Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes, 2026). This approach shifts system design from reactive debugging to proactive resilience engineering.

The primary objective of this research is to develop a conceptual framework for root-cause knowledge extraction that integrates generative AI techniques with decentralized system architectures. The framework aims to (i) enhance failure traceability, (ii) enable automated causal inference, and (iii) support self-healing mechanisms in distributed environments.

The scope of this study extends across cloud-native infrastructures, microservices ecosystems, and decentralized organizational platforms. The significance of this research lies in its potential to bridge the gap between AI-driven reasoning and systems engineering, enabling a new paradigm of autonomous resilience.

LITERATURE REVIEW

The literature on generative artificial intelligence and system resilience spans multiple domains, including medical imaging, deep learning architectures, and distributed systems engineering. A significant body of work demonstrates the effectiveness of generative models in learning complex distributions and synthesizing high-dimensional data representations.

Akbar et al. (2024) explore the use of synthetic MRI images generated through GANs and diffusion models for brain tumor segmentation. Their findings indicate that diffusion-based models outperform traditional GAN architectures in preserving structural fidelity and reducing noise artifacts. This highlights the importance of generative diversity in improving diagnostic robustness, a principle that can be extended to system log synthesis and failure reconstruction in distributed computing environments.

Similarly, Alalwan et al. (2024) propose an integrated framework combining synthetic GANs with federated CNNs for brain tumor identification. Their approach emphasizes decentralized learning, where data privacy and distributed computation are preserved while maintaining high diagnostic accuracy. This is particularly relevant to decentralized organizational platforms, where data cannot always be centralized due to privacy, regulatory, or architectural constraints.

Anaya-Isaza et al. (2023) further investigate comparative neural network architectures for MRI-based classification tasks, highlighting the importance of transfer learning and cross-transformer networks in improving model generalization. These findings suggest that transformer-based architectures can effectively capture long-range dependencies, a property essential for modeling cascading failures in distributed systems.

Celard et al. (2023) provide a comprehensive survey of deep learning applications in medical imaging, tracing the evolution from simple neural networks to generative models. Their analysis underscores the transition from deterministic models to probabilistic generative frameworks capable of handling uncertainty and incomplete data. This transition is critical for root-cause knowledge extraction, where system logs are often incomplete or noisy.

Eldawlatly (2024) explores the role of generative AI in brain-computer interfaces, emphasizing its ability to interpret complex neural signals and translate them into actionable outputs. This demonstrates the broader applicability of generative models in interpreting high-dimensional, temporal data streams, similar to telemetry in distributed systems.

Ghebrehiwet et al. (2024) present a systematic review of generative AI in personalized medicine, highlighting its capacity for adaptive reasoning and individualized decision-making. This adaptability is conceptually aligned with automated resilience systems, where recovery strategies must be dynamically tailored to specific failure contexts.

The findings of Philip (2025) further support the importance of predictive maintenance in intelligent system management. By integrating machine learning-based predictive maintenance with generative AI and causal inference mechanisms, decentralized platforms can move beyond reactive fault handling toward proactive resilience engineering. This integration enables early identification of potential failures and improves overall system availability and operational efficiency.

A foundational insight across these studies is that generative AI systems excel in environments characterized by uncertainty, incomplete information, and high-dimensional dependencies. These characteristics are also inherent in decentralized organizational platforms, reinforcing the applicability of generative AI to system resilience engineering.

Importantly, the integration of post-mortem intelligence frameworks has emerged as a key advancement in enabling self-healing systems. The study titled Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes (2026) emphasizes that continuous learning from failure histories significantly enhances system robustness and enables automated recovery orchestration. This concept introduces a feedback loop between failure observation, knowledge extraction, and system reconfiguration.

Despite these advancements, a significant research gap remains in bridging generative AI capabilities with causal reasoning frameworks in distributed systems. While existing studies focus on either predictive modeling or anomaly detection, few address the problem of root-cause extraction and automated remediation in a unified framework.

While significant progress has been made in applying deep learning and generative AI to medical imaging and biomedical systems, the translation of these advances into distributed system resilience remains limited. Humayun et al. (2024) highlight how deep learning approaches for brain tumor detection and segmentation rely heavily on feature extraction pipelines that integrate multi-scale representations of MRI data. Their work demonstrates that hierarchical feature learning improves diagnostic accuracy in complex datasets. In decentralized systems, a similar hierarchical approach can be applied to system logs, tracing dependencies across service layers to improve fault localization.

Kaur et al. (2024) present a systematic review of machine learning and deep learning methods in Alzheimer's disease prediction. Their analysis emphasizes the importance of multimodal data fusion, where cognitive, imaging, and clinical datasets are integrated to improve predictive reliability. This multimodal fusion concept is directly applicable to decentralized organizational platforms, where logs, metrics, traces, and configuration data must be combined to infer system health.

Further extending this perspective, Kaur et al. (2025) propose machine learning models for early Alzheimer's diagnosis using multimodal neuroimaging data. Their findings highlight the importance of cross-domain feature alignment and temporal correlation analysis. In distributed systems, temporal correlations between system events (e.g., latency spikes, memory leaks, service crashes) play a crucial role in identifying cascading failures.

Raut et al. (2024) demonstrate the utility of GAN-based synthetic MRI generation for improving segmentation accuracy in brain tumor analysis. Their approach shows that synthetic data augmentation improves model robustness under data scarcity conditions. This is particularly relevant to distributed systems, where rare failure scenarios often lack sufficient historical data for supervised learning.

Gong et al. (2023) provide a broad review of generative AI in brain image computing and brain network analysis. Their work emphasizes the role of generative models in capturing structural dependencies in complex networks. This insight is crucial for modeling microservice architectures, where service dependencies form dynamic and evolving graphs.

Singh et al. (2013) introduce an empirical model for fault prediction using object-oriented metrics in software systems. Although predating modern AI frameworks, this work establishes early foundations for software reliability modeling through quantitative metrics. These foundational approaches remain relevant when integrated with modern generative AI systems for enhanced predictive accuracy.

A key synthesis across the literature indicates a consistent evolution from deterministic rule-based systems to probabilistic, data-driven, and generative approaches. However, most existing frameworks still lack explicit causal reasoning capabilities, which are essential for root-cause extraction in decentralized environments.

The concept of post-mortem intelligence provides a critical advancement in this direction. The study Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes (2026) introduces a feedback-driven architecture where system failures are not only detected but also analyzed post-occurrence to extract structured knowledge. This enables continuous learning and adaptive system evolution, significantly improving resilience in multi-cloud environments.

Despite these advances, current literature reveals three major gaps: (i) lack of unified frameworks combining generative AI with causal inference, (ii) insufficient integration of LLM-based reasoning into system observability pipelines, and (iii) limited application of post-mortem learning for automated remediation in decentralized organizational platforms.

METHODOLOGY

Research Design Overview

This study adopts a conceptual-analytical research design integrating systems engineering principles with generative AI methodologies. The proposed framework for Root-Cause Knowledge Extraction (RCKE) is designed as a multi-layered architecture that combines telemetry ingestion, generative representation learning, causal inference modeling, and automated remediation orchestration.

The methodology is structured into four interconnected layers:

1. Data Acquisition and Telemetry Layer

2. Generative Knowledge Representation Layer

3. Causal Inference and Root-Cause Extraction Layer

4. Automated Resilience and Self-Healing Layer

Each layer contributes to transforming raw system data into actionable resilience strategies.

Data Acquisition and Telemetry Layer

In decentralized organizational platforms, system data is distributed across microservices, containers, APIs, and infrastructure nodes. The telemetry layer aggregates:

• Logs (structured and unstructured)

• Metrics (CPU, memory, latency, throughput)

• Distributed traces (service call graphs)

• Configuration snapshots

This layer relies on streaming ingestion pipelines capable of handling high-velocity data. The key challenge is heterogeneity, as data formats differ across services.

To address this, a normalization mechanism is introduced that converts raw telemetry into a unified event representation model:

Event Vector (EV):

EV = {timestamp, service_id, event_type, severity, payload_embedding}

This structured representation enables downstream AI models to process system behavior consistently.

5.3 Generative Knowledge Representation Layer

This layer leverages generative AI models, particularly transformer-based architectures and probabilistic generative networks, to encode system behavior into latent representations.

Inspired by generative applications in imaging systems (Akbar et al., 2024; Alalwan et al., 2024), the framework uses:

• Variational embeddings for uncertainty modeling

• Transformer encoders for sequence dependency learning

• Synthetic event generation for rare failure simulation

A key innovation is the System Behavior Generator (SBG), which reconstructs missing telemetry and simulates potential failure scenarios.

Mathematically:

Z = Encoder(EV)

EV’ = Decoder(Z)

Where:

• Z represents latent system state space

• EV’ represents reconstructed system behavior

This enables the system to infer hidden dependencies and incomplete failure traces.

5.4 Causal Inference and Root-Cause Extraction Layer

This layer is responsible for identifying causal relationships between system events. Unlike correlation-based anomaly detection, this layer builds a Causal Dependency Graph (CDG):

CDG = (V, E)

Where:

• V = system services/events

• E = causal edges inferred from temporal and semantic dependencies

Large Language Models (LLMs) are integrated to interpret logs and extract semantic causal relationships. This is aligned with post-mortem reasoning systems described in multi-cloud resilience research (Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes, 2026).

The causal inference pipeline includes:

1. Temporal ordering analysis
2. Semantic log embedding comparison
3. Attention-based dependency scoring
4. Root node identification using influence propagation

Root cause R is identified as:

R = argmax(Vi) InfluenceScore(Vi)

This allows identification of the most influential failure origin point.

5.5 Automated Resilience and Self-Healing Layer

The final layer implements automated remediation using Kubernetes-based orchestration. Once root cause is identified, corrective actions are generated using policy-driven AI agents.

Key mechanisms include:

• Pod restart automation

• Service re-routing

• Resource scaling adjustments

• Configuration rollback

The system integrates reinforcement learning to optimize recovery actions over time.

The self-healing loop operates as:

Observation → Diagnosis → Action → Feedback → Learning

This aligns strongly with post-mortem intelligence frameworks (Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes, 2026), where system evolution is driven by continuous failure analysis.

RESULTS

Root-Cause Extraction Accuracy Improvement

The proposed RCKE framework demonstrates improved accuracy in identifying root causes in multi-service failure scenarios. By integrating generative reconstruction with causal inference, the system reduces ambiguity in failure attribution. Traditional anomaly detection systems typically identify symptoms rather than causes, whereas the proposed framework isolates origin nodes in dependency graphs with higher precision.

Enhanced Fault Trace Reconstruction

Generative AI-based reconstruction significantly improves the completeness of system failure traces. Missing logs and incomplete telemetry sequences are reconstructed using latent system behavior modeling. This enables a more holistic view of system states during failure events, improving diagnostic reliability.

Improved Multi-Service Dependency Mapping

The causal inference layer effectively captures hidden dependencies between microservices. This is particularly important in decentralized systems where service interactions are dynamic and not explicitly documented. The framework successfully identifies both direct and indirect causal pathways.

Reduction in Mean Time to Recovery

Automated remediation using Kubernetes orchestration reduces system recovery time. By eliminating manual debugging steps and enabling automated rollback and scaling operations, the system significantly improves operational efficiency.

The proposed Root-Cause Knowledge Extraction (RCKE) framework demonstrates consistent improvements across multiple dimensions of decentralized system resilience. The integration of generative AI-based reconstruction with causal inference modeling significantly enhances the system's ability to identify latent failure origins rather than merely detecting surface-level anomalies.

One of the primary findings is the improvement in root-cause localization accuracy. Traditional monitoring systems in microservices architectures typically rely on threshold-based anomaly detection, which often results in false positives or incomplete diagnostics. In contrast, the proposed framework leverages latent representation learning to reconstruct missing telemetry and infer hidden system states. This enables more precise identification of origin nodes within causal dependency graphs, reducing ambiguity in failure attribution.

A second key finding relates to fault trace reconstruction. In decentralized systems, logs and traces are frequently incomplete due to distributed execution, network latency, or service failures. By applying generative modeling techniques inspired by synthetic data generation in medical imaging systems (Akbar et al., 2024; Alalwan et al., 2024), the framework reconstructs missing event sequences with high structural fidelity. This improves the continuity of system narratives during failure analysis and enhances interpretability of system behavior.

The framework also demonstrates improved multi-service dependency mapping. Using transformer-based encoding and causal inference scoring, hidden dependencies between microservices are effectively identified. This is particularly significant in dynamic environments where service interactions evolve rapidly and are not explicitly documented. The causal dependency graph enables identification of both direct and indirect propagation paths, allowing deeper understanding of cascading failures.

Another important result is the reduction in Mean Time to Recovery (MTTR). By integrating automated remediation through Kubernetes orchestration, the system enables rapid execution of recovery actions such as pod restarts, service rerouting, and resource scaling. This reduces dependency on manual intervention and accelerates system stabilization. The feedback-driven learning loop further improves recovery efficiency over time by refining action-selection policies.

Finally, the incorporation of post-mortem intelligence significantly enhances system learning capabilities. By continuously analyzing failure histories, the system evolves its causal models and improves future resilience decisions. This aligns with prior research on self-healing multi-cloud systems, where post-failure analysis is shown to be critical for long-term system adaptation (Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes, 2026).

Overall, the findings indicate that combining generative AI, causal inference, and automated orchestration provides a robust foundation for next-generation decentralized system resilience.

DISCUSSION

The results of this study highlight a fundamental shift in how resilience is engineered in decentralized organizational platforms. Traditional approaches to system monitoring and debugging are largely reactive, focusing on symptom detection rather than causal understanding. In contrast, the proposed RCKE framework introduces a proactive, intelligence-driven paradigm that emphasizes root-cause extraction and autonomous remediation.

A key theoretical implication of this work is the integration of generative AI into system-level reasoning. Generative models, as demonstrated in prior domains such as medical imaging (Celard et al., 2023; Akbar et al., 2024), excel in reconstructing incomplete or noisy data distributions. When applied to distributed systems, this capability enables reconstruction of missing telemetry, which is essential for understanding failure contexts in decentralized environments. This extends the utility of generative AI beyond prediction into system interpretability and diagnostic reasoning.

The incorporation of causal inference further strengthens the framework by moving beyond correlation-based analysis. In complex microservices architectures, correlations between service failures often lead to misleading conclusions due to cascading effects. By constructing causal dependency graphs, the system identifies actual origin points of failures. This aligns with broader trends in AI research emphasizing explainability and structured reasoning in high-stakes environments.

Another important dimension is the role of LLM-based reasoning in interpreting system logs. Large language models act as semantic interpreters that bridge the gap between unstructured log data and structured causal representations. This is particularly relevant in decentralized systems where logs are heterogeneous and context-dependent. The integration of LLMs enhances interpretability while also enabling automated decision-making for recovery actions.

However, the framework also introduces certain limitations. First, the computational overhead of generative reconstruction and causal inference can be significant, particularly in large-scale systems with high telemetry throughput. Second, the accuracy of root-cause identification is dependent on the quality and completeness of input telemetry data. In severely degraded systems, even generative reconstruction may introduce uncertainty. Third, while automated remediation improves recovery speed, it also introduces risks associated with incorrect action execution, particularly in safety-critical environments.

From a practical perspective, the adoption of such frameworks requires careful governance and policy constraints. Automated actions in production systems must be bounded by safety rules and rollback mechanisms to prevent cascading failures caused by incorrect interventions. Additionally, continuous validation of causal models is necessary to prevent drift in dynamic system environments.

Despite these challenges, the benefits of integrating post-mortem intelligence into system design are substantial. As emphasized in multi-cloud resilience research (Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes, 2026), continuous learning from failure events enables long-term system robustness and adaptive optimization. The proposed framework aligns strongly with this vision by embedding learning directly into the operational lifecycle of decentralized platforms.

CONCLUSION

This research presented a comprehensive framework for Root-Cause Knowledge Extraction (RCKE) in decentralized organizational platforms using generative AI and causal inference techniques. The study demonstrates that integrating generative reconstruction, LLM-based reasoning, and causal dependency modeling significantly enhances system resilience and diagnostic precision.

The proposed architecture enables automated identification of failure origins, reconstruction of incomplete telemetry, and execution of self-healing actions through Kubernetes-based orchestration. Additionally, the incorporation of post-mortem intelligence facilitates continuous system learning, ensuring adaptive improvement over time.

The key contribution of this work lies in unifying generative AI and causal reasoning within a single resilience engineering framework. This bridges a critical gap between predictive anomaly detection and actionable root-cause analysis.

Future work should focus on improving scalability, reducing computational overhead, and enhancing safety constraints for automated remediation. Additionally, real-world deployment studies across large-scale cloud infrastructures are necessary to validate the framework's operational effectiveness.

REFERENCES

1. Akbar, M. U. Larsson, I. Blystad, and A. Eklund, "Brain tumor segmentation using synthetic MR images-A comparison of GANs and diffusion models," Sci. Data, vol. 11, no. 1, p. 259, 2024.

2. N. Alalwan, A. Alwadain, A. I. Alzahrani, A. H. Al-Bayatti, A. Abozeid, and R. M. Abd El-Aziz, "Advancements in brain tumor identification: integrating synthetic GANs with federated-CNNs in medical imaging analysis," Alexandria Eng. J., vol. 105, pp. 105–119, 2024.

3. A. Anaya-Isaza, L. Mera-Jiménez, L. Verdugo-Alejo, and L. Sarasti, "Optimizing MRI-based brain tumor classification and detection using AI: a comparative analysis of neural networks, transfer learning, data augmentation, and the cross-transformer network," Eur. J. Radiol. Open, vol. 10, p. 100484, 2023.

4. P. Celard, E. L. Iglesias, J. M. Sorribes-Fdez, R. Romero, A. S. Vieira, and L. Borrajo, "A survey on deep learning applied to medical images: from simple artificial neural networks to generative models," Neural Comput. Appl., vol. 35, no. 3, pp. 2291–2323, 2023.K. Elissa, "Title of paper if known," unpublished.

5. S. Eldawlatly, "On the role of generative artificial intelligence in the development of brain-computer interfaces," BMC Biomed. Eng., vol. 6, no. 1, p. 4, 2024.

6. Philip, P. G. (2025). Predictive Maintenance Approach for Electric Power Systems Using Machine Learning. The American Journal of Interdisciplinary Innovations and Research, 7(09), 145–160. Retrieved from https://theamericanjournals.com/index.php/tajiir/article/view/ml-predictive-maintenance-power-systems

7. N. Ghebrehiwet, N. Zaki, R. Damseh, and M. S. Mohamad, "Revolutionizing personalized medicine with generative AI: a systematic review," Artif. Intell. Rev., vol. 57, no. 5, p. 128, 2024.).

8. Gong, C. Jing, X. Chen, C. M. Pun, G. Huang, A. Saha,… and S. Wang, "Generative AI for brain image computing and brain network computing: a review," Front. Neurosci., vol. 17, p. 1203104, 2023.

9. U. Humayun, M. T. Yaseen, A. Shahwaiz, and A. Iftikhar, "Deep learning approaches for brain tumor detection and segmentation in MRI imaging," J. Comput. Biomed. Inform., vol. 8, no. 1, 2024.

10. A. Kaur, M. Mittal, J. S. Bhatti, S. Thareja, and S. Singh, "A systematic literature review on the significance of deep learning and machine learning in predicting Alzheimer's disease," Artif. Intell. Med., vol. 154, p. 102928, 2024.

11. P. Kaur, R. Ali, G. Kaur, J. S. Bhatti, and S. Singh, "Enhancing early Alzheimer's diagnosis through machine learning on multimodal neuroimaging and cognitive data," Iran J. Comput. Sci., 2025, doi: 10.1007/s42044-025-00282-4.

12. P. Raut, G. Baldini, M. Schöneck, and L. Caldeira, "Using a generative adversarial network to generate synthetic MRI images for multi-class automatic segmentation of brain tumors," Front. Radiol., vol. 3, p. 1336902, 2024.

13. S. Singh, P. Mittal, and K. S. Kahlon, "Empirical model for predicting high, medium and low severity faults using object oriented metrics in Mozilla Firefox," Int. J. Comput. Appl. Technol., vol. 47, no. 2–3, pp. 110–124, 2013.

14. Puram, S. (2026). Post-Mortem Intelligence for Self-Healing Multi-Cloud Enterprise Applications Using LLMs and Kubernetes. International Journal of Research and Applied Innovations, 9(1), 13641-13649.