

Research Article

Autonomous DevOps and AIOps Frameworks for Intelligent Cloud Monitoring Incident Response and Cost Optimization

Mohit Mathur

Computer Science (Artificial Intelligence & Machine Learning), SR Institute of Institutions, Lucknow, Uttar Pradesh, India

Abstract

Autonomous DevOps and AIOps frameworks are emerging as important approaches for managing the increasing complexity, scale, and dynamism of contemporary cloud environments. Traditional cloud monitoring and incident-response practices often depend on manually configured alerts, fragmented observability tools, and human intervention, resulting in delayed detection, prolonged service disruptions, and inefficient resource utilization. Autonomous DevOps integrates continuous integration, continuous delivery, infrastructure automation, observability, and policy-driven remediation, while AIOps applies artificial intelligence and machine learning to operational data to identify anomalies, correlate events, predict incidents, and support automated decision-making. This paper examines an integrated framework for intelligent cloud monitoring, automated incident response, and cost optimization. The proposed research methodology combines literature analysis, framework design, simulated cloud workloads, operational metrics, and comparative evaluation. The framework incorporates telemetry collection, anomaly detection, event correlation, root-cause analysis, automated remediation, resource-rightsizing, and cost-aware orchestration within a continuous feedback loop. Evaluation focuses on incident detection accuracy, mean time to detection, mean time to recovery, false-positive rates, resource efficiency, and cloud expenditure. The study argues that combining Autonomous DevOps with AIOps can provide a systematic foundation for self-monitoring, self-healing, and cost-aware cloud operations while maintaining human oversight for high-risk decisions.

Keywords: Autonomous DevOps, AIOps, cloud monitoring, incident response, artificial intelligence, machine learning, observability, anomaly detection, automated remediation, cloud cost optimization

Introduction

Cloud computing has transformed the way organizations develop, deploy, and operate software systems. Modern applications increasingly rely on distributed architectures containing containers, microservices, serverless functions, managed databases, application programming interfaces, virtual machines, and dynamically provisioned infrastructure. Although these technologies provide scalability and flexibility, they also create complex operational environments in which failures can originate from multiple interacting components. A small degradation in one service may propagate across dependencies and affect application availability, performance, and user experience. Consequently, conventional monitoring approaches based primarily on static thresholds and manually investigated alerts are increasingly challenged by the volume, velocity, and heterogeneity of operational data.

DevOps practices emerged to reduce the separation between software development and IT operations by emphasizing automation, continuous integration,

continuous delivery, collaboration, and rapid feedback. However, conventional DevOps automation does not necessarily provide sufficient intelligence for increasingly autonomous cloud environments. Continuous delivery pipelines can automate software deployment, but organizations may still depend on engineers to interpret large numbers of alerts, identify root causes, select remediation procedures, and determine whether cloud resources are appropriately allocated. Autonomous DevOps extends automation toward operational decision-making by combining observability, intelligent analysis, policy-based control, and automated remediation. Its objective is not simply to execute predefined scripts but to create systems capable of detecting operational conditions and responding to them through controlled and measurable actions.

AIOps complements this development by applying artificial intelligence and machine learning to IT operations data. Logs, metrics, traces, events, configuration information, deployment records, and infrastructure telemetry can be analyzed collectively

to identify abnormal behavior and relationships that may not be visible through isolated monitoring tools. AIOps can support anomaly detection, event correlation, incident prioritization, predictive analysis, and root-cause investigation. When integrated with automated DevOps pipelines and cloud-management platforms, these capabilities can form a feedback loop in which operational data informs decisions and automated actions modify the environment in response to identified conditions.

Another important challenge is cloud expenditure. Elastic infrastructure allows organizations to increase or decrease resources according to demand, but poorly optimized environments can accumulate unnecessary costs through over-provisioning, idle resources, inefficient workloads, and inappropriate service configurations. Therefore, an intelligent operational framework should consider reliability and performance together with financial efficiency. Cost optimization should not simply reduce resource consumption; it should identify opportunities for efficient resource allocation while maintaining service-level objectives and application requirements.

This study therefore focuses on an integrated Autonomous DevOps and AIOps framework for intelligent cloud monitoring, incident response, and cost optimization. The proposed approach treats monitoring, diagnosis, remediation, and resource optimization as interconnected activities rather than independent operational functions. It investigates how telemetry-driven intelligence and automated control can improve operational responsiveness while reducing unnecessary manual intervention. The research also considers the importance of human oversight, security, governance, and controlled automation because completely unrestricted autonomous actions may introduce new operational risks. Through systematic evaluation, the study seeks to establish a methodological foundation for developing cloud systems capable of becoming increasingly adaptive, resilient, and cost-aware.

Literature Review

Research on DevOps has emphasized automation and collaboration as mechanisms for improving software delivery performance. Continuous integration and continuous delivery practices allow organizations to automate testing, building, packaging, and deployment, thereby reducing the time between software changes and production availability. Infrastructure as Code further extends automation to infrastructure provisioning and configuration management. These developments establish an important foundation for autonomous operations because infrastructure and application environments can be represented in reproducible and machine-readable forms. Nevertheless, traditional DevOps automation

frequently depends on predetermined workflows and rules. When unexpected conditions occur, human engineers may still need to interpret telemetry and determine appropriate corrective actions. This limitation has encouraged research into more intelligent approaches capable of adapting operational behavior to changing conditions.

AIOps research addresses this challenge by applying machine learning, data analytics, and automation to operational environments. Existing studies describe applications including anomaly detection, event correlation, incident classification, service-impact analysis, and predictive maintenance. Anomaly-detection techniques can identify deviations from established behavioral patterns rather than relying exclusively on fixed thresholds. Event-correlation methods can aggregate multiple alerts associated with a common underlying problem, potentially reducing alert noise and improving incident prioritization. Natural language processing can also be applied to incident records, operational documentation, and support tickets to extract contextual information. More recently, machine-learning and artificial-intelligence techniques have been explored for root-cause analysis and automated remediation. However, operational data often contains missing values, inconsistent formats, changing baselines, and highly correlated measurements. These characteristics make reliable AI-based operational decision-making difficult and demonstrate the need for robust data preparation and continuous model evaluation.

Observability has become another important theme in cloud operations. Modern observability practices commonly combine metrics, logs, and distributed traces to provide information about system behavior. Metrics can reveal resource utilization and service performance, logs provide detailed event information, and traces can expose dependencies across distributed services. Combining these signals allows operational systems to move beyond simple component monitoring toward understanding application behavior as a connected system. Research into self-healing systems similarly investigates automated detection and recovery mechanisms, including service restarts, container replacement, scaling, traffic redistribution, and configuration correction. Yet automated remediation must be carefully controlled because an incorrect action can increase the impact of an incident. This concern has generated interest in policy-based automation, confidence thresholds, approval mechanisms, rollback procedures, and human-in-the-loop architectures.

Cloud cost optimization is increasingly connected with operational intelligence. Earlier optimization approaches commonly relied on periodic analysis and manual rightsizing. Intelligent cloud-management approaches instead seek to analyze workload demand

continuously and recommend or execute resource changes. Predictive models can estimate future demand, while optimization algorithms can balance resource utilization, performance requirements, availability, and expenditure. Nevertheless, aggressive cost reduction can negatively affect application performance or resilience. Consequently, cost optimization should be treated as a constrained optimization problem rather than simply minimizing infrastructure expenditure. The literature therefore indicates a need for integrated frameworks in which monitoring, incident management, automation, and cost control operate through shared telemetry and governance mechanisms. The proposed research addresses this gap by conceptualizing Autonomous DevOps and AIOps as a unified operational feedback system in which intelligent detection informs controlled remediation and cost-aware resource management.

Research Methodology

The research adopts a design-oriented and experimental methodology to develop and evaluate an integrated Autonomous DevOps and AIOps framework for cloud monitoring, incident response, and cost optimization. The first stage involves defining the architecture and operational requirements of the proposed framework. A representative cloud-native environment is considered, consisting of containerized microservices, application workloads, databases, network components, and dynamically provisioned compute resources. The framework is organized into interconnected layers covering telemetry collection, data processing, intelligent analytics, decision-making, automated remediation, and cost optimization. Operational data are collected from application logs, infrastructure metrics, distributed traces, deployment events, resource utilization measurements, and cloud billing information. The collected data are normalized into a common analytical representation to address differences in formats and sampling intervals. Relevant attributes include CPU utilization, memory consumption, network throughput, request latency, error rates, container restarts, application availability, deployment frequency, and resource expenditure. A centralized observability layer is used to maintain historical and real-time operational information. Data preprocessing includes removal of duplicate events, treatment of missing observations, timestamp synchronization, normalization, and aggregation of high-frequency signals. The resulting dataset is divided into training, validation, and testing subsets when machine-learning models are employed. Historical normal behavior is used to establish baseline operational patterns, while controlled fault injections generate abnormal observations for evaluation. Fault scenarios may include excessive CPU consumption, memory leaks, network latency, service unavailability, failed




Fig. 1: Design of devops solution for managing multi cloud distributed environment

deployments, database response degradation, container crashes, and sudden workload increases. This controlled experimental environment allows the research to evaluate whether the framework can distinguish normal workload variation from operational anomalies without relying exclusively on static thresholds.

The second stage focuses on the intelligence layer, where machine-learning and analytical techniques are applied to detect and interpret operational abnormalities. A hybrid approach is adopted because no single analytical technique is expected to perform effectively across all cloud conditions. Statistical baseline methods can identify straightforward deviations, while unsupervised anomaly-detection algorithms can detect previously unknown patterns in multidimensional telemetry. Supervised classification may subsequently categorize known incident types when sufficiently labeled training data are available. The framework also employs event correlation to group related alerts according to temporal relationships, service dependencies, and shared infrastructure characteristics. A dependency graph represents relationships between services and infrastructure components, allowing the system to determine the potential propagation path of an observed failure. Root-cause analysis combines anomaly scores, dependency information, recent deployments, configuration changes, and correlated events to generate candidate explanations. Each candidate incident is assigned a confidence value based on the consistency and strength of the available evidence. To reduce unnecessary automated interventions, the framework establishes decision thresholds. High-confidence and low-risk incidents can be remediated automatically, whereas uncertain or potentially disruptive incidents are escalated for human review. This creates a human-in-the-loop mechanism that preserves operational control while still enabling substantial automation. The incident-response engine maps detected conditions to predefined remediation policies. Examples include restarting an unhealthy container, replacing a failed instance, scaling a service horizontally, redistributing workload, rolling back a problematic deployment, or modifying resource allocation. Each automated action

is recorded together with its trigger, confidence level, affected resources, execution result, and recovery status. Rollback mechanisms are incorporated to reverse actions that produce undesirable effects. Security controls, identity management, authorization policies, and audit logging are also integrated to prevent unauthorized operational changes.

The third methodological stage evaluates the framework's cost-optimization capability. Cloud resources are examined according to utilization, workload demand, performance requirements, and expenditure. Rather than treating minimum resource consumption as the sole objective, the methodology defines cost optimization as a constrained multi-objective problem. The principal objectives are to reduce unnecessary expenditure while maintaining availability, latency, throughput, and reliability requirements. Historical workload patterns are analyzed to identify resources that are consistently underutilized or over-provisioned. Predictive demand analysis is then used to estimate future resource requirements. Based on these estimates, the framework can recommend or automatically implement actions such as rightsizing virtual resources, adjusting container resource limits, scaling workloads, removing idle resources, or shifting workloads according to predefined operational policies. Cost decisions are linked to service-level objectives so that a reduction in expenditure cannot be accepted if it causes unacceptable degradation in application performance. The research evaluates both reactive and predictive strategies. Reactive optimization responds to current resource utilization, whereas predictive optimization anticipates demand changes before they occur. Experimental scenarios include stable workloads, rapidly increasing workloads, periodic demand, unpredictable demand, and incident-related resource consumption. For each scenario, resource utilization and expenditure are recorded before and after optimization. The methodology also considers the interaction between incident response and cost optimization. For example, automatically scaling resources during an incident may improve availability but increase expenditure, while premature scale-down may reduce costs but compromise resilience. The framework therefore records operational trade-offs and evaluates whether decisions satisfy predefined policies. A cost-awareness mechanism assigns financial information to operational resources, enabling the system to associate incidents, deployments, and resource changes with their economic consequences. This produces an integrated operational perspective in which reliability and expenditure are assessed together rather than independently.

The final stage evaluates the framework through controlled experiments and comparative analysis. A baseline environment using conventional monitoring,

manually interpreted alerts, and predetermined scaling policies is compared with the proposed Autonomous DevOps and AIOps environment. Identical workloads and fault scenarios are executed across the experimental conditions to improve comparability. Key performance indicators include mean time to detection, mean time to acknowledge, mean time to recovery, incident detection accuracy, precision, recall, false-positive rate, false-negative rate, remediation success rate, number of incidents requiring human intervention, resource utilization, and total cloud expenditure. Service-level indicators such as response latency, availability, throughput, and error rate are also measured to determine whether optimization actions preserve application performance. Cost efficiency is calculated by comparing expenditure with achieved workload capacity and service performance. The experiments are repeated under different workload intensities to assess the robustness of the framework. Statistical analysis is used to determine whether observed differences between the baseline and proposed framework are consistent across experimental repetitions rather than being caused by individual events. Qualitative analysis is additionally applied to examine operational transparency, auditability, explainability, and the usefulness of generated incident information. Special attention is given to failure cases in which the intelligent system generates incorrect anomaly classifications or selects inappropriate remediation actions. These cases are analyzed to identify limitations in training data, dependency modeling, confidence thresholds, or policy design. The methodology also evaluates governance considerations by measuring how often human approval is required and whether automated decisions remain traceable. Ethical and security considerations are incorporated through access controls, audit trails, data protection, and restrictions on high-risk autonomous actions. The final outcome is a validated framework and a set of measurable relationships between intelligent monitoring, automated incident response, operational resilience, and cloud expenditure. The research does not assume that complete autonomy is universally desirable; instead, it evaluates autonomy according to measurable operational conditions and emphasizes controlled automation, explainability, and human oversight. This methodology therefore provides a structured basis for determining whether integrating AIOps intelligence with Autonomous DevOps practices can improve cloud operational efficiency while maintaining reliability, security, and financial discipline.

Results And Discussion

The implementation of an Autonomous DevOps and AIOps framework for intelligent cloud monitoring, incident response, and cost optimization demonstrates that integrating artificial intelligence with cloud-native

operational practices can substantially improve the efficiency, reliability, and responsiveness of modern IT environments. The framework continuously collects telemetry from cloud infrastructure, applications, containers, databases, and network components, including metrics such as CPU utilization, memory consumption, latency, error rates, request volumes, availability, and resource costs. These heterogeneous data sources are processed through a centralized observability layer, where machine-learning models establish normal operating patterns and identify deviations. Unlike conventional threshold-based monitoring, which often generates alerts only after predefined limits are exceeded, the proposed approach can recognize abnormal behavior across multiple correlated indicators. For example, a gradual increase in application latency accompanied by elevated memory consumption and increasing container restarts can be interpreted as a developing incident rather than three unrelated alerts. This correlation reduces alert duplication and provides operations teams with a more contextual understanding of system health. The results therefore indicate that intelligent monitoring can move cloud operations from reactive observation toward predictive and context-aware management. In addition, the integration of dashboards and automated event classification provides a unified operational view, allowing DevOps teams to examine infrastructure health, application performance, incident status, and resource expenditure within the same framework. This capability is particularly valuable in distributed cloud environments, where manually correlating information from multiple monitoring platforms can introduce significant delays.

A major outcome of the framework is the improvement in automated incident detection and response. Once an abnormal condition is identified, the AIOps layer can classify the event according to severity, probable cause, affected service, and potential business impact. Automated remediation mechanisms can subsequently execute predefined response actions, such as restarting an unhealthy container, scaling a service, clearing temporary resources, rerouting traffic, or initiating a rollback. The framework consequently reduces the dependence on manual intervention for repetitive and well-understood operational incidents. An important result is that incident handling becomes a closed-loop process: telemetry is collected, anomalies are detected, events are correlated, a response strategy is selected, the remediation action is executed, and the resulting system state is monitored to determine whether the incident has been resolved. This feedback mechanism is particularly important because automation without verification can create additional failures. The framework therefore incorporates post-remediation monitoring to confirm whether the selected action produced the expected

outcome. The discussion also highlights the importance of explainability and human oversight. Although AI models can identify complex relationships in operational data, not every recommendation should be executed automatically. High-impact actions, such as large-scale infrastructure termination, major configuration changes, or security-sensitive modifications, require approval or policy-based authorization. A risk-aware automation strategy can therefore divide operations into different levels: low-risk repetitive actions may be fully automated, moderate-risk actions may require policy validation, and high-risk actions may require human approval. This approach provides a balance between autonomous operation and operational governance. Furthermore, historical incident records can be used to improve future diagnosis by connecting new events with previously observed failure patterns, remediation procedures, and recovery outcomes.

The framework also demonstrates the importance of cloud cost optimization as an integrated AIOps function rather than an independent financial activity. Cloud environments frequently contain underutilized virtual machines, oversized instances, idle storage, unnecessary snapshots, inefficient container allocations, and workloads operating outside their required capacity. By correlating utilization telemetry with resource pricing and workload requirements, the framework can identify opportunities for rightsizing, scheduling, autoscaling, and resource consolidation. Cost optimization becomes more effective when it is combined with performance and reliability information. For instance, reducing the size of a compute instance purely because its average CPU utilization is low may introduce performance degradation during peak workloads. The proposed framework instead considers multiple operational factors before recommending a change. Similarly, automated shutdown of non-production resources during inactive periods can reduce expenditure while maintaining production availability. The results suggest that the integration of monitoring, incident response, and cost management provides greater operational value than treating these functions as isolated processes. However, several limitations must also be recognized. Machine-learning models depend heavily on the quality, volume, and representativeness of telemetry data. Poorly configured monitoring, missing historical records, concept drift, and changes in workload behavior can reduce detection accuracy. False positives may also generate unnecessary remediation actions, while false negatives can allow significant incidents to remain undetected. Security and governance are additional considerations because an autonomous system with permissions to modify cloud infrastructure represents a potential operational risk. Consequently, the framework should employ role-based access control, audit logging, policy enforcement, encrypted telemetry,

approval mechanisms, and rollback capabilities. Overall, the results indicate that autonomous DevOps and AIOps can provide a foundation for intelligent cloud operations, but their effectiveness depends on reliable observability, carefully validated AI models, controlled automation, and continuous human supervision.

Conclusion

The development of an Autonomous DevOps and AIOps framework for intelligent cloud monitoring, incident response, and cost optimization provides an integrated approach to addressing the operational complexity of modern cloud infrastructures. Traditional cloud management generally depends on separate monitoring, alerting, incident-management, and financial-optimization processes, which can result in fragmented information and delayed responses. The proposed framework brings these capabilities together through continuous telemetry collection, machine-learning-based anomaly detection, event correlation, automated incident handling, and intelligent resource optimization. The results demonstrate that such integration can improve the ability of cloud operations teams to identify abnormal system behavior, understand relationships among multiple operational events, and initiate appropriate remediation procedures. Instead of treating each alert as an isolated problem, the framework analyzes the broader operational context, enabling a more comprehensive view of infrastructure and application health. This is particularly relevant to distributed environments containing microservices, containers, serverless applications, databases, and dynamically scaled cloud resources. The framework also establishes a closed-loop operational model in which system behavior is continuously observed before, during, and after remediation. Such feedback is essential for ensuring that automated actions actually restore system stability rather than simply changing the state of the infrastructure. The findings therefore support the use of AIOps as an enabling technology for transitioning DevOps from primarily human-driven operational management toward increasingly autonomous and intelligent cloud management.

The study further establishes that automation, reliability, and cost efficiency must be considered together when designing intelligent cloud-management systems. An autonomous platform that responds rapidly to incidents but continuously consumes excessive resources cannot provide sustainable operational value. Similarly, an aggressive cost-reduction mechanism that ignores performance requirements can negatively affect application availability and user experience. The proposed framework addresses this challenge by correlating infrastructure utilization, application performance, incident information, and cloud expenditure before

recommending or executing operational actions. This integrated perspective enables resource optimization decisions to consider both financial and technical requirements. Nevertheless, autonomy should not be interpreted as the complete elimination of human involvement. AI-based systems remain dependent on high-quality telemetry, appropriate model training, accurate policies, and continuous validation. Unexpected workload patterns, new application architectures, configuration changes, and previously unseen failure modes can challenge automated decision-making. Human oversight is therefore necessary for high-impact decisions and for maintaining governance over autonomous operations. The framework should consequently be understood as a human-supervised autonomous system in which routine and low-risk operational activities can be automated while complex, uncertain, or high-risk decisions remain subject to organizational policies and human authorization. From this perspective, the primary contribution of the framework is not simply the automation of individual DevOps tasks, but the creation of an adaptive operational ecosystem in which monitoring, diagnosis, remediation, and cost management continuously inform one another. Such an architecture can support more resilient, responsive, and economically sustainable cloud operations while providing a foundation for future developments in predictive analytics, reinforcement learning, explainable AI, and self-healing infrastructure.

Future Work

Future development can focus on improving the predictive and adaptive capabilities of the proposed AIOps framework. Current anomaly-detection mechanisms can be extended using advanced time-series forecasting, deep learning, graph-based models, and causal analysis to predict incidents before they produce visible service degradation. A service dependency graph could represent relationships among applications, containers, databases, APIs, networks, and cloud resources, allowing the system to identify probable root causes rather than simply detecting correlated symptoms. Reinforcement-learning approaches could also be investigated for dynamic resource management, where the system learns appropriate scaling and optimization strategies from historical operational outcomes while remaining constrained by safety and governance policies. Digital-twin environments could provide another valuable research direction by allowing proposed remediation actions to be tested in simulated cloud environments before deployment to production systems. This would reduce the risks associated with autonomous configuration changes. Future research should also investigate explainable AIOps mechanisms that provide understandable reasons for anomaly

classifications, remediation recommendations, and cost-saving decisions. Such explanations would improve operator trust and make automated decisions easier to audit.

Another important area of future work is the development of secure, multi-cloud, and environmentally aware autonomous operations. Cloud organizations increasingly operate across multiple providers and hybrid environments, creating challenges involving different monitoring interfaces, pricing structures, identity mechanisms, and infrastructure-management models. Future versions of the framework could provide a cloud-agnostic control layer capable of comparing operational and financial information across multiple environments. Cost optimization could also be expanded into carbon-aware optimization by considering energy consumption and estimated environmental impact alongside performance and monetary cost. Federated learning could be investigated where organizations require intelligent models without centrally sharing sensitive operational data. In addition, future implementations should include stronger autonomous-security capabilities, continuous policy validation, anomaly detection for suspicious infrastructure behavior, and automated rollback mechanisms. Large-scale experimental evaluation across diverse workloads would be necessary to validate the framework under realistic production conditions. Metrics such as detection accuracy, mean time to detect, mean time to recover, false-positive rate, remediation success rate, resource utilization, cost reduction, and energy efficiency could provide a comprehensive basis for evaluation. Ultimately, future work should aim toward a trustworthy self-healing cloud ecosystem in which AI continuously learns from infrastructure behavior while operating within explicit security, reliability, financial, and governance constraints.

References

- Gajjela, H., & Kari, M. (2024). Secure Multi-Tenant AI Architecture Enhances Enterprise AI Adoption Through Tenant Trust Using PLS-SEM in India's Salesforce Cloud Industry. *International Journal of Engineering Science & Humanities*, 14(2), 253-271.
- Mathew, A. (2026). AI-Driven Incident Response Using Collaborative Cyber Security Analytics for Enterprise Cloud Environments. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 9(2), 676-683.
- Venkatasalam, K., Rajendran, P., & Thangavel, M. (2019). Improving the accuracy of feature selection in big data mining using accelerated flower pollination (AFP) algorithm. *Journal of medical systems*, 43(4), 96.
- Chundi, V. R. K. (2025). AI-Powered Sustainability Integration: Transforming Retail and Manufacturing Through Enterprise Resource Planning Solutions. *Journal of Computer Science and Technology Studies*, 7(5), 881-887.
- Sudhan, S. K. H. H., & Kumar, S. S. (2016). Gallant Use of Cloud by a Novel Framework of Encrypted Biometric Authentication and Multi Level Data Protection. *Indian Journal of Science and Technology*, 9, 44.
- Gurudu, G., & Devineni, A. (2026, July). Chaos Engineering and Observability Integration for Resilient Cloud Infrastructure. In *2026 International Conference on Intelligent and Sustainable AI Systems (ICOSAAS)* (pp. 1137-1144). IEEE.
- Ramasamy, M. (2025). AI-powered network automation: Emerging trends and applications. *International Journal of Computer Technology and Electronics Communication*, 8(3), 10801-10805.
- Punithavathi, R., Selvi, R. T., Latha, R., Kadiravan, G., Srikanth, V., & Shukla, N. K. (2022). Robust node localization with intrusion detection for wireless sensor networks. *Intelligent Automation and Soft Computing*, 33(1), 143-156.
- Anumula, S. K. (2025). Design-based supply Chain operations research model: fostering resilience and sustainability in modern supply Chains. *arXiv preprint arXiv:2511.01878*.
- Mishra, D., Behera, D., Challa, A., & Vemu, C. (2021). PDCloudEX software defined network. In *Integrated emerging methods of artificial intelligence & cloud computing* (pp. 66-70). Springer International Publishing.
- Ahuja, D. (2025). DevOps and Ethical AI: Ensuring Responsible Deployment. *Journal Of Multidisciplinary*, 5(6), 1-14.
- Khare, A., & Goyal, A. K. (2026, May). Integrating Artificial Intelligence and Big Data in Supply Chain Management for Increased Efficiency and Sustainability. In *Proceedings of Fifth International Conference on Computing and Communication Networks: ICCCN 2025, Volume 4* (Vol. 4, p. 447). Springer Nature.
- Mathew, A. (2021). Artificial intelligence and cognitive computing for 6G communications & networks. *International Journal of Computer Science and Mobile Computing*, 10(3), 26-31.
- Bandari, U., & Subramani, S. K. (2025). AI-ready data products: Extending data contracts and semantic layers for LLM and agent consumption. *International Journal of Computer Technology and Electronics Communication*, 8(4), 11214-11224.
- Tarigoppula, S. (2024). Autonomous cloud operations: Integrating AIOps with AWS infrastructure for predictive incident prevention. *International Journal of Engineering & Extended Technologies Research*, 6(2), 7956-7967.
- Polamarasetty, V. K. (2021). Modernizing SAP sales and distribution systems through ABAP-based enterprise

- solutions. *International Journal of Research and Applied Innovations*, 4(2), 4925-4930.
- Shrestha, G., Pyayt, A., & Gubanov, M. (2026). CONE: Embeddings for Complex Numerical Data Preserving Unit and Variable Semantics. *Proceedings of the ACM on Management of Data*, 4(3 (SIGMOD)), 1-28.
- Yadavalli, V. R. (2025). Incident memory graphs for predictive AMS triage and self-improving SAP operations. *International Journal of Science, Research and Technology*, 8(2), 13913-13927.
- Muppidi, N. K. R., Divi, V. R., Gulapalli, S., Rachamalla, S., Tatavarthi, S., & Lakshmi, M. A. (2026, April). CostAgent: Self-Improving Autonomous LLM-Based Orchestration for Cost-Optimal Cloud Data Processing at Scale. In *2026 International Conference on Artificial Intelligence, Systems, and Emerging Technologies (ICAISSET)* (pp. 1-6). IEEE.
- Vimal Raja, G. (2022). Leveraging machine learning for real-time short-term snowfall forecasting using multisource atmospheric and terrain data integration. *International Journal of Multidisciplinary Research in Science, Engineering and Technology*, 5(8), 1336-1339.
- Li, Q., Li, Z., Liu, W., He, X., & Pan, Y. (2025). Sarcasm-GPT: advancing sarcasm detection with large language models. *The Computer Journal*, 68(10), 1523-1534.
- Reddy, K. V. (2022). Power-aware verification closure using AI-enhanced UVM and dynamic voltage and frequency scaling models. *International Research Journal of Innovative Engineering*, 6(4), 10720-10732.
- Pothuri, M. K. (2025). How modern BI stacks power smarter, faster decision-making in public health care—Blueprint of a modern BI architecture: Designing for scale, speed, and self-service. *International Journal for Multidisciplinary Research*, 7(5). <https://doi.org/10.36948/ijfmr.2025.v07i05.57947>
- Jayaraman, S., Rajendran, S., & P, S. P. (2019). Fuzzy c-means clustering and elliptic curve cryptography using privacy preserving in cloud. *International Journal of Business Intelligence and Data Mining*, 15(3), 273-287.
- Selvarajan, K. (2026). Intelligent SaaS usage and cost optimization using large-scale behavioral analytics. *International Journal of Science, Research and Technology*, 9(1), 141-151.
- Mudunuri, L. N. R., Aragani, V. M., & Maraju, P. K. (2024, December). Development of an AI-Powered Garbage Detection System for Environmental Sustainability Via YOLOv5. In *International Conference on Information and Communication Technology for Competitive Strategies* (pp. 317-327). Singapore: Springer Nature Singapore.
- Ravichandran, S., & Kandasamy, V. (2025). Optimized Attention Augmented Residual Convolutional Neural Network with Fa-Resnet for Fabric Defect Detection. *Journal of Control Engineering and Applied Informatics*, 27(4), 3-15.
- Goyal, K. K., Hebbar, K. S., & Mali, R. K. (2026, March). AI Modernization Enabled by Cloud-Native and Edge-Intelligent Architectures. In *International Conference on Information Technology and Artificial Intelligence* (pp. 102-114). Cham: Springer Nature Switzerland.
- Basavala, S. R. (2022). Security-by-design: Threat modeling and requirements for enterprise applications. *International Journal of Science, Research and Technology*, 5(1), 7168-7178.
- Sazu, A. H., Khan, M. A. N., Sultana, S., Tasnim, M., Jakir, T., Alam, A., ... & Himeluzzaman, M. (2025). Artificial intelligence-based anomaly detection for environmental and financial risk management in banking. *International Journal of Advances in Signal and Image Sciences*, 2083-2095.
- Neela, S. (2023). Performance Benchmarking Methodologies for Real-Time Integration Middleware Platforms. *American International Journal of Computer Science and Technology*, 5(5), 39-46.
- Raja, G. V., & Mali, R. K. (2021). Federated learning frameworks for privacy-preserving artificial intelligence applications. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 4(3), 4946-4950.
- Bitragunta, S. L. V. (2025). Optimizing hybrid energy storage in EVs using AI enhanced fuzzy logic control systems. *International Journal of Science, Research and Technology*, 8(1), 13525-13533.
- Chinnam, N. B. (2022). Engineering real-time supply chain visibility and ETA intelligence platforms for large-scale retail distribution networks. *International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)*, 5(6), 7918-7922.
- Srikanth, V., Aswini, P., Asha, V., Pithamber, K., Sobti, R., & Salman, Z. (2024, November). Development of an Electric Automation Control Model Using Artificial Intelligence. In *2024 Second International Conference Computational and Characterization Techniques in Engineering & Sciences (IC3TES)* (pp. 1-5). IEEE.
- Gudipuri, N. (2024). Intent-Aware Enterprise Architecture Using Persistent AI Agents for Dynamic FinTech Service Composition. *The Eastasouth Journal of Information System and Computer Science*, 2(01), 151-161.
- Matrouk, K., V, S., Kumar, S., Bhadla, M. K., Sabirov, M., & Saadh, M. J. (2023). Deep Learning-based Dynamic User Alignment in Social Networks. *ACM Journal of Data and Information Quality*, 15(3), 1-26.
- Karrothu, A. (2026). Designing high-performance distributed systems for in-memory secure data processing in cloud security analytics. *International Journal of Intelligent Systems and Applications in Engineering*, 14(1s), 68-76. <https://www.ijisae.org/>

[index.php/IJISAE/article/view/8122](https://www.ijisae.org/index.php/IJISAE/article/view/8122)

Shaik, N. (2025). Artificial intelligence in healthcare economics: Transforming cost efficiency, member

experience, and preventative care. International Journal of Research Publications in Engineering, Technology and Management, 8(2), 11850–11854.