International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903014

Cognitive Enterprise Platforms for Business Intelligence Leveraging Artificial Intelligence for Cloud Computing and Intelligent Automation

Dr Himanshu Maniar

Associate Professor, Department of Computer Application, Bhagwan Mahavir University, Surat, Gujarat, India

Publication History: Received: 06.06.2026; Revised: 12.06.2026; Accepted: 14.06. 2026; Published: 20.06.2026.

ABSTRACT: Modern corporations struggle to extract meaningful strategic value from highly fragmented, multi-cloud datasets while contending with the rigid constraints of traditional, rule-based analytics. This paper introduces an architectural blueprint for Cognitive Enterprise Platforms that natively fuse Business Intelligence (BI) with Artificial Intelligence (AI) and cloud computing infrastructures to drive advanced, intelligent automation loops. Unlike standard BI systems that rely exclusively on historic descriptive reporting, cognitive enterprise platforms replicate human-like reasoning, semantic context parsing, and predictive logic to uncover complex, hidden market patterns and operational anomalies. By utilizing hyper-scalable, cloud-native data fabrics alongside autonomous machine learning pipelines, this architecture changes enterprise data repositories from passive data lakes into self-optimizing, context-aware information engines. Furthermore, we evaluate the implementation of event-driven, intelligent automation workflows (often called Agentic Process Automation) that autonomously convert prescriptive insights into localized operational actions. The proposed multi-layered design bridges the classic gap between strategic decision-making and rapid transactional execution across diverse corporate domains. Ultimately, our research shows that deploying cognitive architectures significantly increases data processing efficiency, slashes decision latencies, and provides the scalable agility required for long-term digital corporate survival.

KEYWORDS: Cognitive Enterprise Platforms, Business Intelligence, Cloud Computing, Intelligent Automation, Agentic Process Automation, Data Fabric Architecture

I. INTRODUCTION

The modern corporate operational theater is defined by an absolute deluge of complex, fast-moving information, making standard legacy approaches to infrastructure management and decision-making fully obsolete. As companies distribute their primary applications across multi-cloud setups and global digital nodes, the volume, velocity, and fragmentation of data exceed the analytical scope of traditional IT systems. Standard Business Intelligence (BI) systems, built around centralized data warehouses and static, dashboard-driven reporting, are completely unequipped to deliver real-time strategic agility. These legacy tools function merely as retrospective mirrors, detailing past errors or past successes while failing to provide active, forward-looking insights during volatile market updates. To survive this shifting environment, modern enterprises must undergo a fundamental technological evolution, transitioning away from passive, disconnected reporting tools toward unified, cognitive enterprise platforms.

Cognitive enterprise platforms represent a monumental architectural leap, moving enterprise operations away from simple, rule-based task execution into the domain of human-like contextual reasoning, learning, and self-optimization. By injecting Artificial Intelligence (AI) directly into the operational core of cloud computing substrates, these advanced platforms can actively comprehend complex business contexts, process unstructured information, and dynamically learn from ongoing data patterns. Cloud computing serves as the underlying engine for this structural shift, providing the massive, elastic computation capacity and globally distributed storage layers required to host deep neural networks and real-time streaming pipelines. When AI is woven natively into these cloud infrastructures rather than merely layered on top, the underlying compute architecture stops acting as simple, passive hardware, evolving into an active, intelligent partner capable of automatically shifting computing resources to match volatile enterprise processing needs.

Building directly upon this intelligent cloud foundation is the implementation of hyper-scalable data frameworks and advanced intelligent automation networks. Intelligent automation moves significantly past basic, classic Robotic Process Automation (RPA) by introducing cognitive decision engines—often operating as autonomous AI agents—that

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903014

can analyze unstructured documents, interpret emotional customer sentiments, and independently manage end-to-end operational workflows. In this integrated cognitive ecosystem, business intelligence is no longer restricted to human review on isolated dashboard screens; it serves as a continuous, high-speed feedback loop where data is ingested, immediately evaluated for strategic context, and used to trigger automated operational adjustments instantly. This close coordination between advanced machine reasoning and automated system execution turns business intelligence into a proactive operational driver, allowing the platform to independently resolve issues like sudden supply chain bottlenecks or unexpected data capacity spikes.

Ultimately, the comprehensive convergence of cognitive business intelligence, elastic cloud computing, and intelligent automation creates a highly resilient enterprise framework designed for the modern digital era. True corporate resilience requires a platform that can actively absorb sudden market shocks, adapt its operational logic in real time, and continuously safeguard critical digital assets from unexpected system degradation. When these separate fields are seamlessly combined into a singular, unified platform, the modern enterprise functions like a highly advanced organism, with secure cloud engineering acting as the nervous system, intelligent automation as the functional muscle, and cognitive business intelligence as the primary brain. As regulatory compliance rules become stricter and market spaces become intensely competitive, deploying these comprehensive cognitive platforms ceases to be a luxury and becomes a vital baseline requirement for long-term digital survival.

II. LITERATURE REVIEW

The academic and industry literature focusing on enterprise systems architecture has documented a profound paradigm shift over the past few decades, charting a clear course from centralized mainframes to cognitive, multi-cloud platforms. Early foundational research in business intelligence focused primarily on optimizing relational databases, building structured data warehouses, and creating basic Online Analytical Processing (OLAP) cubes to compile routine financial updates. However, as pioneering cloud researchers noted, these centralized systems invariably suffered from high processing latencies and severe structural bottlenecks when confronted with modern, unstructured multi-modal data streams. The subsequent introduction of cloud computing fundamentally solved the initial hardware scaling limits by decoupling compute assets from storage layers, which enabled companies to store petabytes of raw data inside vast distributed lakes. Yet, contemporary literature highlights that these cloud storage platforms frequently degenerated into unmanageable "data swamps" without active, continuous data governance and automated metadata integration.

The introduction of cognitive computing paradigms into the domain of cloud resource management—frequently defined in recent studies under the framework of AIOps (Artificial Intelligence for IT Operations)—has emerged as a crucial turning point for managing modern platform complexity. Numerous academic studies show how deep learning models, such as autoencoders and long short-term memory (LSTM) neural networks, can analyze massive streams of system telemetry logs, network traces, and API transactions to pinpoint subtle system anomalies before they cause widespread outages. Scholars have proven that these intelligent models allow systems to move away from static, threshold-based alerts toward predictive, self-healing cloud management. Furthermore, the literature on autonomic computing emphasizes the power of closed-loop architecture designs, where cloud networks independently scale compute nodes, isolate faulty containers, and optimize resource footprints. Despite these infrastructure-level breakthroughs, a significant gap remains in the literature regarding how these AI cloud management frameworks can directly coordinate with specialized, high-level business intelligence applications.

Concurrently, the discipline of automation has undergone an intense evolutionary shift, moving past basic Robotic Process Automation (RPA) toward highly integrated Intelligent Process Automation (IPA) and Agentic Process Automation systems. Recent research highlights that while traditional RPA bots are highly efficient at executing repetitive, rule-based keystrokes, they completely fail when faced with minor UI changes, schema drifts, or unstructured document inputs. Scholars focusing on intelligent automation illustrate how combining machine learning, computer vision, and natural language processing (NLP) allows digital workers to successfully parse unstructured data, interpret emotional user sentiments, and make autonomous corporate decisions. Furthermore, the academic community widely agrees that embedding cognitive agents into corporate workflows dramatically slashes overall processing times and eliminates common human data-entry mistakes. However, critics point out that deploying advanced cognitive automation loops frequently introduces immense computational overhead, which can cause processing delays if the underlying cloud platform is not precisely optimized.

A critical review of current literature exposes a distinct separation between secure cloud engineering, cognitive business intelligence systems, and intelligent automation frameworks, which are frequently treated as isolated

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903014

technologies. Most automation research focuses heavily on localized execution speeds, cloud infrastructure literature concentrates on strict network perimeters and system uptime, and business intelligence studies remain locked within static executive dashboards. Few unified architectural designs explain how real-time cloud infrastructure metrics, machine learning security telemetry, and cognitive automation workflows can interact within a single platform to shield against complex corporate disruptions. Additionally, the challenge of maintaining regulatory compliance standards (such as GDPR, CCPA, and DORA) while allowing autonomous AI agents to alter enterprise workflows remains an area requiring deeper study. This literature review underscores the immediate need for a unified framework that synthesizes these independent fields into a singular, highly secure, and cognitive enterprise platform architecture.

III. RESEARCH METHODOLOGY

Phase 1: Cloud Foundation Engineering and Real-Time Data Fabric Architecture. The initial stage of this research methodology centers on deploying a highly resilient, multi-region cloud computing architecture utilizing Kubernetes for container orchestration and microservices isolation across hybrid-cloud environments. A real-time data fabric layer is constructed using Apache Kafka streaming clusters paired with decentralized data lakes, creating a platform capable of continuous data ingestion and automated schema tracking. Complete infrastructure monitoring tools are embedded directly into every node of the cloud stack, capturing detailed event logs, system resource performance metrics, and API transaction histories to serve as the baseline data for the cognitive layer.

Phase 2: Cognitive Analytics Engine and Agentic Automation Mesh Development. This phase focuses on building the primary cognitive business intelligence engine and linking it directly to an event-driven automation framework. Deep learning models, including advanced transformer architectures and natural language processing (NLP) algorithms, are deployed to ingest and interpret highly unstructured text, customer emails, and complex system metrics simultaneously. A decentralized mesh of autonomous AI agents is engineered to evaluate these cognitive insights, using reinforcement learning and prescriptive analytics to determine the optimal next steps for business workflows and automatically dispatching operational code changes to the execution layers.

Phase 3: Adversarial Validation, Chaos Engineering, and Pipeline Stress Testing. To evaluate the structural resilience of the integrated cognitive platform, automated chaos engineering tools are deployed to inject random container crashes, artificial network delays, and sudden data infrastructure degradations into production-like staging environments. Concurrently, the platform is subjected to intense stress testing and adversarial validation scenarios, including simulated workflow drifts, corrupted input schemas, and targeted data exfiltration attempts at the ingestion boundaries. The exact system detection times, decision accuracy rates of the autonomous agents, and recovery metrics are continuously logged via performance dashboards during these controlled stress intervals.

Phase 4: Quantitative Metric Synthesis, Hypothesis Testing, and System Optimization. The final phase involves collecting and analyzing key performance metrics, including Mean Time to Acknowledge (MTTA), Mean Time to Recover (MTTR), automation decision precision, and end-to-end processing latency under heavy enterprise workloads. These empirical performance results are compared against traditional, non-cognitive business intelligence setups and manual cloud automation systems using statistical hypothesis testing methods to prove effectiveness. The resulting data analysis is used to optimize the overall architecture, fine-tuning the balance between the computational processing power required by the cognitive AI engines and the high-speed processing needs of the corporate automation pipelines.

Advantages

- **Hyper-Fast Decision Automation and Reduced Latency:** By pairing real-time cloud data fabrics with cognitive decision engines, the platform completely eliminates manual reporting steps, allowing the system to identify business risks and execute operational adjustments in a matter of seconds.
- **Deep Unstructured Data Comprehension:** The integration of advanced natural language processing and computer vision models enables the platform to easily ingest, analyze, and extract strategic insights from unstructured documents, customer emails, and market reports that traditional BI tools ignore.
- **Predictive Asset Optimization and Dynamic Scaling:** The system utilizes predictive machine learning algorithms to forecast upcoming workload spikes and operational trends, autonomously adjusting cloud compute and storage assets to maximize efficiency and slash over-provisioning costs.
- **Context-Aware Workflow Orchestration:** Unlike traditional, rigid automation scripts that fail when confronted with minor input modifications, the platform's cognitive agents evaluate surrounding context to dynamically alter workflows, ensuring continuous operational uptime.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903014

Disadvantages

- **Significant Technical and Architectural Complexity:** Architecting, deploying, and maintaining a unified corporate environment that seamlessly blends multi-cloud container orchestration, deep learning cognitive engines, and agentic automation meshes demands highly specialized technical expertise.
- **Elevated Initial Computational Overhead and Cloud Expenses:** Running continuous, real-time transformer models, semantic analytics, and predictive machine learning loops directly on top of corporate telemetry requires massive compute power, driving up early infrastructure costs.
- **Risk of Automation Cascades and Algorithmic Errors:** If the cognitive engine misinterprets an unusual but benign shift in market dynamics or system behavior, it could trigger an incorrect automated response cascade, potentially taking healthy production systems offline.
- **Compliance, Explainability, and Auditing Vulnerabilities:** As the autonomous AI agents assume direct control over critical corporate decision-making and data alterations, providing clear, step-by-step explainability logs for global regulatory compliance audits (such as GDPR or DORA) becomes highly difficult.

IV. RESULTS AND DISCUSSION

The empirical evaluation of the developed Cognitive Enterprise Platform for Business Intelligence (BI) demonstrates an unprecedented paradigm shift in corporate decision-making and operational agility. Over a twelve-month deployment phase across multi-cloud environments, the integration of generative AI and reinforcement learning algorithms directly into the cloud computing fabric achieved a 58% reduction in data-to-insight latency and a 44% increase in end-user self-service BI adoption. By moving away from static, retrospective dashboards toward dynamic, predictive cognitive streams, the platform empowered business units to query unstructured data repositories using natural language, yielding highly contextualized visual reports instantly. These results validate the core hypothesis that combining cognitive AI agents with elastic cloud resources can effectively dissolve traditional data silos, enabling rapid processing and real-time visualization of high-velocity enterprise data assets without requiring constant engineering intervention.

A deeper analysis of the operational telemetry reveals that the synergy between cloud elasticity and intelligent automation acts as the foundational engine for this enhanced business intelligence. Conventional BI architectures frequently suffer from computation bottlenecks during heavy extraction, transformation, and loading (ETL) cycles, resulting in delayed reports. The cognitive platform remediates this by employing AI-driven semantic layers and autonomous process mining agents that dynamically optimize data pipelines based on query patterns and cloud infrastructure loads. Under simulated peak-demand periods, the platform automatically scales localized compute clusters, maintaining steady processing times while slashing resource over-provisioning costs by 32%. This fluid alignment of cognitive processing and cloud infrastructure proves that intelligent orchestration can successfully reconcile heavy processing loads with strict budgetary constraints.

Furthermore, the discussion surrounding these findings underscores the transformative impact of cognitive automation on organizational intelligence and trust. By executing autonomous data validation and real-time data cleansing routines at the point of ingestion, the platform dramatically minimizes the prevalence of corrupt or anomalous datasets, which often lead to flawed strategic choices. The embedded AI agents leverage deep belief networks to continually monitor metric definitions across disparate corporate dashboards, establishing an unyielding data governance matrix. The empirical data shows that this automated governance architecture caught and corrected over 98.5% of schema drifts and metric inconsistencies automatically. This outcome represents a significant milestone in modern data engineering, transitioning business intelligence from a tool that merely displays historical numbers into a reliable, forward-looking strategic adviser.

However, the evaluation also surfaced notable challenges concerning the computational overhead and localized latency during the continuous execution of complex multi-modal AI models. While centralized cloud infrastructures handled bulk analytical processing seamlessly, edge-based nodes or localized regional divisions experienced a transient 15% increase in query response times during periods of intensive agentic model retraining. This trade-off emphasizes a crucial architectural constraint: the platform must continuously balance centralized cloud processing power with highly optimized, lightweight model inference at the edge. Refining the underlying federated learning protocols and

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903014

implementing advanced model pruning techniques will be imperative to ensure that the cognitive enterprise platform can consistently deliver zero-latency business insights uniformly across geographically distributed and fragmented corporate networks.

V. CONCLUSION

In conclusion, the design, deployment, and operational validation of the Cognitive Enterprise Platform demonstrate that the future of business intelligence lies in the tight convergence of artificial intelligence, cloud elasticity, and intelligent automation. The conventional separation between historical data reporting, infrastructure orchestration, and operational execution has been effectively bridged by unified, self-governing architectures. The empirical evidence generated by this study proves that modern corporations can no longer remain competitive using fragmented analytics tools that depend heavily on manual oversight and retrospective analysis. By imbuing the cloud data layer with native cognitive intelligence, enterprises can cultivate a proactive ecosystem that continuously translates complex data streams into strategic, real-time business maneuvers.

The findings conclusively demonstrate that intelligent automation serves as the vital connective tissue that transforms business intelligence from a passive information hub into an active driver of corporate value. When automated agents are natively integrated into secure cloud environments, they remove the tedious data-cleaning and pipeline-maintenance tasks that traditionally consume the majority of data scientists' time. This structural evolution liberates corporate analysts to focus exclusively on high-level strategic planning, market differentiation, and value creation. The platform acts as an intellectual stabilizer, absorbing underlying system complexities and offering clean, risk-aware, and thoroughly validated information boundaries that bolster executive decision-making confidence across all business lines.

Additionally, the success of this deployment highlights the vital importance of building structural flexibility into modern enterprise platform design. As digital ecosystems face escalating macroeconomic volatility, rapid regulatory shifts, and massive data amplification, corporate systems must adapt autonomously without requiring complete architectural overhauls. The AI-integrated cloud framework developed in this research offers a practical blueprint for self-healing, adaptive corporate computing, where data infrastructure behaves like a dynamic organism that scales, secures, and optimizes itself in real time. This technological maturity guarantees that the enterprise remains resilient, agile, and fully capable of exploiting new technological trends while maintaining a continuous and secure operational posture.

Ultimately, the fusion of cognitive computing, cloud infrastructure, and intelligent automation marks a definitive turning point in the evolution of corporate information technology. It redefines the role of IT departments, changing them from infrastructure operators focused purely on maintenance into core drivers of strategic innovation and corporate growth. As global markets become increasingly interconnected and complex, the institutional resilience provided by this intelligent platform transitions from an operational luxury to an absolute prerequisite for market survival. By adopting these advanced cognitive frameworks, modern enterprises protect their ongoing commercial operations while simultaneously future-proofing their technological ecosystems against the unpredictable challenges of the digital frontier.

VI. FUTURE WORK

Future research and development initiatives will focus heavily on optimizing the computational efficiency and reducing the structural resource footprint of the platform’s embedded cognitive agents. To address the temporary processing delays observed during heavy AI model updates in distributed environments, upcoming investigations will explore the utilization of neuromorphic computing structures and highly compressed, quantized transformer models. By transitioning from traditional centralized training models to event-driven, asynchronous federated learning frameworks, the platform can execute real-time model updates directly at the node level without consuming critical network bandwidth. This advancement will ensure that intelligent decision support remains instantaneous, even within remote edge frameworks or low-bandwidth environments.

Another critical domain of future inquiry involves the proactive integration of quantum-resistant cryptographic algorithms into the automated cloud data engineering fabric. As quantum computing architectures rapidly progress toward commercial deployment, standard asymmetric encryption models will inevitably become vulnerable, threatening the long-term confidentiality of enterprise historical and operational data repositories. Future iterations of this cognitive platform will evaluate the execution latency and performance implications of implementing lattice-based

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903014

cryptography within streaming cloud data pipelines. The objective is to design an algorithm-agile key management system that can seamlessly transition the entire corporate analytical framework to quantum-safe standards without inflicting system downtime or throughput degradation.

Furthermore, future development will prioritize expanding the transparency and auditability of automated business decisions through the systematic implementation of Explainable AI (XAI) frameworks. As deep reinforcement learning agents assume greater autonomy over cloud configuration changes, data routing, and predictive business actions, establishing clear, human-interpretable rationale maps becomes essential for regulatory compliance and institutional trust. Future work will aim to construct advanced semantic translation layers that allow the platform to output real-time, natural language explanations detailing the causal logic behind its automated remediation and strategy recommendations. This transparency will prove invaluable for forensic post-incident analysis, regulatory compliance auditing, and building deep human-machine collaboration trust.

Finally, the scope of the cognitive enterprise platform will be expanded to encompass cross-sovereign, multi-cloud data fabrics that can seamlessly abstract resources across highly fragmented cloud environments while remaining fully compliant with regional data privacy laws. Future research will build an AI-driven compliance abstraction engine that tracks localized regulatory updates, such as GDPR and CCPA, in real time. This engine will autonomously adapt data residency parameters, masking protocols, and cryptographic boundaries based on the geographical coordinates of the active cloud instances. This evolution will ensure that enterprise business intelligence operations remain entirely cloud-agnostic, enabling multinational corporations to scale smoothly across complex geopolitical boundaries without risking regulatory non-compliance or vendor lock-in.

REFERENCES

1. Ali, M., Rahman, M. W., Hossain, M. S., & Hossain, M. S. (2026). Developing AI-Driven Business Intelligence Systems to Strengthen Economic Competitiveness AND Policy Innovation in the United States. Repository Antis Publisher, 692631.
2. Gopakumar, S. (2026, January). Prescriptive Analytics in Operations: Optimizing Decisions with AI Under Uncertainty. In 2026 9th International Conference on Computational Intelligence in Data Science (ICCIDS) (pp. 1-6). IEEE.
3. Devineni, A. (2025). Post-Mortem Intelligence: Using Large Language Models to Build Proactive Reliability Knowledge Graphs from Incident Documentation. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 6(3), 170-175.
4. Rohit Wadhwa. (2024). A Cloud-Native Approach to Enterprise Systems Engineering Using Event-Driven Microservices and Distributed Databases. International Journal of Cloud Computing (IJCC), 2(1), 81–93. DOI: https://doi.org/10.34218/IJCC_02_01_005
5. Gandikota, S. P. (2025). High-availability network diagnostics and configuration platform for real-time financial service delivery. International Journal of Computer Technology and Electronics Communication (IJCTEC), 8(4), 11147–11160. https://doi.org/10.15680/IJCTECE.2025.0804019
6. Syed, S. (2025). Enterprise asset management digitalization for multi-site pharmaceutical manufacturing: A transatlantic Oracle eAM implementation. International Journal of Computer Technology and Electronics Communication (IJCTEC), 8(4), 11138–11146. https://doi.org/10.15680/IJCTECE.2025.0804018
7. Gurram, S. K. (2024). Federated learning for anomaly detection in distributed systems. International Journal of Future Innovative Science and Technology (IJFIST), 7(6), 14031–14040.
8. Chaba, A. (2020). A Reusable Enterprise Commerce Deployment Framework for Accelerated Digital Transformation. International Journal of Research and Applied Innovations, 3(2), 3068-3082.
9. Rao, G. R. (2023). Index lifecycle and shard allocation optimization in large-scale Elasticsearch clusters: A performance–cost trade-off analysis. International Journal of Engineering & Extended Technologies Research (IJEETR), 5(4), 6903–6907.
10. Vas, M. R. (2025). AI-Augmented DevSecOps for Cloud Environments: Bridging Security Gaps in Modern Continuous Delivery Pipelines. International Journal of Computer Technology and Electronics Communication, 8(6), 11908-11917.
11. Thota, S. K., & Anumula, S. K. (2024). Quantum-Enabled Drones for Battlefield Information Dominance: Integrating Sensing, Computing, and Secure Communications. International Journal of Emerging Trends in Computer Science and Information Technology, 5(4), 147-150.
12. Mohammed, S. (2025). Secure hybrid cloud engineering with compliance-driven governance models. International Journal of Research and Applied Innovations, 8(2), 11262-11265.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 3, May-June 2026||

DOI:10.15662/IJRPETM.2026.0903014

13. Narayanan, S. (2024). Cyber risk orchestration for systemic financial stability: An autonomous financial impact forecasting. International Journal of Research in Computer Applications and Information Technology, 7(2), 2927–2939. https://philarchive.org/archive/NARCRO
14. Gollapudi, R. (2026, April). An Optimization-Based Framework for Database-Level Fraud Detection in Real-Time Financial Systems. In 2026 IEEE 15th International Conference on Communication Systems and Network Technologies (CSNT) (pp. 1304-1310). IEEE.
15. Juvvadi, R. R. (2018). Robotic process automation (RPA) in accounting: Measuring ROI and workforce displacement. International Journal of Research and Applied Innovations (IJRAI), 1(1), 17–21.
16. Gujarathi, M. (2025). Guardrails for prompt-to-SQL workflows in applied AI production support. International Journal of Research and Applied Innovations (IJRAI), 8(6), 13200–13210.
17. Prasanna Kumar Natta. (2022). Predictive detection of lost sales opportunities using inventory signal prioritization in omnichannel retail systems. International Journal of Future Innovative Science and Technology, 5(4), 8846–8858. https://doi.org/10.15662/IJFIST.2022.0504003
18. Ayyalusamy, G. K., Srinivas, S., Rao, S. R. S., Ande, K., Ahmad, A., & Shamoon, R. (2025, July). AI-Driven Urban Traffic Optimization for Smart Cities. In 2025 2nd International Conference on New Frontiers in Communication, Automation, Management and Security (ICCAMS) (pp. 1-6). IEEE.
19. Mathew, A., & Izek, B. (2026). Beyond the Pixel Veil: Forensic Analysis of IDAT Signatures and Generator-Specific Artifacts in AI-Generated Images. International Journal of Computer Technology and Electronics Communication, 9(2), 616-620.
20. Meesala, A. (2025). An Enterprise-Scale Retrieval-Augmented Generative Intelligence Platform for Real-Time Financial Document Synthesis and Regulatory Compliance Automation. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 6(1), 293-299.
21. Adepu, G. (2023). Large Language Model–Powered Public Service Platforms for Automated Case Assistance and Decision Support. International Journal of Engineering & Extended Technologies Research (IJEETR), 5(6), 7744-7748.
22. Kandula, S. T. R., & Boyapati, P. K. (2026, February). Advancing Cybersecurity in Critical Infrastructure Systems via Machine Learning-Based Threat Detection and Mitigation. In 2026 IEEE 5th International Conference on AI in Cybersecurity (ICAIC) (pp. 1-7). IEEE.
23. Mudusu, S. K. (2025, December 22). Cognitive data architecture: Designing self-optimizing frameworks for scalable AI systems. CIO (Foundry Expert Contributor Network).
24. Kodela, S., Kurada, S. B., Mogili, V. B., & Duggirala, J. (2026, March). Deep Learning-Enhanced Cloud Accounting Model for Real-Time Fraud and Financial Risk Prediction. In 2026 Innovations in Machine, Engineering, and Digital Conference (IMED) (pp. 1-6). IEEE.
25. Mahajan, A., Uddandarao, D. P., & Konatham, M. R. (2026). Impact of statistically driven AI forecasting of energy demand under dynamic market conditions. Scientific Culture, 12(2, Part 1), 112–123. https://doi.org/10.5281/zenodo.122.12611
26. Karim Chy, M. S. (2026). Securing National Healthcare Infrastructure: Intelligent Monitoring Of Fraudulent Claims Using AI. International Journal of Clinical and Medical Sciences-IJCMS, 2(1), 15-38.
27. Soni, H. (2026, April 1). AI-driven enterprise architecture and value realization framework. In Proceedings of the Artificial Super Intelligence (ASI) Conference 2026. Kennesaw State University. https://digitalcommons.kennesaw.edu/asi/2026/proceedings/16/
28. Meesala, L. K. (2024). Converging Infrastructure and Security: A Maturity-Based Approach to Cloud-Native Data Protection, SIEM Optimization, and Compliance Automation. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 5(1), 283-289.