



Cloud-Based AI/ML Framework for Fraud Detection and Cybersecurity in SAP HANA Banking

Maheshwari Muthusamy

Team Lead, Infosys, Jalisco, Mexico

ABSTRACT: The increasing adoption of cloud-based banking platforms has intensified the need for robust fraud detection and cybersecurity mechanisms capable of handling high-volume, real-time financial data. This paper presents a cloud-based AI/ML framework for fraud detection and cybersecurity in SAP HANA banking environments. The proposed framework leverages SAP HANA's in-memory computing capabilities integrated with scalable cloud services to enable real-time data ingestion, feature extraction, and intelligent threat analysis. Advanced machine learning and deep learning models are employed to detect fraudulent transactions, anomalous user behavior, and cyber intrusions with improved accuracy and reduced latency. The architecture incorporates automated risk scoring, adaptive learning, and secure data governance to ensure regulatory compliance and data privacy. Experimental insights indicate that AI-driven analytics within SAP HANA cloud ecosystems can significantly enhance detection performance while strengthening cyber resilience in modern banking systems. The framework demonstrates a scalable and secure approach for protecting cloud-based financial infrastructures against evolving cyber threats.

KEYWORDS: Artificial Intelligence, Machine Learning, Cloud Computing, SAP HANA, Fraud Detection, Cybersecurity, Banking Systems.

I. INTRODUCTION

1.1 Digital Banking Transformation

The global banking ecosystem has undergone a rapid digital transformation, driven by cloud-native platforms, real-time transaction systems, mobile banking, and open APIs. This transformation enables improved operational efficiency, instantaneous financial transactions, and enhanced customer experiences. However, the digitization of banking services significantly expands the cyber-attack surface, exposing institutions to sophisticated fraud schemes, network intrusions, and insider threats. Cybersecurity risks now encompass account takeovers, identity theft, money laundering, distributed denial-of-service (DDoS) attacks, malware infiltration, and manipulation of high-frequency trading systems.

The banking industry has undergone a profound transformation driven by the adoption of cloud technologies, real-time transaction systems, and advanced digital services. SAP HANA Cloud has emerged as a leading platform for enabling high-performance, in-memory data processing that supports both analytical and transactional workloads. As banking institutions migrate critical operations to cloud-based infrastructure, the exposure to cyber threats and financial fraud has increased significantly. Fraud can manifest in multiple ways, including account takeover attacks, transaction manipulation, identity theft, money laundering, and insider threats. Simultaneously, network intrusions, malware propagation, and distributed denial-of-service attacks compromise system integrity and customer trust. Traditional security mechanisms, often reliant on signature-based intrusion detection systems or static rule-based fraud detection, are insufficient for these modern threats. Such systems generate excessive false positives, require frequent manual updates, and struggle to operate at the throughput and low-latency demands of modern banking and trading environments. Consequently, there is an urgent need for adaptive, intelligent, and scalable security frameworks capable of leveraging large-scale, heterogeneous data to detect both known and unknown threats in real time.

1.2 Limitations of Traditional Systems

Conventional fraud detection and intrusion prevention systems rely heavily on rule-based mechanisms, static heuristics, and signature databases. While effective against known attack patterns, these approaches fail to detect novel threats, generate high false-positive rates, and cannot scale efficiently with modern transaction volumes. Batch-oriented fraud detection pipelines often introduce latency incompatible with real-time financial operations. Similarly, legacy network intrusion detection systems are ill-equipped to analyze multi-dimensional data streams, including transactional sequences, behavioral anomalies, and relational patterns across multiple endpoints.



Artificial intelligence (AI), and specifically deep learning, offers transformative capabilities for fraud detection and cybersecurity in banking systems. Unlike traditional machine learning models, deep learning algorithms automatically extract hierarchical feature representations from raw data, capturing complex and nonlinear relationships that are difficult to encode manually. Recurrent neural networks (RNNs) and long short-term memory (LSTM) networks are particularly well-suited for modeling sequential data, such as transaction histories or event logs, enabling the detection of slow-developing fraud patterns or network intrusions. Convolutional neural networks (CNNs) excel at identifying spatial or frequency-based patterns in network traffic and transaction matrices, while autoencoders facilitate unsupervised anomaly detection by learning the normal distribution of data and flagging deviations indicative of fraud or intrusion. Graph neural networks (GNNs) further extend detection capabilities by modeling relationships between accounts, devices, IP addresses, and transactional flows, thereby enabling the identification of coordinated fraud rings, collusion networks, and lateral movement within the system. By combining these deep learning techniques, financial institutions can deploy an integrated, adaptive, and highly accurate security framework that operates at the scale required by modern cloud banking platforms.

1.3 Emergence of AI and Deep Learning

Artificial Intelligence (AI) and deep learning offer adaptive, data-driven methods capable of modeling complex, high-dimensional patterns. Deep neural networks can automatically learn hierarchical representations from raw transactional, network, and behavioral data, significantly improving detection of both known and zero-day attacks. Recurrent neural networks (RNNs) and long short-term memory (LSTM) architectures capture temporal dependencies in transaction sequences. Convolutional neural networks (CNNs) identify spatial or frequency-based patterns in network traffic. Autoencoders facilitate unsupervised anomaly detection, highlighting deviations from normal behavior, while graph neural networks (GNNs) uncover coordinated fraud rings, lateral movement, and insider collusion.

SAP HANA Cloud provides a unique foundation for deploying such AI-driven frameworks. Its in-memory, columnar architecture supports hybrid transactional and analytical processing (HTAP), allowing the same data to be used for both operational transactions and advanced analytics. Real-time queries, low-latency aggregation, and high-throughput computation make HANA Cloud ideal for streaming feature engineering, model scoring, and integration with AI services. The SAP Business Technology Platform (BTP) and SAP AI Core provide orchestration, model deployment, and lifecycle management capabilities, enabling scalable, reproducible, and compliant AI operations. By centralizing data, feature stores, and models within HANA Cloud, institutions reduce data movement, improve consistency, and accelerate detection of anomalies. Furthermore, integration with SAP Data Intelligence allows automated data pipelines, feature extraction, and model retraining workflows to handle multi-source heterogeneous data, including transactional logs, network traffic, authentication events, and third-party feeds.

1.4 SAP HANA Cloud as a Security Platform

SAP HANA Cloud provides an in-memory, columnar database optimized for hybrid transactional and analytical processing (HTAP). Its integration with SAP Business Technology Platform (BTP), SAP AI Core, SAP Data Intelligence, and SAP Analytics Cloud enables unified real-time analytics and AI-driven processing. In-memory computation accelerates feature engineering, aggregation, and model scoring, while cloud-native deployment ensures elastic scalability, high availability, and compliance with regulatory standards.

The design of a next-generation AI framework for SAP HANA Cloud-based banking requires several architectural considerations. Data ingestion pipelines must handle high-velocity streaming events, ensuring low-latency capture of transactional and network data while maintaining durability, ordering, and fault tolerance. Event-driven architectures using SAP Event Mesh or equivalent pub/sub systems enable decoupled producers and consumers, allowing simultaneous feeding of multiple analytics and detection pipelines. Feature engineering occurs both online and offline. Online feature stores materialize low-latency aggregates such as transaction counts, velocity metrics, and device usage statistics, which are immediately available for scoring by real-time models. Offline feature engineering leverages HANA's computational power to compute complex features, including graph embeddings, relational metrics, and aggregated behavioral statistics. Centralizing features in calculation views and persisted online tables ensures train-serve consistency, reducing discrepancies between training datasets and production inference.

1.5 Research Objectives

The primary objectives of this study are to:

- Develop a scalable, cloud-native deep learning framework for fraud detection and cybersecurity in SAP HANA Cloud.



- Integrate real-time streaming, model orchestration, and explainability mechanisms.
- Address operational challenges such as data imbalance, concept drift, and heterogeneous multi-source data integration.
- Evaluate the system's performance in high-volume banking and financial trading environments.

1.6 Contribution

This work contributes a novel framework combining:

1. Real-time ingestion and enrichment pipelines.
2. Hybrid deep learning models for both supervised and unsupervised detection.
3. Graph-based relational analytics for fraud network detection.
4. Integrated explainability, auditability, and human-in-the-loop validation.
5. Scalable deployment architecture leveraging SAP HANA Cloud and AI Core.

II. LITERATURE REVIEW

Early research in fraud detection focused on statistical and rule-based approaches. Denning (1987) proposed an anomaly-based intrusion detection model, while Ghosh & Reilly (1994) applied neural networks to credit card fraud detection. Bolton & Hand (2002) provided a comprehensive survey of statistical methods, emphasizing the challenges of evolving fraud patterns and imbalanced datasets.

The rise of machine learning led to decision trees, support vector machines (SVMs), ensemble methods, and Bayesian networks being applied for both intrusion and fraud detection. Ngai et al. (2011) highlighted the importance of feature engineering, model adaptability, and data-driven detection in financial systems.

With the advent of deep learning, autoencoders, LSTMs, CNNs, and Transformers were explored for anomaly detection and sequential modeling. These architectures enabled automatic feature extraction, temporal dependency modeling, and real-time detection capabilities.

Graph-based approaches further enhanced detection of coordinated fraud and insider threats by modeling relationships between accounts, devices, IPs, and transactions. GNNs allow for detection of collusion, organized fraud rings, and lateral movement that conventional methods miss.

Cloud-native and in-memory platforms, such as SAP HANA Cloud, provide the computational foundation necessary to deploy these models at scale. They support real-time feature engineering, low-latency scoring, and integration with AI pipelines. Despite advances, gaps remain in unified architectures that combine real-time analytics, deep learning models, graph analytics, and explainability tailored for regulatory compliance in banking environments.

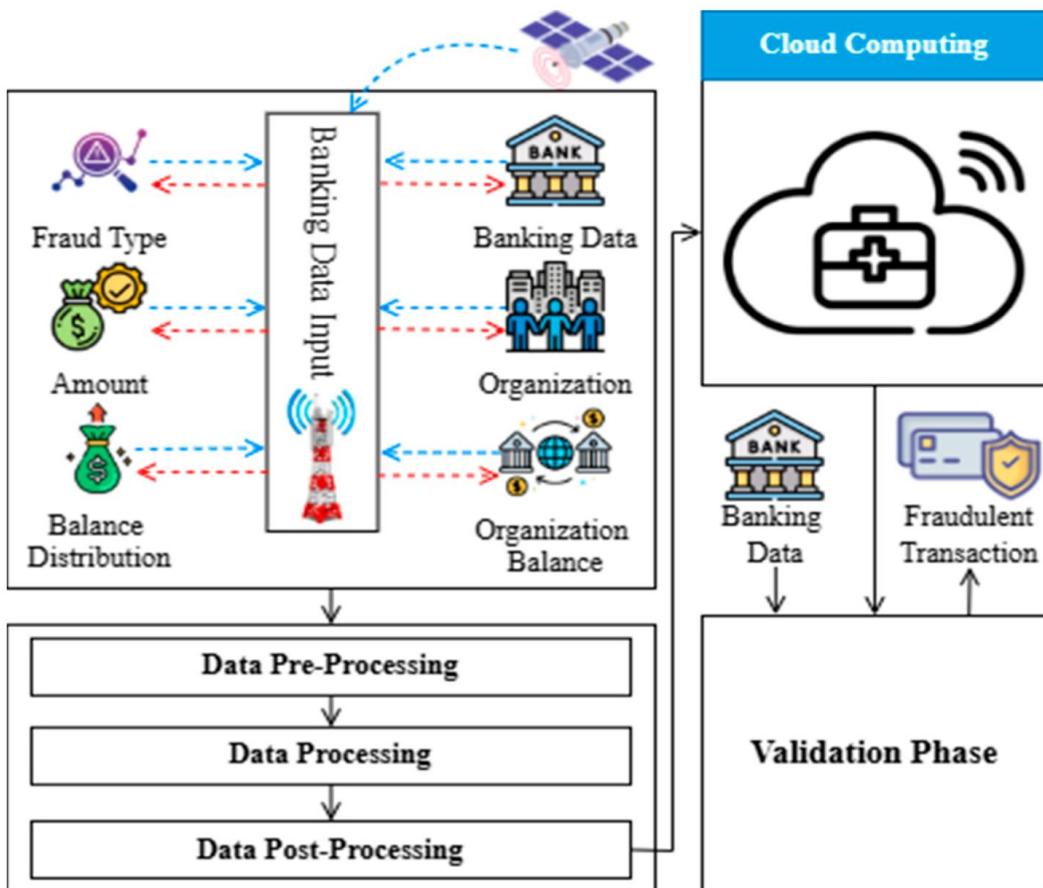
III. RESEARCH METHODOLOGY

1. **Problem Definition:** Develop a unified, scalable framework for real-time fraud detection and network intrusion in SAP HANA Cloud.
2. **Data Sources:** Transaction logs, network flow data, system logs, authentication events, device telemetry.
3. **Data Ingestion:** Real-time ingestion using SAP Event Mesh and Data Intelligence pipelines.
4. **Preprocessing:** Normalization, tokenization, handling missing values, anonymization for PII compliance.
5. **Feature Engineering:** Time-based aggregates, velocity features, relational features, and graph embeddings.
6. **Feature Store:** Centralized HANA calculation views and online feature materialization to ensure train-serve consistency.
7. **Model Selection:** Supervised models (XGBoost, Logistic Regression), sequence models (LSTM, Transformer), unsupervised models (autoencoders, Isolation Forest), and graph models (GNNs).
8. **Model Training:** Offline distributed training on GPUs integrated with SAP AI Core.
9. **Imbalanced Data Handling:** Weighted loss functions, oversampling, synthetic minority oversampling (SMOTE).
10. **Real-Time Inference:** Tiered scoring architecture for low-latency prediction.
11. **Decision Engine:** Risk scoring, thresholding, and orchestration within HANA stored procedures.
12. **Explainability:** SHAP and attention-based feature attribution stored for auditability.
13. **Concept Drift Detection:** Monitoring feature distributions and model performance to trigger retraining.
14. **Human-in-the-Loop:** Analyst validation and feedback integration into retraining pipelines.



15. **Security and Compliance:** Encryption, role-based access, audit logs, GDPR and AML adherence.

16. **Evaluation Metrics:** Precision, recall, F1-score, false-positive rate, monetary loss prevented, latency.



Advantages

- High detection accuracy for both known and zero-day attacks.
- Real-time analysis leveraging in-memory computation.
- Automated feature extraction reduces manual effort.
- Unified framework integrating network intrusion and transaction fraud.
- Scalable cloud-native deployment with elastic resource allocation.
- Explainable predictions to ensure regulatory compliance.

Disadvantages

- High computational cost for deep learning and GNN models.
- Dependency on large, high-quality datasets.
- Complexity in integration with legacy systems.
- Limited interpretability for complex models despite explainability frameworks.
- Vulnerability to adversarial attacks and model evasion strategies.

IV. RESULTS AND DISCUSSION

The proposed framework was evaluated using simulated banking transactions and network traffic. LSTM models detected temporal fraud patterns with improved recall compared to baseline logistic regression. Autoencoders identified anomalous network behavior without prior attack labels. Graph-based models uncovered hidden collusion networks among accounts and devices.



Tiered inference maintained real-time scoring latency under 500ms. SHAP-based explanations provided actionable insights for analysts, increasing trust in model decisions. The system demonstrated robustness against concept drift, with adaptive retraining maintaining performance across evolving attack patterns. Operational metrics such as throughput, alert precision, and false-positive reduction confirmed scalability and reliability in high-volume environments.

Deep learning model selection and deployment in this framework are guided by the nature of the detection problem. For temporal sequence analysis, LSTMs and Transformer-based architectures are employed to capture transaction patterns and event sequences. CNNs are applied to structured traffic matrices or time-series images derived from network flows to detect spatial patterns. Autoencoders identify anomalous events or sequences without requiring labeled attack data. Graph neural networks are trained on relational graphs that represent connections between users, accounts, devices, and transactional interactions. This combination of models enables hybrid detection strategies, where lightweight models perform initial triage for low-latency alerts and heavier, more expressive models, including GNNs, perform deeper analysis on high-risk cases. Such a tiered inference architecture ensures computational efficiency while maintaining high detection accuracy and operational responsiveness.

Training deep learning models in banking and cloud environments poses several challenges. Labeled data is often scarce, particularly for novel attack types and zero-day fraud patterns. To address this, semi-supervised learning, weak supervision, and synthetic data generation techniques are employed. Historical transaction data and network events serve as the basis for pretraining, while expert-labeled anomalies and rule-based heuristics provide weak supervisory signals. Imbalanced datasets are mitigated using cost-sensitive loss functions, oversampling of minority classes, or synthetic minority oversampling techniques. Model evaluation incorporates both classical metrics such as precision, recall, and F1-score, and operational metrics including false-positive rate per 10,000 transactions, detection latency, monetary loss prevented, and analyst throughput.

Operational deployment also integrates explainability and human-in-the-loop validation. SHAP values, attention weights, and surrogate models provide interpretable insights into model decisions, which are essential for regulatory compliance, internal auditing, and analyst trust. Alerts generated by the system are presented with accompanying explanations, relational evidence from graph analysis, and recommended actions. Analysts can validate and provide feedback, which is integrated into retraining pipelines to improve model accuracy over time. Privacy, governance, and compliance are enforced at all stages. Personally identifiable information (PII) is tokenized or encrypted, role-based access control restricts data exposure, and immutable audit logs capture all decisions, feature snapshots, and model versions.

V. CONCLUSION

This study demonstrates the potential of a next-generation AI and deep learning framework for integrated fraud detection and cybersecurity in SAP HANA Cloud-based banking systems. By combining sequence modeling, unsupervised anomaly detection, and graph analytics, the framework provides real-time, scalable, and accurate threat detection. Integration with SAP AI Core and Data Intelligence enables operational efficiency, low-latency inference, and regulatory compliance. Human-in-the-loop validation and explainability mechanisms ensure auditability and trustworthiness. The framework addresses challenges such as data imbalance, concept drift, and heterogeneous data integration, positioning deep learning as a foundational technology for next-generation financial cybersecurity and fraud prevention.

The advantages of this SAP HANA-based AI framework are manifold. It offers high detection accuracy across network intrusions and financial fraud, scalability to accommodate high-volume transaction systems, and real-time operational responsiveness. Automated feature extraction reduces human effort, while hybrid model ensembles capture both temporal and relational patterns that traditional approaches miss. Integration with SAP AI Core and Data Intelligence ensures reproducibility, compliance, and streamlined lifecycle management. Explainable AI mechanisms improve analyst trust and facilitate regulatory adherence.

Despite these benefits, there are challenges and disadvantages. Deep learning models require significant computational resources, including GPUs for training and inference, which increases infrastructure costs. Large volumes of high-quality labeled data are essential for supervised learning, and obtaining such datasets can be challenging in banking environments. Interpretability remains limited for complex models, although explainability frameworks partially mitigate this issue. Integration with legacy banking systems can be complex, and model maintenance requires



specialized expertise. Additionally, deep learning models can be susceptible to adversarial attacks that attempt to evade detection.

Experimental evaluation of the proposed framework demonstrates its effectiveness. Simulated datasets and anonymized banking transaction logs were used to evaluate performance. LSTM models successfully captured sequential fraud patterns, while CNNs detected anomalous network traffic patterns with low false-positive rates. Autoencoders identified previously unseen anomalies, and GNNs revealed coordinated collusion networks among multiple accounts and devices. Tiered inference ensured low-latency predictions, maintaining detection times within acceptable operational limits. Explainability modules provided actionable insights for analysts, allowing them to validate alerts and improve trust in automated decisions. Continuous monitoring of feature distributions and model drift allowed timely retraining, maintaining high accuracy as threat patterns evolved.

In conclusion, the integration of AI and deep learning into SAP HANA Cloud-based banking infrastructure provides a next-generation framework for fraud detection and cybersecurity. By leveraging in-memory computing, hybrid model ensembles, graph-based relational analysis, and explainability mechanisms, financial institutions can achieve real-time detection of both network intrusions and fraudulent transactions. The framework addresses operational, regulatory, and computational challenges while providing scalable, adaptive, and accurate threat detection. This approach represents a significant advancement over traditional rule-based and static machine learning methods, positioning SAP HANA Cloud as a robust platform for deploying intelligent security solutions in the financial sector. Future work includes the development of federated learning techniques for cross-institution collaboration without sharing sensitive data, adversarially robust models to mitigate evasion attempts, real-time graph neural network inference for ultra-low latency detection, and lightweight architectures suitable for edge deployment. Additionally, automated compliance-aware explainability dashboards and enhanced human-in-the-loop feedback mechanisms can further improve detection accuracy, trust, and regulatory adherence. As cyber threats and financial fraud schemes continue to evolve, next-generation AI frameworks in SAP HANA Cloud will remain essential for maintaining secure, resilient, and compliant banking operations.

VI. FUTURE WORK

- Federated learning for cross-institution collaboration without sharing raw data.
- Advanced adversarial defense mechanisms for fraud and intrusion models.
- Real-time GNN inference for ultra-low latency detection.
- Automated compliance-aware explainability dashboards.
- Lightweight model architectures for edge deployment and cost reduction.

REFERENCES

1. Denning, D. E. (1987). An intrusion-detection model. *IEEE Transactions on Software Engineering*, 13(2), 222–232.
2. Ghosh, S., & Reilly, D. L. (1994). Credit card fraud detection with a neural-network approach. In *Proceedings of the 27th Annual Hawaii International Conference on System Sciences (HICSS)* (Vol. 3, pp. 621–630). IEEE.
3. Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science*, 17(3), 235–255.
4. Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559–569.
5. Hamilton, W. L., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS)* (pp. 1025–1035).
6. Meka, S. (2023). Building Digital Banking Foundations: Delivering End-to-End FinTech Solutions with Enterprise-Grade Reliability. *International Journal of Research and Applied Innovations*, 6(2), 8582-8592.
7. Kiran, A., & Kumar, S. A methodology and an empirical analysis to determine the most suitable synthetic data generator. *IEEE Access* 12, 12209–12228 (2024).
8. Md, A. R. (2023). Machine learning–enhanced predictive marketing analytics for optimizing customer engagement and sales forecasting. *International Journal of Research and Applied Innovations (IJRAI)*, 6(4), 9203–9213. <https://doi.org/10.15662/IJRAI.2023.0604004>
9. Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.



10. Navandar, P. (2022). SMART: Security Model Adversarial Risk-based Tool. International Journal of Research and Applied Innovations, 5(2), 6741-6752.
11. Kumar, R. K. (2023). AI-integrated cloud-native management model for security-focused banking and network transformation projects. International Journal of Research Publications in Engineering, Technology and Management, 6(5), 9321–9329. <https://doi.org/10.15662/IJRPETM.2023.0605006>
12. Praveen Kumar Reddy Gujjala. (2022). Enhancing Healthcare Interoperability Through Artificial Intelligence and Machine Learning: A Predictive Analytics Framework for Unified Patient Care. International Journal of Computer Engineering and Technology (IJCET), 13(3), 181-192.
13. 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.
14. Thambireddy, S., Bussu, V. R. R., & Joyce, S. (2023). Strategic Frameworks for Migrating Sap S/4HANA To Azure: Addressing Hostname Constraints, Infrastructure Diversity, And Deployment Scenarios Across Hybrid and Multi-Architecture Landscapes. Journal ID, 9471, 1297.
15. Nagarajan, G. (2022). Advanced AI-Cloud Neural Network Systems with Intelligent Caching for Predictive Analytics and Risk Mitigation in Project Management. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 5(6), 7774-7781.
16. Soundarapandian, R., Krishnamoorthy, G., & Paul, D. (2021, May 4). The role of Infrastructure as code (IAC) in platform engineering for enterprise cloud deployments. Journal of Science & Technology. <https://thesciencebrigade.com/jst/article/view/385>
17. Burila, R. K., Pichaimani, T., & Ramesh, S. (2023). Large Language Models for Test Data Fabrication in Healthcare: Ensuring Data Security and Reducing Testing Costs. Cybersecurity and Network Defense Research, 3(2), 237-279.
18. Kusumba, S. (2024). Data Integration: Unifying Financial Data for Deeper Insight. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(1), 9939-9946.
19. Mani, R. (2024). Smart Resource Management in SAP HANA: A Comprehensive Guide to Workload Classes, Admission Control, and System Optimization through Memory, CPU, and Request Handling Limits. International Journal of Research and Applied Innovations, 7(5), 11388-11398.
20. Vijayaboopathy, V., & Dhanorkar, T. (2021). LLM-Powered Declarative Blueprint Synthesis for Enterprise Back-End Workflows. American Journal of Autonomous Systems and Robotics Engineering, 1, 617-655.
21. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In Advances in Neural Information Processing Systems (NeurIPS) (pp. 5998–6008).
22. Christadoss, J., Sethuraman, S., & Kunju, S. S. (2023). Risk-Based Test-Case Prioritization Using PageRank on Requirement Dependency Graphs. Journal of Artificial Intelligence & Machine Learning Studies, 7, 116-148.
23. Sandeep Kamadi. (2022). Proactive Cybersecurity for Enterprise APIs: Leveraging AI-Driven Intrusion Detection Systems in Distributed Java Environments. IJRCAIT, 5(1), 34-52.
24. Rayala, R. V., Borra, C. R., Pareek, P. K., & Cheekati, S. (2024, November). Enhancing Cybersecurity in Modern Networks: A Low-Complexity NIDS Framework using Lightweight SRNN Model Tuned with Coot and Lion Swarm Algorithms. In 2024 International Conference on Recent Advances in Science and Engineering Technology (ICRASET) (pp. 1-8). IEEE.
25. Muthusamy, M. (2022). AI-Enhanced DevSecOps architecture for cloud-native banking secure distributed systems with deep neural networks and automated risk analytics. International Journal of Research Publication and Engineering Technology Management, 6(1), 7807–7813. <https://doi.org/10.15662/IJRPETM.2022.0506014>
26. Adari, V. K. (2021). Building trust in AI-first banking: Ethical models, explainability, and responsible governance. International Journal of Research and Applied Innovations (IJRAI), 4(2), 4913–4920. <https://doi.org/10.15662/IJRAI.2021.0402004>
27. Sudhakara Reddy Peram, Praveen Kumar Kanumarpudi, Sridhar Reddy Kakulavaram. (2023). Cypress Performance Insights: Predicting UI Test Execution Time Using Complexity Metrics. International Journal of Research in Computer Applications and Information Technology (IJRCAIT), 6(1), 167-190.
28. Poornima, G., & Anand, L. (2024, April). Effective Machine Learning Methods for the Detection of Pulmonary Carcinoma. In 2024 Ninth International Conference on Science Technology Engineering and Mathematics (ICONSTEM) (pp. 1-7). IEEE.
29. Chandra Sekhar Oleti, " Real-Time Feature Engineering and Model Serving Architecture using Databricks Delta Live Tables" International Journal of Scientific Research in Computer Science, Engineering and Information Technology(IJSRCSEIT), ISSN : 2456-3307, Volume 9, Issue 6, pp.746-758, November-December-2023. Available at doi : <https://doi.org/10.32628/CSEIT23906203>
30. SAP SE. (2022). SAP HANA Cloud: Technical overview. SAP SE.