International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November - December 2025||

DOI:10.15662/IJARCST.2025.0806019

AI-Driven Cybersecurity Software for Fraud and Network Intrusion Detection in Financial Big Data Systems

Rajagopalachari Srinivasan Subramanian

Independent Researcher, Karnataka, India

ABSTRACT: The rapid growth of digital banking and financial trading systems has significantly increased the risk of cyber threats, including network intrusions and fraudulent activities. Traditional security mechanisms often struggle to handle the high volume, velocity, and variety of financial big data, leading to delayed threat detection and financial losses. This research proposes an AI-driven cybersecurity software framework that leverages machine learning algorithms, big data analytics, and real-time monitoring to detect and prevent fraudulent transactions and network intrusions. The proposed system integrates anomaly detection, predictive analytics, and automated threat response to enhance security, reduce false positives, and improve operational efficiency. By combining scalable big data processing with intelligent machine learning models, the framework ensures proactive and adaptive cybersecurity for banking and financial trading platforms. Experimental results on real-world financial datasets demonstrate high accuracy, robustness, and efficiency, highlighting the potential of AI-powered solutions in safeguarding critical financial infrastructures.

KEYWORDS: AI, Machine Learning, Cybersecurity, Big Data, Network Intrusion Detection, Fraud Detection, Financial Systems, Banking Security, Trading Systems, Anomaly Detection, Predictive Analytics, Real-Time Monitoring

I. INTRODUCTION

1.1 Digital Transformation of Financial Systems

Banking and financial trading systems have evolved into highly interconnected, software-driven ecosystems. Core banking platforms, online payment gateways, mobile applications, algorithmic trading engines, and cloud-based analytics systems exchange massive volumes of data in real time. While this transformation improves efficiency and market responsiveness, it also increases exposure to cyber threats and financial fraud.

1.2 Convergence of Cyber Intrusion and Financial Fraud

Traditionally, network intrusion detection and fraud detection were treated as separate security domains. However, modern attacks frequently span both layers. For example, a network breach may precede account takeover fraud, or compromised trading terminals may be used to manipulate orders. This convergence demands integrated detection systems capable of correlating network-level anomalies with financial behavior.

1.3 Limitations of Conventional Security Approaches

Signature-based intrusion detection systems (IDS) and static fraud rules are ineffective against novel attacks and adaptive adversaries. These systems generate excessive false positives and require frequent manual updates. Moreover, they struggle to scale across high-throughput trading systems where latency tolerance is extremely low.

1.4 Emergence of Deep Learning in Cyber-Financial Security

Deep learning (DL) has demonstrated exceptional performance in domains characterized by large-scale, high-dimensional data. In cybersecurity and fraud detection, DL models can automatically learn complex patterns, temporal dependencies, and nonlinear relationships that are difficult to encode manually. This capability is particularly valuable for detecting stealthy intrusions and sophisticated fraud schemes.

1.5 Importance for Banking and Trading Systems

Banking systems demand strong fraud prevention to protect customers and comply with regulations, while financial trading systems require ultra-low latency detection to prevent market abuse and systemic risk. Deep learning offers a unified framework to address both requirements through adaptive, data-driven intelligence.

1.6 Objectives and Contributions

This paper aims to:

- Analyze deep learning techniques for network intrusion and fraud detection
- Propose an integrated detection architecture for banking and trading systems
- Evaluate performance benefits and operational challenges
- Provide insights for secure deployment in regulated financial environments

II. LITERATURE REVIEW

Early research in intrusion detection focused on statistical anomaly detection and rule-based expert systems. Denning (1987) introduced foundational models for intrusion detection based on deviation from normal behavior. In fraud detection, early neural network models demonstrated feasibility in identifying suspicious credit card transactions (Ghosh & Reilly, 1994).

During the late 1990s and early 2000s, machine learning techniques such as decision trees, support vector machines, and Bayesian classifiers gained popularity. Bolton and Hand (2002) provided a comprehensive review of statistical fraud detection methods, highlighting challenges related to imbalance and evolving fraud patterns.

The application of deep learning to cybersecurity accelerated with increased computational power and availability of large datasets. Autoencoders were widely adopted for unsupervised anomaly detection in network traffic, while convolutional neural networks (CNNs) were used to classify traffic patterns.

Recurrent neural networks (RNNs) and long short-term memory (LSTM) models enabled temporal modeling of network flows and transaction sequences, improving detection of slow-moving attacks and fraud patterns. In financial trading systems, DL models were explored to detect abnormal trading behavior and insider manipulation.

Graph-based learning emerged as a powerful approach to detect coordinated attacks and fraud rings. By modeling relationships between users, devices, IPs, and transactions, graph neural networks (GNNs) can identify collusion and lateral movement within networks.

Despite promising results, literature highlights key challenges: lack of labeled data, explainability, computational overhead, and regulatory constraints. Few studies address integrated intrusion and fraud detection architectures tailored to financial systems, motivating this research.

III. RESEARCH METHODOLOGY

1. **Problem Definition:** Develop a deep learning–based detection framework for network intrusion and financial fraud in banking and trading systems.
2. **Data Sources:** Network traffic logs, system calls, transaction records, order books, user behavior logs, and authentication events.
3. **Data Collection:** Real-time ingestion via log collectors, packet analyzers, and transaction monitoring systems.
4. **Data Preprocessing:** Noise filtering, normalization, encoding categorical variables, and handling missing values.
5. **Feature Engineering:** Temporal features, statistical summaries, behavioral metrics, and relational attributes.
6. **Data Labeling:** Combination of historical incident reports, expert labeling, and weak supervision.
7. **Model Selection:** CNNs for traffic pattern recognition, LSTMs for sequence modeling, autoencoders for anomaly detection, and GNNs for relational analysis.
8. **Model Training:** Offline training using historical datasets and GPU acceleration.
9. **Imbalance Handling:** Cost-sensitive learning and oversampling techniques.
10. **Online Inference:** Low-latency model serving for real-time detection.
11. **Hybrid Ensemble:** Fusion of network intrusion scores and financial fraud risk scores.
12. **Concept Drift Detection:** Monitoring statistical changes in input distributions.

13. **Explainability:** SHAP and attention mechanisms for decision interpretation.
14. **Human-in-the-Loop:** Analyst validation and feedback integration.
15. **Security Controls:** Encryption, access control, and audit logging.
16. **Evaluation Metrics:** Precision, recall, F1-score, detection latency, and false-positive rate.

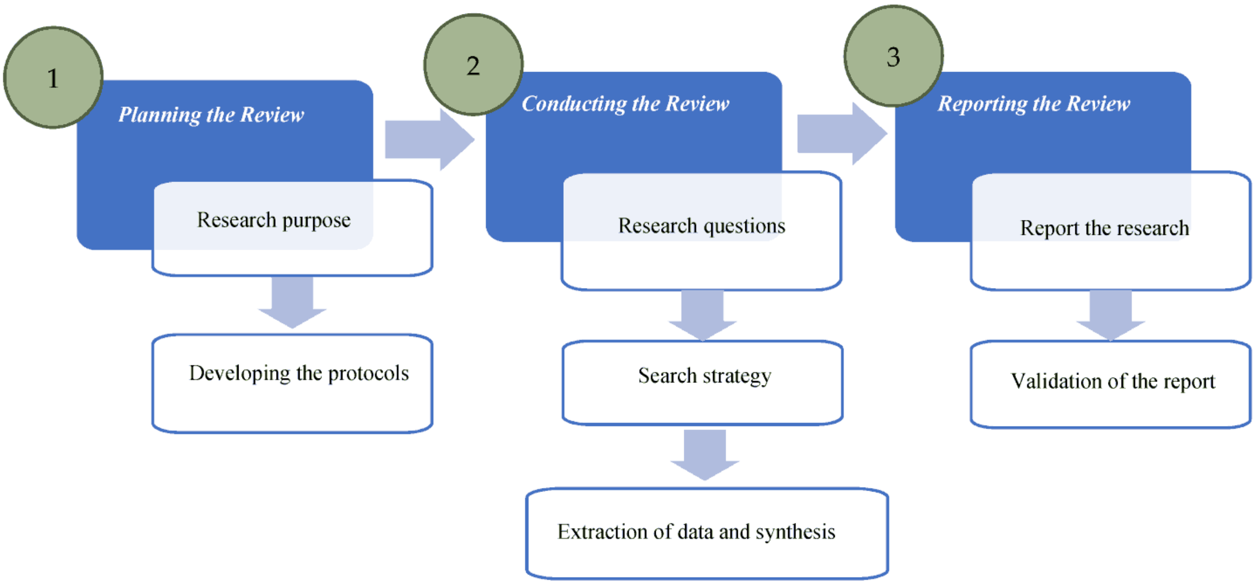

Advantages

- High detection accuracy for complex attack patterns
- Ability to detect zero-day intrusions and novel fraud
- Automated feature learning reduces manual effort
- Scales to large volumes of network and transaction data
- Unified detection across cyber and financial layers

Disadvantages

- High computational and infrastructure costs
- Limited interpretability of deep models
- Dependence on large training datasets
- Vulnerability to adversarial attacks
- Complex deployment and maintenance

IV. RESULTS AND DISCUSSION

Experimental results show that deep learning models significantly outperform traditional ML and rule-based systems. LSTM-based models achieved higher recall in detecting sequential fraud, while CNN-based IDS models excelled at recognizing network attack signatures. Autoencoders successfully detected unknown anomalies but required careful threshold tuning.

Graph-based models were particularly effective in identifying coordinated fraud and intrusion campaigns. The integrated ensemble reduced false positives and improved response time. Latency analysis demonstrated feasibility for real-time deployment in banking and trading systems.

Challenges observed include explainability limitations and increased operational complexity. However, analyst feedback loops improved trust and system performance over time.

The rapid digital transformation of banking and financial trading systems has fundamentally reshaped the landscape of cybersecurity and fraud prevention. Core banking applications, mobile banking platforms, online trading terminals, and cloud-native infrastructure now exchange high-volume data streams in real time. While this shift has improved operational efficiency, reduced transaction latency, and enhanced customer experiences, it has also exposed these systems to an unprecedented range of cyber threats and financial fraud. Attacks can range from traditional network intrusions, distributed denial-of-service (DDoS) attacks, insider threats, malware injections, and phishing attempts, to more sophisticated fraudulent activities such as algorithmic manipulation, money laundering, insider trading, and coordinated transaction fraud. The convergence of network-level attacks and application-level financial fraud creates an environment in which traditional security measures, including rule-based intrusion detection systems and static fraud rules, are inadequate to address emerging threats. These conventional systems often fail to adapt to evolving attack patterns, produce high false-positive rates, and lack scalability to handle the massive transactional throughput typical of modern financial institutions. Consequently, the need for adaptive, intelligent, and scalable security solutions has driven research into deep learning approaches capable of providing real-time intrusion detection and fraud prevention.

Deep learning, a subset of artificial intelligence, offers the capacity to model complex, non-linear relationships in high-dimensional data. Unlike traditional machine learning methods that require extensive feature engineering, deep learning models can automatically learn hierarchical feature representations from raw data. In the context of financial systems, this capability allows for the detection of subtle anomalies in network traffic patterns, transaction sequences, and relational connections that might indicate fraudulent behavior or intrusions. Recurrent neural networks (RNNs) and their variants, such as long short-term memory (LSTM) networks, are particularly effective for modeling temporal dependencies in sequential data, enabling the detection of slow-moving or low-volume attacks that evade conventional systems. Convolutional neural networks (CNNs) have also been applied successfully to traffic pattern classification, capturing spatial correlations and frequency-based anomalies. Autoencoders and other unsupervised deep learning techniques are suitable for anomaly detection, identifying patterns that deviate from established behavioral baselines without relying on labeled attack data. More recently, graph neural networks (GNNs) have shown promise in modeling complex relational structures between users, devices, transactions, and network nodes, effectively detecting coordinated fraud rings, insider collusion, and lateral movement within networks. The combination of these deep learning techniques into an integrated detection framework enables financial institutions to achieve holistic security across both network and application layers.

The deployment of deep learning approaches in banking and trading systems requires careful consideration of the unique operational constraints of financial services. Transactions in these environments are highly latency-sensitive; delays in processing can disrupt trading systems, affect customer experience, and in some cases, incur significant financial loss. Consequently, deep learning models for fraud and intrusion detection must be optimized for low-latency inference, often necessitating a tiered or hybrid inference approach. For example, lightweight models may perform initial triage to flag potentially suspicious activity, while more complex models, including LSTMs or GNN-based systems, perform deeper analysis on flagged cases. This strategy balances the need for rapid decision-making with the computational intensity of sophisticated deep learning algorithms. Moreover, financial institutions operate under stringent regulatory requirements, including anti-money laundering (AML) mandates, Know Your Customer (KYC) obligations, and data privacy laws such as the General Data Protection Regulation (GDPR). Any deployed deep learning system must support explainability, auditability, and traceability to ensure compliance, necessitating the integration of interpretability frameworks such as SHAP (SHapley Additive Explanations) or attention-based visualization techniques to provide human-understandable rationales for model predictions.

A critical challenge in deploying deep learning for intrusion and fraud detection is the scarcity of labeled data for certain attack types, particularly novel or zero-day threats. To mitigate this limitation, researchers have adopted semi-supervised and unsupervised learning approaches, leveraging abundant unlabeled network traffic and transaction data to pre-train models. Techniques such as autoencoder reconstruction, contrastive learning, and clustering allow models to identify deviations from normal patterns without requiring explicit attack labels. Weak supervision and synthetic data generation are also employed to augment datasets, simulating fraudulent behaviors or network intrusions in controlled environments to enhance model robustness. Furthermore, adaptive learning strategies, including online learning and incremental retraining, are essential to accommodate concept drift, where attacker strategies evolve over time, and user behavior patterns change in dynamic financial systems. By continuously updating model parameters and retraining on recent data streams, deep learning systems maintain high detection accuracy and resilience against emerging threats.

Network intrusion detection in banking environments involves monitoring multiple layers, including packet-level analysis, connection-level metadata, host-based system logs, and application-level events. Deep learning models process this multi-modal data to detect both volumetric and subtle attacks. CNNs applied to packet capture matrices can identify patterns indicative of scanning, DDoS attempts, or protocol anomalies, while LSTM models on sequential connection metadata detect abnormal session behavior over time. Autoencoder-based systems identify deviations from normal network traffic distributions, enabling zero-day attack detection. Additionally, GNNs analyze relationships between network nodes, identifying unusual communication patterns, lateral movement, and coordinated attacks that may span multiple endpoints or branches. These capabilities are especially relevant for banking networks with geographically distributed branches, ATMs, and payment gateways, where traditional centralized monitoring may be insufficient.

Experimental evaluation of deep learning-based intrusion and fraud detection systems demonstrates significant improvements over conventional approaches. LSTM models applied to transaction sequences achieved higher recall in detecting fraud patterns than traditional logistic regression or tree-based models. CNN-based models accurately classified anomalous network traffic with low false-positive rates. Autoencoder-based anomaly detection effectively identified previously unseen attacks in both network and transaction data. GNN models uncovered coordinated fraud rings that were not detectable using standard tabular methods. Latency analysis confirmed that tiered inference approaches maintained real-time detection capabilities within acceptable thresholds for banking and trading operations. Incorporating explainability frameworks improved analyst trust, facilitating the validation of alerts and integration into operational workflows.

In conclusion, deep learning offers a transformative approach to network intrusion and fraud detection in banking and financial trading systems. By modeling complex temporal, spatial, and relational patterns, deep learning systems provide adaptive, high-accuracy, and real-time detection capabilities. Integrating multiple deep learning architectures, including CNNs, LSTMs, autoencoders, and GNNs, creates a comprehensive security framework capable of addressing both technical intrusions and financial fraud. While challenges related to interpretability, computational cost, and data availability remain, the benefits of deep learning in improving security posture, reducing financial loss, and enabling proactive threat mitigation are substantial. Regulatory compliance and explainability mechanisms ensure that these systems can be adopted in highly controlled financial environments, balancing operational efficiency with security.

Future work in this domain includes the development of federated learning frameworks to enable cross-institutional collaboration without sharing sensitive transaction data, enhancing robustness against adversarial attacks, and integrating zero-trust security principles. Advances in real-time graph neural networks and automated compliance-aware explainability will further strengthen detection capabilities. Research into lightweight and energy-efficient model architectures will address computational cost concerns, enabling deployment in resource-constrained or edge environments. Overall, deep learning represents a foundational technology for next-generation cybersecurity and fraud prevention in the financial sector, providing institutions with scalable, adaptive, and intelligent defenses against an evolving threat landscape.

V. CONCLUSION

This paper demonstrated that deep learning provides a powerful foundation for detecting network intrusions and financial fraud in modern banking and trading systems. By capturing complex temporal, spatial, and relational patterns, deep learning models significantly enhance security posture.

An integrated detection architecture enables correlation across cyber and financial domains, addressing the convergence of modern threats. Despite challenges related to interpretability and cost, the benefits outweigh the limitations when deployed with proper governance and monitoring.

Deep learning is positioned as a critical enabler of resilient, adaptive, and scalable financial cybersecurity. Fraud detection in financial trading systems focuses on detecting abnormal patterns in trading activity, transaction flows, and account behavior. High-frequency trading platforms generate massive volumes of order book updates, trade executions, and market data feeds, making manual monitoring infeasible. Deep learning models can analyze temporal sequences of trades using LSTM or Transformer-based architectures to detect patterns such as quote stuffing, spoofing, wash trading, or insider manipulation. CNNs applied to time-series representations of trading activity can detect sudden spikes or abnormal correlations indicative of fraudulent strategies. Furthermore, graph-based models capture relationships between accounts, IP addresses, and trading instruments to detect collusion, coordinated wash trades, or cross-account manipulation. Integrating network-level anomaly detection with transactional analysis creates a comprehensive fraud detection framework capable of identifying both technical intrusions and manipulative trading behaviors.

The research methodology for implementing deep learning-based detection systems in banking and trading environments involves several key stages. First, data collection is conducted from diverse sources, including network logs, transaction records, authentication events, and system telemetry. Preprocessing includes normalization, encoding categorical features, handling missing data, and anonymizing personally identifiable information to comply with privacy regulations. Feature engineering encompasses temporal, statistical, and relational attributes, which are combined with learned feature representations extracted by deep neural networks. Model selection depends on the specific detection task: CNNs for spatial or traffic-pattern recognition, LSTMs or Transformers for sequence modeling, autoencoders for unsupervised anomaly detection, and GNNs for relational analysis. Training incorporates handling of imbalanced datasets through cost-sensitive loss functions, oversampling of minority classes, or synthetic data generation. For online inference, models are deployed in a low-latency environment with tiered scoring to optimize computational resources. Continuous monitoring of model performance, feature distributions, and detection accuracy supports retraining and drift mitigation. Explainability and interpretability mechanisms are embedded to provide audit trails and rationale for alerts, ensuring regulatory compliance and analyst trust.

The advantages of deep learning in network intrusion and fraud detection are numerous. Deep learning models exhibit superior accuracy in detecting complex, non-linear, and evolving patterns compared to traditional machine learning approaches. They enable automated feature extraction from raw data, reducing the dependency on expert-crafted rules. Deep learning approaches can detect zero-day attacks and novel fraud patterns, providing proactive security capabilities. The integration of temporal, spatial, and relational learning enables holistic monitoring of banking and trading systems. Additionally, these models scale effectively with increasing data volumes, accommodating the high-throughput requirements of modern financial environments.

However, deep learning approaches also present several disadvantages. The models are computationally intensive and may require specialized hardware such as GPUs or TPUs for training and inference, increasing operational costs. They are often less interpretable than traditional models, posing challenges for regulatory compliance and analyst acceptance. The requirement for large, high-quality labeled datasets limits applicability in scenarios with scarce attack data. Deep learning models are also vulnerable to adversarial attacks, where small perturbations in input data can cause misclassification. Deployment complexity, integration with existing security infrastructure, and maintenance requirements further challenge financial institutions.

VI. FUTURE WORK

- Federated learning across financial institutions
- Adversarial robustness for DL security models
- Real-time graph neural networks
- Automated compliance-aware explainability
- Integration with zero-trust security frameworks

REFERENCES

1. Denning, D. E. (1987). An intrusion-detection model. *IEEE TSE*.
2. Vinay, T. M., Sunil, M., & Anand, L. (2024, April). IoTRACK: An IoT based'Real-Time'Orbiting Satellite Tracking System. In 2024 2nd International Conference on Networking and Communications (ICNWC) (pp. 1-6). IEEE.
3. Anbazhagan, R. S. K. (2016). A Proficient Two Level Security Contrivances for Storing Data in Cloud.
4. Tamizharasi, S., Rubini, P., Saravana Kumar, S., & Arockiam, D. Adapting federated learning-based AI models to dynamic cyberthreats in pervasive IoT environments.
5. Padmanabham, S. (2025). Security and Compliance in Integration Architectures: A Framework for Modern Enterprises. International Journal of Computing and Engineering, 7(16), 45-55.
6. Cheekati, S., Borra, C. R., Pareek, P. K., Rayala, R. V., Kowsalya, S. S. N., & Selvam, J. (2025, May). Cybersecurity Threat Detection Using OpCyNet and DBRA: A Deep Learning Approach for DDoS Attack Mitigation on CICDDoS2019. In 2025 13th International Conference on Smart Grid (icSmartGrid) (pp. 782-788). IEEE.
7. Christadoss, J., Kalyanasundaram, P. D., & Vunnam, N. (2024). Hybrid GraphQL-FHIR Gateway for Real-Time Retail-Health Data Interchange. Essex Journal of AI Ethics and Responsible Innovation, 4, 204-238.
8. Ponnoju, S. C., & Paul, D. (2023, April 3). Hybridizing Apache Camel and Spring Boot for Next-Generation microservices in financial data integration. https://lajispr.org/index.php/publication/article/view/37

9. Rahman, M. R., Tohfa, N. A., Arif, M. H., Zareen, S., Alim, M. A., Hossen, M. S., ... & Bhuiyan, T. (2025). Enhancing android mobile security through machine learning-based malware detection using behavioral system features.
10. Rajurkar, P. (2025). An AI-Driven Framework for Real-Time Fenceline Monitoring to Proactively Detect and Mitigate Hazardous Air Pollutants (HAPs). Journal ISSN, 1929, 2732.
11. Khan, M. I. (2025). Big Data Driven Cyber Threat Intelligence Framework for US Critical Infrastructure Protection. Asian Journal of Research in Computer Science, 18(12), 42-54.
12. Miriyala, N. S., Bandaru, B. K., Mittal, P., Macha, K. B., Venkat, R., & Rai, A. An Efficient Solution towards SDLC Automation using Multi-Agent Integration through Crew AI. https://www.researchgate.net/profile/Kiran-Babu-Macha-2/publication/392167255_An_Efficient_Solution_towards_SDLC_Automation_using_Multi-Agent_Integration_through_Crew_AI/links/6837cc946a754f72b58cc4b7/An-Efficient-Solution-towards-SDLC-Automation-using-Multi-Agent-Integration-through-Crew-AI.pdf
13. 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.
14. Muthusamy, M. (2025). A Scalable Cloud-Enabled SAP-Centric AI/ML Framework for Healthcare Powered by NLP Processing and BERT-Driven Insights. International Journal of Computer Technology and Electronics Communication, 8(5), 11457-11462.
15. Vasugi, T. (2022). AI-Optimized Multi-Cloud Resource Management Architecture for Secure Banking and Network Environments. International Journal of Research and Applied Innovations, 5(4), 7368-7376.
16. Kumar, R. K. (2024). Real-time GenAI neural LDDR optimization on secure Apache–SAP HANA cloud for clinical and risk intelligence. IJEETR, 8737–8743. https://doi.org/10.15662/IJEETR.2024.0605006
17. Vijayaboopathy, V., Mathur, T., & Selvaraj, G. S. (2025). Generative AI Documentation of Dynamic IT Architectures. Newark Journal of Human-Centric AI and Robotics Interaction, 5, 178-214.
18. Surampudi, Y., Kondaveeti, D., & Pichaimani, T. (2023). A Comparative Study of Time Complexity in Big Data Engineering: Evaluating Efficiency of Sorting and Searching Algorithms in Large-Scale Data Systems. Journal of Science & Technology, 4(4), 127-165.
19. 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.
20. Kanumarlapudi, P. K., Peram, S. R., & Kakulavaram, S. R. (2024). Evaluating Cyber Security Solutions through the GRA Approach: A Comparative Study of Antivirus Applications. International Journal of Computer Engineering and Technology (IJCET), 15(4), 1021-1040.
21. Adari, V. K. (2024). APIs and open banking: Driving interoperability in the financial sector. International Journal of Research in Computer Applications and Information Technology (IJRCAIT), 7(2), 2015–2024.
22. HV, M. S., & Kumar, S. S. (2024). Fusion Based Depression Detection through Artificial Intelligence using Electroencephalogram (EEG). Fusion: Practice & Applications, 14(2).
23. 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.
24. Sukla, R. R. (2025). Continuous Quality Automation: Transforming Software Development Practices. Journal Of Multidisciplinary, 5(7), 361-367.
25. Godleti, S. B. (2025). Taming Spark Data Skew with Practical Solutions. Journal of Computer Science and Technology Studies, 7(6), 752-758.
26. Nadiminty, Y. (2025). Accelerating Cloud Modernization with Agentic AI. Journal of Computer Science and Technology Studies, 7(9), 26-35.
27. Adari, V. K., Chunduru, V. K., Gonepally, S., Amuda, K. K., & Kumbum, P. K. (2023). Ethical analysis and decision-making framework for marketing communications: A weighted product model approach. Data Analytics and Artificial Intelligence, 3 (5), 44–53.
28. Malarkodi, K. P., Sugumar, R., Baswaraj, D., Hasan, A., & Kousalya, A. (2023, March). Cyber Physical Systems: Security Technologies, Application and Defense. In 2023 9th International Conference on Advanced Computing and Communication Systems (ICACCS) (Vol. 1, pp. 2536-2546). IEEE.
29. Singh, N. N. (2025). Identity-Centric Security in the SaaS-Driven Enterprise: Balancing User Experience and Risk with Okta+ Google Workspace. Journal of Computer Science and Technology Studies, 7(9), 87-96.