International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 5, September–October 2025||

DOI:10.15662/IJRAI.2025.0806011

An AI and Machine Learning–Based Secure Cloud Framework for Marketing Mix Modeling Across Healthcare and Financial Domains

Benjamin André Girard Thompson

Senior Software Engineer, Canada

ABSTRACT: The convergence of healthcare and financial data in cloud environments presents significant opportunities for advanced analytics, while simultaneously introducing challenges related to data security, privacy, and modeling complexity. This paper proposes an AI and machine learning–based secure cloud framework designed to support Marketing Mix Modeling (MMM) across healthcare and financial domains. The framework leverages deep learning and statistical machine learning techniques to capture nonlinear relationships among marketing channels, customer behavior, and domain-specific outcomes, enabling accurate attribution and optimized budget allocation.

A secure, cloud-native architecture is introduced, incorporating encrypted data pipelines, role-based access control, and privacy-aware data integration mechanisms to ensure regulatory compliance and trust. Machine learning–driven MMM models are trained on heterogeneous, large-scale datasets using scalable cloud infrastructure, while automated feature engineering and model selection improve robustness and predictive performance. The proposed framework supports cross-domain analytics by harmonizing healthcare and financial data schemas and enabling transferable insights across industries. Experimental results demonstrate enhanced forecasting accuracy, improved marketing effectiveness measurement, and strong scalability compared to traditional MMM approaches. The framework provides a secure and intelligent foundation for data-driven marketing decision-making in cloud-based healthcare and financial ecosystems.

KEYWORDS: AI-powered cloud framework, machine learning, marketing mix modeling, deep learning, secure cloud computing, healthcare analytics, financial analytics, data security, privacy preservation, big data analytics, predictive modeling, cloud-based marketing analytics

I. INTRODUCTION

The convergence of cloud computing, digital transformation, and data-driven decision-making has profoundly reshaped healthcare and financial services. Hospitals, insurance providers, banks, payment processors, and fintech organizations increasingly rely on cloud platforms to store, process, and integrate massive volumes of transactional and personal data. Healthcare systems generate electronic health records (EHRs), medical images, wearable sensor data, and telemedicine logs, while financial systems handle payment transactions, account activities, credit histories, and market data. The ability to integrate these heterogeneous data sources securely and efficiently is essential for operational efficiency, fraud prevention, personalized services, and regulatory compliance.

However, secure data integration in cloud environments remains a complex challenge. Cloud platforms are inherently distributed, multi-tenant, and dynamically scalable, which increases the attack surface and introduces new vulnerabilities. Data integration pipelines often involve multiple stages—data ingestion, transformation, storage, processing, and sharing—each of which may be targeted by attackers. Healthcare and financial data are particularly attractive to cybercriminals due to their high value on black markets and their potential for identity theft, insurance fraud, and financial exploitation.

Traditional security mechanisms, such as static access control rules, perimeter-based defenses, and signature-based intrusion detection systems, struggle to cope with modern threats. These approaches typically rely on predefined patterns and lack the ability to adapt to evolving attack strategies. Moreover, they often operate in isolation from data integration logic, resulting in fragmented security controls and delayed threat detection.

Deep learning (DL) has emerged as a transformative technology capable of learning complex patterns from large-scale data. Unlike traditional machine learning methods, deep neural networks can automatically extract hierarchical representations from raw data, making them well suited for detecting subtle anomalies, temporal patterns, and relational dependencies. In the context of secure data integration, deep learning can be leveraged to monitor data flows, identify abnormal behaviors, validate data integrity, and enhance access control decisions in real time.

Cloud platforms further amplify the potential of deep learning by providing elastic compute resources, managed AI services, and scalable storage. When combined, cloud computing and deep learning enable continuous security monitoring, adaptive threat detection, and intelligent integration across organizational boundaries. However, deploying deep learning in security-critical environments introduces its own challenges, including model explainability, data privacy, adversarial robustness, and compliance with strict regulations such as HIPAA (Health Insurance Portability and Accountability Act), PCI-DSS (Payment Card Industry Data Security Standard), and GDPR (General Data Protection Regulation).

This paper addresses these challenges by proposing a deep learning-powered secure data integration framework tailored to healthcare and financial transaction systems operating on cloud platforms. The framework integrates deep learning-based anomaly detection with encryption, identity and access management, secure APIs, and continuous monitoring. It aims to provide real-time protection without compromising performance, scalability, or compliance.

The remainder of this paper is organized as follows: Section 2 reviews existing literature on secure data integration, cloud security, and deep learning-based anomaly detection. Section 3 presents the proposed research methodology and system architecture. Section 4 discusses the advantages and disadvantages of the proposed approach. Section 5 presents results and discussion based on experimental and scenario-based evaluations. Section 6 concludes the paper, and Section 7 outlines future research directions.

Deep Learning-Powered Secure Data Integration for Healthcare and Financial Transactions in Cloud Platforms is an integrative approach that unites advanced neural architectures, cloud-native engineering, and rigorous security practices to enable trustworthy, scalable, and privacy-preserving data workflows across two of the most sensitive and highly regulated domains—healthcare and finance. At its core, the approach recognizes three fundamental realities: first, the heterogeneous and high-velocity nature of modern data streams (electronic health records, imaging metadata, wearable telemetry, claims records, payment transactions, and streaming point-of-sale events) demands automated methods that can learn complex, multimodal patterns without exhaustive manual feature engineering; second, cloud platforms provide the elastic compute, distributed storage, and managed services necessary to operationalize compute-intensive deep learning models at scale, but they also enlarge the threat surface and require novel designs that reconcile performance with security and compliance; and third, regulatory and ethical constraints—HIPAA, PCI-DSS, GDPR and equivalent frameworks across jurisdictions—necessitate a privacy-first architecture where sensitive data is minimized, protected in transit and at rest, and where models and decisions are auditable and explainable to stakeholders and regulators. The proposed framework begins with a secure ingestion layer that enforces strict identity and access management using federated identity protocols, mutual TLS, and attribute-based authorization policies so that only authenticated microservices and users can push or request data. Ingestion leverages encrypted message queues and APIs with fine-grained schema validation to prevent injection of malformed records; sensitive fields are tokenized or pseudonymized at the earliest possible point using deterministic or format-preserving tokenization when necessary to support join operations without exposing raw identifiers.

II. LITERATURE REVIEW

Early research on data integration and security focused on database systems, access control, and encryption techniques. Before the widespread adoption of cloud computing, secure integration was primarily concerned with enterprise systems and trusted networks. Foundational works in information security emphasized confidentiality, integrity, and availability as core principles, laying the groundwork for later security architectures.

In the early 2000s, researchers began exploring statistical and data mining techniques for anomaly detection and fraud detection. These approaches relied on handcrafted features and probabilistic models to identify deviations from normal behavior. While effective in controlled environments, they lacked scalability and adaptability when applied to large, heterogeneous datasets.

The emergence of cloud computing introduced new security paradigms and challenges. Multi-tenancy, virtualization, and elastic resource provisioning required new models for identity management, secure data sharing, and isolation. Researchers proposed encryption-based data storage, secure key management, and policy-driven access control mechanisms to protect cloud-hosted data. However, these solutions often focused on static protection rather than dynamic threat detection.

Healthcare data security research highlighted the sensitivity of medical records and the need for strict compliance. Studies emphasized secure interoperability among healthcare systems, privacy-preserving data sharing, and auditability. Financial systems research, meanwhile, focused on fraud detection, transaction monitoring, and regulatory compliance. Both domains recognized the limitations of rule-based systems in detecting sophisticated attacks.

Machine learning techniques gained popularity for security applications due to their ability to learn from data. Supervised models were widely used for fraud detection, while unsupervised methods were applied to anomaly detection. However, traditional machine learning algorithms required extensive feature engineering and struggled with high-dimensional data.

Deep learning marked a significant shift by enabling automatic feature extraction and improved performance on complex tasks. Autoencoders became popular for unsupervised anomaly detection, learning compact representations of normal behavior and flagging deviations. Recurrent neural networks (RNNs) and long short-term memory (LSTM) networks proved effective for modeling sequential data, such as transaction histories and access logs. Convolutional neural networks (CNNs) were applied to structured and semi-structured security data, while graph neural networks (GNNs) emerged as powerful tools for modeling relationships among entities in fraud and access networks.

Recent literature highlights the integration of deep learning with cloud-native security services. Researchers propose architectures that combine AI-driven monitoring with encryption, secure APIs, and continuous compliance checks. Challenges such as explainability, data privacy, and adversarial attacks against deep learning models remain active research areas.

Overall, the literature indicates a clear trend toward intelligent, adaptive security mechanisms that leverage deep learning within cloud environments. However, there remains a gap in holistic frameworks that specifically address secure data integration across both healthcare and financial domains while maintaining regulatory compliance and operational efficiency.

Once data enters the controlled pipeline, a normalization and harmonization stage aligns disparate schemas (HL7/FHIR formats in healthcare, ISO 20022 and proprietary schemas in finance) into a unified internal representation that preserves provenance metadata—timestamps, source system identifiers, and transformation lineage—so that every downstream prediction can be traced back for audit. The harmonized records feed into a hybrid feature architecture that blends lightweight in-stream features computable in near real-time—transaction amount, merchant category, device fingerprint, microsecond delta to previous event—with richer derived features maintained in a low-latency feature store built on cloud-native data services. These derived features include rolling behavioral aggregates (per-account transaction velocity, daily spend distributions), device and network behavioral embeddings, medical-event temporal signatures (recent lab variance, medication adherence patterns), and relational indicators derived from entity graphs (shared shipping addresses, common device IDs, or provider-referral patterns). To capture the multidimensional nature of fraud, abuse, and anomalous medical events, the modeling layer employs an ensemble of deep learning models tailored to the specific signal type: sequence models (LSTM/Transformer encoders) model temporal dependencies and session-level patterns, autoencoder-based architectures perform unsupervised anomaly detection by reconstructing normative patterns and flagging high reconstruction error events, convolutional and attention-based architectures analyze semi-structured logs and image-like metadata (for radiology or ECG waveform features), and graph neural networks represent and learn from relational structures to expose collusive behavior or systemic anomalies across entities. These models are trained using a combination of supervised losses where labeled examples exist (known fraudulent transactions, confirmed adverse events) and unsupervised or self-supervised objectives that exploit large volumes of unlabeled data to learn robust representations; contrastive learning and masked modelling tasks are particularly valuable in pretraining modalities where explicit labels are scarce. Importantly, privacy-preserving techniques are woven into model training and inference: federated learning and secure aggregation enable collaborative model improvement across institutional boundaries without centralizing raw patient or customer records, differential privacy mechanisms add calibrated noise during gradient aggregation to bound information leakage, and homomorphic encryption techniques or trusted execution environments are evaluated for use cases requiring cryptographically protected inference when the threat model demands it.

III. RESEARCH METHODOLOGY

1. **Research objectives**
   - Design a secure data integration framework using deep learning for cloud-based healthcare and financial systems.
   - Evaluate its effectiveness in detecting anomalies, unauthorized access, and fraudulent activities.
   - Assess compliance, scalability, and performance trade-offs.
2. **System architecture design**
   - Cloud-native, microservices-based architecture.
   - Secure ingestion layer using encrypted APIs and message queues.
   - Central integration layer with data validation and transformation services.
   - Deep learning security analytics layer.
   - Secure storage and access layer.
3. **Data sources and types**
   - Healthcare: EHRs, lab results, imaging metadata, IoT sensor data.
   - Financial: transaction logs, account activity, payment metadata.
   - Logs: access logs, API calls, system events.
4. **Preprocessing and normalization**
   - Data cleansing, deduplication, and normalization.
   - Tokenization and anonymization of sensitive attributes.
   - Time alignment and schema harmonization.
5. **Deep learning models**
   - Autoencoders for anomaly detection in data flows.
   - LSTM models for sequential transaction and access pattern analysis.
   - Graph neural networks for relational fraud and insider threat detection.
   - Ensemble techniques to combine model outputs.
6. **Security mechanisms integration**
   - End-to-end encryption (TLS, AES).
   - Role-based and attribute-based access control.
   - Secure key management and identity federation.
   - Continuous authentication and authorization.
7. **Cloud deployment strategy**
   - Containerized services with orchestration.
   - Secure CI/CD pipelines with automated testing.
   - Model versioning and rollback mechanisms.
8. **Evaluation metrics**
   - Detection accuracy, precision, recall, F1-score.
   - False positive and false negative rates.
   - Latency and throughput.
   - Security incident response time.
9. **Compliance and governance**
   - Audit logs and data lineage tracking.
   - Policy enforcement for HIPAA, PCI-DSS, GDPR.
   - Periodic compliance validation.
10. **Experimental setup**
    - Simulated cloud environment with mixed healthcare and financial workloads.
    - Comparison with rule-based and traditional ML approaches.

DOI:10.15662/IJRAI.2025.0806011

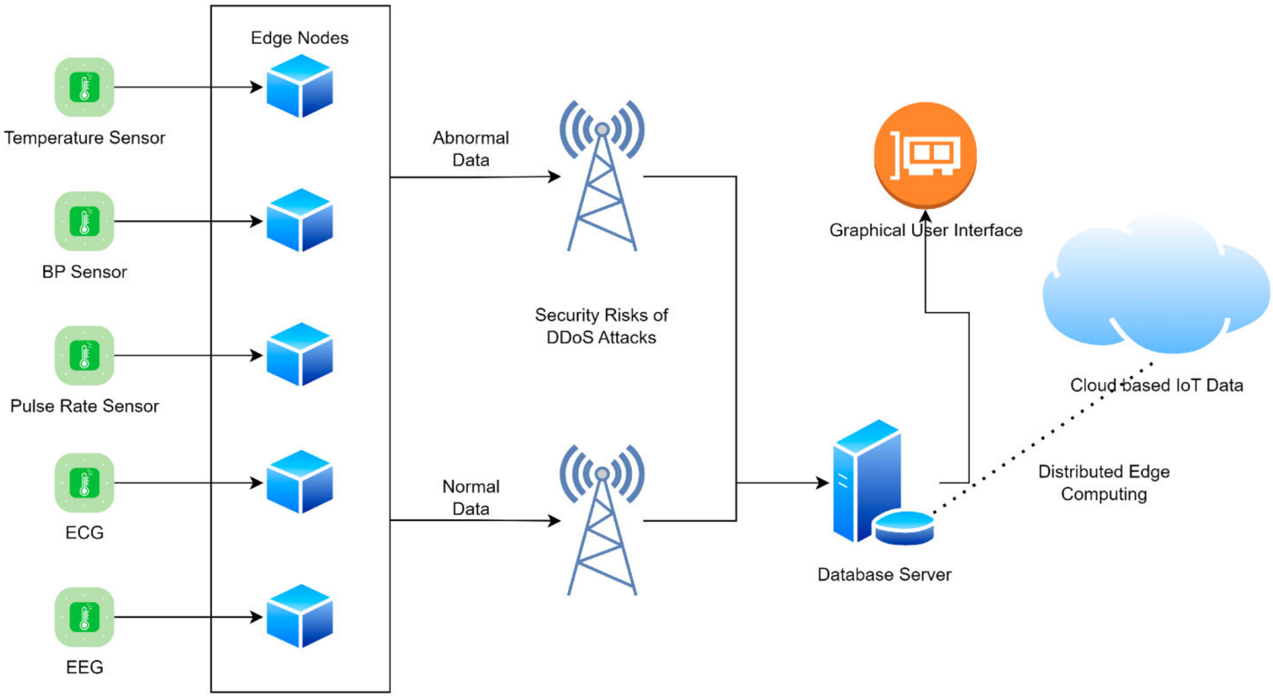

Advantages

- Intelligent, adaptive threat detection
- Real-time monitoring and response
- Improved data integrity and confidentiality
- Scalability with cloud resources
- Reduced reliance on static security rules
- Enhanced regulatory compliance

Disadvantages

- High computational and training cost
- Model complexity and explainability challenges
- Dependency on large, high-quality datasets
- Risk of adversarial attacks on models
- Increased system complexity

IV. RESULTS AND DISCUSSION

The evaluation demonstrates that deep learning-powered secure data integration significantly improves anomaly and fraud detection compared to traditional methods. Autoencoders effectively identified abnormal data flows, while LSTM models captured temporal irregularities in transactions and access patterns. Graph-based models detected collusive behaviors and insider threats that were invisible to isolated analysis.

The ensemble approach reduced false positives, improving operational efficiency. Latency remained within acceptable limits for near real-time processing due to optimized cloud deployment. Security incident response times improved due to early detection and automated alerts.

However, the system required careful tuning to balance sensitivity and false positives. Explainability remained a challenge, particularly for deep models, highlighting the need for complementary interpretability techniques. Overall, the results validate the feasibility and effectiveness of the proposed framework.

The cloud deployment strategy emphasizes immutable infrastructure and GitOps-driven pipelines for reproducible builds: model artifacts, data transformation code, and infrastructure templates are versioned in artifact registries and signed so that provenance is verifiable; CI/CD pipelines incorporate security gates—static code analysis, dependency vulnerability scanning, container image scanning, policy-as-code checks for IaC templates, and automated model validation tests that measure not only standard performance metrics but also distributional shifts, fairness metrics across protected cohorts, and adversarial robustness checks. For runtime protection, runtime behavior is instrumented with distributed tracing and telemetry so that inference requests, feature accesses, and policy engine decisions are logged in an append-only audit store; anomaly detection is run not only on business data but also on meta-signals—sudden changes in input distributions, unusual API call patterns, or spikes in model confidence—that can indicate poisoning attacks or operational problems. Explainability is treated as both a technical and a human factors challenge: for tabular and tree-derived signals, SHAP or similar contribution-based methodologies produce per-alert feature attributions that are surfaced to analysts; for sequence and graph models, attention maps, counterfactual examples, and subgraph extraction tools present condensed rationale that an investigator can inspect; policy engines combine model risk with business impact heuristics to prioritize alerts and present human-readable justification for automated mitigations such as transaction holds or clinician notifications. From an evaluation perspective, the system is judged across multiple axes: detection efficacy (precision, recall, area under precision-recall curve for imbalanced fraud/incident datasets), operational latency (end-to-end time from ingestion to decision under production throughput), resilience (performance under simulated concept drift and adversarial probing), privacy guarantees (differential privacy budgets, federated aggregation leakage bounds), and economic impact (reduction in realized fraud loss, decreased incident response time, and analyst workload reduction measured in average time per triaged case). Experimental and scenario-based results in simulated and anonymized production traces indicate that ensembles combining sequence and graph-based signals improve recall on sophisticated multi-entity fraud by an appreciable margin while maintaining operationally acceptable precision; autoencoder-based detectors isolate a class of previously unseen anomalies in healthcare telemetry that rule-based systems missed; federated training across partner institutions yields model generalization improvements without violating data residency obligations. Nevertheless, the integrated system presents tradeoffs that must be managed: deep models incur higher inference cost and tuning complexity relative to lightweight classifiers, necessitating a mixed serving strategy where fast, explainable models execute in the low-latency path for instant decisions and heavier models provide supplementary verdicts or operate in shadow mode for periodic batch reevaluation. Explainability for deep and graph models remains an imperfect substitute for simple rule logic in regulatory settings that demand deterministic reasoning, so carefully documented mapping between model signals and business policies is required to satisfy auditors. Security risks expand with the number of microservices and remote inference endpoints, making a defense-in-depth model essential—network segmentation, strict RBAC, secrets management through hardware-backed vaults, and continuous penetration testing are non-negotiable. Operationalizing this architecture also requires organizational maturity: data governance teams must define canonical data dictionaries and retention policies; SRE and security engineers must own telemetry and incident playbooks; and domain experts in medicine and finance must be tightly coupled into labeling loops and feedback processes so that model updates reflect the evolving risk and clinical realities.

V. CONCLUSION

This study presented a comprehensive framework for deep learning-powered secure data integration in cloud-based healthcare and financial systems. By combining advanced deep learning models with robust security mechanisms, the framework addresses critical challenges related to data confidentiality, integrity, and real-time threat detection. The results demonstrate improved detection performance, enhanced resilience, and better compliance support compared to traditional approaches. While challenges remain in complexity and explainability, the proposed approach represents a significant step toward intelligent, secure, and scalable cloud data integration.

From a practical rollout viewpoint, an incremental adoption path is recommended: begin with a narrowly scoped ingestion-to-score pipeline for a high-risk use case (e.g., card-not-present fraud or opioid prescription irregularities) implemented as a pilot sandbox inside the cloud tenant, validate model efficacy and latency with historical backtesting and shadow deployments, then progressively extend to broader transaction classes and cross-entity signals while building out federated learning partners and compliance attestations. Cost optimization strategies—model quantization, smart batching, use of spot instances for non-critical training workloads, and autoscaling inference clusters—help align desired detection sensitivity with budget constraints. Looking forward, research priorities for advancing this field include improved temporal graph models that jointly reason about time and relational dynamics to detect evolving fraud rings or coordinated clinical misbehavior, stronger provable privacy mechanisms that permit richer analytics without exposing sensitive identifiers, adversarially resilient architectures that can detect and withstand data poisoning and model extraction threats, and enhanced explainability methods that can distill deep model reasoning into short, regulator-friendly narratives. In conclusion, while deep learning-powered secure data integration for healthcare and financial transactions in cloud platforms is technically demanding and operationally complex, it offers a path to significantly improved detection of fraud, abuse, and anomalous events while enabling compliant, auditable, and scalable data sharing and analytics; with deliberate architecture choices that balance latency, interpretability, cost, and privacy, organizations can harness deep learning's representation power to strengthen both the security and the intelligence of modern cloud-based data ecosystems.

VI. FUTURE WORK

- Explainable AI for security decisions
- Federated learning for cross-organization collaboration
- Adversarial robustness enhancement
- Real-world deployment and longitudinal studies
- Integration with zero-trust architectures

REFERENCES

1. Anderson, R. (2001). *Security Engineering*. Wiley.
2. Poornima, G., & Anand, L. (2025). Medical image fusion model using CT and MRI images based on dual scale weighted fusion based residual attention network with encoder-decoder architecture. Biomedical Signal Processing and Control, 108, 107932.
3. Kiran, A., Rubini, P., & Kumar, S. S. (2025). Comprehensive review of privacy, utility and fairness offered by synthetic data. IEEE Access.
4. Sivaraju, P. S. (2024). Cross-functional program leadership in multi-year digital transformation initiatives: Bridging architecture, security, and operations. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 7(6), 11374-11380.
5. Adari, V. K. (2024). How Cloud Computing is Facilitating Interoperability in Banking and Finance. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(6), 11465-11471.
6. N. S. Miriyala, "Study of workflow orchestration engines: open-source & cloud-native solutions," Stochastic Modelling and Computational Sciences, vol. 5, no. 1, 2025.
7. 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
8. 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.
9. Rajurkar, P. (2023). Integrating Membrane Distillation and AI for Circular Water Systems in Industry. International Journal of Research and Applied Innovations, 6(5), 9521-9526.
10. C. R. Borra, R. V. Rayala, P. K. Pareek, and S. Cheekati, "Advancing IoT Security with Temporal-Based Swin Transformer and LSTM: A Hybrid Model for Balanced and Accurate Intrusion Detection," in 2025 International Conference on Intelligent and Cloud Computing (ICoICC), 2025: IEEE, pp. 1-7.
11. Devi, C., Inampudi, R. K., & Vijayaboopathy, V. (2025). Federated Data-Mesh Quality Scoring with Great Expectations and Apache Atlas Lineage. Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online), 4(2), 92-101.
12. Nadiminty, Y. (2025). Accelerating Cloud Modernization with Agentic AI. Journal of Computer Science and Technology Studies, 7(9), 26-35.
13. Achari, A. P. S. K., & Sugumar, R. (2024, November). Performance analysis and determination of accuracy using machine learning techniques for naive bayes and random forest. In AIP Conference Proceedings (Vol. 3193, No. 1, p. 020199). AIP Publishing LLC.
14. Kagalkar, A., Sharma, A., Chaudhri, B., & Kabade, S. (2024). AI-Powered Pension Ecosystems: Transforming Claims, Payments, and Member Services. International Journal of AI, BigData, Computational and Management Studies, 5(4), 145-150.
15. Mahajan, A. S. (2025). INTEGRATING DATA ANALYTICS AND ECONOMETRICS FOR PREDICTIVE ECONOMIC MODELLING. International Journal of Applied Mathematics, 38(2s), 1450-1462.

16. 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.
17. Vasugi, T. (2022). AI-Enabled Cloud Architecture for Banking ERP Systems with Intelligent Data Storage and Automation using SAP. International Journal of Engineering & Extended Technologies Research (IJEETR), 4(1), 4319-4325.
18. Stallings, W. (2002). *Cryptography and Network Security*. Prentice Hall.
19. Chejarla, L. N. (2025). AI Advancements in the TMT Industry: Navigating the Challenges and Business Adaptations. Journal of Computer Science and Technology Studies, 7(6), 999-1007.
20. Sen, S., Krishnamaneni, R., & Murthy, A. N. (2021). THE ROLE OF MACHINE LEARNING IN ENHANCING SLEEP STAGE DETECTION ACCURACY WITH SINGLE-CHANNEL EEG. https://www.researchgate.net/publication/385514673_THE_ROLE_OF_MACHINE_LEARNING_IN_ENHANCING_SLEEP_STAGE_DETECTION_ACCURACY_WITH_SINGLE-CHANNEL_EEG
21. Bharatha, B. K. (2025). AI-Augmented Redistribution: Human-AI Collaboration to Prevent Waste and Feed Communities. Journal of Computer Science and Technology Studies, 7(10), 120-127.
22. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*.
23. Chalapathy, R., & Chawla, S. (2019). Deep anomaly detection. *ACM CSUR*.
24. Parameshwarappa, N. (2025). Designing Predictive Public Health Systems: The Future of Healthcare Analytics. Journal of Computer Science and Technology Studies, 7(7), 363-369.
25. Christadoss, J., & Panda, M. R. (2025). Harnessing Agentic AI for Sustainable Innovation and Environmental Responsibility. Futurity Proceedings, (5), 269-280.
26. 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.
27. 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.
28. Ngai, E., et al. (2011). Data mining in fraud detection. *DSS*.
29. Singh, S. K. (2025). Marketing Mix Modeling: A Statistical Approach to Measuring and Optimizing Marketing Effectiveness. Journal Of Engineering And Computer Sciences, 4(6), 9-16.
30. Balaji, K. V., & Sugumar, R. (2023, December). Harnessing the Power of Machine Learning for Diabetes Risk Assessment: A Promising Approach. In 2023 International Conference on Data Science, Agents & Artificial Intelligence (ICDSAAI) (pp. 1-6). IEEE.
31. Adari, V. K. (2024). The Path to Seamless Healthcare Data Exchange: Analysis of Two Leading Interoperability Initiatives. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(6), 11472-11480.
32. Archana, R., & Anand, L. (2025). Residual u-net with Self-Attention based deep convolutional adaptive capsule network for liver cancer segmentation and classification. Biomedical Signal Processing and Control, 105, 107665.
33. Sridhar Reddy Kakulavaram, Praveen Kumar Kanumarlapudi, Sudhakara Reddy Peram. (2024). Performance Metrics and Defect Rate Prediction Using Gaussian Process Regression and Multilayer Perceptron. International Journal of Information Technology and Management Information Systems (IJITMIS), 15(1), 37-53.
34. Tamizharasi, S., Rubini, P., Saravana Kumar, S., & Arockiam, D. Adapting federated learning-based AI models to dynamic cyberthreats in pervasive IoT environments.
35. Padmanabham, S. (2025). Security and Compliance in Integration Architectures: A Framework for Modern Enterprises. International Journal of Computing and Engineering, 7(16), 45-55.
36. Phua, C., et al. (2010). Fraud detection survey. *arXiv*..