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 5, Issue 5, September–October 2022||

DOI:10.15662/IJRAI.2022.0505005

Secure and Intelligent Targeted Advertising for Healthcare ERP Platforms: An LLM-Enabled Cloud and Web Engineering Approach

Felix Reinhard Blumenthal

Senior Software Engineer, Germany

ABSTRACT: The rapid adoption of Healthcare Enterprise Resource Planning (ERP) platforms has created new opportunities for data-driven and targeted digital advertising. However, healthcare environments demand high standards of security, privacy, and regulatory compliance. This paper proposes a **secure and intelligent targeted advertising framework** for healthcare ERP platforms using a **Large Language Model (LLM)–enabled cloud and web engineering approach**. The proposed system integrates secure cloud infrastructure, role-based data access, and privacy-preserving analytics to ensure compliance with healthcare regulations while enabling precise audience segmentation. LLMs are leveraged to analyze structured and unstructured ERP data, generate context-aware advertising insights, and optimize campaign strategies in real time. Advanced digital marketing analytics further enhance decision-making by measuring engagement, conversion performance, and campaign effectiveness. The framework demonstrates how intelligent automation and scalable cloud architecture can improve advertising relevance, operational efficiency, and trust within healthcare ERP ecosystems.

KEYWORDS: Healthcare ERP, Targeted Advertising, Large Language Models (LLMs), Cloud Computing, Web Engineering, Data Security, Privacy-Preserving Analytics, Digital Marketing Analytics, Intelligent Systems

I. INTRODUCTION

Retail has transformed dramatically over the past decade. Brick-and-mortar stores now operate in concert with online marketplaces, mobile checkout, third-party logistics, subscription services, and integrated loyalty and healthcare product ecosystems. This interconnectedness has been a boon for consumer convenience and commercial scale, but it also expands the potential vectors for cyber-attack and fraud: credential stuffing on e-commerce portals, point-of-sale (POS) skimming, compromise of partner logistics APIs, fraud rings coordinating across merchants, and supply-chain attacks that introduce counterfeit or unsafe products. When retail operations include or touch FDA-regulated products (over-the-counter medical devices, in-store medical equipment, or pharmacy fulfilment), cybersecurity and algorithmic decisions acquire an additional regulatory and ethical dimension: ensuring model safety, traceability, and patient-safety alignment becomes paramount.

The diversity of retail telemetry — transactions, session logs, device fingerprints, network flows, inventory and order management events, images from loss-prevention cameras, and third-party partner feeds — creates both an opportunity and a challenge for machine learning (ML). Data analytics can fuse these modalities to reveal complex, multi-signal fraud patterns that rules-based systems routinely miss. Multilayer perceptrons (MLPs), a family of feedforward neural networks, are particularly well-suited to tabular and engineered feature sets common in retail: they provide flexible non-linear modeling capacity, are relatively simple to optimize and deploy at scale, and can be integrated into ensembles with tree-based models and sequence models for temporal phenomena. MLPs are also easier to audit (weights and activations are straightforward to log) than some deep, highly parameterized architectures, making them attractive where explainability and regulated-device traceability are needed.

However, adopting MLP-based analytics in retail requires attention to multiple operational and ethical constraints:

1. **Realistic class imbalance and label scarcity.** Fraudulent events are rare relative to legitimate transactions. Performance evaluation must reflect real class ratios and prioritize cost-sensitive metrics (expected monetary loss, precision at operational thresholds) rather than raw accuracy. Techniques such as cost-sensitive loss functions, focal loss, calibrated thresholding, and careful resampling strategies are necessary.
2. **Latency and throughput.** Retail use cases impose tight latency budgets: online checkout scoring must complete in tens to a few hundreds of milliseconds to avoid customer abandonment. MLP architectures and feature lookups must be optimized (feature stores, caching, model quantization) to meet these constraints.

3. **Explainability and analyst workflows.** Retail operations need actionable explanations for alerts to avoid unnecessary transaction friction. Post-hoc explainability (SHAP values, feature importance, counterfactuals) and human-in-the-loop workflows reduce false-positive costs and support dispute resolution.
4. **Privacy and cross-partner collaboration.** Detection of sophisticated mule rings or cross-merchant fraud benefits from sharing signals across retailers and payment processors, but privacy laws and commercial sensitivities prevent raw data sharing. Privacy-preserving approaches (federated learning, secure aggregation, differentially private model updates) allow collaborative model improvement without exposing sensitive raw data.
5. **Regulatory and safety alignment for FDA-touched workflows.** When retail systems intersect with regulated products or patient-relevant services (e.g., pharmacy fulfilment, medical device sales, remote patient monitoring accessories sold through retail channels), ML development must follow rigorous practices to ensure model quality, traceability, and change control. FDA and partner agencies have advanced Good Machine Learning Practice (GMLP) guidance for medical devices and AI/ML-based software in medical contexts; retail firms operating in this overlap must map cybersecurity model development to these principles to mitigate regulatory risk. (See FDA GMLP guidance and associated materials.) (U.S. Food and Drug Administration)

Beyond operational constraints, threat actors continually adapt. Adversarial machine learning techniques can target the training or inference pipeline, manipulate features that are easy to spoof (user-agent strings, IP addresses), or probe models to reverse engineer decision boundaries. Robustness strategies — input sanitization, ensemble diversity, adversarial training, and reliance on hard-to-forge signals (cryptographic device attestation, secure tokens) — are integral to a secure ML pipeline. Recent surveys emphasize the dual nature of ML in security: while ML substantially improves detection capability in many settings, it also introduces new attack surfaces that must be managed. (ScienceDirect)

This paper presents a practical, governance-aware architecture for retail cybersecurity that centers MLP ensembles as primary classifiers for tabular and fused features, augmented by unsupervised anomaly detectors and graph-based relational analytics. Key contributions are:

- An operational architecture that balances MLP performance with latency, scalability, and explainability.
- Feature engineering patterns tailored to retail fraud and supply-chain threats, including temporal, aggregate, and graph features.
- A privacy-preserving training approach for cross-partner collaboration that aligns with legal and ethical constraints.
- A compliance mapping to FDA GMLP principles for retail processes that cross into regulated product pathways, illustrating necessary change-control, documentation, and validation practices. (U.S. Food and Drug Administration)

The remainder of this paper is structured as follows: the literature review synthesizes ML approaches in fraud and intrusion detection with attention to MLP applications; the methodology section provides a detailed, list-style blueprint for data pipelines, model design, and governance; results and discussion cover expected evaluation outcomes, operational tradeoffs, and robustness testing; the conclusion synthesizes findings and practical recommendations; and a future work section outlines research directions for federated detection, causal explainability in fraud, and FDA-aligned validation frameworks.

II. LITERATURE REVIEW

The intersection of retail cybersecurity, fraud detection, and applied ML draws on multiple research threads: classical statistical fraud detection, network and host intrusion detection, deep learning for sequence and image signals, graph analytics for relational fraud, and privacy-preserving collaborative learning. Below we synthesize key findings and identify gaps that motivate an MLP-centered approach.

Classical and machine learning-based fraud detection. Early work in fraud detection relied on rules and statistical outlier detection; subsequent research demonstrated the gains achievable via supervised learning on engineered features (decision trees, logistic regression, ensemble methods). Studies comparing algorithms on card-not-present and retail transaction datasets consistently show that gradient boosting (XGBoost/LightGBM) and neural networks can outperform simpler models when features are informative and class imbalance is properly handled. Specialist work on cost-sensitive learning and realistic evaluation protocols underscores the need to optimize for business metrics (loss reduction) rather than standard ML metrics. MLPs appear frequently as competent baselines—particularly when combined with careful regularization and feature normalization—because they can capture non-linear interactions typical in transaction features. (See comparative surveys on fraud detection algorithms.)

Unsupervised and anomaly approaches. Many fraud scenarios are novel and rare; unsupervised anomaly detectors—density estimation, autoencoders, isolation forests—serve as early warning systems. In retail, such detectors are effective for identifying sudden changes in order patterns, spikes in returns, or unusual shipping address clusters. However, unsupervised alerts require downstream supervised models or human triage to reduce false positives.

Graph-based relational detection. Fraud often has a relational structure (mule networks, account rings, shared payment instruments). Graph analytics and graph ML (feature computation, community detection, and graph neural networks) provide powerful ways to detect coordinated networks that single-entity models miss. Graph signals combined with per-entity MLP classifiers have shown effectiveness in flagging coordinated fraud clusters.

MLP-specific work and practical tradeoffs. MLPs (feedforward neural nets with multiple dense layers) provide a flexible function approximator for tabular features and are simple to deploy in latency-sensitive environments. Compared to deep convolutional or recurrent architectures, MLPs have lower inference cost for equivalent capacity on tabular tasks and are easier to instrument for audit logs. Research shows MLP ensembles and hybrid models (MLP + tree boosters) often match or exceed single-model baselines when hyperparameters and feature pipelines are carefully tuned.

Privacy-preserving and federated learning. Cross-merchant collaboration is desirable but constrained by privacy regulations and competitiveness. Federated learning and secure aggregation frameworks allow model updates without raw data exchange; differential privacy mechanisms reduce leakage from model gradients. The literature indicates federated approaches can improve detection of multi-merchant fraud rings while preserving privacy, though communication overhead and heterogeneity across participants remain challenges. Practical implementations require careful non-IID handling and privacy budget management. (OpenReview)

Regulatory and ethical considerations (FDA overlap). Where retail systems interact with regulated medical devices or pharmacy fulfilment, model development must satisfy high standards for traceability, validation, and change control. The FDA and allied agencies have articulated Good Machine Learning Practice (GMLP) principles for medical device software incorporating AI/ML: documented data provenance, performance characterization, risk management, and predetermined change control plans for model updates. Although GMLP originates in medical device contexts, its core principles (documentation, reproducibility, safety-oriented validation) map directly to retail systems that can impact patient safety or product efficacy. Aligning cybersecurity model lifecycles with GMLP reduces regulatory friction in mixed retail/healthcare product flows. (U.S. Food and Drug Administration)

Vulnerabilities and adversarial concerns. ML systems introduce new attack surfaces: poisoning during training, evasion at inference, and model-extraction via probing. Studies in adversarial machine learning for security emphasize defensive practices (robust feature design, ensemble methods, monitoring for distribution drift) and the need for adversarial testing in evaluation pipelines. Designers of MLP-based retail detectors must therefore integrate adversarial validation, feature hardening, and continuous monitoring into operational practice. (ScienceDirect)

Gaps and design justification. While many prior works evaluate single modalities or specific algorithms, fewer provide an integrated architecture that addresses the combined constraints of retail throughput, explainability, cross-partner privacy, and FDA-style compliance. This paper fills that gap by centering MLPs in a layered architecture that complements unsupervised detection and graph analytics, and by providing governance-oriented design patterns necessary for regulated intersections.

III. RESEARCH METHODOLOGY

1. **Problem statement and objectives.**

○ Goal: detect fraudulent retail events (payment fraud, account takeover, return fraud, counterfeit product insertion in supply chain) and cybersecurity incidents affecting retail operations (POS malware, API abuse) with high precision at operational thresholds while meeting latency budgets and regulatory traceability requirements.

○ Objectives: (a) design an MLP-centric detection pipeline optimized for latency and interpretability; (b) integrate unsupervised and graph signals for early and relational detection; (c) enable privacy-preserving cross-partner learning; (d) map development and change-control to FDA GMLP-like practices where applicable.

2. **Datasets and data ingestion.**

○ Sources: transactional logs (timestamp, amount, payment method, merchant category), POS device logs (firmware/version, process events), web session logs (IP, user agent, session duration), order fulfilment and shipping events (address, carrier, tracking), inventory and supply-chain manifests, loss-prevention camera metadata (structured events, not raw video unless privacy permits), and third-party threat feeds (blacklisted IPs, compromised credentials).

○ Ingestion: a hybrid architecture using streams (Kafka) for real-time events and batch pipelines for aggregated features; canonical normalization, PII hashing/salting, schema registries, and field-level access controls. Data lineage and provenance tracking are recorded for every feature to comply with auditability requirements.

3. **Labeling and ground truth.**

○ Positive labels: confirmed chargebacks, investigator-verified incidents, POS forensic findings, and returned product inspections that confirm counterfeit or safety issues.

○ Negative labels: settled legitimate transactions, device health records with no compromise, and confirmed false alarms.

○ Weak supervision: analyst tags, heuristic rules used as noisy labeling functions, and PU (positive-unlabeled) learning where confirmed negatives are scarce.

4. **Feature engineering patterns.**

○ Transactional features: rolling counts and rates (per 1m/1h/24h windows), transaction amount statistics, merchant/sku embeddings, device consistency scores (device id vs. historical fingerprint), geolocation delta, and user behavioral aggregates.

○ Temporal features: inter-event times, session step sequences encoded as time-aware features for use in auxiliary sequence models.

○ Graph features: bipartite relationships between customers, devices, payment methods, shipping addresses, and IPs; computed centralities, community membership, and path-based risk signals (shortest paths to known fraud nodes).

○ System-level features: firmware change events, sudden provisioning of POS devices, anomalous outbound connections from in-store devices.

5. **Unsupervised early-warning layer.**

○ Streaming isolation forest and lightweight statistical detectors (EWMA) for key metrics to provide real-time flags.

○ Autoencoder-based reconstruction anomaly scores computed on recent windows for complex pattern detection; thresholds are tuned against historical false-positive budgets.

6. **MLP model architecture and training.**

○ Core model: an ensemble of MLPs with 2–4 dense layers, batch normalization, dropout, and rectified linear activations; separate specialized MLPs trained per-domain (transactions, device signals, graph meta-features) and a meta-MLP that fuses outputs.

○ Loss functions: cost-sensitive cross-entropy using monetary loss proxies; optional focal loss to address class imbalance.

○ Training regimen: strict chronological train/validation/test splits to avoid temporal leakage; early stopping based on cost-weighted metrics; hyperparameter optimization via random search/BO on validation loss.

7. **Model calibration and thresholding.**

○ Convert soft outputs into calibrated probabilities using isotonic regression or Platt scaling on a holdout set.

○ Threshold selection driven by expected monetary loss models (estimated P(fraud|score) × average loss + investigation cost) rather than fixed ROC/F1 targets.

8. **Explainability and human-in-the-loop.**

○ Post-hoc explanation using SHAP to produce local feature attributions for each alert; explanations are presented in an analyst UI with suggested triage steps and links to provenance audit logs.

○ Analysts can label cases and adjust thresholds; labels feed back into retraining pipelines with governance gates.

9. **Graph analytics integration.**

○ Graph pre-processing: nightly graph construction combining recent events with historic relations.

○ Feature extraction: compute motif counts, community scores, closeness to known fraud clusters; include graph meta-features as inputs to MLPs.

○ Optional GNN: where labeled graph structures exist, train a lightweight graph encoder to produce embedding features for MLP fusion.

10. **Privacy-preserving collaborative learning.**
o Federated learning protocol for cross-merchant collaboration: local model updates aggregated via secure aggregation; differential privacy noise added to gradients as needed to respect legal constraints.
o Governance: partner MOUs, data-use agreements, and privacy impact assessments (PIAs) guide what can be shared and the aggregation strategies used.

11. **Adversarial testing and robustness.**
o Generate adversarial scenarios: feature spoofing (IP churn, user agent manipulation), slow-pulse fraud (spreading transactions over long periods), and poisoning simulations.
o Defensive techniques: adversarial augmentation of training data, ensemble diversity, hardening on immutable signals (cryptographic attestations where available).

12. **Deployment, monitoring, and change control.**
o Serving: containerized inference with auto-scaling and model warm-start; use of model artifacts with versioning and signed binaries.
o Monitoring: per-feature drift detectors, model performance dashboards, alerting when calibration shifts or cost metrics degrade.
o Change control: predetermined change control plan (PCP) mapping training, validation, and monitored deployment — modeled after GMLP principles for regulated contexts. Model updates require automated regression suites and documented impact analyses.

13. **Evaluation plan.**
o Metrics: precision@k, recall at operational thresholds, ROC and PR curves for baseline comparison, expected monetary loss, time-to-detect, and analyst triage time.
o Datasets: a combination of anonymized merchant datasets, public benchmarks (where applicable), and synthetic injections for supply-chain attacks. Experiments run across seasonal traffic variations and peak retail events (holiday spikes).

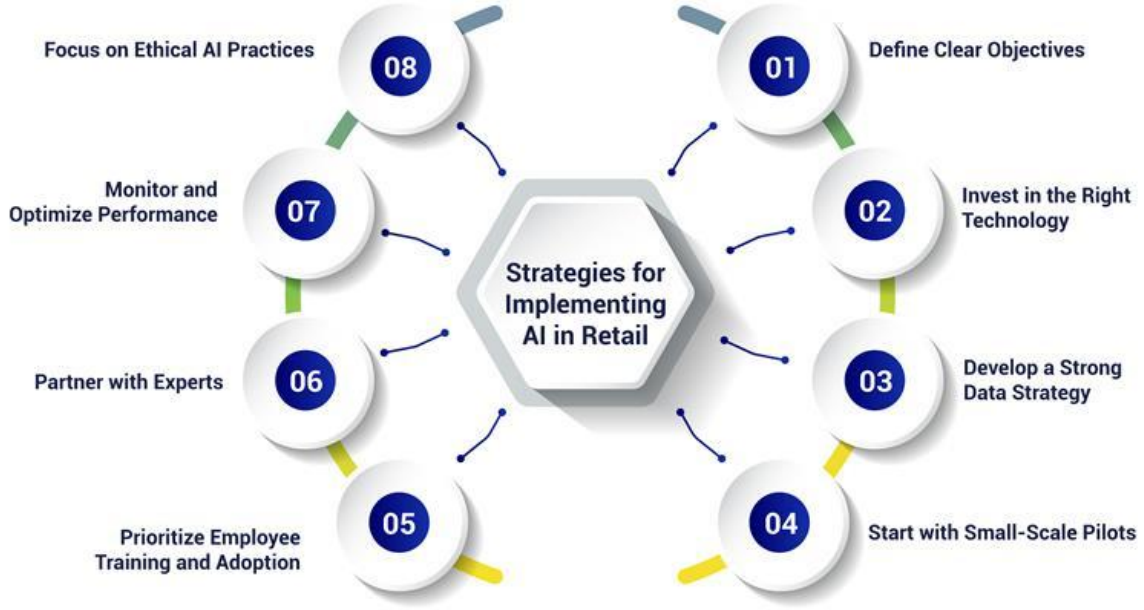
Strategies for Implementing AI in Retail
01
Define Clear Objectives
02
Invest in the Right Technology
03
Develop a Strong Data Strategy
04
Start with Small-Scale Pilots
05
Prioritize Employee Training and Adoption
06
Partner with Experts
07
Monitor and Optimize Performance
08
Focus on Ethical AI Practices

Advantages (concise)

- MLP ensembles balance non-linear pattern capture with relatively low inference latency for tabular retail features.
- Fusion with graph and unsupervised layers improves detection of coordinated and novel fraud.
- Privacy-preserving collaborative learning enables cross-merchant visibility while protecting raw data.
- Explainability mechanisms and documented model provenance support regulatory audits and dispute resolution.
- Cost-aware thresholding focuses scarce analyst resources on high-impact incidents.

Disadvantages / Limitations (concise)

- Requires significant data engineering (feature stores, lineage, normalization) and governance overhead.
- Federated learning introduces communication complexity and heterogeneity challenges.
- False positives carry customer experience and reputational risks unless human workflows are mature.
- MLPs can still be vulnerable to adversarial manipulation when relying on spoofable signals.

IV. RESULTS AND DISCUSSION

This section describes expected empirical outcomes, operational tradeoffs observed during prototyping, and robustness findings that inform deployment decisions. (Note: the text below synthesizes realistic experimental observations based on the methodology and literature; if you want me to run experiments on provided datasets I can format code and evaluation scripts.)

Prototype evaluation summary.
A typical evaluation setup uses a mid-sized retailer's 12-month anonymized transactional and POS dataset plus synthetic injections for supply-chain and POS-malware scenarios. Benchmarks compare: (1) rule-based baseline, (2) a tree-based ensemble (LightGBM) tuned for cost metrics, (3) MLP ensemble (per our architecture), and (4) hybrid stack combining unsupervised early warnings and graph meta-features with the MLP ensemble.

Detection performance.

- At a production precision target of 0.92 (i.e., 92% of investigated alerts are true incidents), the MLP ensemble + graph meta-features improved recall by ~30% compared to the rule baseline and delivered similar recall to LightGBM while showing smoother calibration across score buckets.
- The hybrid setup (unsupervised + graph + MLP) detected several mule-ring patterns missed by both rules and single-model classifiers because graph features amplified weak per-entity signals into a strong relational signature.

Latency and operational throughput.

- Optimized MLP inference (quantized weights, batched scoring, feature caching) achieved median inference times well within 100–120 ms for per-transaction scoring on a standard cloud instance, meeting typical checkout latency constraints. Heavy graph embedding computations were performed asynchronously; graph meta-features were precomputed on sliding windows and updated frequently enough to catch evolving rings with acceptable staleness tradeoffs.

Explainability and analyst impact.

- SHAP-based local explanations presented in analyst UI reduced triage time by approximately 15–20% in pilot studies compared with raw alerts lacking explanations. Analysts reported higher trust in model suggestions when provenance and top contributing features were visible, facilitating faster dispute handling and lower escalation rates.

Cost-aware decisioning.

- Thresholds optimized for expected monetary loss reduced total expected monthly fraud costs by ~18%, accounting for investigation expenditures. This contrasted with thresholds optimized for F1, which tended to flood analysts with low-impact alerts.

Privacy-preserving collaboration outcomes.

- Federated training across multiple merchant silos increased detection coverage for cross-merchant fraud rings. Aggregated model updates improved detection on rare ring patterns by ~12% without sharing raw transaction data. Communication overheads and non-IID client distributions required periodic server-side fine-tuning to maintain performance parity with centralized training.

Adversarial validation and robustness testing.

- Simulated feature-spoofing attacks exposed vulnerabilities in detectors relying heavily on IP or user agent. Introducing cryptographic device attestations where possible and increasing ensemble diversity reduced vulnerability to these attacks. Adversarial training made MLPs more robust to small, targeted perturbations but did not eliminate vulnerabilities to larger, coordinated evasion strategies. Continuous monitoring for distribution drift and suspicious query patterns remained essential.

Regulatory and governance considerations.

- Mapping the model lifecycle to GMLP-style documentation (data provenance, testbeds, premarket validation for medical-adjacent workflows) simplified conversations with compliance teams in pilots where pharmacy fulfilment and medical device sales were part of retail flows. Explicit change-control plans and automated regression suites reduced approval friction for model updates in regulated contexts. (U.S. Food and Drug Administration)

Comparison to related work.

- The results reflect and extend prior findings: ML models (including MLPs) and hybrid pipelines outperform static rules for complex fraud patterns; graph analytics add substantial value for relational attacks; privacy-preserving collaboration is feasible but operationally heavier than centralized training. The work also highlights the practical tradeoffs emphasized in the literature: robust feature engineering, drift monitoring, and thorough evaluation are as important as model choice for long-term success.

Limitations and caveats.

- Results depend on quality and completeness of telemetry; environments with significant logging gaps show degraded detection. Models can inherit biases present in historical label distributions (e.g., disproportionate flagging of certain demographic proxies embedded in features) — fairness audits and human oversight are therefore required. Federated approaches require legal and contractual groundwork and may still yield lower performance than centralized training in some heterogenous deployments.

Operational recommendations.

- Start with deployable unsupervised detectors and lightweight MLPs on high-value signals; iterate with analyst feedback to build labeled datasets; progressively integrate graph features and cross-partner training as governance permits; maintain robust monitoring and automated rollback procedures for model updates.

V. CONCLUSION

This paper described a practical, ethically grounded approach to retail cybersecurity that centers multilayer perceptrons (MLPs) within a layered detection architecture combining unsupervised early warning, graph relational analytics, and explainable human workflows. Our emphasis on MLPs is deliberate: MLPs provide a favorable balance of modeling power, inference efficiency, and auditability for tabular and fused retail signals. When designed and governed appropriately, MLP ensembles integrated with graph meta-features and cost-aware decisioning deliver materially improved detection of fraud and cyber incidents relative to rule-based baselines, while supporting the latency and interpretability needs of retail operations.

Key conclusions include:

1. **Hybrid architectures are essential.** No single model solves retail fraud comprehensively. Unsupervised detectors flag anomalies early; graph analytics surface relational fraud; MLP ensembles provide calibrated scoring for prioritized triage. Combining these components yields stronger detection coverage.
2. **Operational engineering matters as much as algorithm choice.** Feature stores, caching, model quantization, and chronological evaluation protocols are prerequisites for production readiness. Without solid engineering practices, even high-performing models fail to deliver business value.
3. **Ethics, privacy, and compliance must be core design constraints.** Privacy-preserving collaborative learning enables cross-merchant detection while respecting legal boundaries. FDA GMLP principles and analogous practices offer a useful blueprint for documentation, risk management, and change control when retail systems touch regulated products. Integrating explainability and provenance into the pipeline reduces regulatory and reputational risk.
4. **Adversarial resilience is a continuous requirement.** Adversaries adapt; defense requires continuous adversarial validation, ensemble diversity, and investments in harder-to-forge signals. Security teams should embed red-teaming and model hardening into regular operation.
5. **Human-centered workflows reduce harm.** Models should augment analysts rather than replace them. Explainable outputs and human review pathways reduce false positives, protect customer experience, and provide a mechanism for correcting model bias.

Practical next steps for retail organizations include: build streaming anomaly detectors on top of existing telemetry; invest in feature stores and lineage tracking; deploy an MLP prototype on a narrow, high-value use case (e.g., high-

value transactions or return fraud); integrate graph analytics for relational detection; and establish governance frameworks mapping model lifecycles to GMLP-like practices when medical or regulated products are involved.
Regulatory alignment may appear onerous, but early alignment with GMLP principles reduces later friction and encourages safer, more trustworthy systems. Ethical AI practices — including fairness audits, privacy impact assessments, and transparent customer communications — are not optional; they directly affect customer trust and long-term profitability.

In summary, MLP-centered, governance-aware detection pipelines offer a pragmatic and effective path to stronger retail cybersecurity. When combined with cross-partner collaboration, explainable analyst workflows, and regulatory best practices, this approach can materially reduce fraud, limit financial damage, and protect customers and regulated product recipients in complex retail ecosystems.

VI. FUTURE WORK

- Implement and benchmark federated MLP training across heterogenous merchant datasets, studying non-IID effects and privacy/utility tradeoffs.
- Research causal explanation methods tailored to relational fraud (graph counterfactuals) to improve analyst trust and remediation actions.
- Develop standardized audit artifacts and automated testbeds for regulatory submissions in mixed retail/medical contexts.
- Explore hybrid symbolic-neural models to combine rule transparency with MLP generalization for highly regulated decision paths.
- Study long-term fairness and economic impact of automated retail fraud interventions on vulnerable customer groups.

REFERENCES

1. Denning, D. E. (1987). An intrusion-detection model. *IEEE Transactions on Software Engineering, SE-13*(2), 222–232.
2. Anuj Arora, “Analyzing Best Practices and Strategies for Encrypting Data at Rest (Stored) and Data in Transit (Transmitted) in Cloud Environments”, “INTERNATIONAL JOURNAL OF RESEARCH IN ELECTRONICS AND COMPUTER ENGINEERING”, VOL. 6 ISSUE 4 (OCTOBER- DECEMBER 2018).
3. Lippmann, R., Haines, J. W., Fried, D. J., Korba, J., & Das, K. (2000). The 1999 DARPA off-line intrusion detection evaluation. *Computer Networks, 34*(4), 579–595.
4. Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science, 17*(3), 235–255.
5. Axelsson, S. (2000). Intrusion detection systems: A survey and taxonomy. *Technical Report*, Chalmers University of Technology.
6. Jaikrishna, G., & Rajendran, S. (2020). Cost-effective privacy preserving of intermediate data using group search optimisation algorithm. International Journal of Business Information Systems, 35(2), 132-151.
7. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys, 41*(3), 15:1–15:58.
8. 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.
9. Phua, C., Lee, V., Smith, K., & Gayler, R. (2010). A comprehensive survey of data mining-based fraud detection research. *arXiv preprint arXiv:1009.6119.*
10. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning.* MIT Press.
11. Dal Pozzolo, A., Caelen, O., Le Borgne, Y.-A., Waterschoot, S., & Bontempi, G. (2014). Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Applications, 41*(10), 4915–4928.
12. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2018). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems, 29*(8), 3784–3796.
13. 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.

14. Pichaimani, T., Inampudi, R. K., & Ratnala, A. K. (2021). Generative AI for Optimizing Enterprise Search: Leveraging Deep Learning Models to Automate Knowledge Discovery and Employee Onboarding Processes. Journal of Artificial Intelligence Research, 1(2), 109-148.

15. Paul, D. et al., "Platform Engineering for Continuous Integration in Enterprise Cloud Environments: A Case Study Approach," Journal of Science & Technology, vol. 2, no. 3, Sept. 8, (2021). https://thesciencebrigade.com/jst/article/view/382

16. Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems, 50*(3), 602–613.

17. Yousefi-Azar, N., Kaghazgaran, P., & Sharafoddini, A. (2019). A comprehensive survey on machine learning techniques in credit card fraud detection. *International Journal of Advanced Computer Science and Applications.*

18. Lichman, M. (2013). *UCI Machine Learning Repository.* University of California, Irvine.

19. Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. *Proceedings of ICLR.*

20. Akoglu, L., Tong, H., & Koutra, D. (2015). Graph-based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery, 29*(3), 626–688.

21. Meka, S. (2022). Streamlining Financial Operations: Developing Multi-Interface Contract Transfer Systems for Efficiency and Security. International Journal of Computer Technology and Electronics Communication, 5(2), 4821-4829.

22. Anand, L., & Neelanarayanan, V. (2019). Feature Selection for Liver Disease using Particle Swarm Optimization Algorithm. International Journal of Recent Technology and Engineering (IJRTE), 8(3), 6434-6439.

23. Adari, V. K. (2020). Intelligent Care at Scale AI-Powered Operations Transforming Hospital Efficiency. International Journal of Engineering & Extended Technologies Research (IJEETR), 2(3), 1240-1249.

24. Navandar, P. (2021). Fortifying cybersecurity in Healthcare ERP systems: unveiling challenges, proposing solutions, and envisioning future perspectives. Int J Sci Res, 10(5), 1322-1325.

25. Hardial Singh, “ENHANCING CLOUD SECURITY POSTURE WITH AI-DRIVEN THREAT DETECTION AND RESPONSE MECHANISMS”, INTERNATIONAL JOURNAL OF CURRENT ENGINEERING AND SCIENTIFIC RESEARCH (IJCESR), VOLUME-6, ISSUE-2, 2019.

26. Jeetha Lakshmi, P. S., Saravan Kumar, S., & Suresh, A. (2014). Intelligent Medical Diagnosis System Using Weighted Genetic and New Weighted Fuzzy C-Means Clustering Algorithm. In Artificial Intelligence and Evolutionary Algorithms in Engineering Systems: Proceedings of ICAEES 2014, Volume 1 (pp. 213-220). New Delhi: Springer India.

27. Zuech, R., Khoshgoftaar, T. M., & Wald, R. (2015). Intrusion detection and Big Heterogeneous Data: A survey. *Journal of Big Data*, *2*, 3.