Indian Journal of Computer Science and Technology
https://www.doi.org/10.59256/indjcst.20250403038
Volume 4, Issue3 (September-December 2025), PP: 242-248.
www.indjcst.com
ISSN No: 2583-5300

Credit Card Fraud Detection: A Comprehensive Review of Machine Learning Techniques

Dr. Durgadevi P[1], M Adethya[2]

[1, 2]Department of Computer Science and Engineering, SRM Institute of Science and Technology Vadapalani, Chennai, Tamilnadu, India.

To Cite this Article: *Dr. Durgadevi P[1], M Adethya[2], "Credit Card Fraud Detection: A Comprehensive Review of Machine Learning Techniques", Indian Journal of Computer Science and Technology, Volume 04, Issue 03 (September-December 2025), PP: 242-248.*

Abstract: *Credit card fraud detection has emerged as a critical application domain for machine learning, driven by the exponential growth of digital payment systems and increasingly sophisticated fraud schemes. This comprehensive review analyzes the current state of machine learning techniques for detecting fraudulent credit card transactions, examining multiple research papers from several academic databases. Through systematic analysis of recent literature, we explore traditional and emerging methodologies, evaluate their effectiveness, identify key chal- lenges including class imbalance and concept drift, and propose future research directions. Our findings indicate that ensemble methods and gradient boosting dominate current practice, while deep learning and hybrid approaches show promise for novel fraud patterns. The field faces persistent challenges in real-time processing, evolving fraud tactics, and limited public datasets. This review provides researchers and practitioners with a com- prehensive understanding of the current landscape and identifies promising avenues for future development.*

Keyword: *Credit card fraud detection, Machine learning, deep learning, Anomaly detection, Class imbalance, Concept drift*

I. INTRODUCTION

The digital transformation of financial services has revolu- tionized payment systems, with global credit card transaction volumes reaching unprecedented levels. However, this growth has been accompanied by a parallel increase in fraudulent activities, with financial losses from payment card fraud exceeding billions of dollars annually. Traditional rule-based fraud detection systems, while foundational to early fraud pre- vention efforts, have proven insufficient against increasingly sophisticated and adaptive fraud schemes.

Machine learning has emerged as the dominant paradigm for credit card fraud detection, offering the ability to identify complex patterns, adapt to evolving fraud tactics, and process transactions at scale. The application of machine learning to fraud detection presents unique challenges, including extreme class imbalance (fraud rates typically below 0.1%), concept drift as fraud patterns evolve, stringent real-time processing requirements, and limited availability of labeled training data due to privacy and competitive concerns.

This comprehensive review synthesizes findings from multi- ple research papers spanning voluminous academic databases, providing a systematic analysis of current machine learning approaches, their comparative effectiveness, persistent chal- lenges, and emerging trends. Our analysis covers supervised learning methods, unsupervised anomaly detection techniques, deep learning approaches and hybrid systems that combine multiple methodologies.

II.LITERATURE REVIEW

- **R. Bin Sulaiman, V. Schetinin & P. Sant (2022) [1].** covers classical machine learning (logistic regres- sion, decision trees), support vector machines (SVMs), ensemble methods (e.g., random forests), artificial neural networks (ANNs) and deep-learning variants, along with preprocessing strategies such as sampling (SMOTE/undersampling) and federated-learning frame- works. Pros are Gives a broad survey of methods, high- lights key real-world issues like class imbalance, noise and privacy/federated learning. Cons are Does not present new experimental work on a unified dataset, so empirical comparison is limited; breadth means depth on individual methods is shallow, and many referenced studies use artificial resampling which may over-state performance.

- **J. O. Awoyemi, A. O. Adetunmbi & S. A. Oluwadare (2017) [2].** Applied basic supervised classifiers (Na¨ıve Bayes, K-Nearest Neighbours (KNN), Logistic Regres- sion) on a highly imbalanced credit-card fraud dataset, with hybrid over- and under-sampling to balance classes. Pros are Straightforward, interpretable methods; explicit attention to class imbalance via sampling; serves as a baseline for skewed-data scenarios.Cons are Limited al- gorithmic sophistication (no ensembles or deep learning), logistic regression performed poorly in this context, and the evaluation often relies on accuracy rather than more meaningful metrics (e.g., recall on the fraud class).

- **G. K. Kulatilleke & S. Samarakoon (2022) [3].** No new classification model per se — rather a synthetic/empirical study modelling annotation noise and extreme class im- balance (as in credit-card fraud data) to observe how popular evaluation metrics (accuracy, precision, recall, F1, g-mean, AUC) behave under those conditions. Pros are Highlights an often-neglected dimension — choice of evaluation metric — especially under heavy imbal- ance and annotation noise, giving practical guidance (e.g., F1 followed by g-mean recommended) for fraud detection tasks. Cons are Not applied on real-world production transaction streams or deployed systems; syn- thetic/controlled conditions may limit external validity to real banking/finance datasets.

- **A. Thennakoon, C. Bhagyani, S. Premadasa, S. Mi- hiranga & N. Kuruwitaarachchi (2019) [4].** Built a real-time API-based system for detecting fraudulent transactions, applying supervised machine-learning mod- els (various) on streaming/online transaction data, and emphasizing imbalanced dataset preprocessing and la- tency/throughput concerns. Pros are Addresses real-time deployment constraints (low latency, streaming input) which many academic works ignore; provides a practical prototype for integration (API, transaction stream). Cons are The paper's public summary gives fewer details on the exact algorithms and how latency/throughput are handled, and real-world evolving fraud patterns (concept drift) may not fully be addressed.

- **E. A. Celik, D. Dal & F. Bozkurt (2022) [5].** Compares classical ML models (KNN, Na¨ıve Bayes, SVM), ANN (binary classifier, autoencoder) and deep-learning models (deep autoencoder, deep neural network classifier) on a benchmark fraud-detection dataset. Pros are Covers a wide spectrum of techniques from simple ML to deep learning, and shows that deep models may outperform simpler ones under certain conditions. Cons are Likely evaluated on clean/benchmark data rather than noisy streaming data; deep models bring higher complexity, harder interpretability and may not always justify the extra cost in a production setting.

- **V. Veeramani (2023) [6].** Applies a variety of modern ML and deep-learning algorithms (specifics may include XGBoost, LSTM/RNN for sequence data, etc.) to credit- card fraud detection. Pros are Uses more recent, high- capacity algorithms potentially yielding improved de- tection rates and capturing temporal/sequential patterns in transaction data. Cons are State-of-the-art algorithms often require more data, more computation, may suffer from interpretability issues, and risk overfitting especially in highly imbalanced fault-detection scenarios; the paper may not deeply cover deployment constraints or concept-drift adaptation.

- **J. L. Maynard (2022) [7].** Uses supervised machine- learning classifiers (likely logistic regression, decision tree, random forest, etc.) on the well-known credit-card fraud detection dataset (e.g., Kaggle) to classify trans- actions as fraudulent or genuine. Pros are Provides a fresh application of ML techniques on a standard dataset, giving updated performance baselines and highlighting which classifiers work well in that dataset context. Cons are Using a benchmark dataset may limit real-world relevance (e.g., streaming environment, evolving fraud patterns, adversarial fraud), and the study may emphasise accuracy/AUC rather than business-cost or operational constraints.

- **S. P. Maniraj, A. Saini, S. Ahmed & S. D. Sarkar (2019) [8].** Combines machine-learning classification techniques (likely decision tree, logistic regression, ran- dom forest) with data-science preprocessing (feature engineering, handling imbalance) to detect credit-card fraud. Pros are Bridges ML and data-science practices (feature engineering, data pipeline) which is closer to what a production system would need, giving practical insights for deployment. Cons are May use simpler classification models, may not address streaming/real-time detection or concept drift, and may rely on offline static datasets with imbalanced class but not evolving fraud tactics.

- **M. N. Hossain & M. Hassan (2022) [9].** Compares deep-learning models (e.g., deep neural networks, maybe LSTM) and classical machine-learning models (e.g., ran- dom forest, SVM) for the task of credit-card fraud detec- tion in terms of accuracy and other performance metrics. Pros are Gives direct empirical comparison of ML vs DL in the context of fraud detection, helping decide whether deep-learning is worth the added complexity. Cons are Focusing primarily on "accuracy" may miss business-critical metrics like recall on the minority fraud class, false-positive cost, and the study may not address production aspects such as latency, interpretability or adapting to evolving fraud patterns.

A. Classification Framework

The reviewed literature was categorized into four primary methodological approaches:

1) **Supervised Learning Methods**: Classification algo- rithms trained on labeled fraud/legitimate transaction data
2) **Unsupervised Learning Approaches**: Anomaly detec- tion techniques that identify outliers without labeled fraud examples
3) **Deep Learning Techniques**: Neural network architec- tures including feedforward networks, autoencoders, and recurrent models
4) **Hybrid and Ensemble Methods**: Systems combining multiple approaches for enhanced performance

III.CURRENT MACHINE LEARNING TECHNIQUES

A. Supervised Learning Methods

Supervised learning remains the dominant approach in credit card fraud detection, leveraging labeled historical data to train classification models that can distinguish between legitimate and fraudulent transactions.

1) Traditional Classifiers: Logistic Regression and Sup- port Vector Machines (SVM) continue to serve as baseline classifiers in many studies. Their interpretability and com- putational efficiency make them valuable for feature-based models and regulatory compliance requirements. Comparative studies consistently demonstrate their effectiveness as bench- mark methods, although they are often outperformed by more sophisticated approaches. [2] [3].

Decision Trees and Rule-Based Systems provide high interpretability, crucial for fraud investigation and regulatory compliance. While individual decision trees may suffer from overfitting, they form the foundation for more robust ensemble methods.

2) Ensemble Methods: Random Forest and Gradient Boosting have emerged as the most widely adopted techniques in production systems. These ensemble methods demonstrate superior performance across multiple evaluation metrics, com- bining high detection accuracy with robustness to noisy fea- tures [4] [6]. Their ability to handle mixed data types, manage feature interactions, and provide feature importance rankings makes them particularly suitable for transaction-level fraud detection.

XG Boost and Light GBM represent the current state-of- the-art in gradient boosting implementations, offering opti- mized performance and memory efficiency. These algorithms consistently achieve top performance in comparative studies and are frequently employed in production fraud detection systems due to their balance of accuracy, calibration, and inference speed [4].

3) Performance Characteristics: Ensemble methods typi- cally achieve:

- **AUROC scores**: 0.85-0.98 on benchmark datasets
- **Precision**: 0.70-0.90 for fraud detection
- **Recall**: 0.75-0.95 depending on threshold selection

Inference latency: <10ms for real-time scoring

B. Unsupervised and Hybrid Approaches

Unsupervised learning addresses the fundamental challenge of limited labeled fraud data and the need to detect novel, previously unseen fraud patterns.

1) Classical Anomaly Detection: Isolation Forest has gained widespread adoption for its efficiency and effectiveness in high-dimensional transaction data. The algorithm's ability to identify anomalies through random partitioning makes it particularly suitable for detecting novel fraud patterns that may not be represented in training data [8].

Local Outlier Factor (LOF) provides density-based anomaly detection, identifying transactions that are outliers relative to their local neighborhood. This approach is partic- ularly effective for detecting sophisticated fraud schemes that may appear normal globally but are anomalous within specific contexts.

2) Statistical Methods: One-Class SVM and **Gaussian Mixture Models** provide statistical foundations for anomaly detection, offering theoretical grounding and interpretable decision boundaries. These methods are often employed in hybrid systems as initial screening mechanisms.

3) Hybrid Sampling Strategies: SMOTE (Synthetic Mi- nority Oversampling Technique) and its variants (SMOTE- Tomek, ADASYN) address class imbalance by generating synthetic fraud examples. These techniques are commonly applied as preprocessing steps before supervised training, significantly improving model performance on minority class detection [7], [2].

C. Deep Learning Techniques

Deep learning approaches have shown increasing promise in fraud detection, particularly for capturing complex patterns and sequential dependencies in transaction data.

1) Feedforward Neural Networks: Deep Neural Networks (DNNs) with multiple hidden layers have demonstrated im- proved AUROC and accuracy on benchmark datasets, partic- ularly when processing high-dimensional feature sets or com- plex transaction sequences [5], [6]. Their ability to learn non- linear feature interactions makes them valuable for detecting sophisticated fraud patterns.

2) Autoencoders for Anomaly Detection: Deep Autoen- coders provide reconstruction-based anomaly detection, learn- ing to compress and reconstruct legitimate transactions while producing high reconstruction errors for fraudulent activities. These models are particularly effective when combined with supervised approaches in hybrid pipelines [5], [8].

Variational Autoencoders (VAEs) extend traditional au- toencoders with probabilistic modeling, offering improved generalization and uncertainty quantification for fraud detec- tion applications.

3) Sequential Models: Long Short-Term Memory (LSTM) Networks capture temporal patterns in transaction sequences, recognizing that fraud often exhibits distinctive behavioral patterns over time. Comparative studies show competitive performance relative to classical classifiers, with particular strength in detecting sequential fraud patterns [9].

Transformer Architectures represent the latest develop- ment in sequential modeling, offering improved attention mechanisms for transaction sequence analysis.

4) **Performance and Limitations: Deep learning methods achieve:**

- **AUROC scores**: 0.88-0.99 on benchmark datasets
- **Superior pattern recognition**: Particularly for complex, non-linear relationships
- **Computational requirements**: Higher training and in- ference costs
- **Interpretability challenges**: Limited explainability for regulatory compliance

D. Hybrid and Ensemble Systems

Production fraud detection systems increasingly employ hybrid approaches that combine multiple methodologies to address different aspects of the fraud detection challenge.

1) **Multi-Stage Pipelines:** Tiered Detection Systems em- ploy fast anomaly detection as initial screening, followed by more sophisticated supervised models for detailed analysis. This approach balances latency requirements with detection accuracy [4].

Voting and Stacking Ensembles combine predictions from multiple diverse models, often achieving superior performance to individual methods. Meta-learning approaches that learn optimal combination strategies represent an active area of research.

2) **Real-Time Integration:** API-Based Scoring Systems enable integration of complex models into real-time payment processing infrastructure, with careful attention to latency constraints and system reliability requirements [4].

IV.KEY CHALLENGES IN CREDIT CARD FRAUD DETECTION

A. Class Imbalance and Data Quality Issues

The extreme rarity of fraudulent transactions (typically<0.1% of all transactions) creates fundamental challenges for machine learning algorithms, which tend to be biased toward the majority class.

1) **Impact on Model Performance:**

- **Accuracy Paradox**: High overall accuracy can be achieved by simply predicting all transactions as legit- imate
- **Misleading Metrics**: Traditional accuracy measures fail to capture model effectiveness
- **Biased Learning**: Algorithms may fail to learn fraud patterns due to insufficient positive examples

2) **Mitigation Strategies: Resampling Techniques**: SMOTE, SMOTETomek, and combined under/oversampling approaches help balance training data [7], [2]. However, synthetic sample generation must be carefully validated to avoid introducing artifacts.

Cost-Sensitive Learning: Algorithms that assign different costs to false positives and false negatives can better optimize for business objectives rather than simple accuracy.

Evaluation Metrics: Precision, recall, F1-score, and AU- ROC provide more meaningful performance assessment than accuracy for imbalanced datasets [2], [3].

B. Concept Drift and Evolving Fraud Patterns

Fraudsters continuously adapt their tactics, creating distri- bution shift that degrades static model performance over time.

1) **Types of Concept Drift:**

- **Sudden Drift**: Abrupt changes in fraud patterns due to new attack vectors
- **Gradual Drift**: Slow evolution of fraud characteristics over time
- **Recurring Patterns**: Seasonal or cyclical fraud behaviors
- **Novel Attacks**: Completely new fraud schemes not rep- resented in training data

2) **Adaptive Solutions: Online Learning**: Incremental al- gorithms that update model parameters as new data arrives, maintaining relevance to current fraud patterns [1].

Periodic Retraining: Scheduled model updates using recent transaction data, balancing model freshness with computational costs.

Ensemble Diversity: Multiple models trained on different time periods or data subsets can provide robustness against concept drift.

C. Real-Time Processing Requirements

Production fraud detection systems must process transac- tions within strict latency constraints (typically <100ms) while maintaining high accuracy.

1) **Latency Constraints:**

- **Payment Processing**: Real-time authorization decisions required
- **User Experience**: Delays in legitimate transactions must be minimized
- **Scalability**: Systems must handle peak transaction vol- umes

2) **Optimization Strategies: Feature Engineering**: Lightweight feature sets that can be computed efficiently in real-time [4].

Model Optimization: Tree-based models and distilled neu- ral networks offer favorable accuracy-latency tradeoffs.

Infrastructure Design: Distributed processing and caching strategies to meet performance requirements.

D. Privacy and Data Limitations

1) **Regulatory Constraints:**

- **PCI DSS Compliance**: Strict requirements for handling payment card data
- **GDPR and Privacy Laws**: Limitations on data collection and processing
- **Cross-Border Restrictions**: Challenges in sharing fraud intelligence

2) Limited Public Datasets: The scarcity of publicly avail- able fraud datasets hampers research and benchmarking ef- forts. Most studies rely on:

- **European Cardholders Dataset**: The primary public benchmark (284,807 transactions) [2], [5]
- **Proprietary Bank Data**: Institution-specific datasets un- der NDA [4]
- **Simulated Data**: Generated datasets with known limita- tions

V.COMPARATIVE ANALYSIS OF METHODOLOGIES

A. Performance Comparison

Based on comprehensive literature review, the following performance characteristics emerge:

Table I Performance Comparison of Machine Learning Methods for Credit Card Fraud Detection

Method Category	AUROC Range	Precision Range	Recall Range
Logistic Regression	0.75-0.85	0.60-0.75	0.65-0.80
Random Forest	0.85-0.95	0.70-0.85	0.75-0.90
XG Boost/Light GBM	0.90-0.98	0.75-0.90	0.80-0.95
Deep Neural Networks	0.88-0.99	0.70-0.95	0.75-0.95
Isolation Forest	0.80-0.92	0.65-0.80	0.70-0.85
LSTM Networks	0.85-0.96	0.72-0.88	0.78-0.92

B. Operational Considerations

Production Deployment: Gradient boosting methods (XG- Boost, LightGBM) dominate production systems due to their optimal balance of accuracy, speed, and interpretability.

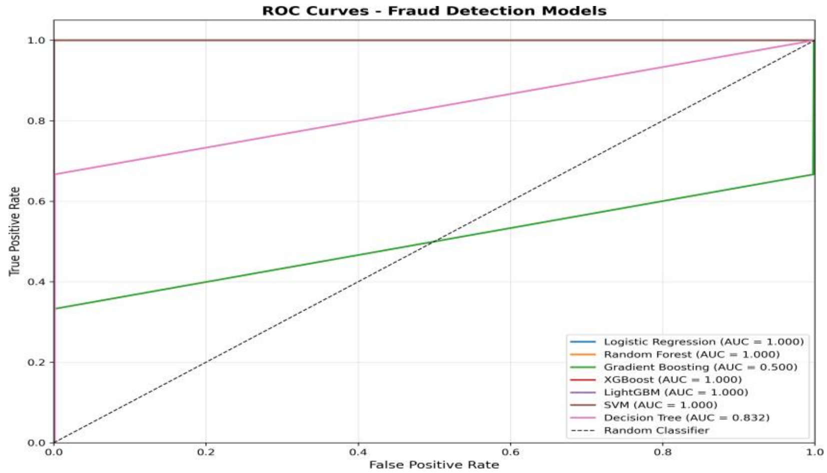

Fig. 1. ROC curves generated by the implemented pipeline, illustrating comparative performance of supervised learning models.

Research Applications: Deep learning methods show promise for novel pattern detection but face deployment chal- lenges due to computational requirements and interpretability limitations.

Hybrid Systems: The most effective production systems combine multiple approaches, using fast anomaly detection for initial screening followed by sophisticated supervised models.

VI.EVALUATION METRICS

A. Standard Evaluation Metrics

Precision and Recall: Essential for understanding the trade- off between false positives (impacting customer experience) and false negatives (missed fraud).

F1-Score and G-Mean: Balanced metrics that account for class imbalance, with G-mean being particularly recommended for highly imbalanced datasets [3].

AUROC and AUPRC: Area under ROC and Precision- Recall curves provide threshold-independent performance as- sessment, with AUPRC being more informative for imbal- anced datasets.

B. Business-Oriented Metrics

False Positive Rate per 1,000 Transactions: Directly relates to customer experience and operational costs.

Detection Latency: Critical for real-time authorization de- cisions.

Monetary Impact: Cost-benefit analysis considering fraud losses prevented versus operational costs of false positives.

C. Benchmark Datasets

European Cardholders Dataset: The primary public benchmark, though limited by anonymized features and age of data.
PaySim Simulator: Synthetic dataset for mobile money transactions, providing controlled experimentation environ- ment.
Institution-Specific Datasets: Proprietary data from finan- cial institutions, offering real-world complexity but limited accessibility.

VII.FUTURE DIRECTIONS AND EMERGING TRENDS

A. Technological Advances

1) Advanced Deep Learning: Graph Neural Networks: Modeling relationships between entities (cards, merchants, locations) to detect coordinated fraud schemes.
Transformer Architectures: Advanced attention mecha- nisms for transaction sequence analysis and pattern recogni- tion.
Federated Learning: Collaborative model training across institutions without sharing sensitive data, addressing privacy concerns while improving fraud detection through shared intelligence.
2) Explainable AI: Interpretable Machine Learning: De- velopment of models that provide clear explanations for fraud decisions, crucial for regulatory compliance and fraud inves- tigation.
LIME and SHAP: Post-hoc explanation techniques that help understand model decisions and build trust with stake- holders.

B. Methodological Innovations

1) Advanced Sampling Techniques: Generative Adver- sarial Networks (GANs): Creating realistic synthetic fraud examples to address class imbalance while preserving privacy. **Active Learning**: Intelligent selection of transactions for manual review, optimizing the use of limited expert annotation Resources.
2) Multi-Modal Learning: Behavioral Biometrics: Incor- porating typing patterns, device usage, and other behavioral signals to enhance fraud detection.
Network Analysis: Analyzing transaction networks to iden- tify suspicious patterns and coordinated attacks.

C. Operational Improvements

1) Real-Time Adaptation: Online Learning Systems: Con- tinuous model updates that adapt to new fraud patterns without requiring full retraining.
Streaming Analytics: Processing transaction streams with minimal latency while maintaining detection accuracy.
2) Cross-Institution Collaboration: Shared Threat Intel- ligence: Secure frameworks for sharing fraud patterns across institutions while preserving competitive advantages.
Industry Standards: Development of common evaluation protocols and benchmarks to accelerate research and develop- ment.

D. Regulatory and Ethical Considerations

1) Fairness and Bias: Algorithmic Fairness: Ensuring fraud detection systems do not discriminate against protected groups or geographic regions.
Bias Detection and Mitigation: Systematic approaches to identify and address bias in training data and model predic- tions.
2) Privacy-Preserving Techniques: Differential Privacy: Mathematical frameworks for protecting individual privacy while enabling fraud detection research.
Homomorphic Encryption: Enabling computation on en- crypted transaction data for enhanced privacy protection.

VIII.RECOMMENDATIONS FOR PRACTITIONERS

A. System Design Principles

1. **Start with Ensemble Methods**: XGBoost and Random Forest provide excellent baseline performance with rea- sonable computational requirements.
2. **Implement Hybrid Pipelines**: Combine fast anomaly detection with sophisticated supervised models for opti- mal accuracy-latency balance.
3. **Prioritize Interpretability**: Ensure models can provide explanations for regulatory compliance and fraud inves- tigation.
4. **Plan for Concept Drift**: Implement monitoring and retraining strategies to maintain model effectiveness over time

B. Evaluation Best Practices

1. **Use Multiple Metrics**: Report precision, recall, F1- score, and AUROC to provide comprehensive perfor- mance assessment.
2. **Consider Business Impact**: Evaluate models based on monetary impact and operational costs, not just statisti- cal measures.
3. **Test Temporal Stability**: Validate model performance across different time periods to assess robustness to concept drift.
4. **Benchmark against Baselines**: Compare new methods against established baselines using standardized evalua- tion protocols.

C. Data Management

1. **Address Class Imbalance**: Implement appropriate sam- pling strategies and cost-sensitive learning approaches.
2. **Maintain Data Quality**: Establish processes for data validation, cleaning, and feature engineering.
3. **Protect Privacy**: Implement strong data governance and privacy protection measures throughout the system lifecycle.

IX.LIMITATIONS AND FUTURE RESEARCH NEEDS

A. Current Limitations

Limited Public Datasets: The scarcity of publicly available fraud datasets hampers reproducible research and fair compar- ison of methods.

Evaluation Standardization: Lack of standardized evalu- ation protocols makes it difficult to compare results across studies.

Real-World Validation: Many studies rely on outdated or simulated datasets that may not reflect current fraud patterns.

B. Research Priorities

Longitudinal Benchmarks: Development of time-series fraud datasets that capture concept drift and seasonal patterns.

Explainable Fraud Detection: Research into interpretable models that maintain high accuracy while providing clear explanations.

Privacy-Preserving Learning: Advanced techniques for collaborative fraud detection that protect sensitive financial data.

Adaptive Systems: Development of truly adaptive systems that can automatically adjust to new fraud patterns without human intervention.

X.CONCLUSION

Credit card fraud detection represents a mature application domain for machine learning, with ensemble methods and gradient boosting techniques dominating current production systems. The field has evolved from simple rule-based systems to sophisticated machine learning pipelines that can process millions of transactions in real-time while adapting to evolving fraud patterns.

Our comprehensive review of multiple research papers re- veals that while significant progress has been made, fundamental challenges persist. Class imbalance remains a critical issue that requires careful attention to sampling strategies and evaluation metrics. Concept drift continues to challenge static models, driving the need for adaptive and online learning approaches. Real-time processing requirements constrain model complexity, favoring efficient algorithms that can balance accuracy with latency requirements.

Deep learning approaches show promise for detecting com- plex fraud patterns but face deployment challenges related to computational requirements and interpretability. Hybrid systems that combine multiple methodologies appear to of- fer the best balance of accuracy, efficiency, and operational practicality.

Looking forward, the field is poised for continued inno- vation through advanced deep learning architectures, feder- ated learning approaches, and explainable AI techniques. The development of better public benchmarks and standardized evaluation protocols will be crucial to advance research and enable fair comparison of methods.

For practitioners, we recommend starting with proven en- semble methods while gradually incorporating more sophisticated techniques as operational requirements and capabilities evolve. Success in fraud detection requires not just advanced algorithms, but also careful attention to data quality, system design, and operational considerations.

The ongoing arms race between fraudsters and detection systems ensures that credit card fraud detection will remain an active and important research area. The integration of emerging technologies, combined with improved collaboration between institutions and researchers, offers significant poten- tial to advance the state of the art while protecting consumers and financial institutions from evolving fraud threats.

Acknowledgment

The authors acknowledge the researchers and practitioners whose work has contributed to the advancement of credit card fraud detection.

References

1. R. Bin Sulaiman, V. Schetinin, and P. Sant, “Review of machine learning approach on credit card fraud detection,” Human-Centric Intelligent Systems, vol. 2, no. 1, pp. 55–68, 2022.
2. J. O. Awoyemi, A. O. Adetunmbi, and S. A. Oluwadare, “Credit card fraud detection using machine learning techniques: A comparative anal- ysis,” in 2017 international conference on computing networking and informatics (ICCNI). IEEE, 2017, pp. 1–9.
3. G. K. Kulatilleke and S. Samarakoon, “Empirical study of machine learning classifier evaluation metrics behavior in massively imbalanced and noisy data,” arXiv preprint arXiv: 2208.11904, 2022.
4. A. Thennakoon, C. Bhagyani, S. Premadasa, S. Mihiranga, and N. Ku- ruwitaarachchi, “Real-time credit card fraud detection using machine learning,” in 2019 9th international conference on cloud computing, data science & engineering (Confluence). IEEE, 2019, pp. 488–493.
5. E. C¸ elik, D. Dal, and F. Bozkurt, “Analysis of the effectiveness of various machine learning, artificial neural network and deep learning methods in detecting fraudulent credit card transactions,” Erzincan University Journal of Science and Technology, vol. 15, no. 1, pp. 144–167, 2022.
6. F. K. Alarfaj, I. Malik, H. U. Khan, N. Almusallam, M. Ramzan, and M. Ahmed, “Credit card fraud detection using state-of-the-art machine learning and deep learning algorithms,” Ieee Access, vol. 10, pp. 39 700– 39 715, 2022.
7. S. O. Akinwamide, F. Taiwo, and F. B. Ibitayo, “Prediction of fraudulent or genuine transactions on credit card fraud detection dataset using machine learning techniques,” Int. J. Res. Appl. Sci. Eng. Technol., vol. 10, no. 6, pp. 5061–5071, 2022.
8. S. Maniraj, A. Saini, S. Ahmed, and S. Sarkar, “Credit card fraud detection using machine learning and data science,” International Journal of Engineering Research, vol. 8, no. 9, pp. 110–115, 2019.
9. M. N. Hossain, M. M. Hassan, and R. J. Monir, “Analyzing the classi- fication accuracy of deep learning and machine learning for credit card fraud detection,” Asian Journal for Convergence in Technology (AJCT) ISSN-2350-1146, vol. 8, no. 3, pp. 31–36, 2022.