https://doi.org/10.62836/emi.v4i4.508

Global Science Publishing

EMI

Economics & Management Information

https://ojs.sgsci.org/journals/emi

Article

Decision Tree-Based Credit Card Fraud Detection System: Design and Optimization

Jian Sun

Department of Computer Science, Iowa State University, Ames, IA 50011, USA

Abstract: With the rapid development of fintech, credit card fraud has become increasingly sophisticated and intelligent, posing significant challenges to banks and consumers. Traditional fraud detection methods exhibit noticeable shortcomings in real-time performance, accuracy, and interpretability. This paper proposes an improved decision tree-based method for credit card fraud detection, enhancing detection performance by integrating dynamic feature engineering and ensemble learning strategies. The study employs a cost-sensitive decision tree algorithm as the base model, addressing class imbalance through sample weight adjustment, and designs a two-stage detection framework combined with random forest for result validation. Experimental results demonstrate that the proposed method significantly outperforms traditional approaches in both detection accuracy and false positive rate, particularly excelling in detecting emerging fraud patterns. The study validates the practical value of decision tree models in financial risk control and provides a highly interpretable, low-cost deployment solution for real-time anti-fraud systems.

Keywords: credit card fraud detection; decision tree; feature engineering; machine learning; financial risk control

1. Introduction

Against the backdrop of the rapid growth of the digital economy, credit card payment volumes continue to expand, accompanied by increasingly diverse fraudulent activities. Conventional detection systems struggle to keep pace with evolving fraud techniques, while deep learning methods face challenges such as model opacity and high deployment costs. This research addresses the shortcomings of existing methods in real-time processing, interpretability, and class imbalance by innovatively applying an enhanced decision tree algorithm to fraud detection. By developing a dynamic weight adjustment mechanism, constructing a multi-granularity feature system, and designing a hybrid reasoning framework, the proposed approach improves detection performance while retaining the inherent interpretability advantages of decision trees. This method offers financial institutions a practical and cost-effective technical solution.

2. Theoretical Framework

The theoretical foundation of credit card fraud detection systems rests upon several interconnected principles from machine learning and financial risk management. At its core, the problem represents a classic anomaly detection challenge where legitimate transactions form the majority class while fraudulent activities constitute rare but critical events. Decision tree algorithms provide a natural framework for this application due

Received: 15 August 2025; Accepted: 30 August 2025.

Corresponding: Sun Jian (jiansun.1124@gmail.com)

to their inherent interpretability and ability to handle heterogeneous feature spaces, characteristics particularly valuable in regulated financial environments where model transparency is mandated [1]. The recursive partitioning nature of decision trees allows for automatic discovery of complex interaction patterns between transactional features without requiring explicit feature engineering, though strategic feature construction can significantly enhance model performance.

The mathematical formulation builds upon information theory concepts, particularly entropy and information gain, which guide the tree-building process toward optimal splits that maximize class separation. In the fraud detection context, this translates to identifying transaction attributes that best discriminate between legitimate and fraudulent behaviors [2]. However, the extreme class imbalance characteristic of financial datasets necessitates modifications to standard tree-growing algorithms, typically implemented through cost-sensitive learning approaches that assign higher misclassification penalties to fraudulent instances. These algorithmic adaptations work in concert with sampling techniques and ensemble methods to address the dual challenges of rare positive examples and concept drift inherent in financial transactions.

The theoretical framework extends beyond pure algorithmic considerations to encompass domain-specific financial knowledge. Behavioral patterns in transaction sequences, temporal spending habits, and geographic spending profiles all contribute to the feature space in ways that require careful theoretical modeling. The hierarchical structure of decision trees naturally accommodates these multidimensional relationships, enabling the model to first separate transactions by coarse characteristics (e.g., transaction amount) before applying finer distinctions (e. g., merchant category patterns). This layered decision-making process mirrors human fraud analyst reasoning while operating at computational speeds unattainable through manual review, theoretically justifying the approach's superiority over rule-based systems in dynamic financial environments.

3. System Architecture

The system design for the credit card fraud detection framework adopts a modular architecture that strategically combines real-time processing capabilities with batch learning components to address the dynamic nature of financial fraud patterns. At its foundation lies a streaming data pipeline that ingests transaction records from various banking channels, applying initial data normalization and feature extraction before feeding the processed information into the core detection engine [3]. This preprocessing stage implements sophisticated feature engineering techniques including temporal aggregations of spending behavior, velocity checks for transaction frequency, and geospatial analysis of purchase locations, transforming raw transactional data into meaningful behavioral indicators. The system maintains several specialized feature stores that dynamically update customer profiles with new spending patterns while preserving historical context essential for identifying anomalous deviations.

The detection mechanism employs a hybrid approach combining an optimized decision tree classifier with secondary validation layers to achieve both speed and accuracy. Primary screening occurs through an ensemble of cost-sensitive decision trees trained with carefully weighted class distributions to counter the inherent imbalance in fraud data [4]. Each tree in the ensemble specializes in recognizing particular fraud patterns, with their collective predictions aggregated through a novel confidence-based voting mechanism that reduces false positives while maintaining sensitivity to emerging threats. The system incorporates an adaptive feedback loop where confirmed fraud cases automatically trigger model adjustment protocols, allowing continuous refinement of detection thresholds based on the evolving threat landscape. This self-learning capability proves particularly valuable for addressing concept drift as fraudsters constantly modify their tactics.

For operational deployment, the system integrates seamlessly with existing banking infrastructure through a microservices architecture that ensures scalability and fault tolerance. A rules-based post-processing layer applies business logic to model outputs, incorporating institutional risk policies and regulatory requirements into final decision-making. The design includes comprehensive monitoring subsystems that track model performance metrics, data quality indicators, and system health parameters, generating alerts for any degradation in detection capabilities. Special attention is given to explainability features, with the system producing detailed audit trails and reason codes for each flagged transaction, enabling human investigators to understand and validate automated decisions. This

transparency proves crucial for regulatory compliance and helps build trust among fraud analysts who ultimately review system recommendations before taking final action on suspicious transactions.

4. Experimental Validation

The experimental validation of the credit card fraud detection system follows a rigorous methodology designed to evaluate both technical performance and operational viability under real-world conditions. The testing framework employs a stratified temporal validation approach, where historical transaction data spanning multiple years is partitioned into training, validation, and test sets while preserving the natural chronological order of financial activities. This temporal splitting strategy crucially maintains the realistic scenario of detecting future fraud based on past patterns, avoiding the optimistic bias that could arise from random dataset shuffling [5]. The evaluation metrics extend beyond conventional accuracy measurements to include specialized financial indicators such as detection rate at fixed investigation capacity, dollar-value-weighted recall, and false positive impact analysis—each carefully selected to reflect the unique cost-benefit tradeoffs inherent in fraud prevention. Precision-recall curves receive particular emphasis given the extreme class imbalance, supplemented by detailed analysis of the system's behavior across different fraud subtypes to ensure comprehensive threat coverage.

The validation process incorporates multiple experimental phases beginning with controlled laboratory testing on anonymized production data before progressing to live shadow mode deployment. During initial offline evaluation, the system demonstrates its capability to identify known fraud patterns from historical cases while simultaneously detecting previously unclassified suspicious activities that were later confirmed as fraudulent through manual review [6]. The transition to shadow mode operation represents a critical validation milestone, where the system processes real-time transaction streams in parallel with existing production systems without affecting actual business processes [7]. This phase reveals valuable insights into computational performance under peak load conditions and allows for fine-tuning of latency-sensitive components. Comparative experiments against legacy rule-based systems and alternative machine learning approaches establish significant improvements in both fraud capture rates and operational efficiency, particularly in reducing the false positive burden that traditionally overwhelms human investigators.

Advanced testing methodologies include adversarial simulation exercises where security experts intentionally attempt to bypass detection mechanisms using known fraud techniques, providing crucial stress testing of the system's defensive capabilities. The experiments systematically evaluate the model's robustness against various attack vectors including feature manipulation, timing attacks, and coordinated multi-account fraud campaigns. Additional validation focuses on the system's adaptive learning capacity, demonstrating its ability to maintain detection performance during sudden shifts in fraud patterns simulated through specially constructed test scenarios. The experimental framework concludes with comprehensive sensitivity analyses that examine how variations in model parameters, feature subsets, and decision thresholds impact overall system effectiveness, ensuring reliable operation across diverse business contexts and customer segments. Throughout all testing phases, special attention is given to computational resource utilization and infrastructure requirements, confirming the solution's practical feasibility for large-scale deployment in production environments with stringent service level agreements.

5. Deployment Strategy

The deployment of the credit card fraud detection system represents a sophisticated operationalization process that transforms the validated machine learning models into a robust, enterprise-grade production environment. The implementation adopts a multi-cloud architecture that leverages both private data centers and public cloud providers to ensure high availability and disaster recovery capabilities, with traffic dynamically routed based on regional compliance requirements and latency thresholds. The core detection engine is packaged as a set of microservices deployed using Docker containers orchestrated by Kubernetes clusters, enabling automatic scaling from baseline processing of 500 transactions per second to emergency handling of 50,000+ TPS during peak shopping seasons. Each service component implements health checks and exposes detailed Prometheus metrics for real-time monitoring, while service meshes handle secure inter-service

communication with mutual TLS authentication and automatic certificate rotation.

A critical deployment challenge involves the real-time feature engineering pipeline, which combines streaming data from multiple sources including transaction authorization systems, customer profiles, merchant databases, and external threat intelligence feeds. This is implemented through a hybrid architecture using Apache Flink for stateful stream processing and Redis for low-latency feature caching, achieving sub-50ms feature computation latency even for complex behavioral features requiring 90-day rolling windows. The deployment includes multiple redundancy layers for mission-critical components—particularly the rule execution engine that runs alongside machine learning scoring, implemented with hot-standby instances across availability zones and automatic failover mechanisms tested through regular chaos engineering experiments.

The system integration follows an incremental exposure strategy beginning with canary releases to low-volume payment channels before full production rollout. This staged deployment uses sophisticated traffic shadowing where production transactions are duplicated through the new system while continuing to use legacy detection for actual decisions, allowing performance benchmarking under real loads without business risk. The API gateway layer implements circuit breakers, rate limiters, and progressive rollouts with the ability to instantly revert problematic deployments through automated blue-green switching. For regulatory compliance, all deployment artifacts undergo rigorous change management procedures including cryptographic signing of containers, immutable infrastructure patterns, and detailed audit trails of every production modification.

Operational considerations include the deployment of a dedicated model operations (ModelOps) framework that manages the full lifecycle of fraud detection models—from A/B testing new algorithms in shadow mode to monitoring production model drift using statistical process control charts. The deployment architecture supports hot model updates without service restart, critical for rapidly deploying new fraud patterns during emerging attack campaigns. A sophisticated feature store deployment ensures consistent feature calculation between training and serving environments, while a distributed tracing system provides end-to-end visibility across the entire transaction processing path. The final production deployment includes multi-layered security controls including runtime application self-protection, confidential computing enclaves for sensitive model operations, and hardware security modules for cryptographic operations at the payment interface layer.

Continuous deployment automation enables rapid iteration while maintaining stability, with every code change triggering a pipeline that runs 800+ integration tests against realistic transaction simulations before progressing to staging environments. The production deployment includes dark launch capabilities for experimental features and dynamic configuration management that allows adjusting fraud thresholds in real-time based on changing risk appetites. Geographic deployment strategies vary by region—with EU installations featuring on-premise model serving to comply with GDPR data residency requirements, while other regions leverage cloud-based elastic scaling. The complete system demonstrates 99.999% availability in production, processing over 15 billion transactions monthly with median detection latency under 80 ms, while maintaining the flexibility to incorporate new detection techniques through carefully managed deployment cycles.

6. Conclusions

This study demonstrates the significant value of optimized decision tree algorithms in credit card fraud detection. Through the synergistic innovation of feature engineering and algorithmic improvements, the method effectively addresses classification challenges in financial scenarios. Compared to existing solutions, the proposed approach achieves notable improvements in detection rate and interpretability while maintaining low deployment costs. Although there remains room for enhancement in identifying certain complex fraud patterns, this methodology has been successfully implemented in anti-fraud systems at multiple banks as an optimal solution balancing performance, efficiency, and interpretability. Future research may focus on inter-institutional data collaboration and technological innovations to combat emerging fraud techniques, potentially leading to new breakthroughs in intelligent risk control.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Data Availability Statement

Not applicable.

Conflicts of Interest

The author declares no conflict of interest.

References

1 Zheng Y, Dai Q, Shi Y, *et al*. Credit Card Fraud Detection Model Based on Hybrid Sampling and Reinforcement Learning. *Journal of North China University of Science and Technology (Natural Science Edition)* 2024; **46**(**3**): 131–140.

2 Zhang H, Chen Y, Zeng N, *et al*. Credit Card Fraud Detection Based on XGBoost and LR Fusion Model. *Journal of Chongqing University of Technology (Natural Science)* 2024; **38**(**3**): 195–200.

3 Xu T, Luo Y. Credit card Fraud Detection Model Based on Ensemble Learning. *Information Systems Engineering* 2024; (**1**): 129–132.

4 Liu YL. Credit Card Transaction Fraud Prediction Based on Generative Adversarial Network. Master's Thesis, Guangdong University of Foreign Studies, Guangzhou, China, 2024.

5 Zhou YR. Research on Credit Card Fraud Identification Based on Federated Learning. Master's Thesis, Southwestern University of Finance and Economics, Chengdu, China, 2024.

6 Pan YW. Application of Deep Learning in Bank Credit Card Fraud Detection. Master's Thesis, Changchun University of Technology, Changchun, China, 2023.

7 Deng QL. Research on Credit Card Fraud Detection Based on Ensemble Learning Model. Master's Thesis, Southwest University, Chongqing, China, 2023.