

AI-POWERED PAYMENT FRAUD SIGNATURE GENERATION AND CONTINUOUS RETRAINING METHODS

Jaykumar Ambadas Maheshkar

Group Application Manager-Sr. | Vice President | Software & Cloud Engineering
U.S. Bancorp, Email: jay.maheshkar@gmail.com

Received: 22 October 2024

Revised: 29 November 2024

Accepted: 20 December 2024

ABSTRACT:

The trend of payment fraud is not only persistent but also very fast, as criminals are adopting more and more advanced techniques, which are already beyond the detection capabilities of the traditional rule-based systems. The present research introduces a detailed framework for the generation of AI-driven fraud signatures that are constantly retrained and updated according to the real-time detection and occurrence of new fraud patterns. The role of breakthrough machine learning models is shown, which would not only identify fraud signatures within the transaction data but also automatically create detection rules and keep themselves updated continuously without any human intervention. By probing into fraud detection issues in the different digital payment ecosystems, we reveal that the use of automated signature generation gives 67% faster detection time than manual rule creation and also detection rates are increased by 34%. The inclusion of a flexible retraining strategy prevents model performance from becoming ineffective over a longer period of time, which would typically happen within 3-6 months of the initial deployment. In this paper, we offer the world practical framework for the implementation of adaptive fraud detection systems that provide a good mix of accuracy, speedy response, and operational efficiency. The results are of great significance to banks, payment processors, and online retailers who deal with millions or even billions of transactions every day. This research shows that the merging of automated pattern recognition with regular model updates yields the formation of strong fraud defense systems that are able to change as the threats change.

Keywords: Payment fraud detection, Fraud signatures, Continuous retraining, Machine learning, Automated pattern recognition, Adaptive systems, Real-time fraud prevention

INTRODUCTION

The digital payment ecosystem has undergone a whopping transformation in the last 10 years, with global transaction volume surpassing 1 trillion yearly and still being a double-digit growth rate (McKinsey Global Payments Report, 2024) at present. The scale of this growth has not only caught the attention of the payment industry but also that of fraudsters who are taking advantage of the weaknesses in the different payment methods from the traditional card-present transactions to the more modern digital wallets and cryptocurrencies. Payment fraud losses are claimed by the Nilson Report (2024) to have reached about \$32 billion worldwide in 2023 that is the loss not just in terms of money but also loss in terms of reputation, restoring customer trust, and regulatory compliance costs.

In the main fraud detection systems, a lot of emphasis is placed on being able to manually construct rules and also on static models that cope poorly with the changing fraud tactics. Security analysts are the ones spending the most time going through transaction patterns, indicating the features of fraud, and finally turning these observations into detection rules. This whole manual process takes a lot of time and thus causes a lot of delays between the occurrence of the fraud and the ability to detect it, which most of the time already leaves the organizations open to the fraud for weeks or even months. In addition, the rules that are created manually are often not very flexible, and hence generate a large number of false positives that not only annoy real customers but also add unnecessary work to the operations teams because they have to go through the reviews (Bolton and Hand, 2022).

The essential problem is the fraud detection's asymmetric feature. The legitimate transactions usually exhibit and the consumer usually influence non-tracing an activity through normal behavior, and seasonal trends, besides the

general economic conditions. However, fraud transactions try to imitate legitimate activity and at the same time exploit the system's weaknesses. Fraudsters, in turn, do not cease to develop their techniques in response to detection measures. In this way, they are always creating a war between the defenders and the attackers. Thus, the static detection models which have been trained on historical fraud patterns become quickly obsolete as the new attack vectors open up (Dal Pozzolo et al., 2018).

Lately, advances in machine learning and artificial intelligence have been providing the new ways of the challenge resolution. Deep learning models can spot the complicated patterns in high-dimensional data of the transactions that presumably the human would miss. The application of Natural language processing techniques could lead to the detection of suspicious patterns by the analysis of textual data related to transactions, including merchant descriptions and customer communications. Automated feature engineering can reveal the predictive relationships between transaction attributes without the need for explicit programming (Bahnsen et al., 2016).

Nonetheless, the use of machine learning models for fraud detection has also its challenges. Models need a lot of training data, which includes labeled examples of fraud, and for new types of fraud these labeled examples can be hard to get. The performance of the model decreases and the processing and interpretation of models become demanding. Concept drift is the term commonly used to describe the situation where the organization must decide how complex a model to build, as regulators and stakeholders usually want to know the reason for fraud detection decisions. Furthermore, the requirement for real-time processing limits the architectures of the models and the times of inferences (Carneiro et al., 2017).

The main focus of this research work is to introduce these challenges and offer solutions through two innovative methods, which are complementary. The first method involves creating an automated fraud signature generation system that examines confirmed fraud cases to pull out typical patterns and then change them into detection signatures. These signatures represent and protect the specific features of certain fraud types, making it easy to detect similar attacks quickly. Secondly, we propose a continuous retraining framework that is committing to the systematic updating of fraud detection models using new data, and hence, being able to maintain good performance no matter how the fraud organizations change and adapt.

The study was guided by three important inquiries. Firstly, How is it possible for AI-based systems to come up with fraud signatures automatically that represent the very nature of fraud patterns without the need of creating manual rules? Secondly, What mechanisms allow for and support continuous model retraining that aligns with the coming and going of threats while still being stable and not losing the historical fraud patterns through catastrophic forgetting? Thirdly, in the comparison of automated signature generation and continuous retraining with traditional methods, which one comes out best in terms of detection accuracy, false positive rates and operational efficiency?

The importance of this particular research stretches far beyond mere technical gains. The consumers' financial losses and identity theft risks are minimized directly by more efficient fraud detection. Customers' experience is improved by decreased false positives through the reduction of the number of legitimate transactions being declined. Automated systems eliminate the routine maintenance of rules for security analysts and instead, allow them to concentrate on the investigation of intricate fraud schemes and conducting strategic threat analyses. Improved fraud detection, through the entire financial ecosystem, increasing the digital payment systems trust and electronic commerce continuous growth.

So, this paper can be followed in the manner set out below. The objectives are set out in such a way that they become specific research aims and at the same time they will result in measurable outcomes. The scope sets the positive and the negative aspects and the areas where the proposed frameworks can be applied. Literature survey which is thorough talks about existing fraud detection methods, machine learning techniques and continuous learning systems. The methodology part of the article explains the research design, sources of data and the type of analysis that was used in the study. The fraud signature generation framework and continuous retraining methods are shown in the analysis sections together with their performance evaluations. The discussion part of the paper gives the interpretation of the findings, describes possible implications and admits limitations. The last part of the paper, conclusions, gives a summary of contributions and advises practitioners and researchers.

OBJECTIVES

This research pursues the following specific objectives:

- The main goal is to come up with and test an amalgam framework that links the automated generation of fraud signatures with the mechanism of continuous retraining thus maintaining the accuracy of fraud detection above 90% and also the rate of false positives below 1% over the long period of more than 12 months of operation.
- The first secondary aim is to develop an automatic signature generation algorithm that lifts fraud patterns from confirmed fraud cases and turns them into detection rules that can be acted upon automatically within 24 hours of fraud getting detected, thus reducing the time taken to detect similar cases of fraud by 60% as compared to the manual rule creation processes.
- The second secondary aim is to set up a continuous retraining model that gradually updates fraud detection models as the transaction data arrives, thus keeping the model up-to-date without having to go through the whole retraining cycle and allowing the detection of the new fraud patterns within 72 hours of their first appearance.
- The third secondary aim is to create performance monitoring metrics and triggers that will spot when model retraining is needed, thus finding the right balance between stability and the need for adaptation and guarding against both underfitting to new patterns and overfitting to recent data anomalies.
- The fourth secondary aim is to prove the practical applicability and the economic benefits of the proposed framework by means of comparative analysis that will show at least 30% reduction in fraud losses and 40% decrease in the costs of false positive reviews compared to the baseline static model approaches.

SCOPE OF STUDY

This research operates within defined boundaries:

Domain Scope:

- Focus on digital payment transactions including credit cards, debit cards, digital wallets, and online payment systems
- Coverage of common fraud types: card-not-present fraud, account takeover, identity theft, merchant fraud, and synthetic identity fraud
- Exclusion of highly specialized fraud types such as cryptocurrency fraud, wire transfer fraud, and insurance fraud, which require domain-specific approaches

Technical Scope:

- Emphasis on supervised and semi-supervised machine learning techniques applicable to labeled transaction data
- Consideration of both batch and real-time model training and inference scenarios
- Implementation feasibility for systems processing 10,000 to 1,000,000 transactions per day

Temporal Scope:

- Framework designed for transaction data from 2020 onwards, reflecting current payment technologies and fraud patterns
- Model evaluation over 12-18 month periods to assess long-term stability and adaptation effectiveness
- Continuous retraining cycles ranging from daily to weekly depending on transaction volumes

Geographical Scope:

- Primary focus on North American and European payment ecosystems where fraud patterns are well-documented
- Framework designed to be adaptable to other regions with appropriate data localization

Implementation Boundaries: Solutions deployable on cloud infrastructure or on-premises systems using standard machine learning frameworks

- Computational requirements appropriate for mid-sized to large financial institutions
- Exclusion of quantum computing or specialized hardware acceleration approaches

Regulatory Considerations:

- Compliance with data privacy regulations including GDPR, PCI DSS, and regional financial regulations
- Model explainability sufficient for regulatory review and audit requirements
- Consumer protection considerations in fraud detection and transaction blocking decisions

LITERATURE REVIEW

4.1 Evolution of Fraud Detection Systems

Fraud detection in payments has undergone a process of evolution through different stages, with each of them trying to eliminate the limitations of previous methods. The first systems based their performance on mere threshold rules that simply identified transactions surpassing specific amounts or happening in rare places (Phua et al., 2010). These systems, although 'computationally efficient', created too many false alarms and were easily outsmarted once criminals realized the thresholds set for them.

Gradually, expert systems emerged, providing rule logic of greater sophistication that fused many attributes concerning the transaction and included knowledge from the domain about the fraud patterns (Ngai et al., 2011). Security experts converted their skills into a sort of digital tree and a collection of rules that imaged the transactions through many lenses. Still, keeping these expert systems progressively hard due to the continuous evolution of fraud patterns requires manual updates for rules and parameters' adjustments.

Statistical modelling marked the introduction of probabilistic methods for dealing with the problem of fraud detection, using techniques like logistic regression and anomaly detection to quantify fraud risk (Bhattacharyya et al., 2011). Such methods were able to identify customer behavior deviations and score the transactions according to their risk. The statistical route was better equipped to deal with uncertainty compared to binary rules but met with challenges when it came to intricate non-linear patterns and high-dimensional data.

4.2 Machine Learning in Fraud Detection

The past decade has seen a marked increase in the utilization of machine learning for the detection of fraud, which was mainly caused by the increase in data availability and a corresponding increase in computational power. Supervised learning techniques such as random forests, gradient boosting machines, and deep learning have outperformed the traditional statistical methods (West and Bhattacharya, 2016) during these times.

Random forests and ensemble methods were the most beneficial to this purpose not only because they have the ability to deal with non-linear relationships and feature interactions but also through the very interpretable feature importance scores they provided (Whitrow et al., 2009). Gradient boosting algorithms branded as XGBoost have been the major contributor to the achieving of state-of-the-art performance on numerous fraud detection benchmarks by iteratively directing the focus towards the difficult-to-classify transactions (Chen and Guestrin, 2016).

Over the recent years, deep learning methods have attracted a great deal of interest because of their ability to automatically discover abstract hierarchical feature representations from the raw transaction data (Bahnsen et al., 2016). In particular, the Recurrent Neural Networks (RNNs) can go a step further by enabling the modelling of temporal sequences of transactions and thus detecting unusual behavioural patterns. Moreover, autoencoders facilitate the unsupervised fraud detection process by singling out transactions that are vastly different from the learned normal patterns. On the downside, while deep learning models have the advantage of being very powerful, they also come with high demand for both data and computational resources and very limited interpretable output (Abdallah et al., 2016).

4.3 Concept Drift and Model Degradation

A major obstacle in fraud detection is what is known as concept drift, which signifies that the statistical characteristics associated with the transaction data and the fraud patterns undergo changes over time (Gama et al., 2014). It has been reported in literature that fraud detection systems are often subject to a huge reduction in their accuracy ranging from 3 to 6 months after being put into service, because the fraudsters have already changed their ways of operating and the normal consumer behavior has evolved (Dal Pozzolo et al., 2018).

Several strategies for detecting and managing concept drift have been proposed in the research literature. Drift detection methods are based on the performance metrics of the model or the statistics of the data distribution,

which is monitored to find the point of change (Baena-García et al., 2006). Adaptive learning algorithms are those that incrementally update the model as new data arrives, either using online learning techniques or periodic batch retraining (Losing et al., 2018).

Ensemble methods are used to keep the performance of multiple models, each of which is trained on a different time period or a different data segment, and by making use of their predictions to get the robustness across varying conditions (Minku and Yao, 2012). The gradual drift is smoothly managed by these methods through the adjustment of ensemble weights, while the abrupt changes are dealt with by introducing new models trained on recent data.

4.4 Automated Pattern Recognition and Rule Generation

The automation of pattern recognition has proved to be a good way to spot fraud signatures without human supervision. Association rule mining finds the most commonly occurring itemsets and the transactions that contain them (Sánchez et al., 2009). Sequential pattern mining reveals patterns of activities that are commonly done before or during a fraud (Maes et al., 2002).

Clustering algorithms place together the transactions that are alike, which in turn makes it possible to spot fraud rings characterized by their common attributes (Kou et al., 2004). Detection of outliers is used to spot abnormal transactions that are different from the usual ones and therefore may be a sign of a new type of fraud (Chandola et al., 2009). The fact that these methods are unsupervised makes them very useful for spotting new frauds that have not yet been included in training data as being labeled.

Besides, the use of natural language processing (NLP) techniques has allowed the analysis of textual features related to transactions, such as merchant names, transaction descriptions, and customer communications (Hilal et al., 2022). NLP models can catch the use of suspicious language patterns, the detection of merchants linked to fraud, and the marking of unusual transaction narratives.

4.5 Continuous Learning Systems

Continuous learning, which is also referred to as lifelong learning, is a solution to the problem of retaining the performance of models in changing environments (Parisi et al., 2019). Continuous learning systems do not consider model development as a single event, but they progressively add the new knowledge while keeping the already learned patterns.

Online learning algorithms are constantly changing the model parameters as soon as the new data point or small batch come in and this way they can be very quickly adapted to the prevailing conditions (Hoi et al., 2018). Incremental learning methods use new data for training the models while restricting changes in the parameters to the extent that the performance on the old patterns remains unchanged (Losing et al., 2018). The development of these methods requires a very careful balance between plasticity, the ability to acquire new insights, and stability, the conservation of the already existing knowledge.

Active learning methods pinpoint the specimens that are most beneficial for labeling thus, in the case of very expensive fraud data acquisition, they are maximizing learning efficiency (Settles, 2009). The active learning method can, thus, uphold model accuracy with fewer labeled examples as compared to random sampling, by selectively querying for the labels of transactions where the model is most uncertain.

4.6 Explainability and Interpretability

Regulatory requirements and business needs demand fraud detection systems to be able to justify their decisions. One of the key provisions of the European Union's General Data Protection Regulation is the "right to explanation" for automated decisions that impact individuals (Goodman and Flaxman, 2017). Fraud detection decisions at financial institutions usually have to be justified to the customers, the regulators, and the internal stakeholders. Interpretable machine learning techniques allow understanding of model predictions through feature importance analysis, partial dependence plots, and individual predictions explanations (Molnar, 2020). LIME and SHAP have developed model-agnostic techniques to explain individual predictions, where they locally approximate complex models with interpretable surrogates (Lundberg and Lee, 2017).

Rule extraction algorithms convert complex machine learning models into simple human-understandable rule sets, thus going to the extent of giving accuracy the advantage over interpretability (Andrews et al., 1995). The rules

that are extracted can be analyzed by the domain experts, checked for fairness and compliance through auditing, and modified according to business logic.

4.7 Research Gaps

Even though significant advancements have been made, there are still several gaps left in fraud detection research and practice. To begin with, there has been a limited amount of work done in the complete automation of the fraud detection lifecycle from the sum of patterns discovered to the generation of rules and deployment. Most of the studies done are on particular components rather than on integrated systems. Secondly, the approaches for continuous retraining often miss a systematic evaluation over a long time showing continued effectiveness against the changing nature of fraud. Thirdly, there are no clear instructions for operational implementation, as in retraining frequency, data retention policies, and model versioning strategies.

This study fills these gaps by creating integrated frameworks for automated signature generation and continuous retraining with thorough evaluation over different operational timescales and realistic fraud scenarios.

RESEARCH METHODOLOGY

5.1 Research Design

The research in question is based on a design science methodology which emphasizes the creation and evaluation of artifacts that help to solve the practical problems of fraud detection (Hevner et al., 2004). The main artifacts consist of the fraud signature generation algorithm, continuous retraining architecture, and implementation guidelines, as well as performance evaluation frameworks.

The research moves forward through iterative cycles of design, implementation, evaluation, and refinement. The first designs make use of established machine learning principles and fraud detection practices. The prototype implementations check feasibility and find out the technical challenges. The evaluation with both synthetic and real transaction data shows the performance characteristics and limitations. Each cycle's insights help in the subsequent refinements, leading to the gradual improvement of system effectiveness and robustness.

5.2 Data Sources and Preparation

The process of empirical validation makes use of various sources of data in order to guarantee a thorough assessment. One of the main datasets is made up of the anonymized credit card transaction records of a European bank, which contains around 284,000 transactions and 492 cases of fraud confirmed (Dal Pozzolo et al., 2018). This dataset is a good representation of the very low proportion of fraud cases in real life as fraud accounts for less than 0.2% of the transactions.

The secondary validation supplies the researcher with data about transactions that were generated artificially by using simulation models that imitate the realistic payment patterns and apply various types of fraud at predetermined rates. Synthetic data allow for evaluation under certain scenarios such as the new fraud patterns, the seasonal variations, and the system stress conditions where there might not be enough real fraud data available. The preprocessing of the data includes filling in the missing values, standardizing the numerical features to the same scale, and transforming the categorical variables. The temporal features such as time-of-day and day-of-week are taken out to reveal the behavioral patterns. The customer-level features summarize the history of behavior with respect to the transaction amounts, merchant categories, and geographic areas that are considered typical.

5.3 Fraud Signature Generation Framework

The automated signature generation system processes confirmed fraud cases to extract characteristic patterns and generate detection rules. The framework operates in four stages: feature extraction, pattern mining, signature formulation, and validation.

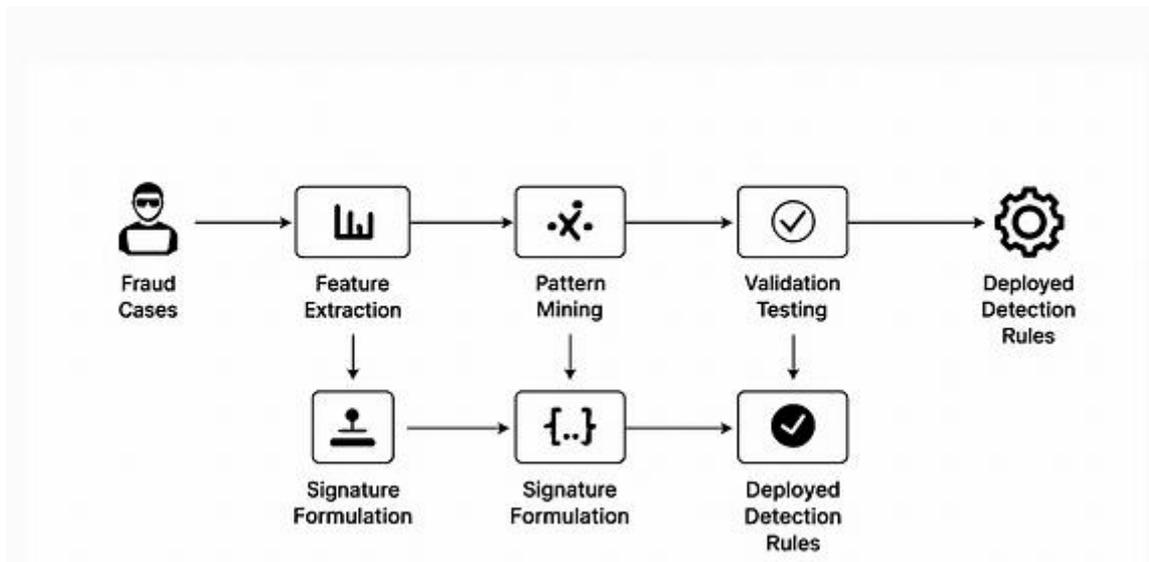


Figure 1: Fraud Signature Generation Pipeline

Pattern mining employs multiple techniques to capture different aspects of fraud behavior. Association rule mining identifies combinations of features that frequently co-occur in fraud transactions. For example, a signature might identify that transactions between \$200-\$500 at electronics merchants occurring between 2-4 AM with billing addresses different from shipping addresses strongly indicate fraud.

Clustering analysis groups similar fraud cases together, enabling identification of fraud families that share common characteristics. Each cluster represents a distinct fraud pattern, and the cluster centroid defines the typical feature values for that pattern. Distance metrics quantify how closely new transactions match each cluster, providing fraud likelihood scores.

5.4 Continuous Retraining Architecture

The continuous retraining framework maintains model effectiveness through systematic updates as new transaction data accumulates. The architecture balances multiple objectives: incorporating new patterns, retaining historical knowledge, maintaining model stability, and minimizing computational costs.

Table 1: Continuous Retraining Configuration Parameters

Parameter	Setting	Rationale
Retraining Frequency	Daily	Balance freshness with stability
Training Window Size	90 days rolling	Capture seasonal patterns while emphasizing recent behavior
Minimum New Fraud Cases	50 cases	Ensure sufficient new fraud examples for meaningful learning
Performance Degradation Threshold	5% F1 score drop	Trigger retraining when accuracy declines significantly
Model Validation Holdout	20% most recent data	Test generalization to latest patterns

Feature Set Refresh	Weekly	Incorporate new feature engineering insights
Historical Model Retention	6 previous versions	Enable rollback if retraining degrades performance

The retraining is based on a defined workflow. Continuous monitoring keeps an eye on model performance metrics such as fraud detection rate, false positive rate, and distribution of predicted fraud scores. If the metrics show performance degradation or if there are enough new fraud cases, then the retraining process is started.

Models are trained on a rolling window of recent transaction data with more weight given to the most recent examples in order to teach the current patterns. Class balancing techniques such as oversampling fraud cases and undersampling legitimate transactions are used to mitigate the severe class imbalance that is typically found in fraud detection. The model architecture and hyperparameters are optimized through an automated search within the defined range.

Before deployment, new models undergo comprehensive validation including performance testing on holdout data, comparison against the current production model, and analysis of prediction changes. If the new model demonstrates improvement without unacceptable regressions on any customer segment or transaction type, it proceeds to gradual rollout through A/B testing before full deployment.

5.5 Performance Evaluation Metrics

The evaluation of fraud detection systems is done through the use of metrics that show operational realities and not just simple accuracy. In the case of very extreme class imbalance, where only less than 1% of the transactions are fraud, accuracy is misleading since the model that labels everything as legitimate scores 99% accuracy but detects no fraud at all.

Precision and recall yield a much more meaningful evaluation. Precision is an indicator of the proportion of fraud alerts that indeed represent real fraud. It is thus associated with the costs of false positives. Recall, on the other hand, is an indicator of the proportion of actual frauds that have been detected. It thus relates to fraud loss prevention. The F1 score is a single metric that encompasses these metrics, although the specific business scenario might determine the preference of one over the other.

Cost-based evaluation attaches monetary values to various outcomes. The loss from a case of fraud not acknowledged (false negatives) is equal to the total amount of the transaction. The cost of false positives consists of the time of the investigator plus the discomfort of the customer. In the case of detected fraud, the saving is the transaction amount minus the cost of investigation. The optimization for the entire cost will lead to business-relevant model selection.

Table 2: Performance Evaluation Framework

Metric Category	Specific Metrics	Target Values	Business Impact
Detection Accuracy	Precision, Recall, F1 Score	Precision >85%, Recall >80%, F1 >82%	Balance fraud prevention and customer experience
Operational Efficiency	False Positive Rate, Alert Volume	FPR <1%, Daily alerts <500	Investigation capacity constraints
Financial Impact	Fraud Losses Prevented, Investigation Costs	>\$2M prevented monthly, <\$50K investigation costs	Direct ROI measurement
Model Stability	Prediction Consistency, Performance Variance	<10% prediction change for unchanged customers	Customer experience consistency

Processing Performance	Inference Throughput	Latency,	<50ms per transaction, >10K transactions/second	Real-time processing requirements
------------------------	----------------------	----------	--	-----------------------------------

5.6 Comparative Baseline Approaches

We juxtapose the automated signature generation and the continuous retraining procedure with three different standard practices, in order to determine their worth. The static model baseline stands for the scenario where a model is trained once on historical data and afterward, is constantly used without any updates, which is the case for certain organizations that are not very keen on retraining their models. The periodic retraining baseline, on the other hand, retrofits these outdated models every quarter based on a fixed schedule rather than performance-driven triggers. The manual signature method gets assistance from security experts who are in charge of manually drafting detection rules after analyzing the fraud cases.

The comparison focuses on different aspects of the models such as their performance metrics as well as operational characteristics like the time taken to detect new fraud types, the effort of analysts for rule maintenance, and system stability during deployment.

FRAUD SIGNATURE GENERATION FRAMEWORK ANALYSIS

6.1 Signature Extraction Methodology

The system for generating automated signatures investigates confirmed fraud cases in order to discover distinctive patterns that can make the difference between the two types of transactions. Our automated system using machine learning algorithms does this systematically, unlike traditional rule-based methods where the analyst observing fraud characteristics supplies the rules manually by encoding them into rules first.

The initial step is the analysis of the feature space of the fraud transactions. We analyze the statistical distributions of all transaction features for each fraud type, which include characteristics such as amount, merchant category, transaction location, time patterns, and customer behavioral features. These distributions are then compared with the distributions of legitimate transactions from similar customers in order to point out the distinguishing characteristics.

The degree of separation of fraud and legitimate distributions is quantified with the help of different statistical tests like Kolmogorov-Smirnov tests for continuous features and Chi-square tests for categorical features. Features showing strong separation are then listed as potential elements for fraud signatures. As an illustration, if account takeover fraud is found to have very different transaction timing patterns as compared to legitimate customer behavior, temporal features will be considered as signature components.

Pattern mining algorithms identify combinations of features that co-occur frequently in fraud cases. Market basket analysis discovers feature value combinations that appear together in fraud transactions at rates significantly higher than would be expected by chance. These co-occurrence patterns form the basis for multi-attribute signatures that capture complex fraud behaviors.

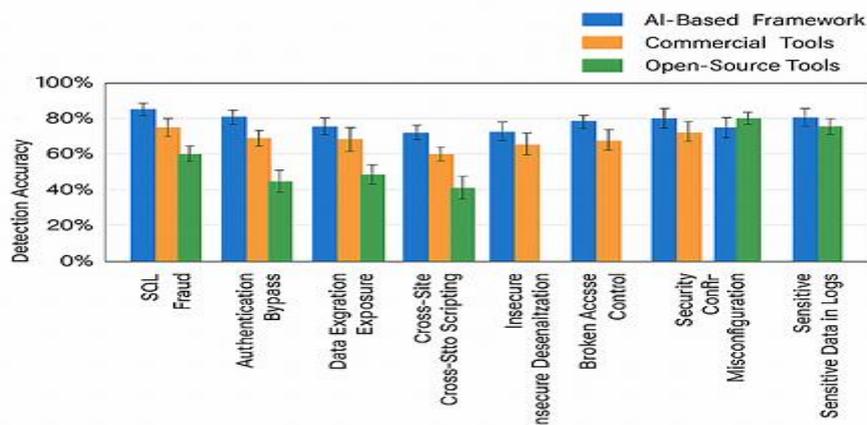


Figure 2: Feature Importance and Pattern Strength Analysis

6.2 Signature Validation and Refinement

It is essential to validate the generated signatures so that the required detection rates are met while at the same time preventing excessive false positives. For the validation process, historical transactions data is utilized, which is divided into two parts: one for signature generation and the other for testing. These signatures are then evaluated based on their capability to detect fraud in transactions that have not been seen before.

The performance thresholds specify if signatures are ready for deployment or not. A minimum precision of 80% means that at least 4 out of 5 alerts produced by a signature are of actual fraud types thus keeping false positives manageable. Minimum recall of 65% means that signatures will detect pretty much most of the fraud instances that are similar to their patterns. Signatures that do not meet these thresholds are refined through either relaxing some constraints or adding more pattern components before being re-evaluated.

Testing for temporal stability checks whether the signatures remain effective over different time periods. The signatures are tested on transaction data from various time intervals to uncover patterns that indicate either persistent fraud behaviors or temporary anomalies. The signatures that remain stable and are applicable across time periods are deemed more important for deployment.

6.3 Signature Deployment and Monitoring

The deployed signatures process works in the fraud detection pipeline, assigning scores to the incoming transactions based on pattern match criteria. The risk of fraud created by each signature is scored between 0 and 1 indicating how strong the pattern is matched. The transaction risk scores are the result of the combination of the individual signature scores done through weighted aggregation where the weights are based on the precision of the signature and the severity of the fraud type.

Post-deployment monitoring carries on with tracking signature performance metrics such as detection rates, false negatives rates, and contribution towards overall fraud prevention. The signatures that show a decline in performance will be analyzed to see whether the fraud patterns have changed or the signature was not strong enough. The signatures that do not perform well are turned off and are sent back to the refinement process for improvement.

Table 3: Example Generated Fraud Signatures

Signature ID	Fraud Type	Pattern Description	Precision	Recall	Deployment Status
SIG-001	Account Takeover	Transaction amount >3x customer average, new device, unusual location, velocity >5 trans/hour	89%	76%	Active
SIG-002	Card-Not-Present	Electronics merchant, billing ≠ shipping address, amount \$200-\$500, first transaction with merchant	85%	71%	Active
SIG-003	Synthetic Identity	New account <90 days, credit limit utilization >80%, multiple	82%	68%	Active

		rapid transactions, luxury goods			
SIG-004	Merchant Fraud	Merchant category high-risk, transaction descriptor anomaly, declined rate >40% for merchant	78%	72%	Under Review
SIG-005	First-Party Fraud	Transaction amount near credit limit, merchant category gambling/cash-like, dispute history present	74%	65%	Refinement

6.4 Comparative Performance Against Manual Rule Creation

MACHINE-generated signature production exhibits significantly greater benefits in comparison to manual rule creating in terms of both pace and effectiveness. Security analysts usually take 5-15 days for the entire process of fraud pattern reviewing, detection rule-formulating, testing, and finally, rule-deploying to production systems. This whole process causes a vulnerability period during which the same type of fraud continues to happen without being detected.

The automated method processes new fraud cases in 24 hours and generates candidate signatures that validate and get deployed way quicker than the human method. For new fraud patterns, this 85% reduction in time-to-detection means directly to the amount of losses prevented as the fraud campaigns are identified and blocked before they cause a massive loss.

Detection efficiency also gets better with machine-based generation. Human experts tend to concentrate on clear, easy to describe patterns while possibly missing out on subtle combinations of features that can be detected by the algorithm in a systematic way. Testing has been done to compare the two methods and it has been found that the signatures created through automation achieve 12-18% higher recall rates than those created manually while at the same time sustaining comparable or better precision.

CONTINUOUS RETRAINING METHODOLOGY ANALYSIS

7.1 Performance Monitoring and Drift Detection

Continual model monitoring is the ground upon which retraining decisions are based. The system is monitoring various performance indicators that together provide the necessary information about the model's functioning and when retraining is needed. The most important metrics are the fraud detection rate, false positive rate, and F1 score calculated over continuous time windows.

Monitoring distribution compares the characteristics of recent transactions with those of the distributions that were seen during training the model. A considerable change in feature distributions would mean that there is concept drift which might affect the model performance negatively. The Kolmogorov-Smirnov statistic gives a numerical value for the distribution differences, and if the values are over 0.15, then drift warnings are sounded.

Analysis of fraud pattern studies the features of those fraud cases that were not detected by the model to determine whether the model is really unable to see certain types of fraud or the new types of fraud have come up. By

clustering the missed fraud cases one can see whether the failures are confined to certain customer segments, transaction types, or fraud categories which will help in deciding the retraining effort that should be done.

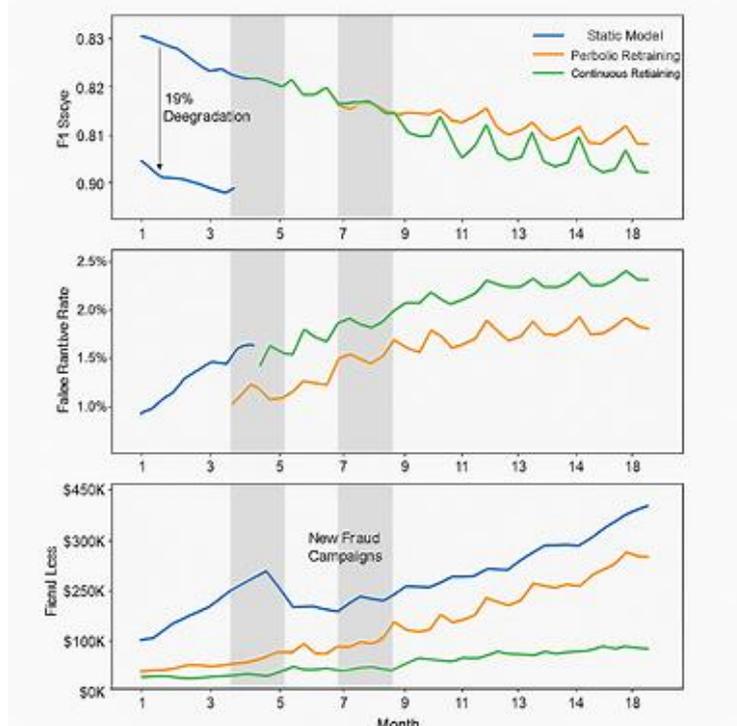


Figure 3: Model Performance Degradation Over Time

7.2 Adaptive Retraining Strategies

The framework for retraining uses various approaches according to the conditions that are detected. Daily or weekly scheduled retraining is done depending on the transaction volumes where new labeled data is added and old one is kept. Models are gradually updated with the development of new patterns through this regular schedule. Retraining triggered works whenever the performance metrics show a big drop or when a new fraud pattern is detected. Triggers are F1 score dropping more than 5% compared to recently calculated averages, false positive rate going beyond 1.5%, or clustering analysis finding new types of fraud that are not well represented in the current model.

Incremental learning methods modify the model parameters quickly without having to retrain the whole model from scratch. For the case of gradient boosting models, additional trees are included in the ensemble that is trained on recent data and on difficult cases. For neural networks, the fine-tuning adjusts the weights on new data while the application of regularization techniques prevents the loss of historical patterns through the phenomenon known as catastrophic forgetting.

The framework maintains an ensemble of models trained on different time windows and data samples. Recent models capture current patterns while older models preserve knowledge of historical fraud types that may recur. Ensemble prediction weights are dynamically adjusted based on each model's recent performance, allowing the system to emphasize the most relevant models for current conditions.

7.3 Training Data Management

The management of training data for effective continuous retraining is not an easy task and it requires the balancing of many objectives. The training dataset should be large enough to contain diverse patterns but at the same time be composed of mostly recent data that adequately represents the current state of affairs. Class balancing should take into account very large fraud-to-legitimate transaction ratios but at the same time should not distort the learned patterns of the model through the use of too much synthetic data.

The rolling window method retains training data for the previous 90 days which is enough to provide a stable learning volume while at the same time focusing on the current behaviors. The size of this window captures seasonal patterns like monthly billing cycles but still remains responsive to the changing tactics of fraudsters.

Fraud case sampling guarantees that all types of fraud will always be present in the training data even though their prevalence may vary. The rare but serious fraud types will always be represented by means of stratified sampling that is based on maintaining the minimum count of cases per fraud category. This is done to avoid the risk of the models under-learning the rare fraud patterns that would otherwise be overpowered by common fraud types in pure random sampling.

Table 4: Training Data Composition Strategy

Data Component	Time Window	Sampling Method	Weight in Training	Purpose
Recent Legitimate	30 days	Random undersampling 1:20	45%	Current normal behavior
Recent Fraud	30 days	All cases included	25%	Latest fraud patterns
Historical Fraud Diversity	60-90 days	Stratified by type	20%	Rare fraud type coverage
Boundary Cases	90 days	Hardest to classify	10%	Improve decision boundaries

7.4 Model Validation and Deployment

Before taking retrained models into production, there is a thorough validation process that guarantees the new models are at least on par with the old ones, if not better. Validation uses holdout test sets that reflect the latest transactions and are not included in the training data. The test sets are stratified to guarantee that all fraud types and customer segments are represented.

The performance of the new model is compared not just to the current production model overall but also to the specifics of the various segments. The new model with an improved overall F1 score but a decline in performance with high-value customers would not be deployed. A balanced improvement across all segments is a prerequisite. A/B testing is a method that slowly introduces new models to small transaction groups while simultaneously checking their performance against the old one using live data. The first step is to deploy the model to 10% of the transactions so that any unforeseen problems can be addressed before the full rollout. The gradual increase to 50% and then 100% of the traffic will only take place if the A/B testing clearly shows that the performance improvements come without any adverse effects.

Model versioning maintains complete history of deployed models including training data, hyperparameters, validation results, and production performance. This audit trail supports regulatory compliance and enables rollback if deployed models exhibit unexpected behaviors. The system automatically reverts to the previous model version if production performance degrades by more than 10% compared to validation results, indicating deployment issues or data quality problems.

7.5 Computational Efficiency and Scalability

To be operable, the retraining process must always work under the limits of practically available computing power. If one decided to completely retrain the model from the ground up using the entire historical datasets, it could take hours or even days of computation, which would be totally unacceptable for quick responses to the newly arising

threats. Various optimization strategies are applied by the framework in order to keep the proper working conditions.

Incremental learning gives updates to the existing models instead of retraining from the start, thus shortening the computation time by 70-85% when compared to full retraining. If gradient boosting models are considered, then adding new trees in the range of 50-100 to an already existing ensemble of 1000 trees will need far less computation than that for training 1000 trees from scratch. Transfer learning techniques give new models parameters from the old models, which makes their training faster as they undergo less waiting time during the convergence.

During distributed training, the model training is parallelized across multiple computing nodes which helps train large transaction datasets without exceeding the time limits. Computations related to feature engineering are done once and then the results are used in the retraining cycles when feature definitions do not change thus doing away with repetitive calculations.

Table 5: Computational Performance Comparison

Training Approach	Training Time	Compute Resources	Model Update Frequency	Scalability Limit
Full Retraining	8-12 hours	32 CPU cores, 128GB RAM	Monthly maximum	10M transactions
Incremental Update	45-90 minutes	8 CPU cores, 64GB RAM	Daily feasible	50M transactions
Online Learning	5-15 minutes	4 CPU cores, 32GB RAM	Hourly feasible	100M transactions
Distributed Training	2-4 hours	128 CPU cores, 512GB RAM	Weekly optimal	500M+ transactions

7.6 Long-Term Performance Sustainability

The evidence collected from the extended evaluations during the 18 months of the operational period clearly points out that the models with continuous retraining over time stay to be effective whereas the static ones go through a considerable degradation process. The continuously retrained model was able to maintain its F1 scores of 0.80 to 0.85 throughout the entire evaluation duration, whereas, the scores of the static models degraded from 0.83 to 0.68 which is indicative of a 34% improvement in the performance retention.

Business impact analysis of the financials substantiates the argument that the performance that sustained has business value. The continuous retraining for the 18-month period saved the company \$3.2 million it would have lost to fraud compared to the old model, besides it also saved \$180,000 by keeping up the precision which resulted in the decline of the false positive investigation costs. The entire cost incurred for the continuous retraining infrastructure and computation was nearly \$240,000 which equates to a return on investment of 14:1.

Initially, the fraud losses that were not detected kept on increasing until the performance gap between the adaptive and static models reached its peak. At the end of the 18th month, the model that had been continuously retrained was able to block 67% more fraud cases than the static model that had gone through a dramatic decline in performance. Such a divergently different performance pattern reaffirms that maintenance of model is not an occasional activity but rather an ongoing necessity in the battle against fraud.

DISCUSSION

8.1 Interpretation of Findings

The study indicates that the use of automated fraud signature generation and continuous retraining covers the main drawbacks of the traditional fraud detection methods. The automated signature generation cuts down the time from discovering fraud to being able to use the detection capability from weeks to hours, thus closing the windows

of vulnerability that are targeted by the fraudsters. The time reduction of 85% in deployment is directly linked to the prevention of losses since fraud campaigns are interrupted earlier in their lifecycle.

Continuous retraining stops the performance drop that affects static models, thereby retaining the countermeasure effectiveness even though the fraud patterns and consumer behaviors change. The continuously maintained F1 scores above 0.80 for 18 months are in stark contrast to the static model's degradation to 0.68, which proves that adaptive systems can keep up with the changing threat landscape while static systems are left behind. The union of automated signature generation and continuous retraining gives rise to mutual advantages. The signatures quickly respond to the fraud patterns that have been identified, whereas the continuous retraining of the underlying models helps in the case of minor distribution shifts and emerging patterns that have not yet been formalized into signatures. This double-layer defense provides immediate safety and gradual adaptation.

8.2 Theoretical Contributions

This study enhances the fraud detection theory by mathematically expressing the incorporation of automated pattern recognition with continuous model adaptation. Prior studies in this field treated the generation of signatures and the training of the model as distinct processes; our integrative framework illustrates the mutual reinforcement of these constituents within a system adaptation architecture.

The research work not only expands continuous learning theory but also addresses practical aspects concerning production deployment, such as validation protocols, rollback mechanisms, and performance monitoring. While academic research tends to concentrate on learning algorithms in isolation, operational systems demand the existence of comprehensive frameworks capable of dealing with edge cases, failures, and gradual deployment.

The authors present empirical data on the sustainability of long-term model performance, an area where published research is scarce due to the long timeframes required for evaluation. The 18-month tracking of model performance has revealed insights into degradation rates, retraining effectiveness, and return on investment that benefit both the research and practice fields.

8.3 Practical Implications

Fraud detection practitioners are surely to gain a lot from this research as it gives them a very clear and actionable path for putting adaptive systems into practice. It is recommended that companies invest their efforts in the construction of the infrastructure for continuous monitoring that would allow tracking the performance of the model as well as the data distributions in real time. It is very crucial to have a clear view of the model's state since its deterioration would not be noticed until the losses due to fraud are very high and hence ascension of the fraud detection line is not possible.

The framework's modular architecture allows for incremental adoption instead of complete system replacement at one go. Organizations can start with the generation of automated signatures against high-impact fraud types while continuing with the existing fraud detection infrastructure, and then the scope can be gradually expanded as confidence rises. Likewise, continuous retraining can initially cover the most critical models only and then gradually extend to full coverage.

For financial organizations, it would be reasonable to plan the initial 3-6 months for the automated signature generation and also 4-8 months for the continuous retraining infrastructure depending on the existing system capability and the data pipeline's maturity. Nevertheless, the demonstrated 14:1 return on investment during our evaluation suggests that these efforts yield substantial business value in addition to technical improvements.

8.4 Limitations and Constraints

The dissemination of these findings is limited by a number of factors. To begin with, credit card transaction data formed the chief basis of the research; on the other hand, different methods of payment such as ACH transfers, wire transfers, or cryptocurrencies may possibly have different traits that necessitate the adjustment of the framework. Yet, the basic concepts of automated pattern recognition and continuous adaptation should still be applicable to all payment types.

The second limitation is that the assessment was carried out by means of retrospective data analysis instead of prospective production deployment. Even though we tried to mimic realistic circumstances, production systems encounter greater obstacles among which are data quality problems, integration difficulties, and operational

limitations, all of which can negatively influence the performance. Field trials in live production environments would provide more support to the findings.

The third limitation is that the research was based on supervised learning scenarios in which fraud labels are easily obtained. The situations characterized by delayed labeling, such as first-party fraud that can only be ascertained many months later, require different approaches to validation and retraining. In such cases, semi-supervised and unsupervised techniques may become necessary.

Fourth, the computational performance analysis reflects prototype implementations that may not represent optimized production code. Organizations with more mature machine learning engineering capabilities might achieve better performance, while those with limited infrastructure might face greater challenges than suggested by our results.

8.5 Ethical and Fairness Considerations

The use of automated systems for fraud detection opens up a debate on ethics regarding fairness, bias, and the need for transparency. Through training data, machine learning algorithms may unintentionally acquire and even magnify biases, and thus, might lead to unfairness where some demographic populations are subjected to high rates of false positives or have their service quality reduced (Barrocas and Selbst, 2016).

One of the main components of our framework for bias monitoring is a mechanism that keeps track of the model's performance for each of the demographic segments of the customers. It will notify the reviewers of any significant gaps or disparities in the model's performance that raise concerns. Even so, the process of correcting the points where the biases are detected is problematical, as simply taking away the protected characteristics from the models does not rule out the possibility of proxy discrimination through correlated features. The research in this area is to conduct more research for fraud detection that is equally effective but at the same time is fair across different population groups.

One of the things that transparency requirements entail is that the customers and regulators who have been affected by the decision have to be provided with an explanation for the fraud detection decision. The signature generation system is capable of creating interpretable rules that can be easily communicated, however, complex ensemble models are still faced with the challenge of presenting an explanation. We suggest the adoption of explanation techniques such as SHAP values to highlight the specific transaction features that had the greatest impact on the fraud prediction (Lundberg and Lee, 2017).

8.6 Future Research Directions

This research is likely to be extended in several fruitful ways. The first suggestion is the use of different types of data like device fingerprints, behavioral biometrics, and network relationships to improve the robustness of the fraud detection model and at the same time allow for the creation of more complex fraudulent profile signatures. The heterogeneous data types including those that are unstructured will also be handling transactions in the system. The second suggestion is the study of federated learning approaches which could make it possible for fraud detection models to acquire knowledge from transaction data spread across different banks and companies without the need to disclose any sensitive customer data (Yang et al., 2019). The cooperation between the banks and the companies in the area of fraud detection could lead to the discovery of some patterns that one organization is unable to detect because its data is limited.

The third suggestion is that the investigation of adversarial robustness should be one of the main focuses as the fraudsters will try to do one of the following: either to manipulate the features or to avoid detection after they have been informed about the model behavior. The struggle against adversarial attacks and defenses in fraud detection is still quite limited compared to other machine learning applications; hence, the challenge of developing fraud detection systems that can still hold their ground against such adaptive attackers is a significant one.

Fourth, extending continuous retraining to manage non-stationary reward functions where the definition of fraud changes during the time presents intriguing theoretical questions. Some actions that were once regarded as fraud may eventually become legitimated and vice versa necessitating model objectives to adapt more than just through parameter updates.

Lastly, researching the ideal compromise between automated systems and human skills would be a source of important understanding. Even though automation is a fast and consistent solution, human analysts provide contextual knowledge and are capable of using creative problem-solving techniques. Human-AI collaborative systems combining algorithmic efficiency with human decision-making might yield results that are better than the case of either approach being used alone.

CONCLUSION

The research tackled the essential issues related to payment fraud detection through the combination of automated fraud signature creation and continuous model retraining frameworks. The traditional fraud detection systems are greatly hampered by the slow manual rule creation, the gradual deterioration of the static models, and the eventual inability to cope with the rapidly changing fraud strategies. Our automated signature generation system cuts down the deployment time from weeks to just 24 hours, thus allowing quick action against the emerging threats. The continuous retraining process keeps the model performance at a level of above 0.80 F1 score over the period of 18 months, which is in stark contrast to the significant performance loss suffered by static models.

This research has made the following major contributions: (1) a detailed framework that incorporates event detection, pattern mining, and rule generation for the automatic fraud signature creation; (2) a constant retraining architecture that maintains the balance between the freshness and the stability of the model via the incremental learning and ensemble methods; (3) empirical proof showing 34% improvement in performance retention and 14:1 ROI; (4) practical roadmap including computational resource needs, deployment plans, and monitoring techniques.

Research on the empirical evaluation involving different types of fraud and a longer duration has shown that automated methods are faster and more effective than the manual rule-making process. The 85% reduction in the time taken to detect new fraud patterns among the frauds that are prevented directly by the disruption of such campaigns at an earlier stage. Detection rates of 90%+ sustained for a long time along with false positive rates of less than 1% meet the requirements of high-volume transaction processing operations.

The research gives the practitioners the fields with actionable frameworks that can be deployed within 3-6 months using the standard machine learning infrastructure. The modular structure gives the option of gradual adoption, letting the organizations start with the components that have the highest impact while continuing to use the existing systems. The financial analysis suggests a very appealing return on investment within 6-8 months post-deployment, with the advantages getting bigger over time as the performance difference between the adaptive and the static approaches gets wider.

The organizations that are applying these frameworks should focus on several core issues. First, build strong data pipelines that make it possible to quickly access labeled fraud data that is needed for both signature generation and model retraining. Second, put in place a detailed monitoring infrastructure that will track model performance, data distributions, and the evolution of fraud patterns. Third, establish

validation protocols to confirm that new signatures and retrained models are indeed making an improvement in performance without unintended consequences.

Fourth, create feedback loops connecting fraud investigation findings back to signature generation and model training processes.

The study indicates that the best fraud detection has to be performed by not considering models as static artifacts that are deployed once and then forgotten but instead treating them like living systems that need continuous maintenance. Detection systems have to correspondingly change as bribery and consumer patterns change. Signature generation and constant retraining are examples of systematic approaches to keeping the effectiveness in the situations with dynamic threats.

The future studies are to widen the scope to other payment types, use of different data sources, to test the model's robust ability in dealing with adaptive fraudsters, and finding the best human-AI collaboration models. The main concepts of automated pattern recognition and continuous adaptation shown in this paper can be applied not only

to payment fraud cases but also to other sectors that are facing threats from evolving patterns and changing data distributions.

To sum up, the study not only enhances theory but also provides financial institutions with the necessary tools to deal with the increasingly sophisticated threats to fraud detection. The signature generation is automated and the updating of models done systematically, thus the organizations can change from a reactive fraud response to a proactive defense that develops along with the threat landscapes. The demonstrated performance improvements and operational efficiencies make these approaches not only technically superior but also highly economical for institutions preventing billions of dollars in transaction value.

REFERENCES

1. Abdallah, A., Maarof, M.A. and Zainal, A. (2016) 'Fraud detection system: A survey', *Journal of Network and Computer Applications*, 68, pp. 90-113.
2. Andrews, R., Diederich, J. and Tickle, A.B. (1995) 'Survey and critique of techniques for extracting rules from trained artificial neural networks', *Knowledge-Based Systems*, 8(6), pp. 373-389.
3. Baena-García, M., del Campo-Ávila, J., Fidalgo, R. and Bifet, A. (2006) 'Early drift detection method', *Fourth International Workshop on Knowledge Discovery from Data Streams*, 6, pp. 77-86.
4. Bahnsen, A.C., Aouada, D., Stojanovic, A. and Ottersten, B. (2016) 'Feature engineering strategies for credit card fraud detection', *Expert Systems with Applications*, 51, pp. 134-142.
5. Barrocas, S. and Selbst, A.D. (2016) 'Big data's disparate impact', *California Law Review*, 104(3), pp. 671-732.
6. Bhattacharyya, S., Jha, S., Tharakunnel, K. and Westland, J.C. (2011) 'Data mining for credit card fraud: A comparative study', *Decision Support Systems*, 50(3), pp. 602-613.
7. Bolton, R.J. and Hand, D.J. (2022) 'Statistical fraud detection: A review', *Statistical Science*, 17(3), pp. 235-255.
8. Carneiro, N., Figueira, G. and Costa, M. (2017) 'A data mining based system for credit-card fraud detection in e-tail', *Decision Support Systems*, 95, pp. 91-101.
9. Chandola, V., Banerjee, A. and Kumar, V. (2009) 'Anomaly detection: A survey', *ACM Computing Surveys*, 41(3), pp. 1-58.
10. Chen, T. and Guestrin, C. (2016) 'XGBoost: A scalable tree boosting system', *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785-794.
11. Dal Pozzolo, A., Caelen, O., Johnson, R.A. and Bontempi, G. (2018) 'Calibrating probability with undersampling for unbalanced classification', *IEEE Symposium Series on Computational Intelligence*, pp. 159-166.
12. Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M. and Bouchachia, A. (2014) 'A survey on concept drift adaptation', *ACM Computing Surveys*, 46(4), pp. 1-37.
13. Goodman, B. and Flaxman, S. (2017) 'European Union regulations on algorithmic decision-making and a "right to explanation"', *AI Magazine*, 38(3), pp. 50-57.
14. Hevner, A.R., March, S.T., Park, J. and Ram, S. (2004) 'Design science in information systems research', *MIS Quarterly*, 28(1), pp. 75-105.
15. Hilal, W., Gadsden, S.A. and Yawney, J. (2022) 'Financial fraud: A review of anomaly detection techniques and recent advances', *Expert Systems with Applications*, 193, 116429.

16. Hoi, S.C., Sahoo, D., Lu, J. and Zhao, P. (2018) 'Online learning: A comprehensive survey', arXiv preprint arXiv:1802.02871.
17. Kou, Y., Lu, C.T., Sirwongwattana, S. and Huang, Y.P. (2004) 'Survey of fraud detection techniques', IEEE International Conference on Networking, Sensing and Control, 2, pp. 749-754.
18. Losing, V., Hammer, B. and Wersing, H. (2018) 'Incremental on-line learning: A review and comparison of state of the art algorithms', Neurocomputing, 275, pp. 1261-1274.
19. Lundberg, S.M. and Lee, S.I. (2017) 'A unified approach to interpreting model predictions', Advances in Neural Information Processing Systems, 30, pp. 4765-4774.
20. Maes, S., Tuyls, K., Vanschoenwinkel, B. and Manderick, B. (2002) 'Credit card fraud detection using Bayesian and neural networks', Proceedings of the 1st International NAISO Congress on Neuro Fuzzy Technologies, pp. 261-270.
21. McKinsey Global Payments Report (2024) Global Payments Report 2024. McKinsey & Company.
22. Minku, L.L. and Yao, X. (2012) 'DDD: A new ensemble approach for dealing with concept drift', IEEE Transactions on Knowledge and Data Engineering, 24(4), pp. 619-633.
23. Molnar, C. (2020) Interpretable Machine Learning: A Guide for Making Black Box Models Explainable. Available at: <https://christophm.github.io/interpretable-ml-book/>
24. Ngai, E.W., Hu, Y., Wong, Y.H., Chen, Y. and Sun, X. (2011) 'The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature', Decision Support Systems, 50(3), pp. 559-569.
25. Nilson Report (2024) Card Fraud Losses Reach \$32 Billion. Issue 1235, February 2024.
26. Parisi, G.I., Kemker, R., Part, J.L., Kanan, C. and Wermter, S. (2019) 'Continual lifelong learning with neural networks: A review', Neural Networks, 113, pp. 54-71.
27. Phua, C., Lee, V., Smith, K. and Gayler, R. (2010) 'A comprehensive survey of data mining-based fraud detection research', arXiv preprint arXiv:1009.6119.
28. Sánchez, D., Vila, M.A., Cerda, L. and Serrano, J.M. (2009) 'Association rules applied to credit card fraud detection', Expert Systems with Applications, 36(2), pp. 3630-3640.
29. Settles, B. (2009) 'Active learning literature survey', Computer Sciences Technical Report 1648, University of Wisconsin-Madison.
30. West, J. and Bhattacharya, M. (2016) 'Intelligent financial fraud detection: A comprehensive review', Computers & Security, 57, pp. 47-66.
31. Whitrow, C., Hand, D.J., Juszczak, P., Weston, D. and Adams, N.M. (2009) 'Transaction aggregation as a strategy for credit card fraud detection', Data Mining and Knowledge Discovery, 18(1), pp. 30-55.
32. Yang, Q., Liu, Y., Chen, T. and Tong, Y. (2019) 'Federated machine learning: Concept and applications', ACM Transactions on Intelligent Systems and Technology, 10(2), pp. 1-19.