Investigating the Impact of Transaction Temporal Patterns on Fraud Detection Models

Yiran Wang *

Department of Art & Science, New York University Shanghai, Shanghai, China

* Corresponding Author Email: yw6901@nyu.edu

Abstract. In the era of digital finance, the growth of transactions has increased the challenge of detecting fraudulent activities. Traditional systems are not enough and need to enhance the ability of capturing the patterns of financial fraud. This study investigates how important the temporal features are in fraud detection. Three algorithms, including Logistic Regression, Random Forest, and XGBoost, are used in this research. They are evaluated both with and without temporal features such as transaction hour and recent transaction counts. Experimental results show that integrating temporal features slightly improves overall accuracy of all used models. But they contribute little to the detection of fraudulent transactions. Static features like account and customer attributes, keep acting as the main predictors of fraud. Among all tested models, XGBoost has the highest and most stable performance. At the same time, Random Forest and Logistic Regression show similar patterns. Feature importance analysis further confirms that temporal features provide moderate but limited benefits. The statement that working out class imbalance problem is more critical when dealing with improving fraud detection could be drawn.

Keywords: Financial fraud detection, Machine learning, Temporal features.

1. Introduction

In the modern digital economic society, financial fraud is seen as an increasingly essential problem. As online and mobile banking services continue expanding, the volume and complexity of financial transactions have grown exponentially. The traditional systems are relatively weak for detecting fraudulent behavior [1]. To work out these challenges, machine learning methods are viewed as useful tools. They could automatically identify the patterns in transaction data [2].

Among all various factors related to fraud detection, temporal information often plays a crucial but often underexplored role. Temporal patterns could reveal dynamic behaviors that static features alone may not able to capture. Incorporating such time-based features into machine learning models can probably enhance their sensitivity to the fraud detection model to some extent [3].

This study tries to develop three machine learning models: Logistic Regression, Random Forest, and XGBoost and then compare them to detect fraudulent transactions, both with and without temporal features. Logistic Regression provides a basic, interpretable baseline. Random Forest improves robustness through ensemble averaging. And XGBoost uses sequential boosting and regularization to improve accuracy. Generally, XGBoost would better compare with other methods [4].

The experimental results show two major findings. Firstly, integrating temporal features significantly improves overall model performance across all algorithms. This shows that time-dependent transaction behavior includes signals in prediction. Secondly, XGBoost consistently outperforms both Random Forest and Logistic Regression.

This research contributes to the literature by stating the importance of temporal dynamics in fraud detection. It also provides empirical evidence for the effectiveness of gradient boosting methods in financial applications.

2. Related Works

In recent financial fraud detection, machine learning techniques are seen as an essential part. They have the ability to capture complex patterns in transaction data [5]. Among these methods, Logistic Regression, Random Forest, and XGBoost are widely applied. Logistic Regression is valued for its simplicity and interpretability. It often serves as a baseline in fraud detection. It could provide insights into the relationship between different features and the fraud [6].

Random Forest constructs multiple decision trees and aggregates their predictions to achieve more stable results as an ensemble learning method. It is very effective in working with large datasets and overfitting [7]. XGBoost represents a more sophisticated ensemble approach. By sequentially building trees that correct the errors of previous ones, XGBoost is more likely to have a high performance. This is one of the reasons that it is popular [8].

Recently, the inclusion of temporal features in fraud detection models has attracted more attention. Temporal patterns such as transaction timing, frequency, and sequences could reveal dynamic behaviors of fraudulent activities. For example, Saldaña-Ulloa et al. proposed a temporal graph network framework modeling event based interactions and time stamped transaction graphs, which leads to the conclusion that integrating temporal relations enhances detection of fraudulent transactions on online payment platforms [9]. Similarly, Aghware et al. emphasized the significance of temporal features, stating that incorporating transaction timestamps and recurrence information can largely enhance the model's detection abilities [10].

Although previous studies have successfully applied traditional machine learning models to fraud detection, few have systematically compared their performance with and without temporal features. The specific impact of temporal information in fraud detection remains underexplored. This research addresses this gap by conducting a comprehensive comparison of Logistic Regression, Random Forest, and XGBoost in detecting fraudulent transactions, both with and without temporal features. Doing these could provide empirical evidence of the importance of temporal information in optimizing fraud detection.

3. Research Design and Methodological Framework

The methodological framework for this research encompasses the data preprocessing pipeline, temporal feature engineering, model training, and evaluation strategies. The overall goal is to examine how temporal features influence the performance of machine learning models in financial fraud detection.

The methodology is divided into two main parts. Firstly, this research focuses on transforming transactional data into a structured format, making it suitable for analysis, as well as the construction of time-aware behavioral features.

Secondly, this research outlines the implementation of three supervised learning algorithms—Logistic Regression, Random Forest, and XGBoost along with detailed descriptions of the training setup and feature importance analysis.

This dual-stage design allows for a direct comparison between models trained on baseline features and those enhanced with temporal dynamics, providing empirical evidence on the contribution of temporal information in aspect of improved fraud detection.

3.1. Data Preprocessing and Temporal Feature Engineering

The dataset used in this study comprises transaction-level records, including features such as transaction amount, date, time, customer ID, and account balance. To maintain temporal consistency, the Transaction_Date and Transaction_Time fields were merged into a single timestamp variable, which is Transaction_DateTime, and transactions were sorted chronologically within each customer. Categorical variables were label encoded, and missing values were imputed either with zeros or with rolling averages. Feature standardization was applied by using z-score normalization.

A series of temporal and behavioral features were then engineered to represent both short-term and long-term user activity.

The number of transactions made in the past 24 hours for each transaction t_i was computed as:

$$\text{txns_last_1d}(t_i) = \sum_{j:t_i - t_j \le 1\ \text{day}} 1 \tag{1}$$

The rolling average of transaction amounts was computed for one-day and seven-day windows, respectively:

$$\begin{aligned} \text{mean_amt_1d}(t_i) &= \frac{1}{N_{1d}} \sum_{j:t_i - t_j \le 1\ \text{day}} A_j, \\ \text{mean_amt_7d}(t_i) &= \frac{1}{N_{7d}} \sum_{j:t_i - t_j \le 7\ \text{days}} A_j \end{aligned} \tag{2}$$

where A_j denotes the transaction amount, and N_{1d}, N_{7d} are the respective counts of transactions in each window.

The rate of balance change between consecutive transactions was defined as:

$$\text{balance_change_rate}(t_i) = \frac{B_i - B_{i-1}}{B_{i-1}} \tag{3}$$

where B_i and B_{i-1} represent the account balances after and before transaction respectively.

The time elapsed since the previous transaction was calculated in seconds as:

$$\text{time_since_last}(t_i) = (t_i - t_{i-1})_{\text{seconds}} \tag{4}$$

Deviations in transaction amounts relative to user-specific averages were quantified by a z-score:

$$\text{amt_zscore}(t_i) = \frac{A_i - \bar{A}}{\sigma_A} \tag{5}$$

where $\bar{A}$ and σ_A denote the mean and standard deviation of all transactions for the same customer.

In addition, categorical temporal indicators were introduced, including a binary flag identifying nighttime transactions:

$$\text{is_night} = \begin{cases} 1, & 0 \le hour(t_i) < 6, \\ 0, & otherwise, \end{cases} \tag{6}$$

and a weekday variable ranging from 0 to 6. Finally, the total number of transactions occurring on each calendar date across all users was added as daily_txn_count.

What's more, class imbalance was addressed through the Synthetic Minority Oversampling Technique (SMOTE) to generate relatively synthetic fraudulent samples since there are a small part of statistics that are fraud, which could be seen from Figure 1.

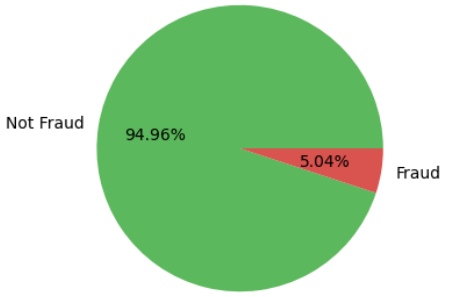

Fig. 1 Fraud vs Non-Fraud Distribution

(Picture credit : Original)

To visually comprehend the dataset, different graphs are plotted.

Fig. 2 Correlation Heatmap

(Picture credit : Original)

Figure 2 reveals that the numeric features share virtually no linear correlation with each other or with the fraud label. This leads to the conclusion that fraud detection may require non-linear modeling approaches to capture the subtle patterns that distinguish fraudulent from legitimate transactions.

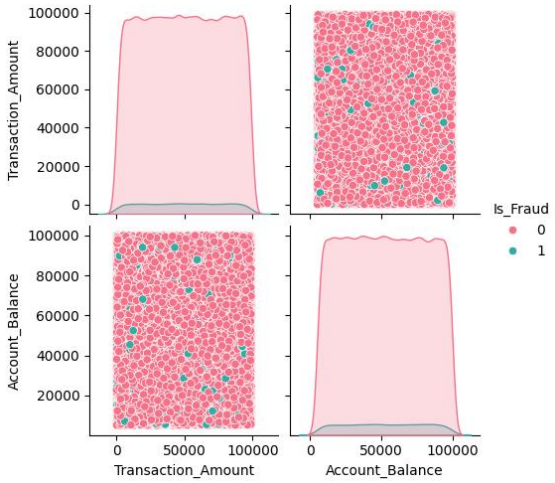

Fig. 3 Feature Distributions: Fraud vs Non-Fraud

(Picture credit : Original)

Figure 3 shows that fraudulent transactions are extremely rare compared to non-fraudulent ones, and a severe class imbalance problem exists.

3.2. Data Preprocessing and Temporal Feature Engineering

In this research, three supervised machine learning algorithms: Logistic Regression, Random Forest, and XGBoost, were trained and compared to get their effectiveness in detecting fraud. Both temporal and non-temporal versions of the dataset were used to evaluate the impact of using time-related features on model performance. And model hyperparameters were tuned to ensure stable training and fair comparison across these algorithms.

Model performance was evaluated using multiple metrics to ensure a balanced assessment, including Accuracy, Precision, Recall, F1-score, and Roc -Auc. All models were trained using identical random seeds and same data splits to ensure experimental reproducibility.

What's more, the evaluation results were systematically compared to identify which classifier achieved the optimal balance between recall and precision.

To interpret model behavior, particularly for the XGBoost classifier, feature importance analysis was conducted to identify the most influential predictors contributing to fraud detection. The feature importance was derived from the model's internal gain-based metric.

This step validates the relevance of temporal features as well as provides practical insights into which behavioral indicators most probably signal potential fraudulent activity.

4. Results

From a holistic perspective in Table 1, after combining temporal features, all models exhibited an increase in overall accuracy, but the recall and F1 scores for the fraudulent transactions actually decreased. This result shows that though temporal features help to improve the correct classification of the majority class (non-fraudulent transactions) and therefore enhance overall predictive accuracy, they do not significantly improve the detection of fraudulent transactions. Generally, XGBoost demonstrates the best overall balance in performance, maintaining stable results among most metrics. However, the use of temporal features does not substantially enhance its ability to identify minority-class instances. Random Forest and Logistic Regression share similar trends: overall accuracy slightly increases, but recall and F1 scores for the fraud decrease, indicating that static features alone are already capturing most of the fraud-predictive signals, and the incremental benefit from temporal features is relatively limited.

Table 1. Model Performance Comparison With and Without Temporal Features

Model	Temporal Features	Accuracy	Precision	Recall	F1-score	ROC-AUC
XGBoost	Yes	0.95	0.04	0.00	0.00	0.4936
Random Forest	Yes	0.79	0.05	0.17	0.08	0.4956
Logistic Regression	Yes	0.76	0.05	0.21	0.08	0.5005
XGBoost	No	0.7896	0.0508	0.1794	0.0792	0.5010
Random Forest	No	0.7404	0.0507	0.2340	0.0833	0.5015
Logistic Regression	No	0.6371	0.0521	0.3605	0.0911	0.5022

Feature importance analysis in Figure 4 could offer further interpretation into this phenomenon. Static features such as Account_Type, Gender, and Device_Type hold the highest importance consistently, showing the result that individual and account-level attributes play an essential role in detecting fraud risk. These features may reflect customer stratification and underlying behavioral differences, then contribute to strong predictive performance even without temporal information. In comparison, temporal and aggregated transaction features, such as daily_txn_count, weekday, and hour, show moderate importance in fraud detection, which is a signal that the model does capture some temporal patterns and those temporal pattern are useful in detection. For instance, accounts with higher daily transaction volumes are more likely to be flagged, and fraudulent activities may be concentrated on certain times. However, the contribution of these features is not that evident and does not substantially improve the identification of fraud instances. Additionally, transaction-related attributes such as Transaction_Type, Transaction_Device, and Merchant_Category also shows moderate importance. The model learns to differentiate fraud likelihood according to transaction type, device used, and merchant category.

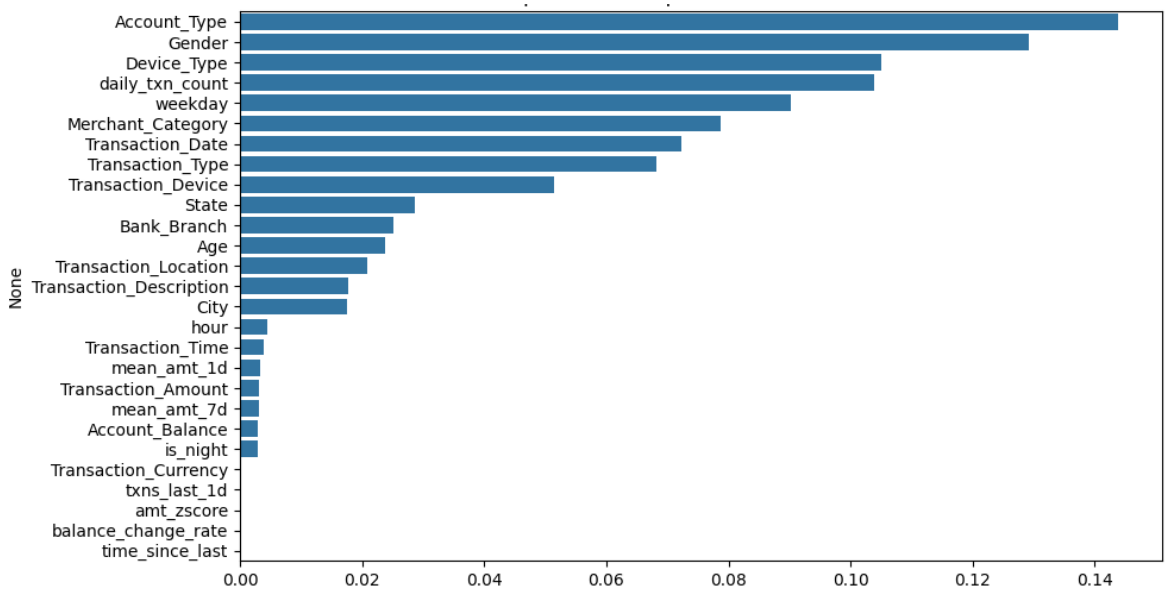

Fig. 4 Feature Importance (Temporal-Enhanced XGBoost Model)

(Picture credit : Original)

To summarize, the introduction of temporal features offers modest improvements in overall predictive performance, primarily benefiting the classification of the majority class. But their impact on detecting rare fraudulent events remains limited. It can be inferred that in highly imbalanced datasets, adding temporal features and basically handling class imbalance may not significantly enhance minor class detection, and further strategies may be necessary to increase the model's sensitivity to fraudulent transactions.

5. Conclusion

This research performed a detailed analysis of machine learning methods for fraud detection, as well as a strict comparison of Logistic Regression, Random Forest, and XGBoost with and without temporal features. Our findings give evidents to conclusion that temporal features hold importance in fraud detection while static features dominate the prediction of fraudulent transactions. In other words, inherent customer and account characteristics already capture most of the predictive signals. At the same time, temporal features, including daily transaction count, transaction hour, and weekday, provide moderate additional information, successfully captured underlying time-dependent patterns in fraudulent behavior. However, their contribution to improving the detection of fraud case is scarce.

XGBoost consistently showed the best overall performance among the evaluated models. While the inclusion of temporal features increased the accuracy, it is also evident that recall and F1 scores for fraudulent transactions decreased for several models. This supports the statement that the challenges of improving minority-class detection is essential in highly imbalanced datasets. Random Forest and Logistic Regression shared similar trends, with patterns that static features alone offer enough predictive ability and temporal features are auxiliary.

Feature importance analysis further confirmed our previous conclusions that both account level and transaction level attributes play critical roles in fraud detection, while temporal features provide complementary interpretations. These results are sufficient to show that temporal information can improve the understanding of transactional behavior. Also, address class imbalance and identify rare fraud events may require more sophisticated approaches, such as anomaly detection techniques, and some hybrid frameworks.

In conclusion, this research indicates the modest benefits of temporal features in fraud detection, and gives us practical insights for developing robust fraud detection systems in imbalanced transaction datasets. In the future, more innovative integrating modeling would be generated to further increase the efficiency of fraud detection.

References

[1] Compagnino A A, Maruccia Y, Cavuoti S, et al. An Introduction to Machine Learning Methods for Fraud Detection. Applied Sciences, 2025, 15(21): 11787.

[2] Chen Y, Zhao C, Xu Y, et al. Deep learning in financial fraud detection: Innovations, challenges, and applications. Data Science and Management, 2025.

[3] Duan Y, Zhang G, Wang S, et al. Cat-gnn: Enhancing credit card fraud detection via causal temporal graph neural networks, arXiv preprint arXiv: 2402.14708. 2024.

[4] Das S R, Sulaiman R B, Butt U. Comparative analysis of machine learning algorithms for credit card fraud detection. FMDB Transactions on Sustainable Computing Systems, 2023, 1(4): 225-244.

[5] Akre Z R. Financial Fraud Detection Based on Machine and Deep Learning: A Review. The Indonesian Journal of Computer Science, 2024, 13(3). https://doi.org/10.33022/ijcs.v13i3.4059

[6] Afriyie J K, Tawiah K, Pels W A, et al. A supervised machine learning algorithm for detecting and predicting fraud in credit card transactions. Decision Analytics Journal, 2023, 6, 100163. Elsevier. https://doi.org/10.1016/j.dajour.2023.100163

[7] Halabaku E, Bytyçi E. Overfitting in machine learning: A comparative analysis of decision trees and random forests. International Journal of Applied Science and Computing, 2024, 39(6): 1–15. https://doi.org/10.32604/iasc.2024.059429

[8] Vaishnavi T, Krishnaveni S, Aravindprakash N, Akilesh S, Hari Prakash A C. Detection of Credit Card Fraud Using Machine Learning. In Proceedings of the 3rd International Conference on Optimization Techniques in the Field of Engineering (ICOFE-2024). Kongu Engineering College, Perundurai, 2025 https://ssrn.com/abstract/4973112

[9] Saldaña-Ulloa D, De Ita Luna G, Marcial-Romero J R. A Temporal Graph Network Algorithm for Detecting Fraudulent Transactions on Online Payment Platforms. Algorithms, 2024, 17(12): 552. https://doi.org/10.3390/a17120552

[10] Aghware F O, Ojugo A A, Adigwe W, et al. Enhancing the random forest model via synthetic minority oversampling technique for credit-card fraud detection. Journal of Computing Theories and Applications, 2024, 1(4): 407–420.