The Eurasia Proceedings of Science, Technology, Engineering & Mathematics (EPSTEM)
ISSN: 2602-3199

The Eurasia Proceedings of Science, Technology, Engineering & Mathematics (EPSTEM), 2024

Volume 31, Pages 1-10

ICATI 2024: International Conference on Advances in Technology and Innovation

Forecasting Fraud Detection Using Data Science Methods

Baris Kavus
Aktif Yatırım Bankası A.S
Istanbul Nisantası University

Negar Sadat Soleimani-Zakeri
Istanbul Nisantası University

Abstract: Fraud detection is critical in various domains, including finance, healthcare, and e-commerce, where fraudulent activities pose significant threats to organizational integrity and financial stability. Traditional fraud detection methods often fail to address the dynamic nature of fraudulent behavior. In response, data science methods have emerged as promising tools for forecasting fraudulent activities by leveraging advanced analytics techniques on large-scale datasets. This research will make significant contributions by focusing on predicting fraud detection through data science methods. The findings will guide on preventing customers from committing fraud. The research questions aimed to be answered in this study are as follows: What are the key factors affecting fraud detection? Which customer behaviors are the strongest predictors of fraud detection? This study will provide a valuable model to the industry, enabling financial institutions to strengthen their risk management strategies and translate innovations in AI into applications.

Keywords: Fraud detection, Logistic regression, XGBoost classifier, CatBoost classifier, Random forest

Introduction

The rapid digitalization of transactions has brought about a parallel increase in credit card fraud, a pressing issue that threatens financial institutions and consumers. The need to detect and prevent fraudulent activities in credit card transactions has thus become a critical area of research, urgently requiring innovative solutions to reduce financial losses and safeguard sensitive information. A key challenge in this field is the significant class imbalance in fraud detection datasets, where legitimate transactions far outnumber fraudulent instances. This imbalance poses a hurdle for traditional machine learning models, which often struggle to accurately identify the minority class (fraud transactions), leading to a high rate of false negatives and overall suboptimal model performance (Gupta et al., 2023; Tran & Dang, 2021).

Researchers have developed various strategies to address the limitations associated with imbalanced datasets, categorized broadly into data-level, algorithm-level, and ensemble-based approaches. Data-level techniques focus on modifying the dataset to achieve a more balanced class distribution. Commonly employed methods include oversampling the minority class (e.g., using the Synthetic Minority Over-sampling Technique (SMOTE) or Adaptive Synthetic Sampling (ADASYN)) and under-sampling the majority class to enhance the performance of the classification algorithms on skewed datasets (Makki et al., 2019; Tran & Dang, 2021). On the other hand, algorithm-level strategies involve adapting the learning process by incorporating cost-sensitive learning mechanisms, which assign a higher penalty to misclassifications of minority class instances. This helps to reduce the bias toward the majority class during training. Additionally, ensemble-based methods such as Random Forest, Boost, and AdaBoost have shown considerable promise by aggregating the results of multiple base classifiers, thus improving the overall accuracy and robustness of fraud detection systems (Awoyemi et al., 2017; Singh et al., 2022).

- Selection and peer-review under responsibility of the Organizing Committee of the Conference

Despite these advancements, significant challenges remain. Even after employing data-level balancing techniques, achieving high precision and recall in fraud detection remains difficult due to the dynamic nature of fraud patterns. Fraud behaviors evolve rapidly, necessitating adaptive models that accommodate new patterns. Consequently, recent research has increasingly focused on integrating multiple techniques to develop comprehensive and robust frameworks for fraud detection. For example, combining data-level balancing with advanced machine learning algorithms, such as deep learning networks or ensemble approaches, has significantly improved classification performance on highly imbalanced datasets (Isangediok & Gajamannage, 2022; Makki et al., 2019).

Moreover, the advent of hybrid approaches, which simultaneously apply multiple strategies (e.g., combining oversampling, cost-sensitive learning, and ensemble techniques), has shown potential to address the limitations of existing methods, such hybrid models improve detection rates and reduce the computational burden of processing large datasets, making them suitable for real-time fraud detection applications. This paper aims to delve into these issues by systematically evaluating the effectiveness of various data science methodologies in forecasting credit card fraud. It will propose an integrated approach that leverages the strengths of different techniques to develop an adaptive, high-performance fraud detection framework capable of operating in a dynamic and data-intensive environment.

This study focuses on fraud detection using machine learning techniques, specifically CatBoost, LightGBM, XGBoost, Logistic Regression, AdaBoost, and Random Forest. First, the performance differences among these methods and their impact on the dataset were analyzed. Then, various hyperparameter tuning and model enhancement strategies were applied to improve the techniques' effectiveness. Finally, the results were compared to determine the most suitable approach for fraud detection.

Literature Review

Recent advancements in credit card fraud detection have explored various strategies to overcome challenges associated with data imbalance, dynamic fraud patterns, and the need for computational efficiency. One promising area is hybrid techniques for balancing datasets and feature selection. Researchers have increasingly employed a combination of under-sampling and oversampling methods to achieve better class distribution. A notable approach is SMOTE-ENN (SMOTE combined with Edited Nearest Neighbors), where synthetic samples are first generated to oversample the minority class, followed by an under-sampling process to remove noisy data points, thus enhancing model robustness (Makki et al., 2019; Warghade et al., 2020). Additionally, feature selection methods such as recursive feature elimination, principal component analysis (PCA), and information gain have been effective in identifying the most relevant features, leading to reduced model complexity and improved generalization performance (Awoyemi et al., 2017; Tyagi & Mittal, 2020).

Deep learning methods, especially Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks, have emerged as powerful tools for detecting complex fraud patterns. LSTM networks, in particular, are well-suited for handling sequential data and capturing temporal dependencies in transaction records, making them effective for identifying subtle changes that could indicate fraudulent behavior. Moreover, autoencoders and Generative Adversarial Networks (GANs) have been utilized for anomaly detection, where the model learns normal behavior patterns and flags deviations as potential fraud instances (Singh et al., 2022; Tran & Dang, 2021).

Due to fraudulent activities' dynamic nature, adaptive and incremental learning models have gained attention. These models can update themselves continuously with new data, enabling them to detect evolving fraud patterns without requiring frequent retraining from scratch. Techniques such as AdaBoost and Streaming Random Forest have demonstrated effectiveness for real-time fraud detection, balancing accuracy and computational efficiency (Isangediok & Gajamannage, 2022; Tyagi & Mittal, 2020)

Traditional metrics like accuracy are often inadequate to evaluate models effectively in imbalanced data. Alternative metrics such as the Area Under the Precision-Recall Curve (PR-AUC), Matthews Correlation Coefficient (MCC), and F-measure have been recommended for better assessment (Makki et al., 2019; Warghade et al., 2020). Cost-sensitive learning approaches have also been adopted to address the imbalanced nature of the data by assigning different costs to misclassifications, which significantly enhances the model's ability to detect minority class instances (Singh et al., 2022).

Comparative studies on different machine learning algorithms have highlighted the superior performance of ensemble techniques, such as Random Forest and Gradient Boosting, compared to individual classifiers like Decision Trees, Support Vector Machines (SVM), and K-nearest Neighbors (KNN). These ensemble models reduce bias and variance, improving the robustness of fraud detection systems (Awoyemi et al., 2017; Gupta et al., 2023; Warghade et al., 2020).

In addition, transfer learning and domain adaptation approaches have been employed to improve fraud detection by leveraging knowledge from related tasks or domains. With labeled fraud data being scarce, transfer learning allows models trained on large datasets from one domain to be adapted for use in another with minimal labeled data, thus enhancing model generalization capabilities. Techniques such as fine-tuning pre-trained models and using domain-specific regularization have shown promise in this area (Singh et al., 2022; Varmedja et al., 2019).

Lastly, explainability and interpretability become crucial as fraud detection models become more complex. Techniques like Gradient Boosting Decision Trees (GBDT) and SHAP (SHapley Additive exPlanations) provide insights into the contribution of each feature to the final prediction, making the decision-making process more transparent. This transparency is essential for regulatory compliance and building trust in automated fraud detection systems (Isangediok & Gajamannage, 2022; Tran & Dang, 2021).

The existing literature indicates that integrating diverse data preprocessing techniques, sophisticated machine learning algorithms, and interpretability methods can markedly improve the efficacy and robustness of credit card fraud detection systems. This study seeks to consolidate these advancements by developing a comprehensive framework incorporating deep learning approaches, ensemble techniques, and adaptive strategies, thereby addressing the multifaceted challenges of forecasting fraud detection.

Method

Dataset and Preprocessing

The dataset used in this study comprises 284,807, with 31 and a highly imbalanced class distribution. Most data represent non-fraudulent transactions, while a small fraction constitutes fraudulent cases. Due to this imbalance, initial preprocessing steps were necessary to ensure a more balanced dataset and to mitigate the risk of model overfitting. These steps included checking missing values. No null values were detected. Robust scale algorithms were used for the two off-scale variables. The data was divided into training, validation, and test sets with a 20% ratio. FeatureWiz was used as the feature selection method. FeatureWiz is an automated feature engineering and selection tool designed to enhance model accuracy by identifying the most relevant features within a dataset. It employs sophisticated statistical techniques and machine learning algorithms to evaluate the contribution of each feature to the target variable. The process encompasses several key steps: data cleaning and preprocessing to address issues such as missing values and outliers; feature engineering, which involves generating new features through transformations, interactions, and polynomial functions; and feature selection, using methods, LightGBM's feature importance. Additionally, FeatureWiz addresses multicollinearity by detecting and removing highly correlated features, thereby improving the model's interpretability and reducing overfitting. It is particularly effective in handling large and imbalanced datasets, streamlining the feature selection process while maintaining or enhancing predictive performance.

A "Random Under-Sampling" technique was implemented to address the class imbalance, which involved reducing the number of majority class instances to match the minority class. Specifically, the non-fraudulent transactions were down-sampled to 492 cases, equaling the fraudulent cases to achieve a 50/50 ratio. This approach resulted in a balanced sub-sample of the original dataset, with an equal representation of fraud and non-fraud transactions. Following the under-sampling, the data was shuffled to eliminate any potential ordering biases and ensure that the models' performance remained consistent across multiple runs. While random Under-Sampling can effectively address class imbalance, it also carries the drawback of potential information loss.

Machine Learning Models

This study used six machine learning algorithms to detect fraud: CatBoost, LightGBM, XGBoost, Logistic Regression, AdaBoost, and Random Forest. These algorithms were chosen for their demonstrated effectiveness in handling classification tasks, particularly in imbalanced datasets. Each model brings unique strengths to the problem. CatBoost, a gradient-boosting algorithm based on decision trees, excels in handling categorical data through its native feature encoding and is known for its fast-training speed. LightGBM, another gradient-boosting framework, employs a leaf-wise tree growth strategy, which enhances speed and memory efficiency, particularly with large datasets. XGBoost, also a gradient boosting method, is distinguished by its regularization capabilities and level-wise tree growth, often yielding superior performance in predictive modeling competitions. Logistic Regression, a linear model, was included for its simplicity and utility as a baseline approach for binary classification tasks. AdaBoost, an ensemble learning technique, constructs a strong classifier by iteratively combining multiple weak classifiers, focusing on correcting misclassified samples. Lastly, Random Forest, an ensemble of decision trees, enhances classification performance by aggregating the predictions from numerous trees, thus reducing the likelihood of overfitting. To optimize the performance of each algorithm, hyperparameter tuning was conducted using techniques such as random search to determine the most suitable parameter configurations for the task.

Threshold Optimization, Model Training and Evaluation

In addition to training the models, threshold optimization was applied to enhance the classification performance, especially given the imbalanced nature of the data. The default decision threshold of 0.5 was adjusted to optimize precision, recall, and F1-score metrics. The optimal threshold was determined by evaluating the trade-offs using metrics like the Precision-Recall curve and Area Under the Receiver Operating Characteristic Curve (ROC-AUC), selecting the threshold that provided the best balance for fraud detection.

The models were trained using the training set, with hyperparameter tuning performed on the validation set to identify the optimal configurations for each algorithm. The final evaluation of the models was carried out on the test set, employing various performance metrics to assess their effectiveness. The metrics included accuracy, which measures the overall correctness of the model's predictions, and precision, representing the proportion of true positive predictions out of all predicted positive instances, indicating the accuracy of fraud detection among the predicted cases. Recall, or sensitivity, was used to evaluate the model's ability to identify fraudulent cases, calculated as the ratio of true positive predictions to the total number of fraud cases. The F1-Score, which is the harmonic mean of precision and recall, was considered to provide a balance between these two metrics. Additionally, the ROC-AUC metric was utilized to assess the model's capability to distinguish between fraudulent and non-fraudulent cases, with higher values indicating better discrimination.

Hyperparameter Optimization

To ensure optimal performance for all models, hyperparameter tuning was conducted using techniques such as random search. These methods allowed for an extensive search across various hyperparameter combinations to identify the configurations that yielded the best results for each algorithm. These optimized settings were intended to maximize the predictive capabilities of the models in detecting fraudulent transactions.

Following hyperparameter tuning, a comprehensive comparison and analysis of the models were performed to identify the most effective technique for fraud detection. The evaluation considered the impact of threshold optimization, alongside the influence of hyperparameter tuning, on the models’ overall performance. The models were assessed using various metrics, including accuracy, precision, recall, F1-score, and ROC-AUC, to understand their strengths and weaknesses in different scenarios thoroughly. The results were then analyzed to determine which model achieved the best performance, highlighting the conditions under which each algorithm excelled.

Results and Discussion

This study aimed to identify the best methods for achieving accurate results in detecting fraud within highly imbalanced datasets. Various techniques were applied to Kaggle-Credit Card Fraud Detection datasets (https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud), including under-sampling, which incorporates machine learning algorithms such as CatBoost, LightGBM, XGBoost, Logistic Regression, AdaBoost, and Random Forest. These datasets, characterized by a significant class imbalance between fraud and non-fraud instances, necessitate specialized methods to enhance model performance. Examining the distributions provides insight into the degree of skewness in the features, allowing us to assess the balance of the dataset. Additional distributions of other features offer further insights into potential patterns. In the future, techniques to reduce skewness in these distributions will be applied within this notebook.

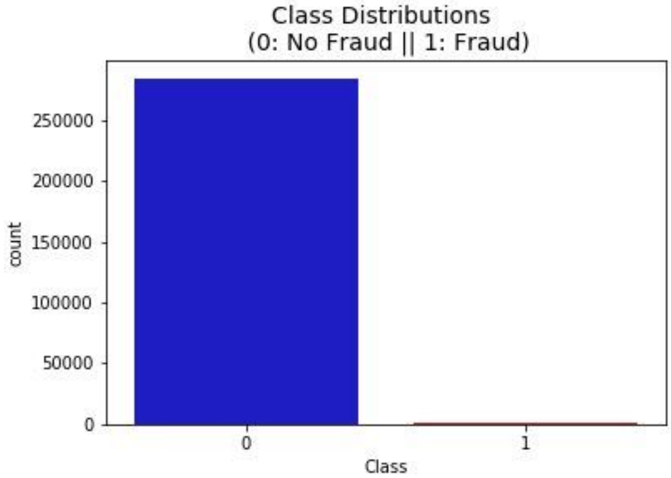

Figure 1. Distribution of Kaggle dataset

Figure 1 shows the distribution of classes in a fraud detection dataset, specifically indicating the imbalance between non-fraud (class 0) and fraud (class 1) transactions. The non-fraud class (0) has a significantly higher count, represented by the large blue bar on the left, while the fraud class (1) has a minimal count, shown as a nearly invisible red bar on the right. This extreme imbalance (where most cases are non-fraud) highlights the challenge in fraud detection, as standard models may struggle to identify the minority (fraud) cases accurately. To address this, methods like resampling or algorithmic adjustments will likely be needed to improve model performance on the minority class.

The variables Amount and Time were scaled using the robust scaling algorithm. The robust scaler algorithm scales data based on the median and interquartile range, specifically to reduce the impact of outliers. This adjustment ensures that outliers in the data distribution do not negatively impact model performance. Next, we balanced the imbalanced data using random under-sampling. This technique reduces the number of instances in the majority class to match the minority class, creating a more balanced dataset for improved model performance on both classes. Figure 2 shows the distribution of classes balanced dataset.

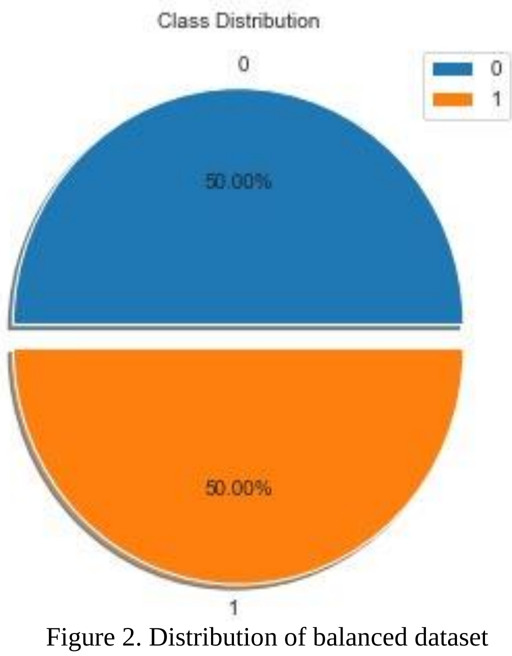

Figure 2. Distribution of balanced dataset

Figure 2 displays a balanced class distribution in the dataset, with both classes (0 for non-fraud and 1 for fraud) representing 50% each. This balanced distribution is likely the result of applying a technique like random under-sampling to address the initial imbalance. By equalizing the number of instances for both classes, the model is expected to have improved performance in identifying fraud cases without being biased toward the majority (non-fraud) class.

Table 1. Model performance metrics

	Accuracy	Precision	Recall	f1	Roc_Auc
Logistic Regression	0.959391	0.989247	0.929293	0.958333	0.977118
Random Forest	0.944162	0.988889	0.898990	0.941799	0.985003
ADA	0.939086	0.939394	0.939394	0.939394	0.983509
XGBoost	0.923858	0.946809	0.898990	0.922280	0.981653
LightGBM	0.934010	0.957447	0.909091	0.932642	0.987838
CatBoost	0.949239	0.978495	0.919192	0.947917	0.989487

Table 1 presents the performance metrics for six machine learning modelsLogistic Regression, Random Forest, AdaBoost, XGBoost, LightGBM, and CatBoost—evaluated on a fraud detection dataset. Each model's effectiveness is assessed using accuracy, precision, recall, F1-score, and ROC-AUC metrics. These metrics offer a comprehensive view of each model's ability to correctly identify instances of fraud while maintaining a low false positive rate. Using F1-score as the main metric, CatBoost and Logistic Regression stand out as the top-performing models for fraud detection, with F1-scores of 0.947 and 0.953, respectively. These scores indicate their strong balance between precision and recall, making them reliable for identifying fraud without excessive misclassification. Random Forest, AdaBoost, and LightGBM also perform well, though slightly lower, while XGBoost shows a lower F1-score, suggesting it may prioritize precision over recall. Overall, CatBoost and Logistic Regression are this dataset's most effective models for balanced fraud detection.

Since the F1-score was highest for Logistic Regression, threshold optimization was performed on this model. This approach allows for fine-tuning the decision threshold to achieve an optimal balance between precision and recall, further enhancing the model's performance in fraud detection by adjusting it to best capture fraud cases without excessively misclassifying non-fraud instances.

Figure 3. Threshold tunning curve

Figure 3 shows the F1-score for different threshold values, helping identify the optimal threshold for balancing precision and recall in fraud detection. In this case, a threshold of 0.484 achieves the highest F1-score of 0.9637, indicated by the red marker on the curve. This threshold represents the point at which the model best captures fraud cases without over-penalizing non-fraud instances. Lowering or raising the threshold from this optimal point would decrease the F1-score, highlighting the importance of fine-tuning to achieve balanced and effective performance.

Next, feature selection was performed using FeatureWiz, resulting in the selection of 10 variables: ['V14', 'V4', 'V10', 'scaled_amount', 'V19', 'V13', 'V20', 'V21', 'V25', 'V2']. The primary goal of feature selection is to improve the model's score by identifying the most relevant variables; however, in some cases, an increase in score may not be observed, which is considered normal. This outcome can occur when the removed variables have minimal impact on performance, indicating that the model is already optimized with a smaller, more relevant subset of features.

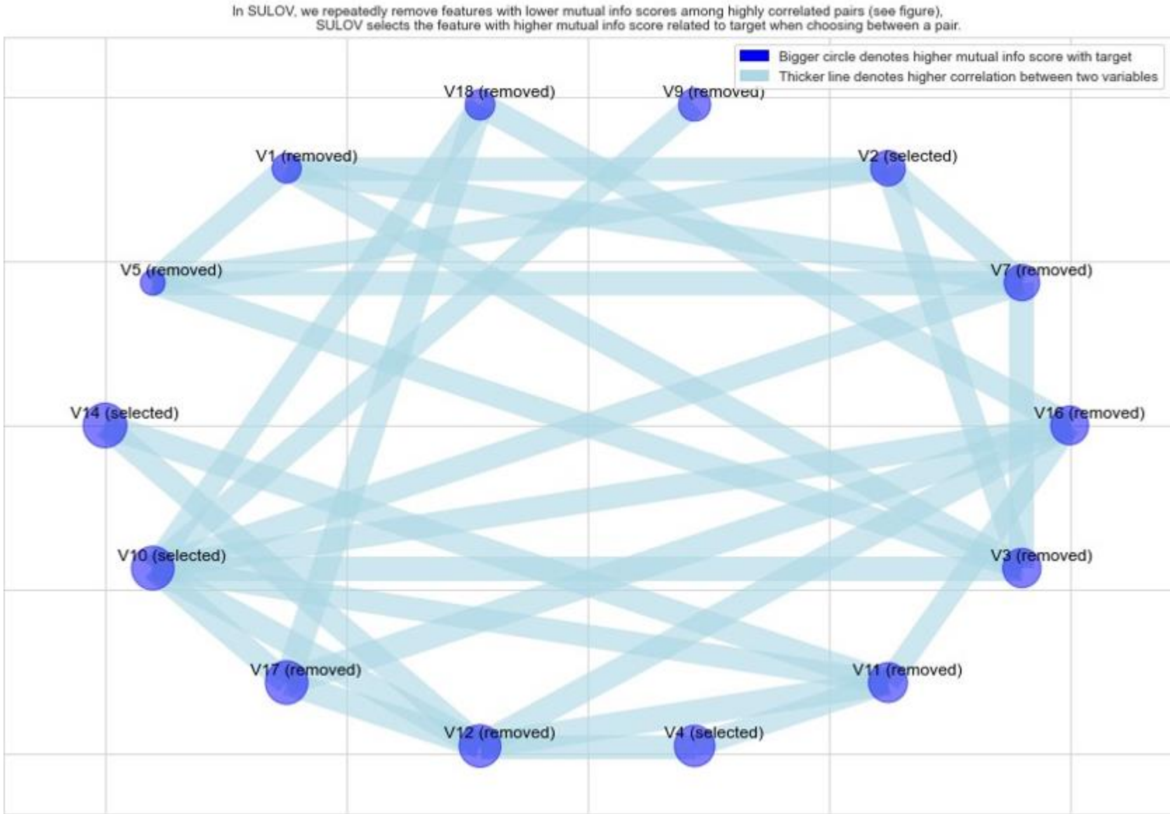

Figure 4. SULOV method for removing highly correlated features

Figure 4 illustrates how the Search for Uncorrelated List of Variables (SULOV) method works by removing highly correlated features. In SULOV, pairs of highly correlated features are compared, and the feature with the lower mutual information score related to the target variable is removed, while the one with the higher score is retained. The larger circles represent features with higher mutual information scores with the target variable, indicating higher relevance. Thicker lines denote stronger correlations between feature pairs, guiding the removal of redundant features. This method helps retain the most informative features and reduce multicollinearity, improving model performance. In this example, features like V4 and V10 are selected, while others like V1, V5, and V18 are removed due to high correlation with selected features.

Table 2. Model performance metrics after feature selection

	Accuracy	Precision	Recall	f1	Roc_Auc
Logistic Regression	0.939086	0.978022	0.898990	0.936842	0.985261
Random Forest	0.944162	0.958333	0.929293	0.943590	0.983766
ADA	0.908629	0.926316	0.888889	0.907216	0.974232
XGBoost	0.939086	0.957895	0.919192	0.938144	0.977840
LightGBM	0.923858	0.928571	0.919192	0.923858	0.981550
CatBoost	0.949239	0.958763	0.939394	0.948980	0.989384

After applying feature selection, Table 2 shows the performance metrics for six machine learning models: Logistic Regression, Random Forest, ADA, XGBoost, LightGBM, and CatBoost. The models are evaluated on accuracy, precision, recall, F1-score, and ROC-AUC. Feature selection led to varying impacts on the models' F1-scores. XGBoost and CatBoost showed improved F1 scores, indicating a better balance between precision and recall after reducing features. Random Forest also saw a slight improvement, suggesting it benefitted marginally from feature selection. However, Logistic Regression and ADA experienced a drop in F1-score, indicating that these models may rely on a broader set of features for optimal performance. LightGBM's F1-score slightly decreased, showing a minor impact. Overall, feature selection helped optimize some models, particularly tree-based methods like XGBoost and CatBoost, but was less beneficial for simpler models like Logistic Regression and ADA.

Threshold optimization was performed again. This step allows further fine-tuning of the decision threshold for each model to enhance the balance between precision and recall after feature selection. Re-optimizing the threshold can help adjust the models to the updated feature set, ensuring optimal performance in detecting fraud cases while minimizing misclassification.

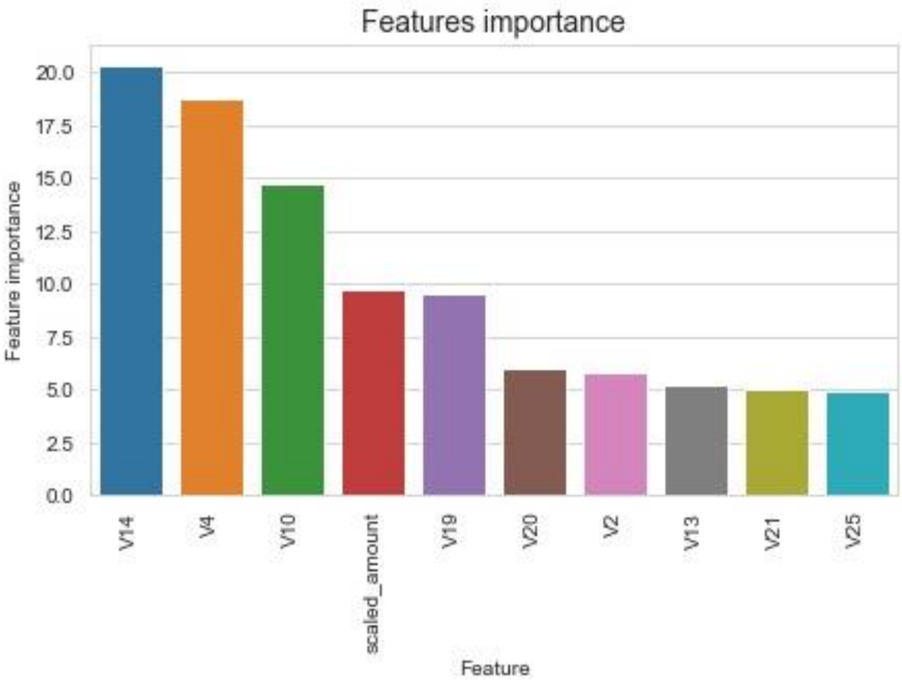

Figure 5. Feature importance

Figure 5 illustrates the relative contribution of each selected variable in predicting fraud cases. The analysis identifies V14, V4, and V10 as the most influential features, with V14 contributing the highest importance, followed closely by V4. These top features are critical in enhancing the model's predictive capacity and distinguishing between fraud and non-fraud cases. scaled_amount and V19 also demonstrate moderate importance, suggesting they provide substantial, though less critical, information for the model. In contrast, features such as V2, V13, V21, and V25 display relatively lower importance, indicating they contribute minimally to the model's performance. This distribution of feature importance suggests that a select group of variables, particularly V14, V4, and V10, serve as the primary predictors in the model, significantly influencing its ability to accurately detect fraud within the dataset.

Table 3. Model performance metrics after hyper parameter

	f1	Roc_Auc	CV Score Mean Before Optimization	CV Score Mean After Optimization
Logistic Regression	0.958333	0.979180	0.9397	0.9409
Random Forest	0.918033	0.983509	0.9309	0.9239
ADA	0.936170	0.988868	0.9207	0.9302
XGBoost	0.938144	0.984642	0.9349	0.9449
LightGBM	0.938776	0.986601	0.9306	0.9402
CatBoost	0.947917	0.987941	0.9374	0.9377

Further hyperparameter optimization was conducted; however, the results indicate no significant improvement compared to the previous table. The F1-score, ROC-AUC, and cross-validation (CV) score remains largely unchanged across models, suggesting that the initial settings were already close to optimal for this dataset. This outcome implies that additional tuning of hyperparameters did not yield substantial performance gains, and the models may have reached their maximum efficiency given the current feature set and data characteristics.

Conclusion

This study evaluated various machine learning techniques for fraud detection within a highly imbalanced dataset. Logistic Regression, Random Forest, AdaBoost, XGBoost, LightGBM, and CatBoost were analyzed through feature selection, hyperparameter tuning, and threshold optimization to enhance performance. CatBoost and Logistic Regression emerged as the top-performing models, exhibiting strong F1-scores, indicative of their effective balance between precision and recall. However, further hyperparameter tuning yielded minimal improvements, suggesting that the initial configurations were near-optimal. The study underscores the importance of tailored feature selection and threshold tuning in handling class imbalances, though certain models may reach performance limits even with optimization.

Recommendations

Future research should explore hybrid models combining multiple techniques, such as ensemble and deep learning, to adapt to evolving fraud patterns. Additionally, incorporating real-time incremental learning could enhance the models' adaptability to dynamic fraud behaviors. Investigating advanced data augmentation methods, such as synthetic sample generation, may improve the models' ability to handle imbalanced datasets. Finally, explainability tools like SHAP could be integrated to enhance model interpretability, allowing for better regulatory compliance and trust in automated fraud detection systems.

Scientific Ethics Declaration

The authors declare that the scientific ethical and legal responsibility of this article published in EPSTEM Journal belongs to the authors.

Acknowledgements or Notes

* This article was presented as an oral presentation at the International Conference on Advances in Technology and Innovation (www.icati.net) held in Antalya/Turkey on November 14-17, 2024.

References

Awoyemi, J. O., Adetunmbi, A. O., & Oluwadare, S. A. (2017). Credit card fraud detection using machine learning techniques: A comparative analysis. *2017 International Conference on Computing Networking and Informatics (ICCNI),* 1–9.

Gupta, P., Varshney, A., Khan, M. R., Ahmed, R., Shuaib, M., & Alam, S. (2023). Unbalanced credit card fraud detection data: A machine learning-oriented comparative study of balancing techniques. P*rocedia Computer Science, 218*, 2575–2584.

Isangediok, M., & Gajamannage, K. (2022). Fraud detection using optimized machine learning tools under imbalance classes. *2022 IEEE International Conference on Big Data (Big Data)*, 4275–4284.

Makki, S., Assaghir, Z., Taher, Y., Haque, R., Hacid, M.-S., & Zeineddine, H. (2019). An experimental study with imbalanced classification approaches for credit card fraud detection. *IEEE Access, 7,* 93010–93022.

Singh, A., Ranjan, R. K., & Tiwari, A. (2022). Credit card fraud detection under extreme imbalanced data: A comparative study of data-level algorithms. *Journal of Experimental & Theoretical Artificial Intelligence, 34*(4), 571–598.

Tran, T. C., & Dang, T. K. (2021). Machine learning for prediction of imbalanced data: Credit fraud detection. *15th International Conference on Ubiquitous Information Management and Communication (IMCOM),* 1–7.

Tyagi, S., & Mittal, S. (2020). *Sampling approaches for imbalanced data classification problem in machine learning.* In P. K. Singh, A. K. Kar, Y. Singh, M. H. Kolekar, & S. Tanwar (Eds.), *Proceedings of ICRIC 2019* (Vol. 597, pp. 209–221). Springer International Publishing.

Varmedja, D., Karanovic, M., Sladojevic, S., Arsenovic, M., & Anderla, A. (2019). Credit card fraud detection—Machine learning methods. *18th International Symposium (INFOTEH),* 1–5.

Warghade, S., Desai, S., & Patil, V. (2020). Credit card fraud detection from imbalanced dataset using machine learning algorithm. *International Journal of Computer Trends and Technology, 68*(3), 22–28.

Author Information

Baris Kavus
Aktif Yatirim Bankasi A.S
Esentepe Mahallesi Kore Sehitleri Caddesi No: 8/1 Sisli / Istanbul/Türkiye
Istanbul Nişantaşı University
Maslak Mahalesi, Taşyoncası Sokak, No: 1V ve No:1Y
Sarıyer-Istanbul/Türkiye
Contact e-mail: baris.kavus@aktifbank.com.tr

Negar Sadat Soleimani-Zakeri
Istanbul Nişantaşı University
Maslak Mahalesi, Taşyoncası Sokak, No: 1V ve No:1Y
Sarıyer-Istanbul/Türkiye

To cite this article:

Kavus, B., & Soleimani-Zakeri, N.S. (2024). Forecasting fraud detection using data science methods. *The Eurasia Proceedings of Science, Technology, Engineering & Mathematics (EPSTEM), 31,* 1-10.