

ALLEGHENY COLLEGE
DEPARTMENT OF COMPUTER SCIENCE

Senior Thesis

Python-Powered Machine Learning for Credit Card Fraud Detection: An Analytical Tool Approach

by

Bill Ochieng

ALLEGHENY COLLEGE

DEPARTMENT OF COMPUTER AND
INFORMATION SCIENCE

Project Supervisor: **Dr. Janyl Jumadinova**
Co-Supervisor: **Dr. Francisco Navarro-Sanchez**

03 May 2024

Abstract

The prevalence of digital payments has escalated the complexity and frequency of credit card fraud, presenting a significant challenge in financial security. This research focuses on the application of Logistic Regression to predict fraudulent transactions, utilizing a dataset of one million transactions with an 8.74% fraud incidence rate. Key features such as transaction distance from the cardholder's home, the transaction type (online or chip-based), and the use of PIN numbers are considered to determine the likelihood of fraud.

This study initially explores the data's highly unbalanced nature, where fraudulent cases are significantly underrepresented. To address this, Logistic Regression is applied, and model performance is evaluated, revealing limitations in handling the imbalanced data. Subsequently, an oversampling technique is employed to balance the dataset, significantly improving model precision and recall metrics for fraudulent transactions from 0.61 and 0.72 to 0.93 and 0.95, respectively.

The effectiveness of the model after oversampling shows a robust capability to identify fraudulent transactions, which is crucial for real-time fraud detection systems used by financial institutions. The results indicate that while the accuracy marginally decreased from 95.9% to 94.1%, the adjusted model offers a more reliable solution for predicting fraud, effectively reducing the risk and potential financial losses due to fraud.

Keywords: Credit Card Fraud, Logistic Regression, Machine Learning, Fraud Detection, Oversampling, Data Imbalance, Transaction Analysis, Financial Security, Real-time Fraud Detection

Table of Contents

Introduction	4
Related Works	9
Method of Approach	12
Experimental Evaluation	23
Model Evaluation Metrics Summary	29
Conclusion	31
Future Work	33
References	35

List of Figures

1	background	5
2	Fraud statistics	6
3	Architecture Of The Proposed System	12
4	correlation values	14
5	correlation heatmap	15
6	correlation bargraph	15
7	Logistic Regression	16
8	KNN	17
9	balanced dataset bar graph	18
10	ROC unbalanced dataset	18
11	dashboard	22
12	ROC/AUC unbalanced dataset	26
13	ROC/AUC balanced dataset	26

List of Tables

1	The unbalanced nature of the dataset	13
2	Scaled Features for Fraud Prediction	24
3	Initial classification_report	25
4	Final Classification_report	25
5	Initial confusion metrics	27
6	Final Confusion metrics	28

Introduction

Fraudulent activities have been a significant concern across various sectors of the economy, manifesting a substantial financial burden on global markets. The Association of Certified Fraud Examiners (ACFE) in 2022 estimated that a typical organization loses 5% of its annual revenue to fraud [17]. In terms of credit card fraud, which remains a prominent threat within the financial sector, losses amounted to \$22.8 billion in 2016 alone, marking a 4.4 percent increase from the previous year according to Nilson Report[6]. These numbers reflect the critical need for effective fraud detection and prevention mechanisms, which have evolved from traditional expert-based systems to sophisticated analytical techniques enabled by advances in technology.




Figure 1: background




Figure 2: Fraud statistics

This introduction delineates the domain of fraud analytics, initially focusing on the significant impact of fraud as evidenced by global statistics. The term “fraud” can be comprehensively defined as wrongful or criminal deception intended to result in financial or personal gain. Throughout history and even in nature, deception practices akin to fraud—such as camouflage in chameleons—demonstrate the intrinsic behavior to mislead for gain. Analogously, fraudsters employ similar tactics in the financial world, meticulously planning their strategies to remain undetected within transaction data, akin to a chameleon blending into its environment.

The multifaceted nature of fraud necessitates a diversified approach to its detection and prevention[5]. Traditional methods that use rule-based systems are increasingly complemented by sophisticated machine learning techniques due to the dynamic and evolving nature of fraud. Predictive analytics (supervised learning), descriptive analytics (unsupervised learning), and social network analytics are pivotal in detecting complex fraud patterns that are imperceptibly concealed within large datasets.

A fundamental component in developing effective fraud analytics is fea-

ture engineering, which enhances the dataset’s utility for detecting fraudulent patterns. This process involves creating, selecting, and transforming variables to better capture the nuances of fraud. However, the rarity of fraud incidents within datasets poses a significant challenge—fraud is an uncommon event, with fraudulent transactions typically constituting less than 1% of all transactions. This imbalance makes it difficult for machine learning models to learn and identify fraudulent patterns effectively due to their scarcity, necessitating specialized techniques to manage this issue.

In addressing the complexities of fraud detection in the digital era, this paper explores the use of Logistic Regression and K-Nearest Neighbor (KNN) algorithms to analyze and predict fraudulent activities in credit card transactions. The dataset employed in this study, sourced from Kaggle, comprises over one million transactions, providing a robust foundation for applying and evaluating these machine learning techniques. The data includes features such as the distance from home, the transaction type, and whether the transaction was authenticated using a PIN, among others, all of which are critical in profiling typical and atypical transaction behaviors.

The evolution of digital payments and the corresponding rise in cyber criminalities make the task of fraud detection more challenging yet essential for maintaining the integrity of financial transactions[9]. According to the Data Breach Index, over 5 million records are stolen daily, emphasizing the ongoing battle against fraud in both card-present and card-not-present payment types. Despite existing fraud prevention methods, the persistence and evolution of fraud tactics necessitate advanced analytical strategies to effectively combat these illicit activities.

By integrating economic frameworks such as Cost-Benefit Analysis and the Risk Management Model, this study not only assesses the technical effectiveness of the applied machine learning models but also their economic impacts. This dual approach provides a holistic view of the potential and limitations of current fraud detection methodologies and highlights the necessity for continuous innovation in this field.

This research aims to provide a comprehensive understanding of fraud analytics, from its definition and fundamental characteristics to the advanced analytical methods used to detect and prevent it. Through a detailed examination of Logistic Regression and K-Nearest Neighbor methodologies, this study contributes to the ongoing efforts to secure digital transactions against the pervasive threat of credit card fraud.

Ethical Implications

The ethical implications of fraud detection are multifaceted and significant. The use of data in fraud detection raises questions about privacy, consent, and the potential for bias in predictive models. Ensuring that data collection and analytics methods respect individual privacy rights and comply with regulatory standards is paramount. Additionally, the deployment of these systems must consider the potential for false positives, where innocent individuals could be wrongly accused of fraud, leading to social and financial consequences. This paper discusses these ethical considerations and the measures taken to mitigate such risks, ensuring that fraud detection algorithms are not only effective but also fair and just.

Threats to Validity

This study faces several threats to validity that could affect the conclusions drawn from the data analysis. The primary concern is the representativeness of the dataset, which comprises over one million transactions sourced from Kaggle. While extensive, this dataset may not fully capture the diversity of fraudulent activities globally or the evolving tactics of fraudsters. Additionally, the imbalance in the dataset, with fraudulent transactions constituting less than 1% of all transactions, poses challenges to model accuracy and generalizability. This section explores these threats in detail and discusses the strategies employed to address them, such as advanced sampling techniques and model validation processes.

This research aims to provide a comprehensive understanding of fraud analytics, from its definition and fundamental characteristics to the advanced analytical methods used to detect and prevent it. Through a detailed examination of Logistic Regression and K-Nearest Neighbor methodologies, this study contributes to the ongoing efforts to secure digital transactions against the pervasive threat of credit card fraud.

Related Works

In developing advanced algorithms for credit card fraud detection, I embarked on a comprehensive exploration of existing scholarly works, particularly focusing on the evolution and application of machine learning in this field. My literature review, grounded in seminal studies, documents the dynamic nature of fraud and the technological advancements that both enable and combat these illicit activities.

The foundational work by Barker, provides a critical overview of the history and methodologies of credit card fraud, underscoring the adaptive strategies of fraudsters in response to evolving technological landscapes [3]. This paper emphasizes the constant need for evolving countermeasures in fraud detection and has significantly shaped the conceptual framework of my research¹.

Further, the analytical study by Baker, M. J., & Miceli, T. explores the economics of crime, focusing on the credibility of threats in law enforcement and its implications for fraud prevention [2]. Their work delves into the strategic interplay between criminal activities and policy enforcement, revealing how the credibility of law enforcement actions influences criminal behavior. This research is instrumental in understanding the broader economic and strategic elements that underpin fraud prevention mechanisms².

I also reviewed the significant contributions of Ajeet Singh and Anurag Jain, who address the challenges of application-level fraud detection using machine learning techniques such as J48 decision tree, AdaBoost, and Random Forest [15]. Their comparative analysis provides insights into the effectiveness of various methods in detecting credit card fraud, particularly highlighting the importance of feature selection in enhancing model performance³.

The work by C. Victoria Priscilla and D. Padma Prabha introduces a novel two-phase feature selection approach that combines filter and wrapper

methods [12]. This methodology significantly improves the classification performance and computational efficiency of machine learning models, offering a promising solution to the challenge of high-dimensional and imbalanced datasets prevalent in fraud detection⁴.

The papers by Naoufal Rtayli and Nourddine Enneya, and N. Sivakumar and Dr. R. Balasubramanian provide a deep dive into the use of Random Forest Classifier and Logistic Regression for feature selection and fraud risk identification [13]. Their methodologies underscore the importance of precise feature selection and the strategic application of machine learning techniques to enhance the detection and prevention of credit card fraud⁵.

Additional relevant research includes studies on the application of predictive analytics technology in Canada by Kosemani Temitayo Hafiz, Dr. Shaun Aghili, and Dr. Pavol Zavorsky, which discusses the utilization of predictive analytics to detect credit card fraud [8].

Navanshu Khare and Saad Yunus Sait present an insightful analysis using machine learning models to detect credit card fraud, emphasizing the collation of different models to improve detection accuracy [7]. Dushyant Singh, Saubhagya Vardhan, and Dr. Neha Agrawal present a comprehensive study on the effectiveness of various machine learning algorithms in identifying fraudulent transactions[16].

Priyanka Yadav and colleagues propose a novel framework for credit card fraud detection using distributed data mining techniques to enhance the scalability and efficiency of fraud detection systems [1]. Anuruddha Thennakoon and his team focus on real-time fraud detection using machine learning, emphasizing the immediacy of detecting fraudulent activities[4].

Franscisca, Nonyelum, and Ogwueleka explore various data mining methods in credit card fraud detection systems [10]. Nutan Suman provides a comprehensive review of the methodologies employed in credit card fraud detection. Sushmito Ghosh and Douglas L. Reilly utilize a neural network for credit card fraud detection, illustrating the adaptability and efficiency of neural networks [?].

Aihua Shen, Rencheng Tong, and Yaochen Deng examine the strengths and weaknesses of various classification models in credit card fraud detection [8]. Lastly, Abhinav Srivastava and colleagues use a Hidden Markov Model for detecting credit card fraud, focusing on the sequential aspect of transactions [11]. These studies collectively enrich the understanding of credit card fraud detection, providing a robust theoretical foundation and practical insights into employing advanced technologies to tackle this pervasive issue

effectively[14].

Method of Approach




Figure 3: Architecture Of The Proposed System

My tool uses predictive analytics which is supervised learning, that means that certain types of transactions will be labeled as suspicious or fraud, and then runs through the logic done via feature engineering and states this transaction is fraudulent yes or no? It gives response in accordance to the training and the set features logic that have patterns that are able to separate fraudulent transaction from legitimate transactions. And after that has been determined using my tool then the fraud investigation team comes in and look into it in more detail and see how they can combat the concern and intercept that transaction before it is complete. The tool is mostly useful

to financial investigators who are trying to combat transactions that illegitimate and it helps them be able to flag and analyze fraudulent transaction by having the feature data analysed and categorised as whether fraud or not and lead to them creating and having fraud prevention strategies thus saving the company a lot of money and losses.

The proliferation of digital transactions has concurrently escalated the incidences of credit card fraud, presenting a formidable challenge to the financial sector. The advent of machine learning offers promising avenues for the development of advanced fraud detection systems. This study embarks on creating a logistic regression model tailored for predicting credit card fraud, navigating through an elaborate process involving data collection, preprocessing, exploratory analysis, and model training and evaluation. The purpose of this sequential methodology is to meticulously address the intrinsic and extrinsic challenges posed by fraud detection tasks, notably the class imbalance problem, and to enhance the model's predictive accuracy and reliability. The initial stage in building the fraud detection model involves collecting a comprehensive dataset embodying transactions with a dichotomous classification: fraudulent and legitimate. This dataset exhibits a pronounced class imbalance, with fraudulent transactions significantly outnumbered, a common trait in fraud detection datasets due to the relative rarity of fraudulent activities. Observe below that out of 1 million transactions in the dataset, only 87,403 (8.74%) transactions are fraudulent

```
df[['fraud']].value_counts()
```

Table 1: The unbalanced nature of the dataset

Fraud	Count
0.0	912597
1.0	87403

The inherent class imbalance poses a substantial challenge, as it can lead to model bias towards the majority class, reducing the predictive accuracy for the minority class (fraudulent transactions). To address this, we explored several strategies, including oversampling the minority class, undersampling the majority class, and employing synthetic data generation techniques like SMOTE (Synthetic Minority Over-sampling Technique).

```

#oversampling using the imbalanced-learn python library
from imblearn.over_sampling import RandomOverSampler
oversample = RandomOverSampler(sampling_strategy=1)
X_over, y_over = oversample.fit_resample(X, y)

#re-merge for EDA
df_over = pd.concat([X_over,y_over],axis=1)

```

These methods aim to balance the class distribution, thereby enhancing the model's ability to learn from both classes equally. Exploratory Data Analysis (EDA) plays a crucial role in understanding the dataset's characteristics. Through EDA, we identify patterns, anomalies, and correlations between features that could influence the model's fraud detection capabilities. This process involves statistical summaries, visualization techniques, and correlation analyses, providing insights into the data's structure and guiding the preprocessing steps. Visualizing Distribution and Correlations: Visual tools like histograms, box plots, and heatmaps allow us to observe the distribution of individual features and their relationship with the fraud outcome. For instance, box plots comparing transaction amounts between fraudulent and legitimate transactions can reveal significant differences in their distributions, suggesting transaction amount as a potentially indicative feature of fraud. Correlation Analysis: Calculating the correlation coefficients between features and the target variable helps in identifying the features most relevant to predicting fraud. Features with high correlation to the fraud variable are prioritized in the model training phase.

```

[91]: used_pin_number      -0.100293
      used_chip           -0.060975
      repeat_retailer     -0.001357
      distance_from_last_transaction  0.091917
      distance_from_home   0.187571
      online_order         0.191973
      ratio_to_median_purchase_price  0.462305
      fraud                1.000000
      Name: fraud, dtype: float64

```

Figure 4: correlation values




Figure 5: correlation heatmap




Figure 6: correlation bargraph

Logistic Regression

Logistic Regression is a statistical method for binary classification. It estimates the probability that a given input belongs to a particular category. The probability output is transformed using the logistic function (sigmoid), ensuring that the output lies between 0 and 1. The model prediction is typically made by setting a threshold of 0.5, above which the transaction is classified as fraudulent.




Figure 7: Logistic Regression

```
#as part of preprocessing, I shall transform the  
#independent variable to make it appropriate to be passed  
#into logistic regression model  
from sklearn.preprocessing import StandardScaler  
columns = X_over.columns  
scaler = StandardScaler()  
X_train_transformed = scaler.fit_transform(X_train)  
X_train_transformed = pd.DataFrame(X_train_transformed, columns = columns)  
X_train_transformed.head()
```

Logistic regression is chosen for its efficiency and effectiveness in binary classification tasks. The model is developed to estimate the probability that a given transaction is fraudulent, based on the input features. This probabilistic approach is particularly suited for scenarios where understanding the uncertainty of predictions is crucial.

K-Nearest Neighbors (KNN)




Figure 8: KNN

```
fraud_detection.train_and_evaluate_model(  
    "KNN", KNeighborsClassifier(n_neighbors=5)  
)  
fraud_detection.save_model("KNN", "credit_card_fraud_knn_model.sav")
```

The KNN algorithm is employed as an auxiliary analysis tool to further validate the logistic regression model's findings and to explore the data's structure from a different perspective. KNN's reliance on the proximity of data points to make predictions makes it a valuable tool for identifying non-linear patterns that logistic regression might miss. This analysis involves selecting an optimal k value, which is crucial for balancing the bias-variance trade-off inherent in KNN.

Each step in the methodology is designed to build upon the previous, ensuring a holistic and rigorous approach to model development: Preprocessing: Prepares the dataset for analysis, addressing class imbalance and ensuring the model learns from a balanced representation of fraud and legitimate transactions.




Figure 9: balanced dataset bar graph

EDA: Provides deep insights into the dataset, guiding the feature selection and preprocessing strategies by highlighting key patterns and anomalies.




Figure 10: ROC unbalanced dataset

Model Development: The core phase where the predictive model is built, trained, and evaluated. The choice of logistic regression is justified by its suitability for binary classification and its interpretability, which is crucial for trust and adoption in financial applications. **KNN Analysis:** Serves as a complementary analysis, offering a different perspective on the data's structure

and the model's performance, thereby enriching the overall understanding and robustness of the fraud detection capabilities. Conclusion The development of a credit card fraud detection tool via logistic regression and supplementary KNN analysis represents a comprehensive approach to tackling the challenges posed by digital transaction fraud. By systematically addressing class imbalance, leveraging EDA for informed feature selection, and employing a rigorously evaluated predictive model, this study contributes significantly to the ongoing efforts in financial security. The methodology not only enhances the predictive accuracy of the fraud detection model but also provides a replicable framework for future research in this domain.

Integrating the Flask application into the broader narrative of the credit card fraud detection project involves connecting the back-end analytical process with the front-end user interaction facilitated by Flask. This section aims to seamlessly incorporate the Flask app's role within the project's methodological framework, detailing its integration from data processing to user interaction and prediction delivery.

Integrating Flask into the Project's Methodological Framework Overview

The Flask application represents the culmination of our project, serving as the interface between the user and the machine learning model developed through rigorous data analysis. It is the bridge that translates user input into actionable data, feeding it to the back-end model for fraud prediction and presenting the results in an understandable format. This section explores how the Flask app is intricately linked to each phase of the project, from initial data analysis to model deployment and user interaction.

From Data Analysis to Model Deployment Model and Scaler Loading

At the core of our Flask app are the logistic regression model and the scaler, both crucial components developed in the preceding stages of our research. The model, trained to predict fraudulent transactions based on a set of features, and the scaler, used to normalize input data, are loaded into the app. This direct use of the model and scaler bridges the gap between the analytical work conducted in Python and the real-world application of these tools through the Flask app.

```
model = joblib.load("bill_credit_card_fraud_prediction.sav")
scaler = joblib.load("bill_credit_card_fraud_prediction_scaler.pkl")
```

These lines of code symbolize the transition from the development and evaluation phase, where the model's efficacy is rigorously tested, to the deployment phase, where the model is put into practice.

Feature Importances The Flask application also incorporates the knowledge gained from the Exploratory Data Analysis (EDA) and model refinement stages, particularly regarding feature importances. The definition of feature importances within the app:

```
feature_importances = {  
    # Feature importance values based on analysis  
}
```

This inclusion is pivotal, as it reflects the analytical process's findings on which features significantly influence the model's fraud prediction capabilities. This not only informs the user of the relative importance of each feature but also enhances transparency regarding how the model makes its predictions.

User Interaction and Prediction The Transaction Form

The Flask app's transaction form allows users to input transaction data, representing the features identified as significant during the EDA and model development phases:

```
class TransactionForm(FlaskForm):  
    # Fields for each feature
```

This form is the user's entry point to interact with the model, designed based on the project's analytical insights. Each field corresponds to a feature that the logistic regression model uses to predict transaction fraudulence, underscoring the form's design's direct connection to the analytical findings.

Prediction and Visualization

Upon form submission, the Flask app processes the user input, scales it using the previously loaded scaler, and feeds it into the model for prediction:

```
def return_prediction(model, scaler, sample_json):  
    # Process user input and predict
```

The app then presents the prediction results to the user in a user-friendly manner, alongside visualizations of the inputted features and their importances. This not only delivers the model's prediction but also provides a visual representation of the data that led to the prediction, enhancing user understanding and trust in the model.

```
@app.route("/prediction")
def prediction():
    # Display prediction results and feature importances
```

The use of Plotly for visualization further bridges the gap between analytical processes and user experience, presenting the feature importances and user inputs in an interactive format. This choice reflects a thoughtful integration of back-end analysis with front-end user interaction, ensuring that the insights gained from the project's methodological approach are communicated effectively to the user.

The Flask app is the practical manifestation of the project's methodological journey, from initial data collection and analysis to model development and evaluation. It encapsulates the project's analytical essence, transforming it into a user-oriented tool that leverages sophisticated machine learning techniques to predict credit card fraud. This seamless integration of Flask into the project's workflow exemplifies how advanced data analysis can be translated into real-world applications, bringing valuable insights directly into the hands of users.

Thank You. Here is the Information You Gave:

- Distance From Home: 15
- Distance From Location of Last Transaction: 20
- Ratio to Median Purchase Price: 2
- Repeat Retailer?: 1
- Used Chip?: 1
- Used Pin Number?: 1
- Online Transaction?: 1

Prediction: This transaction is not fraudulent




Figure 11: dashboard

Experimental Evaluation

To demonstrate the utility and accuracy of our model, we conducted a test to predict a single transaction's class, specifically targeting the detection of a fraudulent transaction. This test not only exemplifies the model's predictive capabilities but also showcases the preprocessing steps required to prepare data for prediction. First, we selected a transaction previously labeled as fraudulent from our dataset. This selection was done by filtering the dataset for instances with the 'fraud' label set to 1 and taking the first entry:

```
df_1fraud = df[df['fraud'] == 1].head(1)
```

Next, we prepared the transaction for prediction. This preparation involved removing the 'fraud' column, as our objective was to predict this label based on other features:

```
df_1fraud = df_1fraud.drop('fraud', axis=1)
```

Subsequently, we scaled the features using a previously fitted scaler, transforming the data to match the scale of the model's training data:

```
df_1fraud, _ = scaler.transform(df_1fraud)
```

The features considered for this specific transaction were distance from home, distance from the last transaction, ratio to median purchase price, whether the retailer was repeated, use of chip, use of pin number, and whether it was an online order. Below is the table of scaled feature values used for predicting a fraudulent transaction:

Table 2: Scaled Features for Fraud Prediction

Feature	Scaled Value
Distance from home	-0.40851561
Distance from last transaction	1.29597531
Ratio to median purchase price	0.55667687
Repeat retailer	0.36800856
Used chip	-0.6671046
Used pin number	-0.24467067
Online order	0.52500085

Following scaling, we converted the scaled features back into a DataFrame, ready for prediction:

```
df_1fraud = pd.DataFrame(df_1fraud, columns=columns)
```

Model Prediction Utilizing the logistic regression model (logmodel), we predicted the class of the transaction:

```
logmodel.predict(df_1fraud)
```

The model returned an array([1.]), correctly identifying the transaction as fraudulent. This outcome not only validates the model’s capability to discern fraudulent transactions but also illustrates the importance of accurate feature scaling and data preparation in the prediction process. The culmination of this research project is marked by the comprehensive evaluation of the logistic regression model developed for credit card fraud detection. This evaluation is a critical component of the research, providing insights into the model’s performance and its practical applicability in identifying fraudulent transactions. The process involves analyzing the model’s predictive accuracy through various metrics, including precision, recall, F1-score, and the area under the Receiver Operating Characteristic (ROC) curve (AUC). The utilization of these metrics offers a multi-dimensional view of the model’s effectiveness, addressing both its strengths and areas for improvement.

Model Performance Analysis

The precision metric, which indicates the model’s accuracy in predicting fraudulent transactions, shows a remarkable improvement from 0.61 to 0.95. This significant enhancement suggests that the model has become highly

effective in identifying true positive cases without being overwhelmed by false positives. This improvement is crucial in the context of fraud detection, where the cost of false alarms can be high, both in terms of operational efficiency and customer trust.

Table 3: Initial classification_report

	precision	recall	f1-score	support
0.0	0.96	0.99	0.98	273810
1.0	0.89	0.61	0.72	26190
accuracy			0.96	300000
macro avg	0.93	0.80	0.85	300000
weighted avg	0.96	0.96	0.96	300000

Similarly, the recall metric, which measures the model’s ability to detect all actual fraudulent transactions, increased from 0.72 to 0.94. This improvement demonstrates the model’s enhanced sensitivity in identifying fraud cases among the transactions it analyses. High recall is essential in minimizing the risk of overlooking fraudulent activities, thereby reducing the potential financial and reputational damage to the institution employing this model.

Table 4: Final Classification_report

	Precision	Recall	F1-Score	Support
0.0	0.95	0.93	0.94	273968
1.0	0.93	0.95	0.94	273591
Accuracy			0.94	547559
Macro avg	0.94	0.94	0.94	547559
Weighted avg	0.94	0.94	0.94	547559

The F1-score, a harmonic mean of precision and recall, stands at 0.94, reflecting a balanced performance between the model’s precision and recall. This balance is indicative of the model’s robustness, affirming its reliability in fraud detection across various scenarios without disproportionately favoring

one metric over the other. Moreover, the model's AUC score, representing its ability to distinguish between fraudulent and non-fraudulent transactions across different thresholds, further confirms its efficacy.




Figure 12: ROC/AUC unbalanced dataset




Figure 13: ROC/AUC balanced dataset

A high AUC score near 1.0 is indicative of the model's excellent classification ability, reinforcing its utility as a powerful tool in the fight against credit card fraud. The evaluation results underscore the logistic regression model's potential in real-world applications. The significant improvements

in precision and recall, coupled with a high AUC score, validate the model’s capacity to serve as a reliable fraud detection tool within financial systems.

Table 5: Initial confusion metrics

	Predicted Negative	Predicted Positive
Actual Negative	271789	2021
Actual Positive	10233	15957

The confusion matrix is used to represents the performance of a classification model in the context of a binary classification task, such as fraud detection. It’s a powerful tool for understanding not just the model’s accuracy but the nature of its errors. True Negatives (TN): 271,789 instances were correctly predicted as non-fraudulent (Predicted Negative and are Actual Negative). This represents the bulk of predictions and shows that the model is highly effective at identifying legitimate transactions. False Positives (FP): 2,021 instances were incorrectly predicted as fraudulent (Predicted Positive) despite being non-fraudulent (Actual Negative). These are Type I errors, leading to unnecessary investigations or alerts but are relatively low, indicating the model’s high precision. False Negatives (FN): 10,233 instances were incorrectly predicted as non-fraudulent (Predicted Negative) despite being fraudulent (Actual Positive). These are Type II errors, representing missed fraudulent transactions, which are critical in fraud detection contexts due to the potential for financial loss. True Positives (TP): 15,957 instances were correctly identified as fraudulent (Predicted Positive and are Actual Positive). This number, while lower than the True Negatives, is crucial for evaluating the model’s ability to catch fraudulent activities. From this confusion matrix, several important performance metrics can be calculated: Precision (for the Positive class): The ratio of True Positives to the sum of True Positives and False Positives. It measures the model’s accuracy in predicting positive instances. High precision indicates a low rate of false alarms. Recall (or Sensitivity for the Positive class): The ratio of True Positives to the sum of True Positives and False Negatives. It measures the model’s ability to detect all relevant instances. High recall is especially important in fraud detection, as it reflects the model’s effectiveness at identifying fraudulent transactions. F1-Score: The harmonic mean of precision and recall. It balances both metrics, providing a single measure to evaluate the model’s overall performance when precision and recall are equally important. Accuracy: The

ratio of correct predictions (both True Positives and True Negatives) to all predictions. While not explicitly calculated from the confusion matrix, it can be easily derived and offers a quick glance at overall performance. However, in imbalanced datasets, like fraud detection, accuracy can be misleading if the model is biased towards the majority class. The confusion matrix highlights the model's strengths in identifying non-fraudulent transactions while maintaining a reasonable capability to detect fraudulent ones. The relatively low False Negatives compared to True Positives is promising but still suggests room for improvement, especially since missing fraudulent transactions can be costly. Future model enhancements might focus on reducing False Negatives to improve recall, possibly through more sophisticated modeling techniques or feature engineering.

Table 6: Final Confusion metrics

	Predicted Negative	Predicted Positive
Actual Negative	255532	18436
Actual Positive	13860	259731

The following unit tests were designed and executed, each addressing a different aspect of the system:

- Data Loading Tests:**
 - Success Scenario:** Tested the successful loading of data, ensuring the DataFrame is not None and not empty. This verifies that the model can correctly load and recognize the dataset.
 - File Not Found Scenario:** Tested loading with a non-existing file, expecting a `FileNotFoundException`. This checks the model's error handling capabilities when provided with incorrect file paths.
- Data Preprocessing Test:** Tested the preprocessing functionality to ensure that feature scaling is applied correctly. The test verifies if the mean of the scaled features is close to 0 and the standard deviation is close to 1, within an acceptable tolerance. Additionally, it confirms that the target variable `y` is properly extracted and not null.
- Exploratory Data Analysis (EDA) Test:** Conducted a test to ensure the EDA process completes without raising exceptions. This test assumes that the EDA process either modifies the dataset or generates plots directly, and its success indicates the stability and error-free execution of the EDA functions.
- Model Saving and Loading Test:** Tested the functionality for saving and loading a trained model using `joblib`. The process involves training a Logistic Regression model, saving it to a file, and then loading it to verify if the loaded model is not None. This test ensures that models can be

persisted and retrieved successfully for future use. Model Training and Evaluation Test; Focused on testing the training and evaluation process of a Logistic Regression model. This test verifies that the model is added to the model registry (model.models) and is not None post-training. It ensures that the model training and evaluation process is correctly implemented and that metrics are calculated. These unit tests collectively validate the robustness and reliability of the Fraud Detection Model across various critical functionalities: data handling, preprocessing, exploratory analysis, model persistence, and evaluation. By passing these tests, the model demonstrates its preparedness for real-world application, highlighting its capabilities in accurately detecting fraudulent transactions while ensuring flexibility and reliability through systematic testing.

This capacity to accurately identify fraudulent transactions can significantly enhance security measures, safeguarding both the financial institutions and their customers from the repercussions of fraud. Looking ahead, while the current model demonstrates promising results, the domain of fraud detection is ever-evolving, necessitating continuous model refinement. Future research could explore the integration of more complex machine learning algorithms, the inclusion of additional predictive features, and the adoption of adaptive learning techniques to further improve the model's accuracy and adaptability. Moreover, gathering and incorporating feedback from the model's practical deployment will be invaluable in identifying real-world challenges and opportunities for enhancement. The logistic regression model's performance evaluation highlights its effectiveness and potential in credit card fraud detection. The significant improvements in key performance metrics, achieved through meticulous model development and evaluation processes, set a solid foundation for its application in safeguarding financial transactions. As the model moves from a research setting to practical deployment, it opens avenues for continuous improvement and adaptation, aiming to stay ahead in the dynamic landscape of fraud detection.

Model Evaluation Metrics Summary

Accuracy: This is the ratio of correctly predicted transactions (both fraudulent and non-fraudulent). However, due to the imbalanced nature of fraud datasets, accuracy is not the best metric as it can be misleading. High accuracy might be achieved if the model simply predicts the majority class (non-

fraudulent transactions). Precision: This measures the accuracy of positive predictions. Formally, it is the ratio of correctly predicted positive observations to the total predicted positives. High precision relates to a low false positive rate. Recall (Sensitivity): This is the ability of the model to find all the relevant cases (fraudulent transactions). It is the ratio of correctly predicted positive events to the actual positives. F1-Score: This is the weighted average of Precision and Recall. Therefore, this score takes both false positives and false negatives into account. ROC Curve and AUC: The ROC curve is a graphical plot that illustrates the diagnostic ability of a binary classifier system as its discrimination threshold is varied. The area under the curve (AUC) measures the entire two-dimensional area underneath the entire ROC curve and provides an aggregate measure of performance across all possible classification thresholds.

From the initial model (before oversampling), we observed the following metrics:

Precision for fraudulent transactions (1.0): 0.89 Recall for fraudulent transactions (1.0): 0.61 F1-score for fraudulent transactions: 0.72 Overall accuracy: 95.917% These metrics indicated a model that was quite conservative about labeling transactions as fraudulent, which is why the recall was relatively low. Given the imbalance in the dataset, this model might not be effective in a real-world scenario where missing fraudulent transactions can be very costly.

After applying oversampling to balance the dataset, the model performance showed improvement:

Precision for fraudulent transactions: 0.93 Recall for fraudulent transactions: 0.95 F1-score for fraudulent transactions: 0.94 Overall accuracy: 94.132% These metrics demonstrate a much better balance between precision and recall, indicating a robust model that can effectively identify fraudulent transactions without excessively misclassifying non-fraudulent transactions.

The improved model after addressing the imbalance via oversampling provides a more reliable tool for fraud detection. It demonstrates a substantial improvement in identifying fraudulent transactions (high recall) while maintaining a strong accuracy in its predictions (high precision). This balance is crucial for practical applications in fraud detection, where both missing actual frauds and falsely labeling transactions as fraudulent can be very costly.

Conclusion

The integration of machine learning algorithms, specifically Logistic Regression and K-Nearest Neighbors (KNN), into the realm of credit card fraud detection represents a significant advancement in combating this persistent issue. This research examined the effectiveness of these models not just in terms of technical accuracy but also through their economic impacts, employing Cost-Benefit Analysis (CBA) and the Risk Management Model for a comprehensive evaluation.

Logistic Regression proved valuable for its ability to estimate the probability of fraudulent transactions with a clear probabilistic output, enhancing interpretability and enabling stakeholders to understand how various factors contribute to fraud risk. Despite its strengths, the linear assumptions inherent in Logistic Regression may limit its ability to capture more complex or nuanced fraud patterns without continuous updates and feature engineering.

KNN, on the other hand, capitalizes on the similarity of new transactions to historical data, offering adaptability to new and evolving fraud tactics. This model's effectiveness, however, is heavily dependent on the quality and representativeness of the data used, requiring meticulous tuning to ensure it accurately distinguishes between fraudulent and legitimate transactions.

The economic analysis anchored by the CBA outlined the costs related to the adoption and maintenance of these models against their potential benefits, which include reduction in fraud instances, operational efficiencies, and enhanced consumer confidence. The Risk Management Model further helped in assessing how these algorithms mitigate transactional risks, safeguarding financial assets and organizational reputation effectively.

The findings reveal that while Logistic Regression is highly effective in reducing fraud incidence due to its precision and recall, KNN excels in adapting to complex fraud patterns, making it invaluable in dynamic fraud environments. Economically, the benefits of implementing these models are sub-

stantial, encompassing not only direct reductions in fraud-related losses but also improvements in operational efficiency and customer trust, crucial for long-term business sustainability.

This study illustrates the critical role of advanced machine learning techniques in not only detecting and preventing fraud but also in enhancing the economic and strategic frameworks within which financial institutions operate. The substantial initial investment in these systems is justified by significant long-term gains—highlighting a proactive approach to fraud management that aligns with broader risk management and economic objectives.

The application of Logistic Regression and KNN in credit card fraud detection is not merely a technical solution but a strategic decision that offers extensive economic benefits and aligns with comprehensive risk management strategies. This research underscores the importance of evaluating such technologies through both a technical lens and an economic perspective, ensuring that they provide not only high predictive accuracy but also contribute positively to the overall financial health of the implementing organization.

Here is the link to my artifact tool link

Here is the link to my presentation Presentation

Future Work

My exploration into developing a logistic regression model for credit card fraud detection has been an enlightening journey, highlighting both the potential and challenges of leveraging machine learning in combating financial fraud. This initial success has set the groundwork for further exploration into more sophisticated machine learning algorithms, acknowledging the necessity for continuous innovation in fraud detection methodologies due to the ever-evolving tactics of fraudsters.

Advanced Machine Learning Algorithms: While logistic regression has proven effective due to its simplicity and interpretability, the complex and adaptive nature of fraud patterns suggests significant benefits from exploring more advanced algorithms. Deep learning models, for instance, have the potential to discern intricate patterns within extensive datasets, capturing complex, non-linear relationships that simpler models may miss. Additionally, ensemble methods such as Gradient Boosting Machines (GBMs) and Random Forests could enhance prediction accuracy and resilience by harnessing the collective strength of multiple models.

Data Diversification and Expansion: The effectiveness of fraud detection systems greatly depends on the diversity and volume of the data they train on. Expanding the dataset to include a wider array of transactions across various regions, currencies, and fraud scenarios will not only enhance the model's applicability but also improve its responsiveness to new and emerging fraud trends through real-time data utilization.

Long-term Model Effectiveness: Ensuring the long-term effectiveness of these models is critical in the dynamic field of fraud detection. I envision a future where continuous model evaluation and iterative updates keep pace with changing fraud patterns. This approach may include automated re-training mechanisms and feedback loops that integrate new data, essential for maintaining the accuracy and relevance of the models over time.

Model Interpretability and Fairness: As machine learning models increase in complexity, the importance of interpretability and fairness becomes more pronounced. Future work will focus on methods that enhance model transparency to foster trust and accountability. It will also address potential biases, ensuring equitable treatment across all user transactions and that advancements in fraud detection are achieved without compromising fairness.

Integration with Other Fraud Prevention Technologies: The next frontier involves integrating machine learning models with other fraud prevention technologies, such as behavioral biometrics, device fingerprinting, and anomaly detection systems. This holistic fraud management ecosystem will create a multi-layered defense that enhances fraud detection capabilities without compromising user experience.

Economic and Social Impact Considerations: Moving forward, it is essential to consider the economic impacts and the balance of cost-benefit aspects of enhanced fraud detection systems. My commitment to continuous innovation and adaptation is fueled by a vision to significantly advance our capabilities in detecting and preventing fraud, thus safeguarding the integrity of financial systems and enhancing consumer trust globally.

References

- [1] Amir Aziz, Hamid Ghous, Multan Punjab, and Pakistan. 2021. Fraudulent transactions detection in credit card by using data mining methods: A review. 79, (January 2021), 18.
- [2] Matthew J Baker and Thomas J Miceli. 2021. Crime, credible enforcement, and multiple equilibria. *International Review of Law and Economics* 68, (2021), 106030.
- [3] Katherine J Barker, Jackie D’amato, and Paul Sheridon. 2008. Credit card fraud: Awareness and prevention. *Journal of financial crime* 15, 4 (2008), 398–410.
- [4] Peter Burns and Anne Stanley. 2002. Fraud management in the credit card industry. *Federal Reserve Bank of Philla Payment Cards Center Discussion Paper* 02-05 (2002).
- [5] Anastasia Cheliatsidou, Nikolaos Sariannidis, Alexandros Garefalakis, Jamel Azibi, and Paschalis Kagias. 2023. The international fraud triangle. *Journal of Money Laundering Control* 26, 1 (2023), 106–132.
- [6] David S Evans. 2004. The growth and diffusion of credit cards in society. *Payment Card Economics Review* 2, (2004), 59–76.
- [7] Sandhya Gangadharan, M. Abishek, S. Kumar, and R. Kumar. 2022. Credit card fraud detection using machine learning algorithms. 313–320. https://doi.org/10.1007/978-981-19-5221-0_30
- [8] Temitayo Kosemani, Shaun Aghili, and Pavol Zavorsky. 2016. The use of predictive analytics technology to detect credit card fraud in canada. June 2016. 1–6. <https://doi.org/10.1109/CISTI.2016.7521522>
- [9] Dennis Ng. 2018. Evolution of digital payments: Early learnings from singapore’s cashless payment drive. *Journal of Payments Strategy & Systems* 11, 4 (2018), 306–312.

- [10] Francisca Ogwueleka. 2011. Data mining application in credit card fraud detection system. *Journal of Engineering Science and Technology* 6, (June 2011), 311–322.
- [11] Amit Pandey and Achin Jain. 2017. Comparative analysis of KNN algorithm using various normalization techniques. *International Journal of Computer Network and Information Security* 11, 11 (2017), 36.
- [12] C Victoria Priscilla and D Padma Prabha. 2021. A two-phase feature selection technique using mutual information and XGB-RFE for credit card fraud detection. *Int. J. Adv. Technol. Eng. Explor* 8, (2021), 1656–1668.
- [13] Naoufal Rtayli and Nourddine Enneya. 2020. Selection features and support vector machine for credit card risk identification. *Procedia Manufacturing* 46, (2020), 941–948.
- [14] Chennai Selaiyur and Tamil Nadu. A study on consumer approach towards credit cards-with reference to salem.
- [15] Ajeet Singh and Anurag Jain. 2019. Adaptive credit card fraud detection techniques based on feature selection method. In *Advances in computer communication and computational sciences: Proceedings of IC4S 2018*, 2019. Springer, 167–178.
- [16] N Sivakumar and R Balasubramanian. 2015. Fraud detection in credit card transactions: Classification, risks and prevention techniques. *International Journal of Computer Science and Information Technologies* 6, 2 (2015), 1379–1386.
- [17] Stephen Trouard, Mark Morgan, and Billy Morehead. 2023. Value of certified fraud examiner certification-perceptions of government financial managers. *Business Education Innovation Journal* 15, 1 (2023).