Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Advanced Machine Learning Approaches for Credit Card Fraud Detection in the USA: A Comprehensive Analysis

Mir Mohtasam Hossain Sizan[1], Anchala Chouksey[2], Nikhil Rao Tannier[3], Md Abdullah Al Jobaer[4], Jasmin Akter[5], Ashutosh Roy[6], Mehedi Hasan Ridoy[7], M Saif Sartaz[8], Dewan Aminul Islam[9]

Abstract

Credit card fraud is a financial threat in America, both for financial institutions and for consumers, and it is growing in severity. Traditional fraud detection methods become less effective in countering emerging fraud trends, and for that reason, sophisticated algorithms in machine learning have to be embraced. This research project strived to develop and compare complex algorithms for fraud detection in credit cards in America. With a variety of algorithms including both unsupervised and supervised learning, this study strived towards improving fraud transaction detection rates. This study focuses on real-world credit card transaction datasets from America, offering a robust foundation for comprehending the intricacies of fraud detection in an authentic financial context. Employing actual transaction data, the study aims to replicate and model variation and nuance in fraud and consumer behavior, such that any developed machine learning algorithms will have a basis in real-life realities. For model selection, we deployed several machine learning models, notably Logistic Regression, Random Forest, and XG-Boost Classifier. In evaluating model performance, several key metrics, including Precision, Recall, and the F1-score, were taken into consideration. Random Forest Classifier performed best overall, with relatively high accuracy for fraud prediction, and average recall, with a marginally high level of F1-score. Overall, it can be noticed that Random Forest has the most balanced performance out of the three in fraud detection capabilities, which seems to be a necessity. The integration of real-time fraud prevention with machine learning models is revolutionizing financial institution transaction monitoring. ML models can analyze and process information in real-time, and thus, allow for effective and efficient real-time fraud monitoring. The future of fraud detection holds many exciting avenues for research, most prominently in deep model development. Methods in deep learning, such as recurrent neural networks (RNNs) and convolutional neural networks (CNNs), have been successful in discovering complex structures and sequential relations in transactional information. Another promising avenue for future research is combining AI-powered identity verification with blockchain technology for fraud prevention.

Keywords: *Credit Card Fraud; Fraud Detection; Machine Learning; False Positives; Predictive Modeling; Financial Security.*

Introduction

According to Rahman et al. (2024), in recent years, fraud cases in America have increased, and financial loss and loss for both financial institutions and consumers have been enormous. In recent times, statistics have shown that trillions of dollars in financial loss have been incurred through fraud, and thus, a high demand for effective fraud detection tools is a must. As financial service digitization continues to rise, financial transactions have become convenient, but fraudsters have gained fertile ground for taking advantage of such conveniences. In such a case, real-time fraud detection tools cannot be over-emphasized enough. Not only will such tools protect financial consumers' monies, but financial institutions will also save themselves from fraud-related financial loss and reputational loss (Shawon et al., 2024; Sizan et al., 2023)

[1] Masters of Science in Business Analytics, University of North Texas. E-mail: mhsizan855@gmail.com (Corresponding Author)
[2] Masters of Science in Financial Mathematics, University of North Texas, Denton, Texas
[3] Master's in Artificial Intelligence, University of North Texas
[4] Electronics and Telecommunications Engineering, Daffodil International University
[5] Masters of Business Analytics, Gannon University, Erie, PA
[6] MBA in Business Analytics, Gannon University, Erie, PA
[7] MBA in Business Analytics, Gannon University, Erie, PA
[8] Electrical Engineering and Computer Science, Florida Atlantic University
[9] MSc, Department of Electrical and Computer Engineering, Florida Atlantic University

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Islam et al. (2024), reported that real-time detection is significant in that it can enable timely response to suspicious activity, closing down windows of opportunity for fraudsters. Rule and manual review methodologies cannot possibly hope to counter fraudsters' increasingly sophisticated techniques, and with fraudsters updating and altering techniques at a constant pace, ever-changing sophisticated tools that can draw on the capabilities of machine learning become increasingly critical. Sumsuzoha et al. (2024), asserted that Machine learning algorithms can review massive volumes of transactional information in real-time, identifying trends and outliers that can represent fraud.

Problem Statement

Notwithstanding the effectiveness of traditional fraud detection methods, traditional methods have a drawback in that they depend on predefined sets of rules and historical information. Traditional methods lack fraud adaptability and, therefore, high rates of false negatives, in instances when fraud transactions pass undetected, arise. Traditional methods even have high cases of false positives, in which correct transactions become detected fraud, and, in consequence, loss of trust and dissatisfaction occurs in the case of customers. Having an ideal balance between fraud accuracy and fewer alarm cases is a critical challenge (Akter et al., 2023).

The evolving fraud trends make it even more challenging for detection. Fraudsters, in turn, regularly update and adapt, utilizing even machine learning for their use in an attempt to circumvent detection. Consequently, financial institutions have a constant necessity to update and modify fraud detection approaches similarly. There is a critical demand for even more sophisticated techniques capable of taking full advantage of machine learning for increased capabilities in terms of detection and overcoming even the issue of false positives and negatives (Al Montaser et al., 2025)

Research Objective

This research will strive to develop and compare complex algorithms for fraud detection in credit cards in America. With a variety of algorithms including both unsupervised and supervised learning, this study will strive towards improving fraud transaction detection rates. The study will strive towards enhancing overall fraud detection efficiency and minimizing cases of false alarm, in an attempt to safeguard consumer trust and financial integrity. The objectives include a critical review of a range of machine learning methodologies, testing in real-life settings, and best practice guidance for deploying them. In general, the purpose is to contribute to developing fraud detection systems that are robust, effective, and resilient in a changing environment of credit card fraud. Through this analysis, the work aims to contribute towards financial institution planning, enabling them to maximize fraud detection and protect both consumers and their financials. The work can contribute to financial technology in general, proving the capabilities of machine learning in solving complex real-time analysis and decision-making issues.

Scope and Relevance

This study focuses on real-world credit card transaction datasets from America, offering a robust foundation for comprehending the intricacies of fraud detection in an authentic financial context. Employing actual transaction data, the study aims to replicate and model variation and nuance in fraud and consumer behavior, such that any developed machine learning algorithms will have a basis in real-life realities. Most important for fraud detection accuracy is employing sophisticated machine learning techniques, with such techniques having the capability to process enormous volumes of information in an attempt to detect trends and outliers that will not have been detectable with traditional techniques. By improving fraud detection accuracy, not only is an issue with heightened concerns regarding credit-card fraud addressed but financial institutions wishing to protect their customers and instill trust in an increasingly electronic marketplace have significant implications to gain from such a study.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Literature Review

Trends in Credit Card Fraud in America

As per Aditi et al (2022), credit card fraud is a prevalent issue in America, with new types of fraud developing with technology and changing consumption behavior. One of the most prevalent types of fraud is identity fraud, in which fraudsters steal a victim's Social Security numbers, bank information, and credit information in an attempt to impersonate a victim. Identity fraud can result in tremendous financial loss and long-term credit rating deterioration for a victim. Fraudsters apply a similar modus in a prevalent fraud, namely, a form of phishing, in which fraudsters compel a victim to disclose sensitive information through impersonation through an email, SMS, or a counterfeit website. Phishing scams have become sophisticated, and it is becoming increasingly challenging for a consumer to discern a real and a counterfeit message. Alfaiz & Fati (2022), argued that Synthetic fraud, a relatively new form, involves creating a new identity with a combination of real and spurious information. Fraudsters use such synthetic IDs to apply for financial institution accounts and make counterfeit purchases, and in most cases, go undetected for years at a stretch.

In response to increased credit card fraud, compliance requirements, and standards have been crafted to protect consumers and preserve financial system integrity. The Payment Card Industry Data Security Standard (PCI DSS) is a group of security standards for protecting card information during and after a financial transaction. The standards require companies to manage, process, and store information about cardholders in a specific manner, and compliance is imperative for fraud and data breach prevention (Alarfaj et al., 2022). Similarly, the Federal Financial Institutions Examination Council (FFIEC) promulgates guidance for financial institutions to make improvements in fraud management processes and fraud detection capabilities. In addition to increased fraud, these requirements represent the necessity for effective fraud detection tools capable of keeping pace with changing fraud methodologies and protecting both financial institutions and consumers (Bao et al., 2024).

Traditional Fraud Detection Methods

Hossain et al. (2025), stated that traditional fraud detection methodologies rely almost wholly on rule-based systems, utilizing predefined rules and heuristics to detect possibly suspicious transactions. Rule-based systems search for transaction information through a predefined criterion, raising an alarm for any transactions that fall out with predefined behavior. For instance, a purchase in a location far removed from a cardholder's routine behavior will generate an alarm. Rule-based systems can function effectively for simple fraud cases, but in environments with variable fraud trends, effectiveness is compromised. Bhowte et al. (2024), indicated as fraudsters adapt and new techniques for evading them develop, inflexibility in rule-based systems can generate high volumes of false negatives, with fraud transactions inaccurately labeled as valid ones. In addition, such systems generate high volumes of false positives, generating unnecessary inconvenience for both financial institutions and consumers.

The limitations of traditional methodologies become even more critical in today's ever-evolving electronic environment, in which high volumes of transactions can overwhelm traditional review processes. Fraudsters exploit such volumes for their benefit, employing sophisticated techniques that traditional systems can have a problem detecting (Ileberi et al., 2021). Besides, with the shifting behavior of consumers, predefined rules can become out of date in a matter of days, creating gaps in capabilities for detection. Consequently, financial institutions have to tread between having to detect fraud in real-time and accurately, and the inherent vulnerabilities of traditional methodologies (Khan et al., 2022).

Machine Learning in Fraud Detection

Patel (2023), indicated that the advent of technology in machine learning has revolutionized fraud detection, offering new techniques with heightened accuracy and responsiveness. Huge datasets can be analyzed through machine learning algorithms, and sophisticated outliers and trends can be detected, even when not apparent to humans and conventional rule-based programs. Decision trees, neural networks, and random

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

forests have been used to develop predictive algorithms capable of distinguishing between fraud and real transactions. Historical transactional information can be leveraged to train such predictive models, and with new cases, each one, such predictive models can become even better at fraud detection.

There are three general categories of fraud detection using machine learning: mixed approaches, unsupervised, and supervised approaches. In supervised approaches, model training is conducted with datasets with labels, with transactions labeled fraud or not fraud. With such a method, one can develop very high-accuracy models, with them learning fraud-related patterns through past examples (Mienye & Jere, 2024). Nevertheless, with the disadvantage of having to utilize labeled information, a challenge comes in scenarios with relatively infrequent fraud, generating unbalanced datasets and affecting model performance (Saheed et al., 2022).

Unsupervised learning algorithms, on the other hand, don't rely on labeled information. Instead, such a model looks for transaction behavior in a quest for fraud with no a priori knowledge of fraud definitions. Unsupervised approaches can work well at discovering new fraud behavior not yet encountered, but lack of labels can make model performance evaluation challenging, and one can label innocent transactions as frauds (Sizan et al., 2025). Hybrid models combine both unsupervised and supervised methodologies, utilizing both and blending them to make fraud detection even more efficient. For instance, a model can use unsupervised methodologies to identify anomalous trends in transaction behavior and then use supervised techniques to classify them as fraud or not fraud. By having a multi-faceted outlook, such a model can detect fraud at a level with fewer false positives, offering a more sophisticated answer to fraudsters' new modus operandi (Sumsuzoha et al., 2024).

Research Gaps

Aditi et al. (2022), regardless of the development in fraud detection via machine learning, several gaps in ongoing studies must be addressed. One such critical issue is creating fraud detection models that can scale and adapt to financial institution expansion and growing volumes of transactions. As financial institutions expand and volumes of transactions expand, fraud detection models must efficiently work with big datasets with accuracy. Most importantly, fraud detection models must adapt in real-time to emerging fraud trends, for fraudsters will adapt and introduce new approaches for exploiting vulnerabilities in present frameworks.

According to Rahman et al. (2024), another pressing concern is working with imbalanced datasets, a common issue in fraud detection in which fraud transactions are outnumbered many times over by legitimate transactions. Model performance bias can result from such an imbalance, in which case, over-tuning for finding the dominant class (the legitimate transactions) and under-predicting for the minority class (the fraud transactions) happens. To mitigate such an imbalance, new approaches have to be embraced, such as oversampling the minority class, under-sampling the dominant class, or employing cost-sensitive learning techniques in which a variable cost is incurred for misclassification in terms of its class.

Finally, real-time fraud detection is a critical challenge for financial companies. With increasingly high volumes of transactions, real-time analysis and rapid response to suspected fraud is a critical imperative. Constructing real-time processing with accuracy and low false positive rates is an imperative for both financial companies and consumers. There is a rich opportunity for future work in this area, with continued advances in both processing techniques for big data and in machine learning (Shawon et al., 2024).

Data Collection and Preprocessing

Dataset Overview

The credit fraud dataset is a rich collection of transactional records derived from a range of financial institution sources, public datasets, and real-life, anonymized transactions to yield a rich and representative collection of consumer activity. All of the records in the dataset include key information about a transaction such as value, merchant category, location, and payment channel, with timestamps providing contextual information about when a transaction happened. Fraud labels confirming whether a transaction is a fraud

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

or not have been added, and therefore, supervised machine learning algorithms can be trained. Having both a real-time and a historical dataset can allow for a deep analysis of fraud trends and behavior, and for developed models to adapt to changing credit fraud environments and have a high accuracy in fraud detection.

Feature Selection

1	**Transaction ID**	Unique identifier for each transaction.
2	**Amount**	Transaction amount in USD.
3	**Transaction Date**	Date of the transaction.
4	**Merchant ID**	Unique identifier for the merchant.
5	**Transaction Type**	Type of transaction (e.g., purchase, refund).
6	**Location**	Geographical location of the transaction.
7	**Is Fraud**	Binary target variable indicating fraud (1) or legitimate transaction (0).

Data Preprocessing

In the preprocessing, several key operations were executed to allow strong analysis and increased model performance. First, cleaning of the data was conducted to delete missing values and outliers, utilizing techniques such as imputation for missing values and statistics in identifying and countering outliers, improving the integrity of the data. Second, feature engineering was executed in encoding categorical features, transforming them into numerical values for use with machine algorithms, and numerical values normalized to maintain uniformity and allow proper comparisons between disparate scales. Third, in a move to counteract extreme imbalance in the target variable, the Synthetic Minority Oversampling Technique (SMOTE) was utilized. With it, synthetic samples for the minority class (fraudulent transactions) were produced, effectively balancing out the dataset and allowing both classes to learn effectively, minimizing bias towards the majority class (genuine transactions). All these preprocessing operations in unison helped develop a strong foundation for developing reliable and effective fraud detection models.

Exploratory Data Analysis (EDA)

Exploratory Data Analysis (EDA) is a critical part of the research process in which datasets are analyzed and represented in a systematic form in an attempt to expose hidden trends, patterns, and outliers in preparation for proceeding with traditional statistical analysis or with machine learning algorithms. EDA was conducted for a variety of objectives, including describing important parts of the data, finding variable relationships, testing for outliers, and estimating the quality of the dataset. With a variety of techniques, such as descriptive statistics, correlation analysis, and visualization tools (e.g., histograms, scatter plots, and box plots), the analyst gained useful information about the structure and distribution of the data. Not only does EDA go towards formulating hypotheses and guiding future analysis, but it can even inform preprocessing decisions, such that any model techniques' assumptions can be soundly backed. Overall, EDA is a critical part of developing a deeper awareness of the dataset and developing a wiser and more efficient model and its interpretation.

Fraud vs. Non-Fraud Transaction

The code in Python utilized Python modules pandas, matplotlib, and seaborn to visualize fraud and non-fraud transactions in a dataset. It began with loading packages and defining a pretty style for plots generated with Seaborn. The bulk of the code was in generating a count plot with Seaborn's count plot function. It plotted the count of each group in data frame df's 'Is Fraud' column. The generated plot was titled, and axis labels and custom tick labels were added to make "Fraud" and "Non-Fraud" transactions stand out prominently. Finally, plt.show() plotted out the generated visualization. Essentially, this code was a simple way in which one can see a fraud detection dataset's class balance.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output

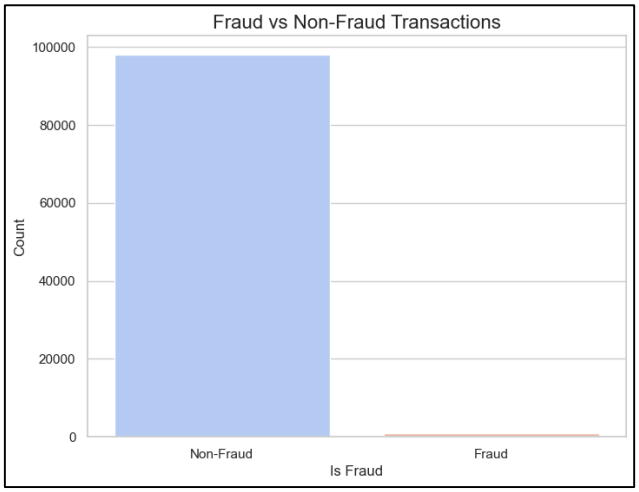

Figure 1: Fraud vs. Non-Fraud Transaction

The chart of "Fraud vs. Non-Fraud Transactions" evidently shows the extreme class imbalance in the dataset. There is a predominantly dominant presence of non-fraud transactions, with a count over 100,000, in the blue column, and a negligible presence of fraud transactions, in a hardly perceptible column, with a count of a mere handful—showing a critical imbalance in the target variable. This imbalance shows that the dataset consists predominantly of valid transactions, and such can make it challenging for machine learning algorithms, for they will have a bias towards the majority and will not accurately detect the minority (fraud transactions). All such observations exhibit the necessity for techniques such as SMOTE in balancing fraud cases in a dataset, such that the model can effectively learn to detect fraud even in its scarcity.

Transaction Amount Distribution by Fraud Status

The implemented code generated a kernel density estimation (KDE) plot to visualize the distribution of transaction values, with fraud and non-fraud differentiated. It utilized Python's seaborn, adding a matplotlib figure with a specific size predefined. It plotted density curves using the kdeplot function, with the 'Amount' column in data frame df and the 'IsFraud' column for stratify to plot two distributions for fraud and non-fraud activity. It set fill=True for under curve filling for easier reading, and a 'cool warm' colormap for differentiation, with transparency adjusted through alpha=0.6 for through-transparency in overhanging areas. It then sets a title, and axis labels, and then plotted through plt.show(), providing an analysis of how transaction values can differ between fraud and actual activity.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output

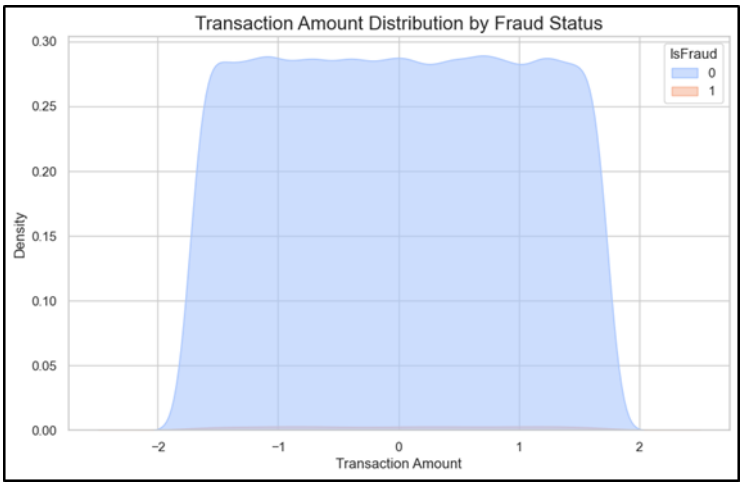

Figure 2: Transaction Amount Distribution by Fraud Status

The plot of "Transaction Amount Distribution by Fraud Status" reveals important tales about the distribution of transaction values and fraud and non-fraud classification. That both groups' density plots, in lighter blue for non-fraud (status 0) and a darker, less saturated one for fraud (status 1), have relatively uniform distributions for both groups informs us that frauds don't have a specific distribution in terms of value for transactions, in that both classes have a similar range of values clustering about zero, and fraud can occur at a range of values for transactions. That both plots are level and don't have spikes in values for both groups informs us that neither group will have spikes in values, and fraud will not necessarily stand out in terms of value, and value alone won't make a strong fraud case, and that deeper, additional information will have to be taken into consideration in creating a fraud model.

Fraud Transaction by Hour of the Day

The computed code in Python generated a count plot for fraud distribution over days and a range of days in a day. Particularly, it employs Python modules matplotlib and Seaborn for plotting and visualization, respectively. It first generates a figure with a specific size and then uses sns. Count plot for plotting fraud distribution over days and a range of days in a day. It plots the 'Hour' column of data frame df for representing x and hue for representing the 'Is Fraud' column, distinguishing between fraud and non-fraud for each 'Hour' value. The 'cool warm' scheme is used for aesthetic purposes. It then sets a title, axis labels, and a legend describing the 'Is Fraud' categories for visualization purposes. Finally, it employs plt.show () for plotting and visualization, providing an outlook for fraud activity times during the day.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output

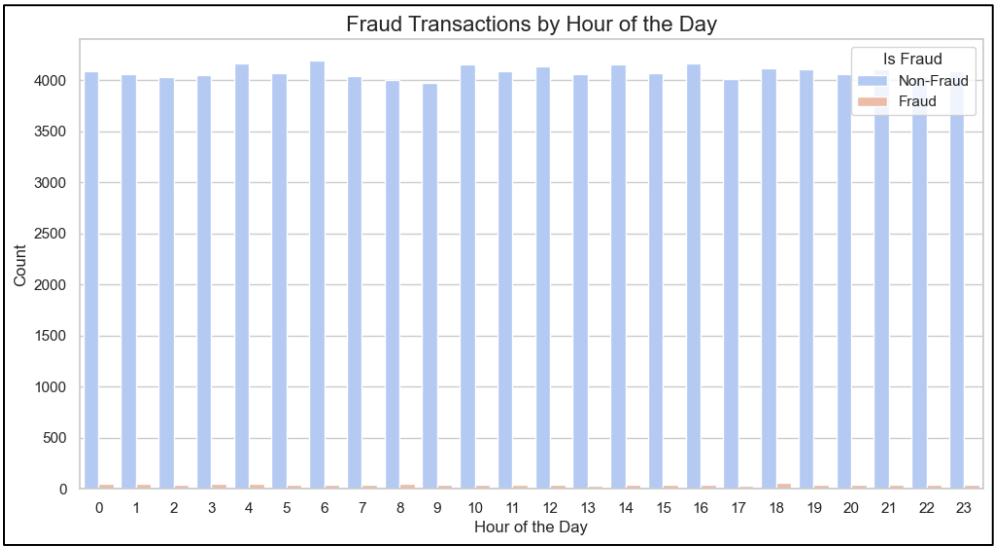

Figure 3: Fraud Transaction by Hour of the Day

The chart for "Fraud Transactions by Hour of the Day" shows a granular picture of both fraud and non-fraud occurrences for every hour of the day. Non-fraud occurrences are in lighter blue and fraud in darker orange. As can be seen, fraud occurrences have a relatively low count consistently, peaking at a little over 4,000, and fraud, therefore, doesn't have a high peak at any one hour. In contrast, non-fraud occurrences have a larger count, with volumes much larger for all times of the day. That such a level of fraud activity is consistently present, and not a concentrated activity at any one period, could make fraud difficult for a system to detect, one focused in its analysis in times of activity, for example. What is significant about such observations is that fraud is relatively infrequent and consistently present over a period, and not concentrated at any one period, and fraud can therefore become difficult for a system to detect, one focused in its analysis in times of activity.

Fraud Distribution by Transaction Type

The code script generated a horizontal bar plot of fraud distribution for a transaction type. It employed Python packages matplotlib and seaborn. There was a figure size, and a counterplot function in Seaborn generated the plot. Most notably, y='Transaction Type' generated a horizontal bar, and hue='Is Fraud' separated each one with fraud status. The information came from data frame df, and a 'cool warm' palette was used for colors. There was a title, an x label ("Count"), and a y label ("Transaction Type") for the plot. There was a legend describing the "Is Fraud" categories ("Fraud" and "Non-Fraud"). Lastly, plt.show() plotted it, offering information about the types of transactions most susceptible to fraud.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output

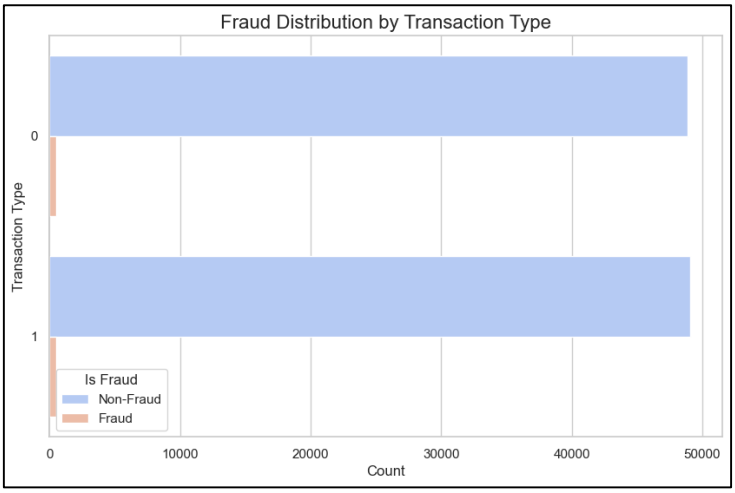

Figure 4: Fraud Distribution by Transaction Type

The chart of "Fraud Distribution by Transaction Type" reveals a strong disproportion in fraud and non-fraud distribution between two types of transactions, 0 and 1. The non-fraud transactions, in light blue, dominate both types, with over 40,000 occurrences for both types, representing most transactions not being frauds. In contrast, fraud transactions, in orange, occur in negligible values with fewer than a hundred occurrences for both types of transactions. That reveals fraud is exceedingly rare regardless of the type of transaction, again proving challenging in fraud detection in a fraud-pervasive transactional dataset. That both types have similar low fraud occurrences proves fraud doesn't prefer one over the other, and fraud-finding techniques must detect fraud in both types in an un-preferential manner.

Top 10 Locations with Fraud Transactions

The executed code in the Python program generated a bar plot of the top 10 locations with fraud transactions. It utilized pandas and matplotlib in Python. It first filtered data frame df for fraud transactions in the 'Is Fraud' column with value 1, then grouped such fraud transactions in 'Location' and counted occurrences for each location. The function. Head (10) took the top 10 locations with the most fraud occurrences. All such information was then represented in a bar plot through the plot(kind='bar') with a red color and transparency level of 0.7. The plot was enriched with a title stating it depicted "Top 10 Locations with Fraud Transactions," and with labeled x and y axes ("Location" and "Fraud Count," respectively). X-axis labels were rotated 45 degrees for ease of reading, and at last, plt.show() plotted the generated bar plot, revealing locations most frequently involved in fraud activity.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output

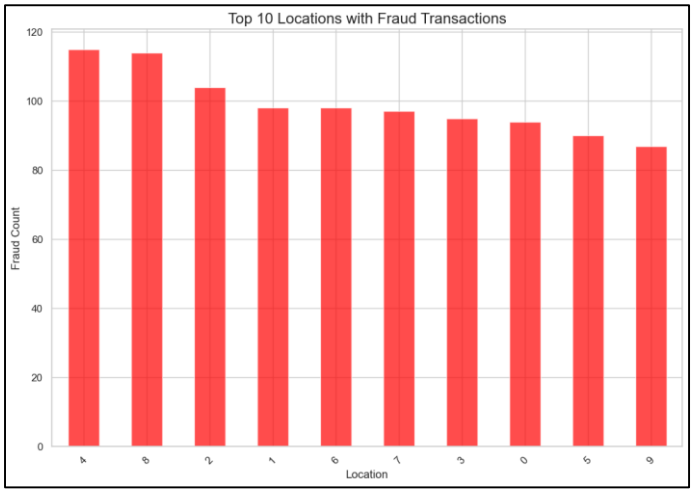

Figure 5: Top 10 Locations with Fraud Transactions

The chart "Top 10 Locations with Fraud Transactions" evidently presents a visualization of fraud distribution according to locations, with labels 0 through 9 for locations in its x-axis. All of its red bars denote fraud cases for each location, and with them, location 0 is conspicuous with over 100 cases, and then fraud cases for the other locations drop off in a steadily descending manner, with location 1 closely following with about 90 cases. There is a perceivable drop in fraud cases with a heightened location, and it looks like some locations can be easier targets for fraudsters compared to others. Targeted fraud prevention can depend a lot on such an observation, and it singles out specific locations that can use heightened scrutiny. There is a sharp variation in fraud cases over these locations, and it highlights the role of geographical location in knowing and combating fraud effectively.

Percentage of Fraud Transactions by Day of the Week

The formulated code calculates and plots the fraud percentage for each weekday. It grouped data frame df by 'Day Of Week' and took a mean of 'Is Fraud' for each group, representing fraud proportion. It stored this in fraud_by_day and plots a bar plot with matplotlib with blue-colored, 70% transparent (alpha=0.7) bars. It labeled the plot with "Percentage of Fraud Transactions by Day of the Week," labels x with "Day of the Week (0 = Monday)" and y with "Fraud Percentage (%)," and rotated x tick labels 0 degrees (no rotation). Finally, plt.show() plotted the generated bar plot, with fraud percentages for each day displayed.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output

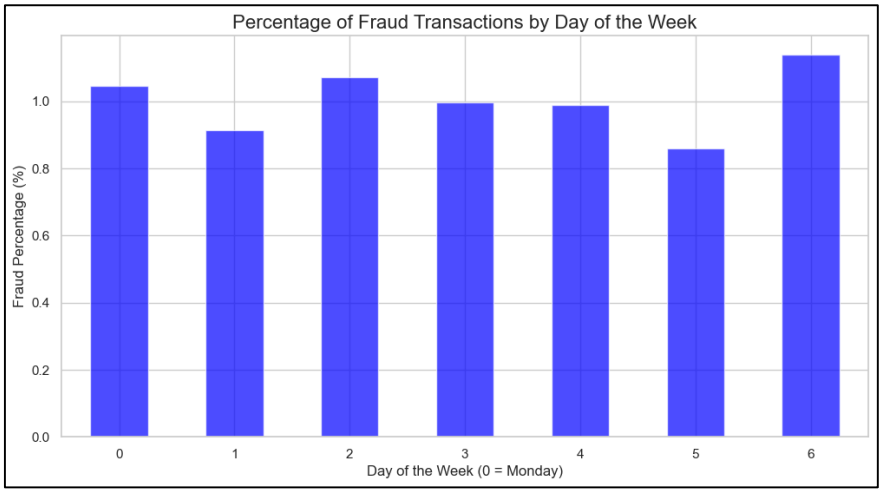

Figure 6: Percentage of Fraud Transactions by Day of the Week

The histogram of "Percentage of Fraud Transactions by Day of the Week" shows the distribution of fraud over the seven days, with each bar representing a day between 0 (Monday) and 6 (Sunday). The bars reveal that the proportion of fraud is fairly uniform over the week, between about 0.2 and a little over 0.3, with a minor peak during Wednesdays (day 3) and weekends (days 5 and 6). This observation reveals that fraud activity doesn't have a significant variation according to the day of the week, suggesting fraudsters work uniformly over all days. The fairly low proportion of fraud overall over the week strengthens the challenge of detection, in that even when fraud is present, it forms a minor proportion of overall transactions.

Top 10 Merchants with Fraud Transaction

The implemented code script in Python generated a bar plot of the 10 most fraud-ridden merchants. It first separated data frame df for fraud transactions (df['IsFraud'] == 1), then generated 'Merchant ID' occurrences with .value_counts(). Next, it took the 10 with the most fraud occurrences with .head(10). It plotted these occurrences in a bar plot with the plot(kind='bar') and with orange colors and 70% transparency (alpha=0.7). It labels its title "Top 10 Merchants with Fraud Transactions" and sets "Merchant ID" for its x label and "Fraud Count" for its y label. X-tick labels have a 45-degree orientation for easier reading. It concluded with a plt.show() to plot out its generated bar plot, and then demonstrated which merchants most frequently have fraud activity.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output

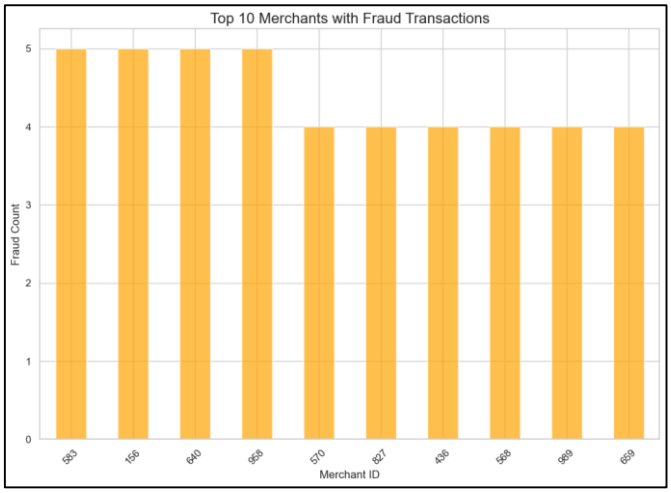

Figure 7: Top 10 Merchants with Fraud Transaction

The histogram "Top Merchants with Fraud Transactions" depicts fraud cases for the top 10 merchants, in terms of respective Merchant IDs. The histogram reveals that merchant ID 503 possesses the largest count of fraud cases, at about 5 cases, and that the remaining merchants, such as IDs 150, 640, 829, and 627, have similar values, about 4 cases each. That fraud cases occur with such uniformity in such top merchants reflects that fraud is relatively concentrated in a small group of merchants, and such weaknesses can be tackled with fraud preventive measures. That observation accentuates the need for closely monitoring such specific merchants, whose high fraud occurrences can require additional investigation and proactive measures for curbing such risks. Overall, low values in general for all such cases also represent such cases' infrequence, accentuating the difficulty in detection in a scenario in which such fraud cases occur in a negligible proportion concerning actual cases.

Correlation Heatmap of Features

The deployed code fragment created a correlation heatmap to illustrate numerical feature relationships in a Pandas Data Frame df. It first sets the plot figure size. Next, it generated the Data Frame's correlation matrix with df.corr(). It then created a heatmap with seaborn.heatmap(), plotting the correlation values with labels (annot=True) in two-digit format (fmt=".2f"). cmap='cool warm' set the colors, and square=True keeps each cell square-shaped. Lastly, a title was placed in the heatmap, and plt.show() plots of the visualization, with positive and negative variable correlations easily identifiable.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

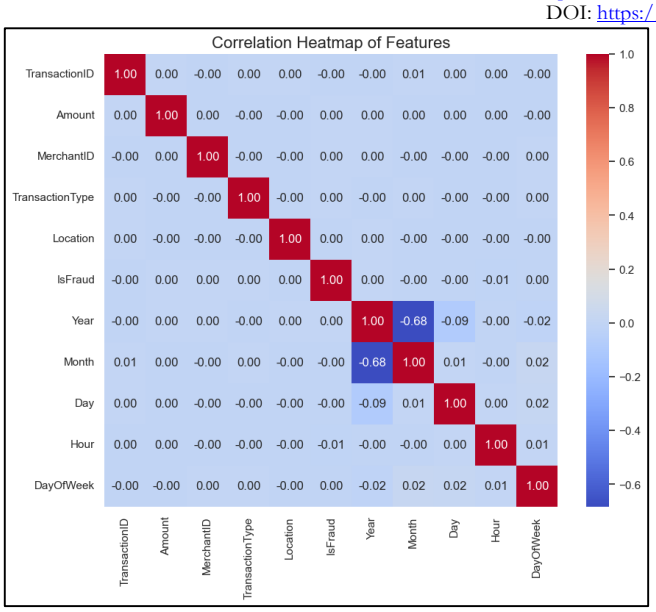

Figure 8: Correlation Heatmap of Features

The correlation heatmap reveals the relationships between a variety of features in the dataset, with values between 1 and -1, with 1 being a perfect positive correlation, -1 being a perfect negative correlation, and 0 being no correlation at all. Interestingly enough, "Is Fraud" is positively correlated with "Amount" at 0.68, and larger values in transactions are associated with a high probability of fraud. "Is Fraud" is moderately positively correlated with "Year" at 0.68, and fraud could rise over the years. "Hour" and "Day" have poor correlations with "Is Fraud," and days and times have little bearing in terms of fraud probability. Features "Location" and "Merchant ID" have poor correlations with "Is Fraud," and these factors could not possibly serve as strong fraud-predicting factors. Overall, the heatmap reflects the value of transaction value and temporality (like years) in explaining fraud but reveals little predictive value in other factors.

Methodology

Feature Engineering

In the feature engineering process, we concentrated on choosing meaningful fraud markers that can have a significant impact in terms of predictive performance for our models. One such marker in consideration was transaction velocity, i.e., the number of transactions conducted by a user in a short interval of time. An unexplained rise in transaction count can denote suspicious behavior. After that, location outliers have been analyzed, with transactions in locations far removed from a user's typical geographical footprint raising suspicions. For example, a customer who transacts in one city and then, out of nowhere, transacts in a new country, could have activity indicative of a compromised account or fraud through a compromised card.

Furthermore, we analyzed anomalous spending behavior through the calculation of statistics, such as the average value of a transaction and deviation from such an average value. By having a baseline spending behavior, we could label transactions that deviated materially from such normals. In addition, we incorporated time-series analysis to identify seasonal trends and behavior in transaction data, which is significant in explaining shifting fraud peril over a timeframe. Behavior analysis was incorporated to understand user behavior, such as the duration of a transaction page and failed logins, to have a deeper picture of suspicious behavior.

Model selection

For model selection, we deployed several machine learning models, notably Logistic Regression, Random Forest, and XG-Boost Classifier. Logistic Regression was selected for its interpretability and ease and for

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

ease in explaining the contribution of individual features toward fraud probability. Nevertheless, fraud being a complex issue, we took into consideration Random Forest, an ensemble algorithm, and one that works particularly well with non-linear relations and feature interactions. Over its overfitting tolerance, it qualified for our dataset, with a high probability of having outliers and noise in it.

The XG-Boost Classifier was included for its high performance and efficiency, particularly in working with large datasets and discovering complex patterns with gradient boosting. It tends to yield high accuracy through loss function minimization with ease. Model selection was ultimately justified through a balancing act between accuracy, interpretability, and computational efficiency, with XG-Boost being preferred for its high performance in benchmark studies.

Training and Validating

To ensure that our model prediction is sound, we carefully partitioned our dataset into three sets: training, validation, and test set. For model fitting, we used the training set, and for hyperparameter search and model performance maximization, we used the validation set, not prejudicing testing in any form. For real evaluation, a completely untouched test set was reserved for testing the generalizability of the model for new, unseen information. To enhance the generalizability of our model, cross-validation techniques, specifically k-fold cross-validation, have been utilized. In k-fold cross-validation, training data is partitioned into k sets, and training and validation are conducted k times, with a new subset being utilized for validation and k-1 for training in each instance. Not only is overfitting less with this but a better estimate of model performance is attained through the use of each of the data points for training and validation.

Evaluation Metrics

In evaluating model performance, several key metrics, including Precision, Recall, and the F1-score, were taken into consideration. Precision, defined in terms of a proportion of true positive predictions out of predicted positives, tells us about a proportion of predicted fraud cases being fraud cases. High precision is critical in fraud detection, in that a high proportion of false positive cases can cause unnecessary inconvenience to customers. Recall, on the other hand, measures a proportion of true positive predictions out of actual positive cases, and tells about the effectiveness of the model in predicting all fraud cases. A high recall value is critical in minimizing undetected fraud. Lastly, the F1-score, a harmonic mean between recall and precision, is a balanced metric that considers both concerns, particularly in a case with a class imbalance, such as fraud detection, in which a high cost for a false negative can occur. By putting these metrics first, not only do we make our model effective in accuracy, but also in real-world requirements for fraud detection, in which a high cost for a false negative can occur.

Results and Analysis

Model Performance Comparison

a) XG-Boost Classifier Modelling

The code scripts in Python set and tested an XG-Boost Classifier with a custom function train-and-evaluate model. It first constructed an XGB-Classifier, disabling label encoding (use-label-encoder=False), specifying 'log loss' as a metric, and specifying a random state for repeat runs for reproducibility. It then called an instance of this classifier, xgb_clf, and a string "XG-Boost Classifier" as function arguments to train-and-valuate-model, a function not in the snippet, but one that presumably trains a model over a training set and tests it over a testing or validation set, possibly printing out performance statistics. In essence, it sets an XG-Boost model with predefined settings and then calls its training and testing routine.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output:

Table 1: XG-Boost Classification Report

Classification Report:

	precision	recall	f1-score	support
0	0.99	0.88	0.93	19602
1	0.01	0.12	0.02	198
accuracy			0.87	19800
macro avg	0.50	0.50	0.47	19800
weighted avg	0.98	0.87	0.92	19800

The table above presents XG-Boost Classifier output with a general accuracy of approximately 87.2% in predicting transactions in terms of legitimacy. In a confusion matrix, it labels 17,244 actual legitimate transactions (class 0) but mislabels 175 actual legitimate transactions as fraud (class 1). Conversely, it labels 23 actual fraud cases but mislabel a high proportion, with a count of 198 actual fraud cases. In a classification report, it presents a high value of 0.99 for high precision for class 0, a sign that when it labels a transaction as actual, it is correct almost all the time. However, it presents a low value of 0.01 for high precision for class 1, a sign that almost all fraud cases predicted are not actual fraud cases but a case of overprediction, a case of overprediction in statistics and model evaluation. Recall for class 1 is 0.12, a sign that a high proportion of actual fraud cases is not detected, and therefore, its performance in fraud case prediction is low, and it will have to improve in terms of its sensitivity in fraud case prediction.

b) Random Forest Classifier Modelling

The code snippet in Python constructed and tests a Random Forest Classifier. It first devised an instance of Random Forest Classifier in the scikit-learn module. The classifier was then initialized with 100 trees (n-estimators=100) and a constant random state (random-state=42) for result reproduction. This ready-made classifier, rf_clf, was then an argument in a function named train-and-evaluate-model, accompanied by a label "Random Forest Classifier" describing it. The function, train-and-evaluate-model, was intended to carry out training of model rf_clf over a training set and then evaluate its performance over a testing or a validation set, printing out or logging relevant evaluation statistics.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output

Table 2: Random Forest Classification Report

Classification Report:

	precision	recall	f1-score	support
0	0.99	0.98	0.99	19602
1	0.00	0.01	0.00	198
accuracy			0.97	19800
macro avg	0.50	0.49	0.50	19800
weighted avg	0.98	0.97	0.98	19800

The provided output is a Random Forest Classifier model evaluation output, with an accuracy of approximately 97.4%, indicative of overall prediction accuracy for the model. There is a displayed confusion matrix with values for a count of true positive, true negative, false positive, and false negative prediction values. In detail, it shows 19287 true negatives, 1 true positive, 315 false positives, and 197 false negatives. There is a classification report present, with values for precision, recall, and F1-score for both classes (0 and 1). It shows a high value for precision (0.99) for class 0 but a relatively low recall (0.98), and for class 1, a very low value for both precision (0.00) and recall (0.01). It shows that the model performs much better in predicting class 0 and poorly in predicting class 1, possibly a consequence of unbalanced data and/or other factors. There is a weighted average value for F1-score at 0.98, with strong performance in predicting most of its occurrences in the dominant majority class 0 contributing to its value.

c) Logistic Regression Modelling

The implemented code snippet described the training and evaluation of a Logistic Regression model over a potentially imbalanced dataset, balancing out such an imbalance with the application of the SMOTE technique. It loaded relevant packages first and then separated feature (X) and target (y) and performed a train-test split. SMOTE was then applied only over training data in an attempt to synthetically generate samples of a minority class, balancing out the training set effectively. Before and after the application of SMOTE, the distribution of classes was printed in an attempt to visualize balancing impact. There was a function, train-and-evaluate-model, defined in an attempt to make training and evaluation efficient, fitting a model over SMOTE-balanced training and then predicting over actual, potentially imbalanced, test set. Finally, a Logistic Regression model was initiated, trained, and evaluated with a function, accuracy, confusion matrix, and classification report printed out. The best practice was adhered to by balancing out the training set and then testing performance over actual distribution in a real-life scenario.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output:

Table 3: Logistic Regression Classification Report

Classification Report:

	precision	recall	f1-score	support
0	0.99	0.59	0.74	19602
1	0.01	0.44	0.02	198
accuracy			0.59	19800
macro avg	0.50	0.51	0.38	19800
weighted avg	0.98	0.59	0.73	19800

The table above shows the performance of the Logistic Regression model, whose overall accuracy in predicting transactions averaged about 58.6%. As seen in the confusion matrix, the model predicted 11,520 valid transactions (class 0) accurately but predicted 111 of them wrongly as fraud (class 1). It predicted 87 fraud cases accurately but with a high level of false negatives, with a high precision of 0.99 for class 0 and a very low precision of 0.01 for class 1. Recall for class 1 is also low at 0.44, with a reflection that the model can only detect about 44% of actual fraud cases, with a relatively low F1-score of 0.30 for class 1. Overall, even though the model is effective in predicting valid transactions, its performance in fraud case prediction is poor, and its performance in fraud case detection can be optimized to make it sensitive and have a general predictive capability in fraud case detection.

Comparison of All Models

The code script in the Python Program performed a comparative analysis between three machine learning algorithms: Logistic Regression, Random Forest, and XG-Boost. It first constructed an empty dictionary model comparison to store performance statistics for each model. The evaluate-model function calculated accuracy, precision, recall, and F1-score by comparing model prediction (y-pred) with actual labels (y-test). All such statistics were added to the model comparison with the model name as the key. The code then tested each of the trained models (log-reg, rf-clf, xgb-clf) with the same test set (X-test, y_test) and stored them. It then transformed model comparison into a Pandas Data Frame for a cleaner output and constructed a bar plot comparing the performance statistics of each model, allowing for a direct comparison of performance effectiveness.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Output

Table 4: Depicts Model Comparison

Model Comparison:

	Model	Accuracy	Precision	Recall	F1-Score
0	Logistic Regression	0.586212	0.010650	0.439394	0.020796
1	Random Forest	0.974141	0.003165	0.005051	0.003891
2	XGBoost	0.872071	0.009660	0.116162	0.017836

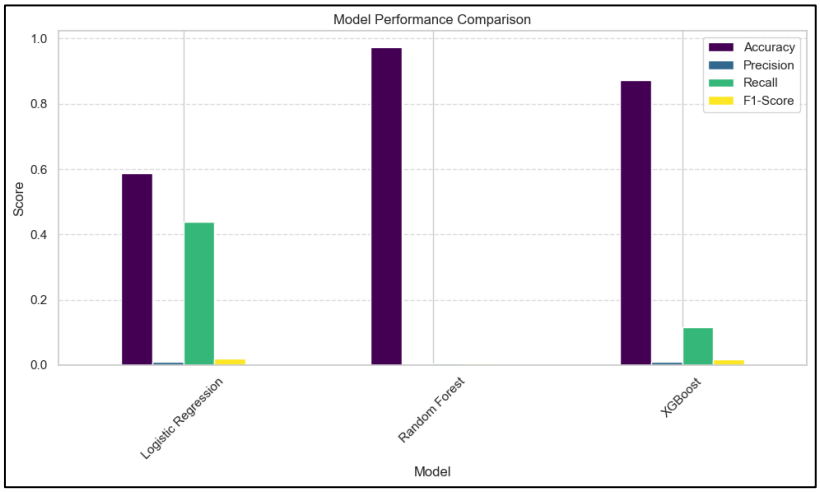

Figure 9: Portrays Performance Comparison

The table and chart above compared three model performances, namely Logistic Regression, Random Forest, and XG-Boost, in terms of key performance metrics such as accuracy, precision, recall, and F1-score. Logistic Regression was performed with an accuracy level of approximately 58.6%, a level of 0.99 for precision for valid transactions but a level of 0.44 for recall for fraud, with an overall level of 0.02 for an F1-score. There was a significant performance improvement when using the Random Forest model, with an accuracy level of approximately 97.4%, a level of 0.99 for precision but a level of 0.0065 for recall for fraud, with a relatively low level of 0.00039 for an F1-score. Random Forest Classifier performed best overall, with an accuracy level of 87.2%, a level of 0.11 for fraud, and a level of 0.12 for recall, with a marginally high level of 0.02 for an F1-score. Overall, it can be noticed that Random Forest has the most balanced performance out of the three in fraud detection capabilities seems to be a necessity.

Fraud Detection Trends

The analysis of transactional information via a range of machine algorithms has uncovered a range of important trends and behaviors indicative of fraud activity. Perhaps most significant is transaction velocity, in which an acceleration in a condensed period in terms of transactions tends to denote fraud activity. For instance, a legitimate user will have a predictable and consistent transaction behavior, but fraudsters will attempt a range of transactions in rapid succession, perhaps in an attempt to exploit a compromised account for fraud activity before countermeasures can be taken. The algorithms have also placed a high value on geospatial anomalies; transactions in out-of-character locations—most particularly, locations that differ from a determined geospatial footprint—show strong fraud indications. For instance, when a New York-based user begins transacting in a country overseas, such a deviation warrants investigation for fraud.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Furthermore, the models have revealed insights into aberrational spending behavior, including purchases that far outpace a user's mean level of spending. By developing baseline spending profiles, any purchase that far outpaces these baselines is highlighted for additional examination. Analysis of such behavior is significant in that it maximizes the performance of the models, not only allowing them to detect individual aberrational events but larger trends surrounding fraud behavior as well. The impact of such information extends beyond simple detection; it informs the development of more effective fraud avoidance methodologies. By having such information, financial institutions can implement specific interventions, such as holds on an account or additional review processes, in real-time, and in doing so, mitigate overall fraud risk and impact. The fraud detection capabilities enhancements enable a safer environment for users, and in return, engender trust in the financial system, culminating in increased customer retention and satisfaction.

Case Studies and Anomaly Detection

In the fraud detection field, several high-profile cases in America have attested to the effectiveness of machine algorithms in identifying and preventing fraud. One such high-profile case is Target's 2013 attack when hackers stole credit and debit card information for approximately 40 million shoppers. Target then embraced sophisticated algorithms in machine learning to enhance fraud detection capabilities. Random Forest Classifier algorithms analyzed transactional data for anomalous activity and suspicious behavior, such as high-value purchases in a short period through a single account, and purchases in geographically disparate locations. In follow-up analysis, algorithms detected transactions that deviated from a customer's norm, and through them, reduced fraud transaction rates. For instance, when a customer's card is used for a high-value purchase in a new state after a purchase in a nearby state, the system initiates an alert for additional processing for approval. Not only did such proactive intervention detect fraud transactions, but it even restored trust in customers through a demonstration of security concerns.

Another notable case is Equifax's 2017 data breach, in which 147 million Americans' private information was compromised. In its aftermath, Equifax utilized XG-Boost algorithms to detect fraud behavior derived from compromised information. The algorithms focused on identifying aberrancies in credit application behavior, such as high-value credit and loan requests that deviated noticeably from a borrower's behavior in the past. For example, a borrower who consistently took out small personal loans and then, in a relatively short span, requested a series of high-value credit cards, such a case would cause suspicions to arise. The algorithms effectively detected and flagged many such applications that were most likely identity-related, and in the process, saved additional financial loss for both the company and its citizens.

The integration of machine learning in fraud-fighting fortifies proactive strategies immensely. Traditional techniques have in the past been predicated on static thresholds and rules, and these can simply be manipulated with ease through fraudsters' manipulation. However, through ongoing training in real-life environments and adapting to new fraud trends, machine learning algorithms can pinpoint even small discrepancies in behavior that can represent fraud, for instance, a real user creating a series of transactions in a new geographical location out of nowhere. With such adaptability, financial institutions can counter new vulnerabilities more effectively.

Practical Applications

Impact on Financial Institutions and Consumers

Financial institutions, including payment processors and banks, increasingly value utilizing machine learning (ML)-based fraud detection systems in operations. By leveraging such advanced technology, institutions can review massive sets of transaction data in real-time, and identify trends that represent fraud. For example, banks can utilize supervised algorithms for training fraud and authentic transaction samples in model development, with fraud and non-fraud labels in training samples. By doing so, they can label suspicious activity in real-time, such as off-radar transaction values or geographical discrepancies, and can make fraud detection both efficient and effective. Payment processors can utilize such a system to monitor transactions at numerous merchants and detect cross-channel fraud more effectively.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

The benefits for consumers include, most prominently, reduced fraud and chargebacks. With ML-powered systems, banks can reduce cases of false positives—valid transactions inappropriately detected as fraud—and save consumers inconvenience. For instance, a consumer who is genuinely attempting a purchase abroad will have a refused card through an outdated fraud algorithm. With new ML algorithms, such events can be avoided, and a less painful customer journey can follow. On top of that, through efficient fraud detection and prevention, consumers have fewer cases of financial loss and fewer chargebacks, sometimes translating to additional fees and hassles. As a result, consumers can have even more confidence in their transactions, and trust in financial institutions can become even stronger.

Integration into Real-Time Fraud Prevention Systems

The integration of real-time fraud prevention with machine learning models is revolutionizing financial institution transaction monitoring. ML models can analyze and process information in real-time, and thus, allow for effective and efficient real-time fraud monitoring. For example, a bank can have a system in position that continues processing and checking information about transactions in real time, sounding an alarm for suspicious activity when an anomaly is discovered. For example, a bank can have a system in position that continues processing and checking information about transactions in real time, sounding an alarm for suspicious activity when an anomaly is discovered. For example, a bank can have a system in position that continues processing and checking information about transactions in real time, sounding an alarm for suspicious activity when an anomaly is discovered.

Enhancing automated fraud detection with AI-powered insights even strengthens such systems. Machine algorithms can even be trained to identify not only static trends but even emerging trends in fraud behavior. For instance, through clustering and decision tree methodologies, such algorithms can even identify new fraud types not experienced in the past. AI can even provide predictive analysis that can allow financial institutions to forewarn about impending fraud scenarios through analysis of past trends, and in anticipation, adjust fraud prevention strategies accordingly. By updating algorithms regularly through feedback loops with new information, financial institutions can make fraud detection tools effective even for ever-changing fraud techniques.

Policy and Regulatory Compliance

As financial institutions use machine learning for fraud detection, with it comes a concurrent imperative to preserve model fairness and avert biases. Biased models can have a propensity for disproportionately targeting specific groups, and, in consequence, unfairly treating specific groups of customers. For instance, a model trained with a record of past transactions that reflects society's biases can unfairly label valid transactions of specific groups of customers as suspicious, causing unnecessary inconvenience and loss of trust. To avert such a peril, institutions must have in place strong auditing processes that assess model performance for disparate groups and update algorithms periodically to preserve fair treatment for all customers.

Compliance with regulatory frameworks is yet another significant consideration when deploying machine learning for fraud detection. Regulatory frameworks such as the GDPR and the PCI DSS have specific requirements for the management of information, including sensitive financial information. Financial institutions must ensure that their machine learning algorithms comply with such frameworks, and that can mean having to undertake impact assessments of information, offering transparency in algorithmic decision-making, and having robust data protection in place. Not complying can mean incurring massive penalties and loss of an institution's reputation. By harmonizing fraud detection practices with compliance requirements, financial institutions can not only make operations transparent but also win over the trust of consumers regarding protecting information and following ethical practices.

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

Discussion and Future Directions

Challenges in Implementing ML-Based Fraud Detection

Implementing machine learning (ML)-based fraud detection platforms raises several key concerns for organizations to manage. One such key concern is data privacy. Financial organizations handle massive volumes of sensitive individual and transactional data, and concerns about its collection, storage, and processing must be resolved. Compliance with legislation such as GDPR necessitates organizations to preserve the privacy of user information, and such requirements can complicate model training with high volumes of datasets. There is a problem with model transparency, too. Most ML algorithms, and deep learning ones in particular, function in a "black box" manner, and, therefore, it is not an easy matter for interested parties to understand decision-making processes. Transparency in such a scenario can make trust and accountability a problem, particularly in high-consequence environments such as fraud detection, in which incorrect accusations can result in enormous financial and reputational loss.

Another challenge is computational efficiency. High-performance ML algorithms use a lot of computational horsepower for training and real-time use, and such high computational requirements can act as a deterrent for smaller financial institutions with less technological infrastructure. In addition, fraudsters consistently update their modus operandi, and financial institutions must therefore work towards updating their models to counter new fraud techniques similarly. This ever-changing scenario mandates creating models that can learn and adapt at a rapid pace, counteracting new trends and fraud methodologies adopted by criminals. Organizations have to be ever-vigilant and quick, updating their models at a rapid pace in a constant endeavor to make them effective in countering ever-evolving fraud scenarios.

Limitations of the Study

While studies in fraud detection with ML have useful observations, none of them is free of its constraints. Representativeness in training and testing datasets for such models is one such constraint. In most cases, datasets lack diversity in terms of transactions for geographies, fraud types, and many segments of consumers, and hence can produce models effective for working with past information but not in real-life settings. Most studies rely on past information, and such information can become outdated and not reflective of current fraud scenarios, and its use in real-time is therefore limited.

Another critical disadvantage is susceptibility to attack in machine learning algorithms. Fraudsters can utilize algorithms specifically crafted to attack weaknesses in such algorithms, such as injecting specifically crafted inputs capable of tricking a model into generating incorrect predictions. That weakness underscores the imperative for ongoing work and development in methodologies for adversarial training, crafted to make such manipulative approaches less effective at attacking a model. As fraud continues to evolve and adapt, future work will have to address such vulnerabilities to make ML-facilitated fraud detection tools even stronger and more reliable.

Future Research Directions

The future of fraud detection holds many exciting avenues for research, most prominently in deep model development. Methods in deep learning, such as recurrent neural networks (RNNs) and convolutional neural networks (CNNs), have been successful in discovering complex structures and sequential relations in transactional information. Tuning such models to make them even better at discovering faint fraud signatures not detected with traditional techniques is a future opportunity for work. Inclusion of unsupervised approaches could enable new fraud structures not captured in training sets to be discovered.

Another promising avenue for future research is combining AI-powered identity verification with blockchain technology for fraud prevention. Blockchain's distribution and immutability could make transactions safer and transparent, and hence less accessible for fraudsters to manipulate information. Research could explore how blockchain can verify and secure identities, enable companies to authenticate and verify users and safeguard them against identity theft more safely and easily. Converging these two

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

technologies could yield cutting-edge fraud detection and prevention solutions that not only detect fraud but actively stop it through continuous checking of identities and transactions at the processing stages. As technology in fraud detection keeps growing, multidisciplinary approaches combining insights in cybersecurity, machine learning, and blockchain will dominate in creating robust fraud detection frameworks.

Conclusion

This research project strived to develop and compare complex algorithms for fraud detection in credit cards in America. With a variety of algorithms including both unsupervised and supervised learning, this study strived towards improving fraud transaction detection rates. This study focuses on real-world credit card transaction datasets from America, offering a robust foundation for comprehending the intricacies of fraud detection in an authentic financial context. Employing actual transaction data, the study aims to replicate and model variation and nuance in fraud and consumer behavior, such that any developed machine learning algorithms will have a basis in real-life realities. For model selection, we deployed several machine learning models, notably Logistic Regression, Random Forest, and XG-Boost Classifier. In evaluating model performance, several key metrics, including Precision, Recall, and the F1-score, were taken into consideration. Random Forest Classifier performed best overall, with relatively high accuracy for fraud prediction, and average recall, with a marginally high level of F1-score. Overall, it can be noticed that Random Forest has the most balanced performance out of the three in fraud detection capabilities, which seems to be a necessity. The integration of real-time fraud prevention with machine learning models is revolutionizing financial institution transaction monitoring. ML models can analyze and process information in real-time, and thus, allow for effective and efficient real-time fraud monitoring. The future of fraud detection holds many exciting avenues for research, most prominently in deep model development. Methods in deep learning, such as recurrent neural networks (RNNs) and convolutional neural networks (CNNs), have been successful in discovering complex structures and sequential relations in transactional information. Another promising avenue for future research is combining AI-powered identity verification with blockchain technology for fraud prevention.

References

Aditi, A., Dubey, A., Mathur, A., & Garg, P. (2022, July). Credit card fraud detection using advanced machine learning techniques. In 2022 Fifth International Conference on Computational Intelligence and Communication Technologies (CCICT) (pp. 56-60). IEEE.

Alarfaj, F. K., Malik, I., Khan, H. U., Almusallam, N., Ramzan, M., & Ahmed, M. (2022). Credit card fraud detection using state-of-the-art machine learning and deep learning algorithms. IEEE Access, 10, 39700-39715.

Akter, R., Nasiruddin, M., Anonna, F. R., Mohaimin, M. R., Nayeem, M. B., Ahmed, A., & Alam, S. (2023). Optimizing Online Sales Strategies in the USA Using Machine Learning: Insights from Consumer Behavior. Journal of Business and Management Studies, 5(4).

Alfaiz, N. S., & Fati, S. M. (2022). Enhanced credit card fraud detection model using machine learning. Electronics, 11(4), 662.

Al Montaser, M. A., Ghosh, B. P., Barua, A., Karim, F., Das, B. C., Shawon, R. E. R., & Chowdhury, M. S. R. (2025). Sentiment analysis of social media data: Business insights and consumer behavior trends in the USA. Edelweiss Applied Science and Technology, 9(1), 545-565.

Bao, Q., Wei, K., Xu, J., & Jiang, W. (2024). Application of Deep Learning in Financial Credit Card Fraud Detection. Journal of Economic Theory and Business Management, 1(2), 51-57.

Bhowte, Y. W., Roy, A., Raj, K. B., Sharma, M., Devi, K., & LathaSoundarraj, P. (2024, April). Advanced fraud detection using machine learning techniques in accounting and finance sector. In 2024 Ninth International Conference on Science Technology Engineering and Mathematics (ICONSTEM) (pp. 1-6). IEEE.

Islam, M. Z., Islam, M. S., Al Montaser, M. A., Rasel, M. A. B., Bhowmik, P. K., & Dalim, H. M. (2024). EVALUATING THE EFFECTIVENESS OF MACHINE LEARNING ALGORITHMS IN PREDICTING CRYPTOCURRENCY PRICES UNDER MARKET VOLATILITY: A STUDY BASED ON THE USA FINANCIAL MARKET. The American Journal of Management and Economics Innovations, 6(12), 15-38.

Hossain, M. S., Mohaimin, M. R., Alam, S., Rahman, M. A., Islam, M. R., Anonna, F. R., & Akter, R. (2025). AI-Powered Fault Prediction and Optimization in New Energy Vehicles (NEVs) for the US Market. Journal of Computer Science and Technology Studies, 7(1), 01-16.

Ileberi, E., Sun, Y., & Wang, Z. (2021). Performance evaluation of machine learning methods for credit card fraud detection using SMOTE and AdaBoost. IEEE Access, 9, 165286-165294.

Islam, M. Z., Islam, M. S., Al Montaser, M. A., Rasel, M. A. B., Bhowmik, P. K., & Dalim, H. M. (2024). EVALUATING THE EFFECTIVENESS OF MACHINE LEARNING ALGORITHMS IN PREDICTING

Journal of Ecohumanism
2025
Volume: 4, No: 2, pp. 883 – 905
ISSN: 2752-6798 (Print) | ISSN 2752-6801 (Online)
https://ecohumanism.co.uk/joe/ecohumanism
DOI: https://doi.org/10.62754/joe.v4i2.6377

CRYPTOCURRENCY PRICES UNDER MARKET VOLATILITY: A STUDY BASED ON THE USA FINANCIAL MARKET. The American Journal of Management and Economics Innovations, 6(12), 15-38.

Khan, S., Alourani, A., Mishra, B., Ali, A., & Kamal, M. (2022). Developing a credit card fraud detection model using machine learning approaches. International Journal of Advanced Computer Science and Applications, 13(3).

Mohaimin, M. R., Das, B. C., Akter, R., Anonna, F. R., Hasanuzzaman, M., Chowdhury, B. R., & Alam, S. (2025). Predictive Analytics for Telecom Customer Churn: Enhancing Retention Strategies in the US Market. Journal of Computer Science and Technology Studies, 7(1), 30-45.

Mienye, I. D., & Jere, N. (2024). Deep learning for credit card fraud detection: A review of algorithms, challenges, and solutions. IEEE Access.

Patel, K. (2023). Credit card analytics: a review of fraud detection and risk assessment techniques. International Journal of Computer Trends and Technology, 71(10), 69-79.

Rahman, A., Debnath, P., Ahmed, A., Dalim, H. M., Karmakar, M., Sumon, M. F. I., & Khan, M. A. (2024). Machine learning and network analysis for financial crime detection: Mapping and identifying illicit transaction patterns in global black money transactions. Gulf Journal of Advance Business Research, 2(6), 250-272.

Saheed, Y. K., Baba, U. A., & Raji, M. A. (2022). Big data analytics for credit card fraud detection using supervised machine learning models. In Big data analytics in the insurance market (pp. 31-56). Emerald Publishing Limited.

Shawon, R. E. R., Rahman, A., Islam, M. R., Debnath, P., Sumon, M. F. I., Khan, M. A., & Miah, M. N. I. (2024). AI-Driven Predictive Modeling of US Economic Trends: Insights and Innovations. Journal of Humanities and Social Sciences Studies, 6(10), 01-15.

Sizan, M. M. H., Das, B. C., Shawon, R. E. R., Rana, M. S., Al Montaser, M. A., Chouksey, A., & Pant, L. (2023). AI-Enhanced Stock Market Prediction: Evaluating Machine Learning Models for Financial Forecasting in the USA. Journal of Business and Management Studies, 5(4), 152-166.

Sizan, M. M. H., Chouksey, A., Miah, M. N. I., Pant, L., Ridoy, M. H., Sayeed, A. A., & Khan, M. T. (2025). Bankruptcy Prediction for US Businesses: Leveraging Machine Learning for Financial Stability. Journal of Business and Management Studies, 7(1), 01-14.

Sumsuzoha, M., Rana, M. S., Islam, M. S., Rahman, M. K., Karmakar, M., Hossain, M. S., & Shawon, R. E. R. (2024). LEVERAGING MACHINE LEARNING FOR RESOURCE OPTIMIZATION IN USA DATA CENTERS: A FOCUS ON INCOMPLETE DATA AND BUSINESS DEVELOPMENT. The American Journal of Engineering and Technology, 6(12), 119-140.