Journal of Information Systems Engineering and Management
2025,10(4)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

Leveraging Unsupervised Machine Learning and Deep Learning for Enhanced Card Fraud Detection

[1]Krishiv Garg, [2]Dr Umang Soni

[1]Bhupindra International Public School, Patiala

krishivgarg2008@gmail.com

[2]Assistant Professor, Department of Mechanical Engineering, Netaji Subhas University of Technology, New Delhi

Umang.soni@nsut.ac.in

ARTICLE INFO

Received: 11 Nov 2024

Revised: 29 Dec 2024

Accepted: 20 Jan 2025

ABSTRACT

The modern financial ecosystem is highly reliant on digital transactions, and despite the convenience, they have a significant flaw with high rates of fraudulent activities. Card fraud poses a severe risk to consumers and financial institutions, leading to economic losses, data breaches, and trust erosion. Traditional fraud detection methods mainly use supervised learning models or heuristics that fail to keep up with evolving fraud patterns and often have too many false positives and negatives. This research study aims to develop a standard detection framework to identify fraudulent transactions without relying on labelled fraud data. This is done by leveraging unsupervised learning techniques, namely isolation forest, One-Class SVM and deep learning-based Autoencoder using the OPENML ID: 45955 dataset referred to as 'dataset.csv'. We show that unsupervised methods can be effective when fraudulent labels are sparse, while deep learning approaches offer notable improvements in balancing detection (recall) and correctness (precision). Empirical results reveal that Autoencoders yield a superior F1-score compared to Isolation Forest and OneClass SVM, with an Accuracy of up to 92.51%, Precision of 56.01%, Recall of 66.48%, and ROCAUC of 0.8074.

Keywords: Fraud Detection, Machine Learning, Cybersecurity, Anomaly Detection, Financial Security

1. INTRODUCTION

Cybersecurity has been at the forefront of recent developments in financial infrastructure and is extremely important to safeguard the digital infrastructure employed by the financial institutions with whom we entrust our assets. One major issue they face is bad actors trying to commit fraud; one of the most common types of fraud is credit card fraud. In the cycle of using a credit card, there are multiple vulnerabilities that can be exploited both on the user end and the merchant end. Hence, it is essential for financial institutions to be able to leverage artificial intelligence to detect fraud and address it before it can cause major issues. For this research paper, we have defined fraud in accordance with Black's Law Dictionary- Fraud consists of some deceitful practice or wilful device resorted to with intent to deprive another of his right or in some manner to do him an injury[1]. With digital payments rising, digital fraud has become a significant issue for all stakeholders, namely, customers and financial institutions like banks. Fraud not only causes monetary damage to individuals but also strains merchant relationships with customers and key commercial partners, like card issuers and fulfilment vendors.

2. LITERATURE REVIEW

2.1 Existing Approaches

Credit card fraud detection methods have evolved from simple rule-based systems to complex Artificial Intelligence models. Key approaches include:

2.1.1 Rule-Based Systems: Early fraud detection methods relied on developer-defined rules and heuristics (e.g. blacklists of fraudulent card numbers, flags for transactions over certain amounts or in rapid succession). These systems are easy to implement and interpret, but are static and often overly simplistic. Fraudsters learn and adapt to the rules, rendering them ineffective over time [2]. For example, if a fixed spending limit triggers an alert, criminals may keep transactions below the threshold. While rule-based frameworks formed the backbone of many legacy fraud systems, they tend to produce high false alarm rates and miss novel fraud patterns[3]. Bolton and Hand (2002) observed that such static rules eventually only catch “incompetent” fraudsters who don’t adapt, highlighting the need for more flexible approaches [4].

2.1.2 Statistical Methods: Banks have been using these methods for a very long time to spot fraud. Statistical methods search data for deviations from the expected behaviour. Finance firms often use thresholding, clustering and peer group analysis and probability models, including Bayesian networks or Hidden Markov models, to predict the likelihood of a transaction turning out to be fraudulent [5]. Rather than operating under direction, they flag any irregular transaction. For example, models could mark a sudden high-value purchase in a foreign country as an outlier and label a fast series of high-value international transactions as fraud without any confirmation, which is counterintuitive if, for example, the user is travelling. Another difficulty they face is spotting fraud trends since criminals constantly modify their techniques, and the fact that they are using fraud detection datasets with often severe class imbalances.

2.1.3 Machine Learning: Compared to heuristics and rule-based methods, Machine learning has significantly improved fraud detection rates. Another advantage of using machine learning is that if it is regularly trained on new data, it continues to improve in almost all metrics. But it requires a large volume of labelled examples for training. This is a problem, especially for our use case, where it is difficult to accurately label huge amounts of data[6].

2.1.4 Deep Learning: Deep learning models achieve higher overall accuracy and recall than methods we’ve previously discussed because they can automatically learn intricate features from raw data. For example, a 2024 systematic review noted the effectiveness of deep models like LSTMs and transformers across credit card fraud datasets. An important advantage of deep learning is its ability to process massive datasets and discover subtle patterns which other machine learning models often miss; however, it typically acts as a “black box,” making explaining its decisions more difficult (a very important consideration in financial industries) [7].

2.2 Limitations of current models

Despite significant advances in model development, current fraud detection models face several challenges that limit their effectiveness and deployment. Some of the important ones are as follows:

2.2.1 High Rates of False Positives: A major problem in detecting card fraud is the high rates of false positives. Genuine transactions vastly outnumber fraudulent transactions, which causes models to generate false positives even with high accuracy. Studies report false favourable rates in credit card fraud screening as high as *30–70%*, meaning a large fraction of transactions flagged for investigation are innocent [8]. False alarms of this kind can erode consumer confidence, which might cause them to reduce usage or completely give up a card following a large number of declined transactions due to this error. Reducing false positives is challenging since fraud patterns mostly coincide with normal behaviour. To identify the few fraud cases, the system usually sweeps broadly and falsely flags some genuine activity. Even current advanced models still find great difficulty juggling low false positives with high fraud detection rates.

2.2.2 Imbalanced Datasets: Fraud detection datasets show that a small minority of all transactions are fraudulent, which is because they model real life, where only a small chunk of all transactions are false. Although a naive classifier would miss all actual fraud cases, ironically, it can achieve over 99% accuracy by assuming every transaction is not fraud, which ruins the user experience . Conversely, conventional machine learning algorithms can fail to sufficiently learn the traits of fraud since they are biased toward the majority class. As mentioned earlier, this also contributes to high false positive and negative rates [9]. Beyond the ones listed above, other difficulties include concept drift—that is, the constant evolution of fraud patterns requiring models to be flexible. Second, as the section on deep learning notes, a lack of interpretability means that banks frequently require justifications for why a transaction was flagged, which is also called a “black box”. Although crucial, these problems fall outside the main focus of this paper, but are nonetheless an active research subject.

2.3 Recent advances

To address the limitations of current models, recent research has introduced more sophisticated techniques and novel paradigms in card fraud detection. Some of the important developments include:

2.3.1 Ensemble and Hybrid Learning: Ensemble learning is a very popular strategy used in fraud detection. Ensembles combine the predictions of multiple models, allowing them to achieve better accuracy and robustness than a single model working alone. Recent work done by experts has explored hybrid ensembles that mix different algorithm types (e.g., combining a supervised classifier with an unsupervised anomaly detector) to capture complementary perspectives of fraud. For example, Wallny (2022) reports that an ensemble of diverse classifiers significantly reduced the overall cost of fraud—cutting fraud-related losses by about 30% compared to the best individual model [10].

2.3.2 Graph-Based Fraud Detection: To find relational fraud trends, one developing trend is graphically modelling credit card transactions and entities. Nodes in a graph show things like credit card accounts, stores, devices, or IP addresses; edges show interactions—that is, two cards sharing a device or a card used at a merchant. Often working in networks or rings, fraudsters may use many cards at one merchant or a stolen card used at a cluster of related merchant IDs. Graph-based techniques can reveal these connections of interest hidden from view in isolated transaction analysis. Graph neural networks (GNNs) have been recently used to this challenge. For example, Xiang et al. (2023) constructed a temporal transaction graph, where every transaction is represented as a node connected by edges if they share specific traits or occur in sequence. They then utilised a GNN with temporal attention to distribute risk among linked transactions. Even with only a small fraction of transactions labelled as fraud, this semi-supervised approach outperformed other state-of-the-art fraud detectors [11].

3. MOTIVATION

The Association of Certified Fraud Examiners reports that organisations lose approximately 5% of their annual revenue to fraud[12]. As the world rapidly digitises and more people use cards, card acceptance increases, increasing the chance of being defrauded. A card's payment processing cycle has multiple vulnerabilities; the perpetrator can exploit any of them, whether it be the customer or the payment processor[13]. Fraud detection is critical for financial institutions, e-commerce platforms, and other industries to safeguard assets, maintain trust, and comply with Anti-Money Laundering Laws and Know-Your-Customer regulations, which require companies to prevent fraud-related crimes[14].

4. AIMS OF THIS PAPER

1. Three anomaly detection approaches—Isolation Forest, One-Class Support Vector Machine, and Autoencoders—applied to a real-world-like dataset, OPENML ID: 45955, referred to as

   'dataset.csv'

2. Comparison of model performance using accuracy, precision, recall, F1-score, and ROC-AUC, discussing the strengths and limitations of each technique.

3. In-depth insights into how tuning hyperparameters and altering training dynamics can dramatically improve fraud detection rates, particularly in Autoencoder models.

5. RESEARCH OBJECTIVES

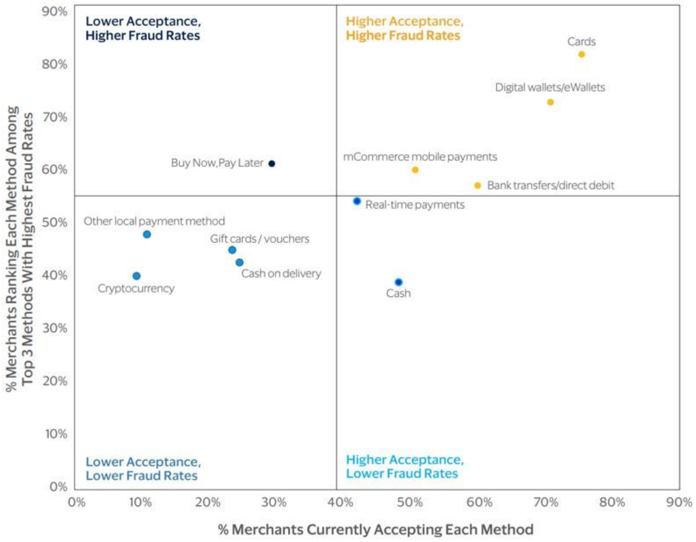

Figure 1:Payment method acceptance versus fraud rate(Global e-commerce payments and fraud report 2024) [15]

The Global Commerce Payments and Fraud Report strongly correlates with a payment method's acceptance and card fraud rates.[16] According to the report, cards, one of the most highly accepted payment methods, have the highest fraud rate; hence, financial institutions need to be able to detect fraudulent card payments accurately, ensure consumer protection, consumer trust, and protect themselves.

The primary objective of this research is to develop an effective and scalable fraud detection model that can accurately identify fraudulent transactions in real-time financial systems. Given the challenges posed by highly imbalanced datasets and evolving fraud patterns, this study aims to explore unsupervised and deep learning approaches to detect anomalies in credit card transactions without relying heavily on labelled fraud data. Specifically, the research seeks to: (1) evaluate the effectiveness of Isolation Forest, One-Class Support Vector Machine, and Autoencoders in detecting fraudulent activities, (2) compare their performance based on key metrics such as accuracy, precision, recall, and F1-score, and (3) propose a robust anomaly detection framework that financial institutions can deploy for real-time fraud monitoring. The ultimate goal of this research paper is to evaluate and compare algorithms and create models to reduce financial losses and improve customer transaction security while maintaining high accuracy and efficiency.

6. METHODOLGY

This section details the approach to developing fraud detection models, including dataset selection, feature engineering, model architecture, and evaluation metrics.

6.1 Dataset Description

A robust real-world or real-world-like dataset that includes both fraudulent and non-fraudulent transactions is the most essential part of creating a fraud detection model. This study used the real-world-like opensource dataset *'card_transdata.csv'* from OPENML, ID: 45955, by Iwo Godzwon [17]. This dataset captures transaction patterns and behaviours that could indicate potential fraud in card transactions. The data comprises several features designed to reflect the transactional context, such as geographical location, transaction medium, and spending behaviour relative to the user's history [18]. The dataset applied was used to create an Isolation Forest Model, a One-Class Support Vector Machine, and a Deep Learning Autoencoder. The dataset consists of one million simulated transactions.

Attribute Description:

1. **distance_from_home**: This is a numerical feature representing the geographical distance in kilometers between the transaction location and the cardholder's home address.
2. **distance_from_last_transaction**: This numerical attribute measures the distance in kilometres from the location of the last transaction to the current transaction location.
3. **ratio_to_median_purchase_price**: A numeric ratio that compares the transaction's price to the median purchase price of the user's transaction history.
4. **repeat_retailer**: A binary attribute where '1' signifies that the transaction was conducted at a retailer previously used by the cardholder, and '0' indicates a new retailer.
5. **used_chip**: This binary feature indicates whether the transaction was made using a chip (1) or not (0).
6. **used_pin_number**: Another binary feature, where '1' signifies using a PIN for the transaction, and '0' shows no PIN was used.
7. **online_order**: This binary attribute identifies whether the purchase was made online ('1') or offline ('0').
8. **fraud**: A binary target variable indicating whether the transaction was fraudulent ('1') or not ('0').

6.2 Data Preprocessing

Before training our models, we performed essential data preprocessing steps to ensure optimal model performance. The dataset (card_transdata.csv) was checked for missing values, but none were found. Regardless, the code to drop missing values was included. Next, the outliers were examined but not explicitly removed since they could represent fraudulent transactions. Since the dataset contains numerical attributes with varying scales, for feature scaling in every model, we applied MinMax Scaling to the following numerical features to normalise them between 0 and 1 to ensure equal weighting during model training:

"distance_from_home", "distance_from_last_transaction", "ratio_to_median_purchase_price", "repeat_retailer", "used_chip", "used_pin_number", "online_order".

6.3.1 Proposed Models and Reasoning

In this study, we intentionally chose two unsupervised learning algorithms, Isolation Forest and One-Class Support Vector Machine (SVM), and one deep learning algorithm, Autoencoder. While shortlisting the algorithms, the first criterion was that we did not choose any supervised learning models. This might seem unconventional, as traditional financial systems have used supervised learning, but this criterion is critical to the accuracy of a fraud detection model. It was established due to the nature of the problem we are trying to solve; in the cases of financial fraud, especially card fraud, there are two key things that we need to keep in mind before choosing our algorithm:

1) The dataset from the real world is usually unlabelled.

2) Fraudsters continually adapt their methods to evade detection.

3) Fraud detection data sets are highly imbalanced, meaning there are far fewer fraudulent transactions than normal ones.

Supervised learning has been used for many applications involving machine learning. But in card fraud detection, it is inappropriate because supervised learning needs a lot of labelled data. But the information we gathered for our study contains this. However, banks typically require hand labelling of the raw data, which is costly, time-consuming, and prone to errors—qualities that can significantly compromise the accuracy of our model. Second, con artists constantly change their strategies to evade detection; thus, historical data's reflections of fraud trends might not be present in the future. Although supervised learning models can learn from past data, their fundamental weakness is usually their inability to identify more recent anomalies [19]. The third and most important weakness is that highly imbalanced fraud detection data sets cause the supervised learning models to become biased towards non-fraud transactions in our use case, which may lead to high accuracy but low recall, causing inconvenience to consumers.

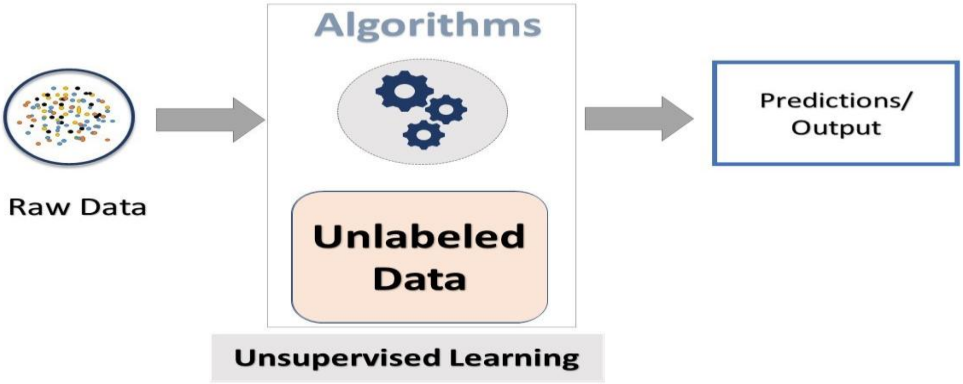

Figure 2: Working of Unsupervised Learning Algorithms [20]

Using unsupervised learning, one of our most troubling problems of having a well-labelled data set is solved because unsupervised learning algorithms don't require labelled data like supervised learning algorithms. Similarly, Autoencoders also don't require labelled data.

6.3.2 Reasoning for choosing Isolation Forest, One-Class Support Vector Machine, and Deep Learning-Based Autoencoders

Considering the inefficiency of supervised learning algorithms, we chose unsupervised and deep learning models. This addresses our primary concern of a lack of labelled data. Isolation Forest was our first algorithm, which we tested. It was chosen because it is a fast and scalable model that efficiently isolates outliers by recursively partitioning the data, making it a decent choice for detecting rare fraudulent transactions in large datasets, keeping in mind that you will have hundreds of millions of transactions for fraud detection. Second, we chose One-Class SVM, which learns a decision boundary around normal transactions and flags deviations as potential fraud. Although One-Class SVM can capture nonlinear patterns, it is computationally expensive and often performs poorly in high-dimensional datasets. We implemented Deep Learning-based Autoencoders to overcome these limitations, which learn a compressed representation of normal transactions and detect fraud based on reconstruction errors.

6.4 Implementation and Software Tools

We require data processing, machine learning, and deep learning libraries to implement the Fraud detection models' code.

The following tools and Libraries have been used:

- Python 3: The programming language used for data analysis and model development
- Pandas: Library for data manipulation and preprocessing
- NumPy: Library used for Numerical Computations
- Scikit-Learn: Library for machine learning models(i.e. Isolation Forest and One-Class SVM)
- TensorFlow: Used for Deep Learning based Encoder
- Matplotlib: Data Visualisation, for example, in a heatmap and a confusion matrix

6.5 Data Splitting Strategy

- To maintain data integrity and generalisation:
- 80% of the dataset was allocated for training.
- 20% of the dataset was used for testing.
- We used stratified sampling to preserve the original fraud-to-non-fraud ratio, ensuring a representative split.
- Training sets were further divided:
- Autoencoder and One-Class SVM were trained only on normal transactions to learn normal behaviour.
- The test set contained both normal and fraudulent transactions to evaluate model performance.

6.6 Autoencoder for Fraud Detection

Autoencoders are deep learning-based anomaly detection models that learn to reconstruct normal transaction patterns. Fraudulent transactions exhibit high reconstruction error, allowing them to be flagged as anomalies.

6.6.1 Model Architecture The Autoencoder consists of:

- Encoder:
- 16 neurons (ReLU activation) - Extracts meaningful transaction patterns.
- 8 neurons (ReLU activation) - Bottleneck representation, reducing dimensionality.
- Decoder: o 16 neurons (ReLU activation) - Attempts to reconstruct the input data. o 7 neurons (Sigmoid activation) - Produces final reconstruction values.

6.6.2 Model Training

- Training Objective: Minimise Mean Squared Error (MSE), ensuring that the model learns typical transaction structures.
- Optimiser: Adaptive Moment Estimation (Adam) to facilitate stable learning.
- Batch Size: 64 transactions per batch for efficient training.
- Epochs: 50 to allow convergence.
- Validation Set: Normal transactions from the test set were used to monitor performance and avoid overfitting.

6.6.3 Fraud Detection Using Reconstruction Error

- After training, reconstruction errors were computed for both normal and fraudulent transactions.
- A threshold of normal transaction errors was set at the 95th percentile.
- Transactions exceeding this threshold were classified as fraud.

6.6.4 Model Evaluation

- Computed precision, recall, F1-score, and ROC-AUC.
- A Confusion Matrix was plotted to assess fraud detection efficiency.
- A Histogram of Reconstruction Errors helped visualise anomaly separation.

6.7 One-Class SVM for Fraud Detection

One-class SVM is an unsupervised anomaly detection technique that models the normal distribution of transactions and flags outliers.

6.7.1 Model Training

- Kernel: Radial Basis Function (RBF) for non-linear separation.
- Gamma: Scale-adjusted for feature normalization.
- Nu (ν-parameter): 0.02 to regulate anomaly proportion.

6.7.2 Anomaly Detection Strategy

- One-Class SVM assigns each transaction a decision function score.
- Scores below a certain threshold indicate fraud.
- Threshold Determination: Percentile-based analysis on normal transaction scores.

6.7.3 Model Evaluation

- Classification report (precision, recall, F1-score, ROC-AUC).

- Confusion Matrix plotted to assess fraud detection effectiveness.
- Decision Score Histogram analysed for fraud score distribution.

6.8 Isolation Forest for Fraud Detection

Isolation Forest is an unsupervised model that isolates anomalies by randomly partitioning data features.

6.8.1 Model Training

- Number of Trees: 100 trees for better anomaly separation.
- Contamination Level: 2% (estimated fraud rate in the dataset).

6.8.2 Anomaly Detection

- Anomaly Score Computation: Transactions with shorter path lengths in the decision tree were flagged as fraud.
- Threshold Selection: Dynamically adjusted based on contamination parameter.

6.8.3 Model Evaluation

- Confusion Matrix, ROC-AUC, and F1-score were used to measure fraud detection accuracy.
- Histogram of Anomaly Scores analysed to visualise fraud detection boundaries.

7. EVALUATION OF FRAUD DETECTION MODELS

This section comprehensively evaluates the fraud detection models implemented in this study. The evaluation is based on classification performance metrics, confusion matrices, anomaly score distributions, and comparative analysis of the models.

7.1 Evaluation Metrics

To assess the performance of the models, the following evaluation metrics were used:

- Accuracy: Measures overall correctness, but is misleading in imbalanced datasets.
- Precision: Measures how many of the detected fraud cases were actual fraud.
- Recall (Sensitivity): Measures how well the model captures fraudulent transactions.
- F1-Score: The harmonic mean of precision and recall.
- ROC-AUC Score: Measures how well the model distinguishes between fraud and non-fraud transactions.

Results in Table Form:

Model	**Accuracy**	**Precision**	**Recall**	**F1-Score**	**ROC-AUC Score**
Autoencoder	92.51%	56.01%	66.48%	60.80%	0.8074
One-Class SVM	92.79%	64.82%	38.22%	48.09%	0.6812
Isolation Forest	90.23%	24.44%	5.62%	9.14%	0.5198

Figure 3: Experimental Results

7.2 Evaluation of Autoencoder Model

- Highest Recall (66.48%), meaning it captures more fraud cases than other models.
- Balanced F1-Score (60.80%), ensuring a good trade-off between precision and recall.
- Highest ROC-AUC Score (0.8074), proving its superior ability to separate fraud from non-fraud.
- The accuracy is at 92.51%, showing high overall correctness.
- Moderate Precision (56.01%), meaning some normal transactions are misclassified as fraud.

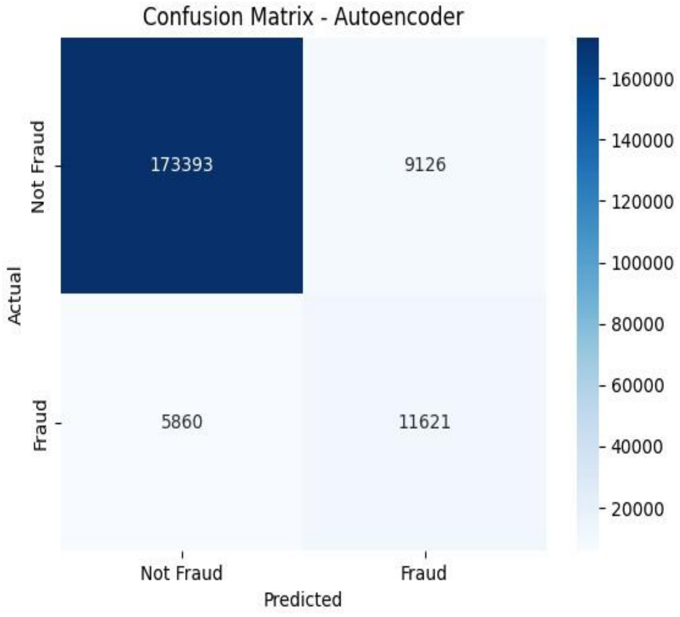

Figure 4: Confusion Matrix (Autoencoder)

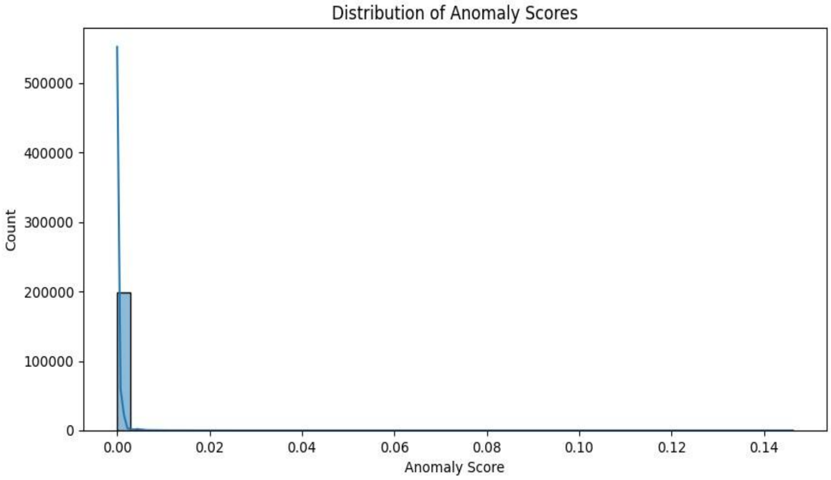

Figure 5: Distribution of Scores (Autoencoder)

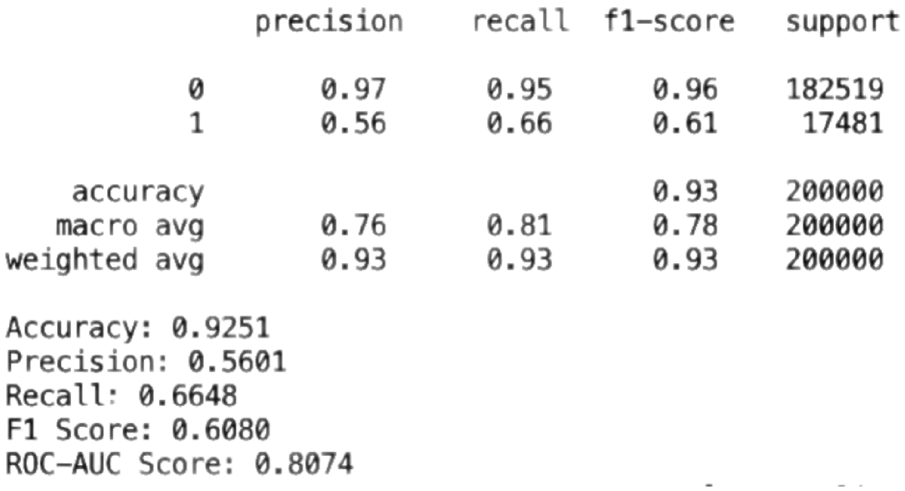

```
              precision    recall  f1-score   support

           0       0.97      0.95      0.96    182519
           1       0.56      0.66      0.61     17481

    accuracy                           0.93    200000
   macro avg       0.76      0.81      0.78    200000
weighted avg       0.93      0.93      0.93    200000

Accuracy: 0.9251
Precision: 0.5601
Recall: 0.6648
F1 Score: 0.6080
ROC-AUC Score: 0.8074
```

Figure 6: Classification Report (Autoencoder)

7.3 Evaluation of One-Class SVM

- Moderate Precision (64.82%) but low recall (38.22%) means it misses many fraud cases.
- Lower ROC-AUC Score (0.6812) compared to Autoencoder (50 Epochs).
- High Accuracy (92.79%), but misleading due to poor fraud recall.

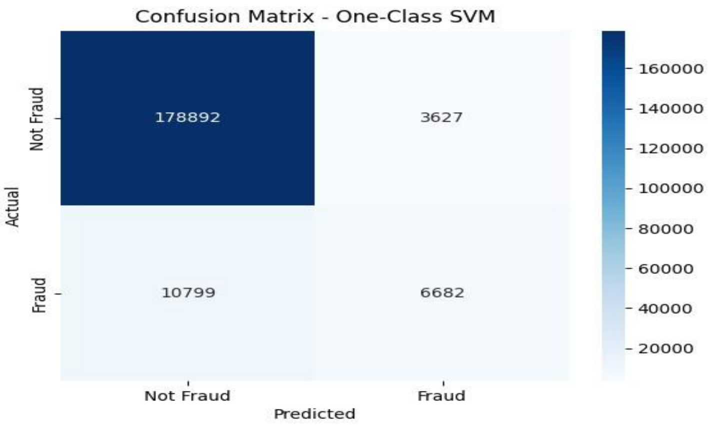

Figure 7: Confusion Matrix (One-Class SVM)

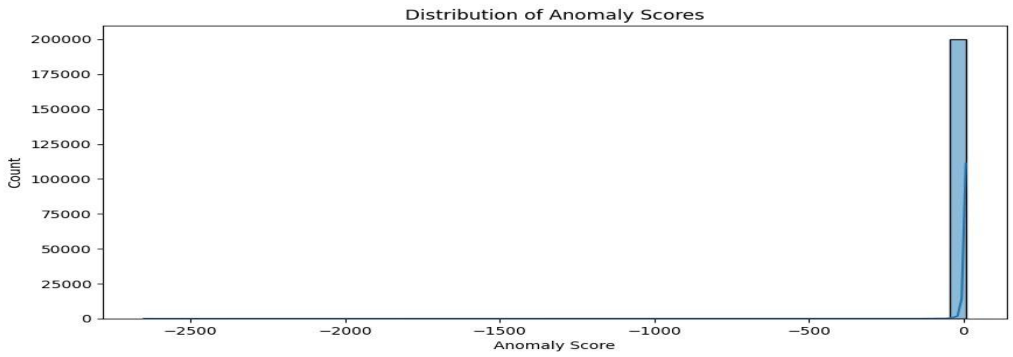

Figure 8: Distribution of Anomaly Scores (One-Class SVM)

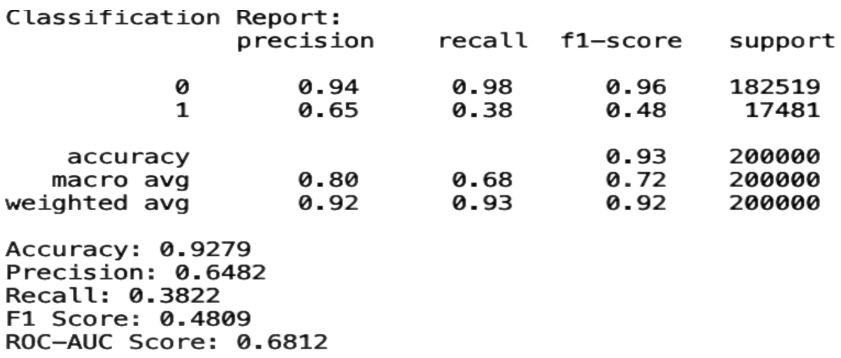

```
Classification Report:
              precision    recall  f1-score   support

           0       0.94      0.98      0.96    182519
           1       0.65      0.38      0.48     17481

    accuracy                           0.93    200000
   macro avg       0.80      0.68      0.72    200000
weighted avg       0.92      0.93      0.92    200000

Accuracy: 0.9279
Precision: 0.6482
Recall: 0.3822
F1 Score: 0.4809
ROC-AUC Score: 0.6812
```

Figure 9: Classification Report (One-Class SVM)

7.4 Evaluation of Isolation Forest

- Worst Recall (5.62%), meaning it fails to detect most fraud cases.
- ROC-AUC Score (0.5198) is close to random guessing.
- Low Precision (24.44%), meaning most flagged fraud cases are actually normal transactions.
- Accuracy (90.23%) is misleading since it does not account for fraud detection effectiveness. Overall, it is not a viable fraud detection method.

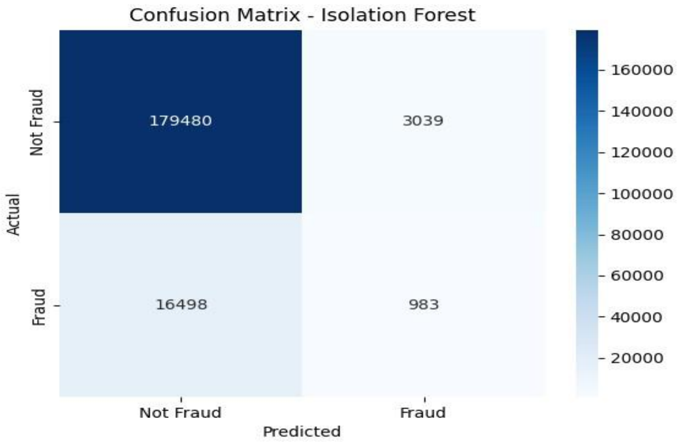

Figure 10: Confusion Matrix (Isolation Forest)

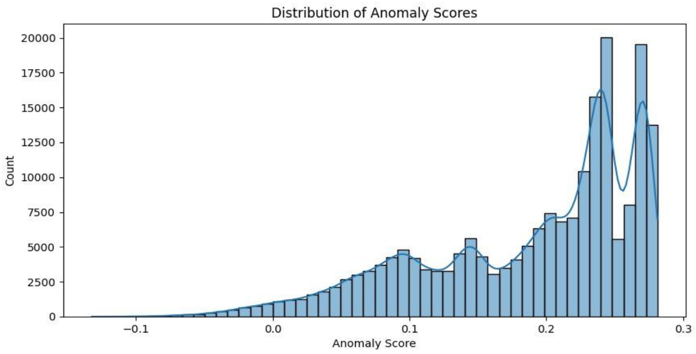

Figure 11: Distribution of Anomaly Scores (Isolation Forest)

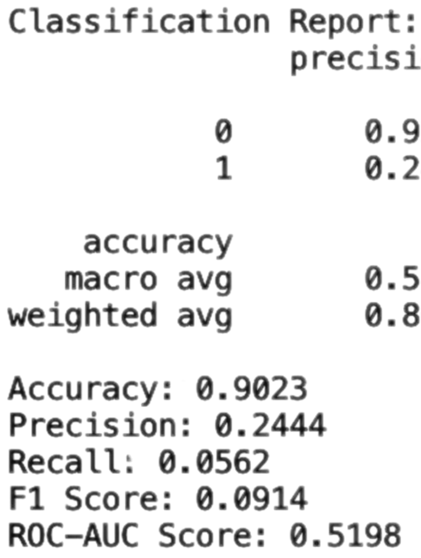

```
Classification Report:
               precision    recall  f1-score   support

           0       0.92      0.98      0.95    182519
           1       0.24      0.06      0.09     17481

    accuracy                           0.90    200000
   macro avg       0.58      0.52      0.52    200000
weighted avg       0.86      0.90      0.87    200000

Accuracy: 0.9023
Precision: 0.2444
Recall: 0.0562
F1 Score: 0.0914
ROC-AUC Score: 0.5198
```

Figure 12: Classification Reports (Isolation Forest)

7.5 Comparative Analysis and Best Model Selection

Figure 13: Comparative Analysis

Criteria	Best Performing Model	Reason
Fraud Recall	Autoencoder	Best at detecting fraud cases
Precision	One-Class SVM	Higher Precision but poor recall
Overall Balance	Autoencoder	Best F1-Score and Highest ROC-AUC
Accuracy	One-Class SVM	Highest accuracy but misleading due to low fraud recall

8. CONCLUSION

Autoencoder is the best model due to its superior recall, F1-Score, and ROC-AUC performance. However, One-Class SVM may be considered an alternative; it greatly sacrifices recall and requires many computing resources and time. Isolation Forest is unsuitable for fraud detection as even though accuracy is important, recall and precision are more important in our case of fraud detection. The conclusion is that the deep learning models may be better than supervised and unsupervised learning models for our use case. Future research in this area should primarily focus on improving feature engineering and trying to use hybrid models to enhance detection further. Combining an Autoencoder with a One-Class SVM can significantly improve precision and recall.

REFERENCES

[1] Staff, T. (2014, January 18). *FRAUD*. The Law Dictionary. https://thelawdictionary.org/fraud/

[2] Bolton, R. J., & Hand, D. J. (2002). Statistical Fraud Detection: A review. *Statistical Science*, *17*(3). https://doi.org/10.1214/ss/1042727940

[3] Xiang, S., Zhu, M., Cheng, D., Li, E., Zhao, R., Ouyang, Y., Chen, L., Australian Artificial Intelligence

[4] Institute, University of Technology Sydney, Department of Computer Science and Technology, Tongji University, Shanghai Artificial Intelligence Laboratory, & Tencent Jarvis Laboratory. (2023). Semi-supervised credit card fraud detection via Attribute-Driven Graph representation. In *The Thirty-Seventh*

[5] *AAAI Conference on Artificial Intelligence (AAAI-23)*. https://www.xiangshengcloud.top/publication/semisupervised-credit-card-fraud-detection-via-attribute-driven-graph-representation/Sheng-AAAI2023.pdf

[6] Richard J. Bolton, David J. Hand "Statistical Fraud Detection: A Review," Statistical Science, Statist. Sci. 17(3), 235-255, (August 2002)

[7] Garg, K. (2024). How can artificial intelligence be used to detect and mitigate zero-day vulnerabilities? *Scholarly Review ., SR Online: Showcase*(Equinox 2024). https://doi.org/10.70121/001c.124880

[8] Vanschoren, J. (n.d.). *OpenML*. OpenML: Exploring Machine Learning Better, Together. https://api.openml.org/d/45955

[9] *Year-over-Year Developments in Financial Fraud Detection via Deep Learning: A Systematic Literature review*. (n.d.). https://arxiv.org/html/2502.00201v1

[10] Wallny, F. (2022). False Positives in Credit Card Fraud Detection: Measurement and Mitigation. *Proceedings of the 55th Hawaii International Conference on System Sciences*, 1572. https://hdl.handle.net/10125/79527

[11] Tao, X., Zheng, Y., Chen, W., Zhang, X., Qi, L., Fan, Z., & Huang, S. (2021). SVDD-based weighted oversampling technique for imbalanced and overlapped dataset learning. *Information Sciences*, *588*, 13-51. https://doi.org/10.1016/j.ins.2021.12.066

[12] Xiang, S., Zhu, M., Cheng, D., Li, E., Zhao, R., Ouyang, Y., Chen, L., Australian Artificial Intelligence

[13] Institute, University of Technology Sydney, Department of Computer Science and Technology, Tongji University, Shanghai Artificial Intelligence Laboratory, & Tencent Jarvis Laboratory. (2023). Semisupervised credit card fraud detection via Attribute-Driven Graph representation. In The Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI-23). https://www.xiangshengcloud.top/publication/semisupervised-credit-card-fraud-detection-via-attribute-driven-graph-representation/Sheng-AAAI2023.pdf

[14] Godzwon, I. (2024, May 7). Credit_Card_Fraud. OPENML.org. Retrieved March 2, 2025, from --------https://openml.org/search?type=data&status=active&id=45955

[15] Association of Certified Fraud Examiners. (n.d.). *ACFE press release*.

[16] https://www.acfe.com/aboutthe-acfe/newsroom-for-media/press-releases/press-release-detail?s=2024- Report-to-the-Nations

[17] PricewaterhouseCoopers. (n.d.). *Combating fraud in the era of digital payments*. PwC.

[18] https://www.pwc.in/industries/financial-services/fintech/dp/combating-fraud-in-the-era-of-digitalpayments.html

[19] *Anti-Money laundering (AML)*. (n.d.). FINRA.org. https://www.finra.org/rules-guidance/keytopics/aml

[20] Merchant Risk Council, Visa Acceptance Solutions, Verifi, & B2B International. (2024). *2024 Global eCommerce Payments & Fraud Report*. https://www.cybersource.com/content/dam/documents/campaign/fraud-report/global-fraud-report2024.pdf

[21] Merchant Risk Council, Visa Acceptance Solutions, Verifi, & B2B International. (2024). *2024 Global eCommerce Payments & Fraud Report*. https://www.cybersource.com/content/dam/documents/campaign/fraud-report/global-fraud-report-2024.pdf

[22] Vanschoren, J. (n.d.). *OpenML*. OpenML: Exploring Machine Learning Better, Together. https://api.openml.org/d/45955

[23] Almuqati, Mohammed & Sidi, Fatimah & Mohd Rum, Siti Nurulain & Zolkepli, Maslina & Ishak, Iskandar. (2024). Challenges in Supervised and Unsupervised Learning: A Comprehensive Overview. International Journal on Advanced Science Engineering and Information Technology. 14. 1449-1455. 10.18517/ijaseit.14.4.20191.

[24] Sarker, I. H. (2021). Machine learning: algorithms, Real-World applications and research directions. *SN Computer Science*, *2*(3). https://doi.org/10.1007/s42979-021-00592-x

[25] Bairwa, V. (n.d.). *Unsupervised Machine learning: definition, working, types, pros & cons and applications*.

[26] https://www.edushots.com/Machine-Learning/unsupervised-machine-learning-overview