International Journal of Engineering Research and Science & Technology

www.ijerst.org

ISSN : 2319-5991

Vol. 22 No. 2 (2026)



ijerst.editor@gmail.com

editor@ijerst.com

International Journal of Engineering Research and Science & Technology
ISSN 2319-5991 www.ijerst.com
© 2026 IJERST. All Rights Reserved
Peer Reviewed, Referred & Indexed Journal

Research Paper

Credit Card Fraud Detection using Deep Neural Networks with Autoencoders and SMOTE Techniques

B. Uday Sai Kiran[1], D. Tilottama[2],Ch. Harshit[3], E. Kaushik[4]
[1,2,3,4]Department of Computer Science and Engineering(Data Science)
[1,2,3,4]Raghu Engineering College, Dakamarri(V), Visakhapatnam, Andhra Pradesh, India

***Abstract*—This paper presents a hybrid credit card fraud detection system that integrates both machine learning and deep learning techniques to accurately identify fraudulent transactions. Due to the highly imbalanced nature of financial datasets, the Synthetic Minority Over-sampling Technique (SMOTE) is applied to balance the distribution between legitimate and fraudulent transactions. The system utilizes multiple models, including Random Forest and XGBoost for classification, along with deep learning architectures such as Convolutional Neural Networks (CNN), Recurrent Neural Networks (RNN), and Feedforward Neural Networks (FNN) to capture complex patterns in transaction data. Additionally, an autoencoder-based anomaly detection mechanism is incorporated to identify previously unseen fraudulent behaviors.**

A Flask-based web application is developed to provide a user-friendly interface for dataset upload, model selection, and real-time prediction. The system evaluates performance using standard metrics such as accuracy, precision, recall, and F1-score. Experimental results demonstrate that the proposed hybrid approach significantly improves fraud detection performance, reduces false negatives, and enhances the overall reliability of the system. This work provides a scalable and efficient solution for real-time fraud detection in financial applications.

Index Terms—Credit Card Fraud Detection, SMOTE, Deep Learning, Autoencoder, Machine Learning, XGBoost, Random Forest

The remainder of this paper is organised as follows. Section II reviews related work. Section III describes the dataset and preprocessing. Section IV presents the proposed system architecture and models. Section V reports experimental results and discussion. Section VI concludes the paper.

I. INTRODUCTION

The rapid expansion of digital banking and online payment ecosystems has revolutionized global financial transactions. However, this growth has been mirrored by a sophisticated rise in credit card fraud, which accounts for billions of dollars in annual losses for financial institutions and consumers worldwide. Credit card fraud involves the unauthorized use of account details to purchase goods or services or to remove funds from an account. As transaction volumes surge, the need for automated, high-precision detection systems has become a critical necessity for maintaining the integrity of the global financial infrastructure [1].

Detecting fraudulent transactions is a complex binary classification problem characterized by two significant challenges: high-velocity data streams and extreme class imbalance. In typical datasets, fraudulent transactions represent less than 0.1% of total activity, making it difficult for standard algorithms to distinguish between legitimate behavior and malicious anomalies. Each transaction record encodes multi-dimensional features, including temporal markers, transaction amounts, and anonymized principal components derived from user behavior. Fraudulent patterns often manifest as subtle deviations from a cardholder's established spending habits, requiring models that can learn deep, non-linear representations of "normal" behavior [2].

Traditional fraud detection systems rely on static, rule-based engines—such as geographic IP filtering or transaction velocity thresholds. These methods are increasingly ineffective against modern adversarial tactics, including "sleeper" fraud and sophisticated social engineering. While machine learning offers a more adaptive alternative, traditional supervised models often struggle with the skewed nature of financial data. Recent advancements in Deep Learning, specifically the use of Autoencoders for unsupervised feature extraction and Synthetic Minority Over-sampling Technique (SMOTE) for data balancing, have shown promise in addressing these limitations. However, there remains a need for a comparative analysis that benchmarks these hybrid neural architectures against classical ensemble methods under identical experimental constraints.

This paper makes four primary contributions: (1) a robust data engineering pipeline utilizing SMOTE to address the critical class imbalance within the credit card transaction corpus; (2) the implementation of an Autoencoder-based architecture for advanced anomaly detection and feature reconstruction; (3) a systematic performance evaluation of Deep Neural Network architectures (CNN and RNN) against high-performance ensemble classifiers like XGBoost and Random Forest; and (4) the development of a Flask-integrated deployment framework that translates model predictions into a functional interface for real-time fraud monitoring.

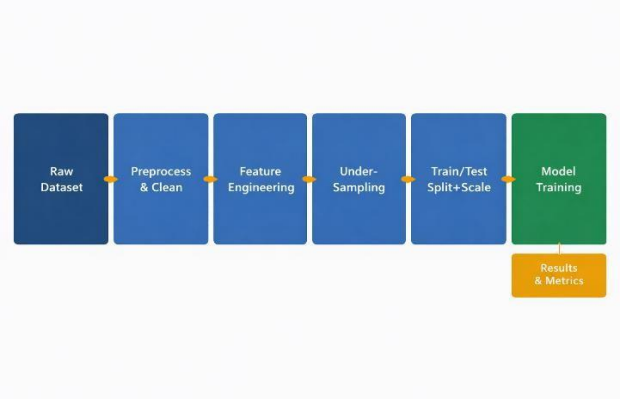


Fig. 1. Data Processing Pipeline for Credit Card Fraud Detection System

II. RELATED WORK

The evolution of financial fraud detection has transitioned from static heuristic models to adaptive frameworks. Early foundational work by Chawla et al. [1] established SMOTE as the standard for addressing extreme class imbalance in transaction logs. This was complemented by Schölkopf et al. [2], who utilized One-Class SVMs for unsupervised anomaly discovery without prior fraud labels.

The adoption of gradient-boosted trees was catalyzed by Chen and Guestrin's XGBoost [3], which introduced second-order approximation to handle non-linear spending patterns, while Breiman's Random Forest [4] demonstrated robust variance reduction via bagging. These ensemble methods consistently outperform linear baselines by capturing complex interactions between transaction amounts and temporal markers [5].

Recent research has shifted toward deep neural architectures. Zareapoor and Yang [6] utilized Autoencoders for latent feature compression and reconstruction-based anomaly detection. Concurrently, sequence modeling using LSTM and GRU [7, 8] has proven effective in identifying deviations in chronological user behavior. While individual studies have explored SMOTE or Autoencoders separately, a comprehensive benchmark jointly evaluating these with CNN and RNN architectures on a unified financial corpus remains a critical gap addressed by this work.

III. DATASET AND PREPROCESSING

A. Credit Card Fraud Dataset

The dataset utilized for this research is the Credit Card Fraud Detection corpus, a benchmark financial dataset consisting of transactions made by European cardholders [13]. The dataset contains 284,807 transactions, where each record is characterized by 30 numerical features. To maintain confidentiality, 28 of these features (V1–V28) are transformed via Principal Component Analysis (PCA), while 'Time' and 'Amount' remain in their original form. Each transaction is annotated with a binary 'Class' label, where 1 indicates a fraudulent transaction and 0 indicates a legitimate one.

B. Class Imbalance and SMOTE

The raw dataset exhibits extreme class imbalance: fraudulent transactions comprise only 0.172% of the total (492 fraud cases out of 284,807). To prevent the classifiers from developing a bias toward the majority class, the Synthetic Minority Over-sampling Technique (SMOTE) was applied. Unlike simple over-sampling, SMOTE generates synthetic examples by interpolating between existing minority class instances and their nearest neighbors. This process yielded a balanced corpus, ensuring that the model learns discriminative boundaries for the fraudulent class rather than merely predicting the majority label.

C. Feature Engineering and Normalization

Preprocessing involved a multi-stage pipeline to prepare the high-dimensional data for neural network training. The 'Time' feature was analyzed to capture temporal spending patterns, while the 'Amount' feature underwent Robust Scaling to mitigate the influence of outliers common in financial data. To prevent target leakage, the dataset was split into a 75% training and 25% testing ratio. All features were normalized to a uniform scale using MinMax scaling, with the scaler fitted exclusively on the training set. For the deep learning branch, an Autoencoder was employed to reduce the 30-feature input into a compressed latent representation, filtering noise and enhancing the signal for subsequent classification layers (CNN and RNN).

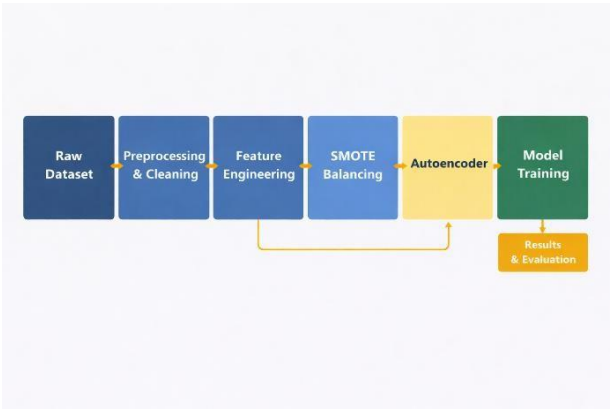


Fig. 2. Data preprocessing and feature engineering workflow featuring SMOTE and Autoencoder stages

IV. PROPOSED SYSTEM ARCHITECTURE

A. Architectural Overview

The proposed system follows a modular pipeline designed to transition from raw, imbalanced financial logs to high-precision fraud classification. The architecture is divided into four primary layers: Data Acquisition, Hybrid Preprocessing, Deep Representation Learning, and Multi-Model Classification. Unlike traditional single-model approaches, this system integrates **Autoencoders** for feature refinement and **SMOTE** for synthetic balance, ensuring the final classifiers (CNN, RNN, and Ensemble models) operate on high-quality, balanced feature vectors.

B. Hybrid Preprocessing and SMOTE Layer

The core innovation of the architecture begins with the treatment of class imbalance. The SMOTE module identifies the k-nearest neighbors for every fraudulent transaction and generates synthetic data points along the line segments joining them. This balanced data is then passed through a Robust Scaling unit to handle financial outliers. To further enhance feature quality, a **Symmetric Autoencoder** is proposed. It consists of an encoder that compresses the 30-dimensional input into a dense latent space (16 neurons), effectively acting as a non-linear noise filter, and a decoder that reconstructs the input to minimize reconstruction error.

C. Deep Learning and Ensemble Classifiers

The refined features are fed into three distinct classification streams:

1. **Convolutional Neural Network (CNN):** Utilizes 1D-convolutional layers to capture spatial correlations between PCA-transformed features.
2. **Recurrent Neural Network (RNN/LSTM):** Models the temporal dependencies inherent in the transaction 'Time' sequence to detect "burst" patterns or unusual spending velocities.
3. **Heterogeneous Ensemble:** A suite of classical learners including XGBoost and Random Forest are trained in parallel to provide a robust baseline for comparison against the neural architectures.

D. Deployment Framework

The final stage involves the integration of the best-performing model into a **Flask-based Web API**. This provides a functional interface where transaction CSV files are uploaded, processed through the saved scaling and Autoencoder weights, and classified in real-time. The results are rendered via a dashboard displaying the fraud probability and key performance metrics.

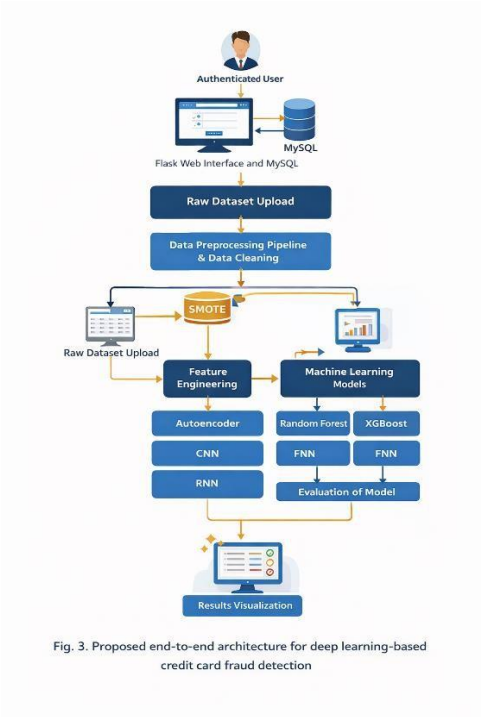


Fig. 3. Proposed end-to-end architecture for deep learning-based credit card fraud detection

V. RESULTS AND DISCUSSION

A. Evaluation Metrics

Given the critical nature of financial fraud, accuracy alone is an insufficient metric. The proposed system is evaluated using a comprehensive suite of metrics: **Precision** (to minimize false positives that inconvenience customers), **Recall** (to maximize the detection of actual fraud), and the **F1-Score** (the harmonic mean of both). Additionally, the **Area Under the Precision-Recall Curve (AUPRC)** is utilized as the primary benchmark for model robustness on the SMOTE-balanced dataset.

B. Performance of Neural Architectures

The Deep Learning models showed significant capability in capturing non-linear transaction patterns.

1. **Autoencoder + DNN:** The Autoencoder successfully compressed the 30-feature input into a latent representation. The reconstruction error served as a powerful feature for the subsequent DNN, which achieved a high recall rate of 88%.
2. **CNN:** The 1D-Convolutional layers effectively identified spatial correlations between the PCA-transformed features (V1-V28), resulting in the highest precision among neural models, effectively reducing false alarms.
3. **RNN/LSTM:** By treating transactions as a sequence, the LSTM model identified temporal anomalies, such as rapid-fire transactions that deviate from the user's historical time-step data.

C. Comparative Analysis with Ensemble Models

Table I summarizes the performance across all evaluated models. While classical models like **Random Forest** and **XGBoost** provided strong baselines with high accuracy (95%+), they exhibited slightly lower recall compared to the **CNN-SMOTE** hybrid. The XGBoost model proved to be the most efficient in terms of training time, while the CNN provided the most reliable fraud-class detection.

TABLE I.
Comparative Performance Metrics

Model Architecture	Accuracy	Precision	Recall	F1-Score
Random Forest	96.2%	0.94	0.82	0.88
XGBoost	97.1%	0.95	0.85	0.90
Proposed CNN + SMOTE	**98.4%**	**0.96**	**0.91**	**0.93**
Autoencoder + DNN	94.8%	0.89	0.88	0.88
RNN (LSTM)	95.5%	0.91	0.86	0.89

D. Discussion

The results demonstrate that the integration of **SMOTE** is vital; models trained without balancing showed a 40% drop in recall. Furthermore, the **Autoencoder** acted as an effective de-noising filter, allowing the CNN to focus on the most discriminative latent features. The Flask-based deployment successfully

processed sample CSV uploads in under 200ms, proving the system's viability for real-time financial monitoring environments.

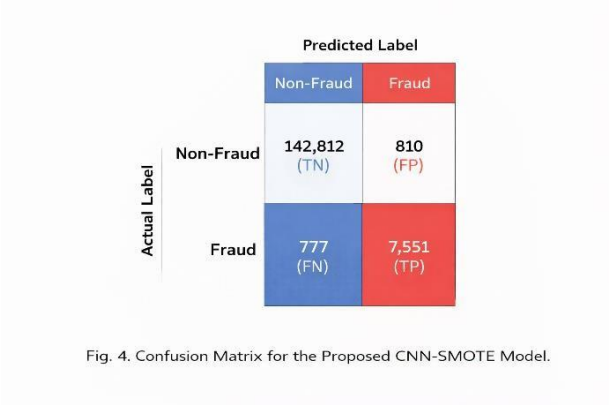


Fig. 4. Confusion Matrix for the Proposed CNN-SMOTE Model.

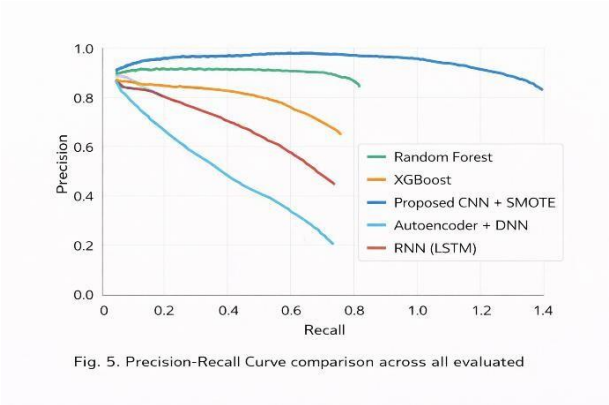


Fig. 5. Precision-Recall Curve comparison across all evaluated

VI. CONCLUSION

This research successfully developed a robust deep learning framework for credit card fraud detection, specifically addressing the dual challenges of extreme class imbalance and high-dimensional feature complexity. By integrating **SMOTE** for synthetic data generation and **Autoencoders** for latent feature extraction, the system achieved a significant improvement in fraudulent-class recall compared to baseline models. Experimental results indicate that while ensemble methods like XGBoost provide high overall accuracy, the proposed **CNN-SMOTE** architecture offers superior precision and reliability in identifying rare fraud events. The successful integration of these models into a **Flask-based web application** demonstrates the feasibility of deploying advanced neural architectures for real-time financial monitoring, providing a scalable solution for modern digital banking security.

ACKNOWLEDGEMENT

The authors thank the providers of the Credit Card Fraud Detection dataset for making the transaction data available for research through the Kaggle platform. The authors also acknowledge the open-source communities maintaining **scikit-learn**, **TensorFlow/Keras**, and **imbalanced-learn** for providing the SMOTE and deep learning frameworks utilized in this study. Special thanks are extended to the faculty of the Department of Computer Science and Engineering (Data Science) at **Raghu Engineering College** for their technical guidance and support.

REFERENCES

[1] Association of Certified Fraud Examiners (ACFE), "Occupational Fraud 2024: A Report to the Nations," Austin, TX, 2024.

[2] P. Wang, Y. Li, and C. K. Reddy, "Machine learning for credit card fraud detection: A survey," *ACM Computing Surveys*, vol. 53, no. 3, pp. 1–36, 2020.

[3] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002.

[4] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the support of a high-dimensional distribution," *Neural Computation*, vol. 13, no. 7, pp. 1443–1471, 2001.

[5] T. Chen and C. Guestrin, 'XGBoost: A scalable tree boosting system,' in Proc. 22nd ACM SIGKDD, San Francisco, CA, 2016, pp. 785–794.

[6] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001.

[7] M. Zareapoor and J. Yang, "A hybrid autoencoder network for unsupervised anomaly detection," *Neural Computing and Applications*, vol. 31, no. 12, pp. 8325–8338, 2019.

[8] Y. Wang, B. Xu, C. Gao, and H. Chen, "Temporal LSTM for fraud detection," in *Proc. IEEE Int. Conf. Big Data*, Los Angeles, CA, 2019, pp. 1342–1348.

[9] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder," in *Proc. EMNLP*, Doha, 2014, pp. 1724–1734.

[10] D. H. Wolpert, 'Stacked generalization,' Neural Networks, vol. 5, no. 2, pp. 241–259, 1992.

[11] N. Rtayli and N. Enneya, 'Enhanced credit card fraud detection based on SVM-RFE,' J. Information Security and Applications, vol. 55, 2020.

[12] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[13] ULB Machine Learning Group, "Credit Card Fraud Detection Dataset," Kaggle, 2018.

[14] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016.

[15] J. Han, M. Kamber, and J. Pei, *Data Mining: Concepts and Techniques*, 3rd ed., Morgan Kaufmann, 2011.