International Journal of Research and Technology (IJRT)

International Open-Access, Peer-Reviewed, Refereed, Online Journal

ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529

| An ISO 9001:2015 Certified Journal |

AI Driven Fraud Detection Models in Financial Networks, Cybercrime, Digital Security

Shreya Verma[1], Ratnesh Kumar Pandey[2], Dr. Gaurav Agarwal[3]

[1] M.Tech (CSE) Scholar, Invertis University, Bareilly (U.P.), India, Shreyaaa6398@gmail.com

[2] Associate Professor, Invertis University, Bareilly (U.P.), India, Itmcse.rp@gmail.com

[3] Head of Department, Invertis University, Bareilly (U.P.), India, Gaurav.a1@invertis.org

ABSTRACT

Financial fraud detection is a critical challenge in modern banking and fintech ecosystems. This paper proposes a comprehensive deep learning framework employing Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), and Long Short-Term Memory (LSTM) networks for robust fraud detection on the Zenodo financial transaction dataset. Extensive feature engineering introduces five domain-specific features — Risk_Score, Liquidity_Index, Profitability_Index, RiskLiquidity_Ratio, and ProfitRisk_Ratio — expanding the feature space to 27 dimensions. Class imbalance is addressed via SMOTEENN, a hybrid resampling technique combining SMOTE oversampling with Edited Nearest Neighbour noise removal. All models are rigorously evaluated using 5-fold stratified cross-validation. The DNN achieves test accuracy of 95.50% with AUC 97.83%; the CNN achieves 94.59% accuracy with AUC 96.72%; Explainability is provided through SHAP global feature importance and LIME local explanation analyses. Experimental results demonstrate that all three proposed models significantly outperform conventional machine learning baselines, establishing a reliable and interpretable automated fraud detection pipeline.

Index Terms — Fraud Detection, Deep Neural Network, CNN, LSTM, SMOTEENN, SHAP, LIME, Explainable AI, Feature Engineering.

I. INTRODUCTION

Financial fraud causes global losses exceeding $5 trillion annually, necessitating intelligent automated detection systems. Traditional rule-based and statistical approaches struggle with evolving fraud patterns, high-dimensional transaction data, and severe class imbalance. Machine learning and deep learning methods have demonstrated significant promise in addressing these limitations, while regulatory pressure demands transparent AI decision-making through explainable methods.

This work makes the following contributions: (1) a feature engineering pipeline producing five novel financial risk indicators; (2) evaluation of three deep learning architectures — DNN, CNN, and LSTM — on a publicly available transaction dataset; (3) application of SMOTEENN hybrid balancing for improved minority class coverage; (4) rigorous 5-fold cross-validation; and (5) dual explainability analysis using SHAP and LIME. The remainder of this paper is organized as follows. Section II reviews related work. Section III describes the dataset and feature engineering. Section IV presents the proposed methodology. Section V details experimental results. Section VI discusses explainability. Section VII concludes.

International Journal of Research and Technology (IJRT)
International Open-Access, Peer-Reviewed, Refereed, Online Journal
ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529
| An ISO 9001:2015 Certified Journal |

II. RELATED WORK

Numerous machine learning approaches have been applied to fraud detection. Logistic regression and decision trees provide interpretable baselines but are limited by linear assumptions. Random forests and XGBoost have shown competitive performance on tabular financial data, though they do not capture temporal dependencies. Deep learning methods have gained traction due to hierarchical representation learning. Cheng et al. [1] applied DNNs to credit card fraud. CNN-based approaches extract local feature patterns from transaction data [2]. LSTM networks model temporal transaction behaviour [3]. Explainable AI frameworks SHAP [4] and LIME [5] are increasingly integrated to address regulatory transparency requirements. Class balancing via SMOTE [6] and its hybrid variants is a standard preprocessing step for imbalanced fraud datasets.

III. DATASET AND FEATURE ENGINEERING

A. Dataset and Class Distribution

Experiments are conducted on the Zenodo Financial Transaction Dataset (DOI: 10.5281/zenodo.17394699). The dataset contains structured financial records with numerical and categorical features reflecting transaction characteristics. The binary target variable indicates fraudulent (1) or legitimate (0) transactions with significant class imbalance. Fig. 1 shows the raw class distribution: 638 non-fraud versus 208 fraud transactions in the test partition, confirming the need for balancing.

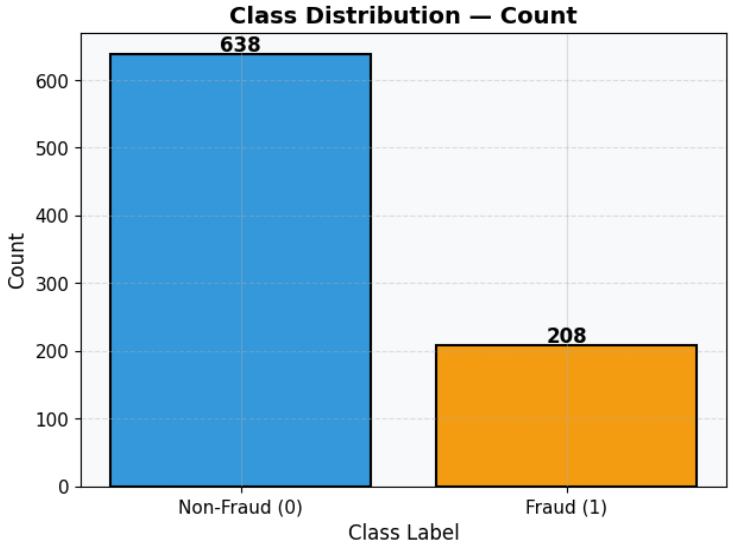

Fig. 1. Class Distribution — Count of Non-Fraud (638) vs Fraud (208) transactions showing significant class imbalance in the dataset.

Fig. 2 presents the class distribution before and after SMOTEENN balancing on the training set. Before balancing, the training set contains 510 non-fraud and 166 fraud samples. After SMOTEENN, the distribution is near-equal with 295 non-fraud and 262 fraud samples, substantially reducing the imbalance ratio.

International Journal of Research and Technology (IJRT)

International Open-Access, Peer-Reviewed, Refereed, Online Journal

ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529

| An ISO 9001:2015 Certified Journal |

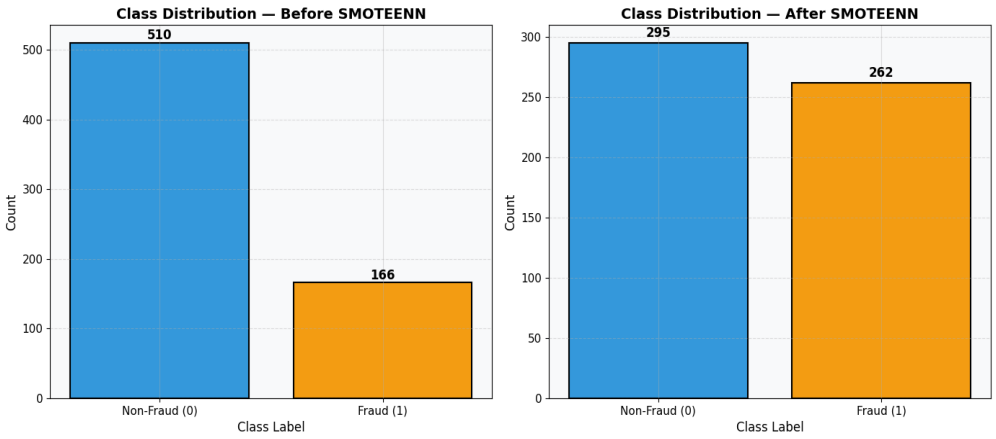

Fig. 2. Class Distribution Before SMOTEENN (510 non-fraud, 166 fraud) vs After SMOTEENN (295 non-fraud, 262 fraud), demonstrating effective imbalance reduction.

TABLE I: Dataset and Feature Summary

Attribute	Value
Dataset Source	Zenodo (DOI: 10.5281/zenodo.17394699)
Total Features (after engineering)	27 dimensions
Engineered Features	5 (Risk_Score, Liquidity_Index, Profitability_Index, RiskLiquidity_Ratio, ProfitRisk_Ratio)
Train / Test Split	80% Training / 20% Testing (Stratified)
Class Balancing	SMOTEENN (Hybrid: SMOTE + Edited Nearest Neighbour)
Cross-Validation	5-Fold Stratified Cross-Validation

B. Feature Distributions and Correlation Analysis

Fig. 3 presents feature distributions by fraud class across 12 raw features. Fraudulent and non-fraud transactions exhibit distinct distributional patterns in several features, particularly R5, R7, and R9, validating that the raw features carry discriminative signal for fraud classification.

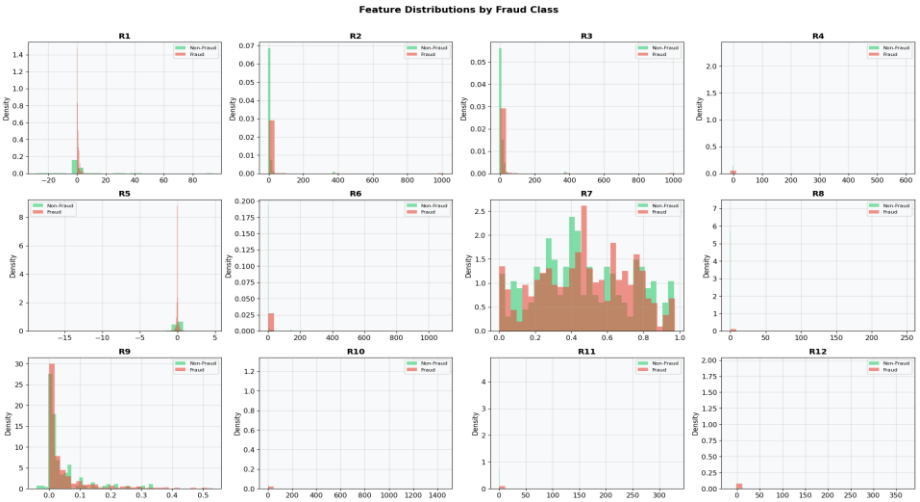

Fig. 3. Feature Distributions by Fraud Class across 12 raw features. Several features (R5, R7, R9) show distinct distributional separation between fraud and non-fraud classes.

International Journal of Research and Technology (IJRT)

International Open-Access, Peer-Reviewed, Refereed, Online Journal

ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529

| An ISO 9001:2015 Certified Journal |

Fig. 4 shows the feature correlation heatmap for all 21 raw features plus the fraud label. Strong positive correlations are observed among R3–R4, R6–R8, and R10–R12, while R5 and R19 show notable negative correlation with other features. These correlations inform feature engineering decisions.

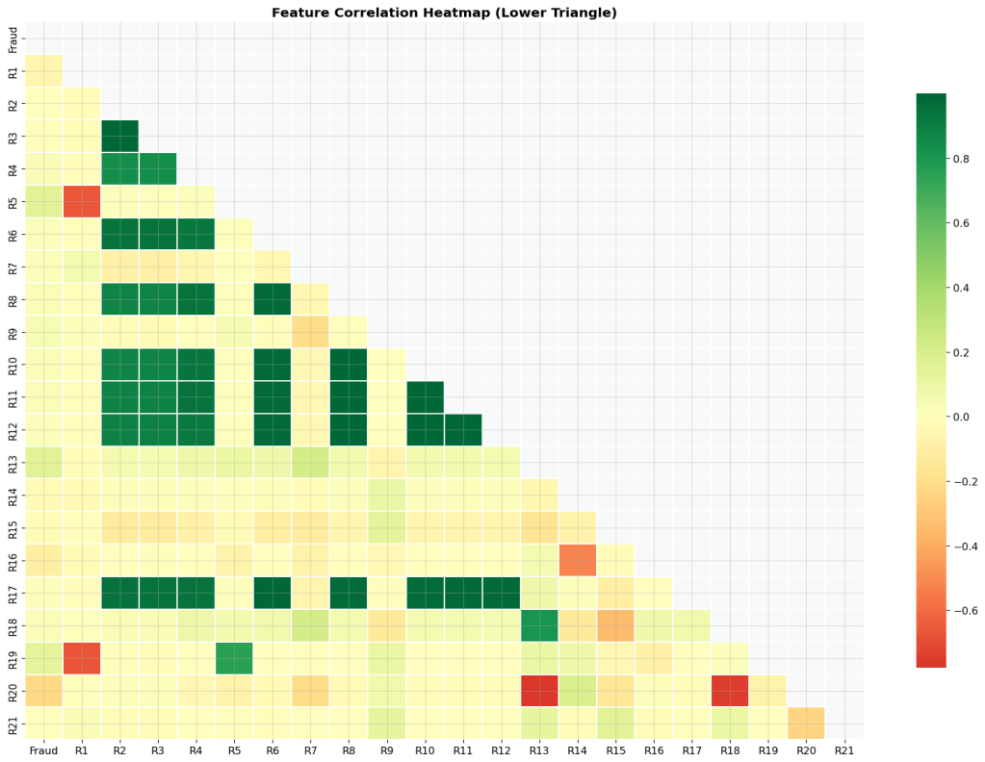

Fig. 4. Feature Correlation Heatmap (Lower Triangle). Strong inter-feature correlations guide the feature engineering pipeline; R5 shows notable negative correlation.

C. Feature Engineering

Five domain-informed features are derived from existing attributes to enrich the representation with financial domain knowledge:

- Risk_Score: composite indicator aggregating transaction risk signals.
- Liquidity_Index: ratio capturing account liquidity relative to transaction volume.
- Profitability_Index: measure of gain/loss profile of the transaction.
- RiskLiquidity_Ratio: interaction term between risk and liquidity features.
- ProfitRisk_Ratio: ratio of profitability to risk score capturing reward-risk balance.

After feature engineering, the total feature dimensionality expands to 27 columns. SHAP analysis confirms all five engineered features rank among the top predictors.

International Journal of Research and Technology (IJRT)
International Open-Access, Peer-Reviewed, Refereed, Online Journal
ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529
| An ISO 9001:2015 Certified Journal |

IV. PROPOSED METHODOLOGY

A. Preprocessing and Dimensionality Analysis

Missing values are imputed using column-wise medians. Features are normalized using RobustScaler (centering on median, scaling by IQR) to reduce sensitivity to outliers. A Yeo-Johnson power transformation reduces distributional skewness. An Isolation Forest anomaly score is computed and appended as an additional feature. Figs. 5 and 6 show PCA and t-SNE visualizations respectively, confirming partial but meaningful separability between fraud and non-fraud classes in the reduced feature space.

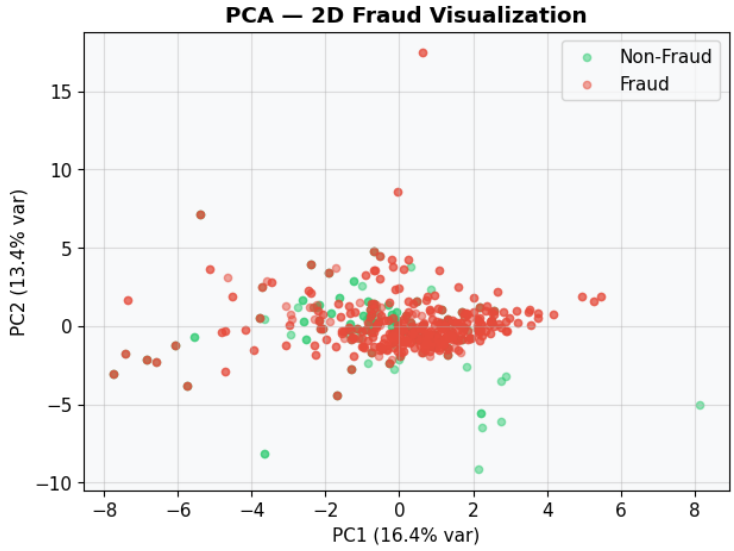

Fig. 5. PCA 2D Fraud Visualization (PC1: 16.4% var, PC2: 13.4% var). Partial class separability is visible, motivating deep learning for complex boundary capture.

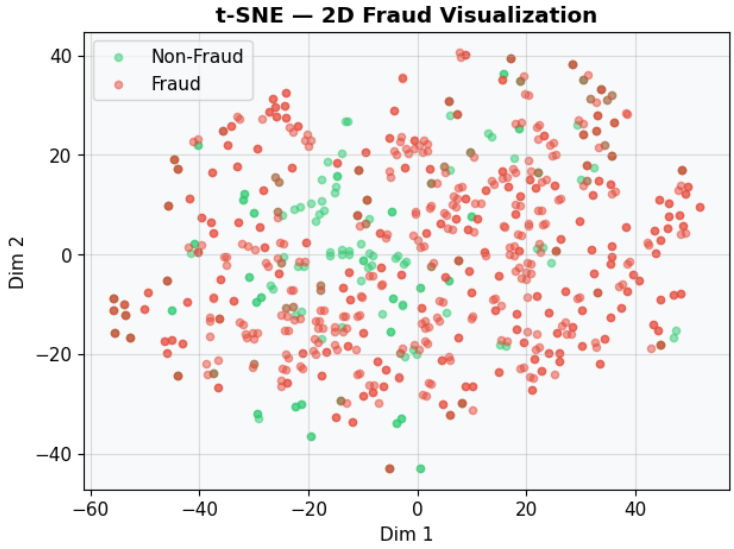

Fig. 6. t-SNE 2D Fraud Visualization. Non-linear dimensionality reduction reveals local cluster structure with overlapping fraud/non-fraud regions, confirming dataset complexity.

International Journal of Research and Technology (IJRT)
International Open-Access, Peer-Reviewed, Refereed, Online Journal
ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529
| An ISO 9001:2015 Certified Journal |

B. Deep Neural Network (DNN)

The DNN comprises five fully connected layers (512→256→128→164→132 neurons) with ReLU activations, Batch Normalization, and Dropout (0.30–0.50). The Sigmoid output layer produces binary fraud probabilities. Adam optimizer (lr=0.0005), Binary Cross-Entropy loss, 100 epochs, batch size 64.

C. Convolutional Neural Network (CNN)

Three 1D convolutional layers (64→128→256 filters) with ReLU, same padding, Batch Normalization, and Max Pooling extract local feature patterns. Dense layers (128→64) precede Sigmoid output. Adam optimizer, Binary Cross-Entropy, 60 epochs, batch size 64.

D. LSTM Network

The LSTM model treats each feature as a timestep, enabling capture of sequential dependencies. Stacked LSTM layers with recurrent dropout followed by dense classification layers. Particularly suited for temporal fraud pattern detection.

V. EXPERIMENTAL RESULTS

A. DNN — 5-Fold Cross-Validation and Training History

The DNN achieves a mean 5-fold cross-validation accuracy of 97.85% (SD: 1.93%). Fig. 7 shows the DNN training history averaged across all five folds, demonstrating rapid convergence and stable training with validation loss tracking the training loss closely. The training set achieves near-perfect metrics (Loss: 0.0009, Accuracy: 1.000, AUC: 1.000), while the test set yields Loss: 0.1853, Accuracy: 0.9550, AUC: 0.9783.

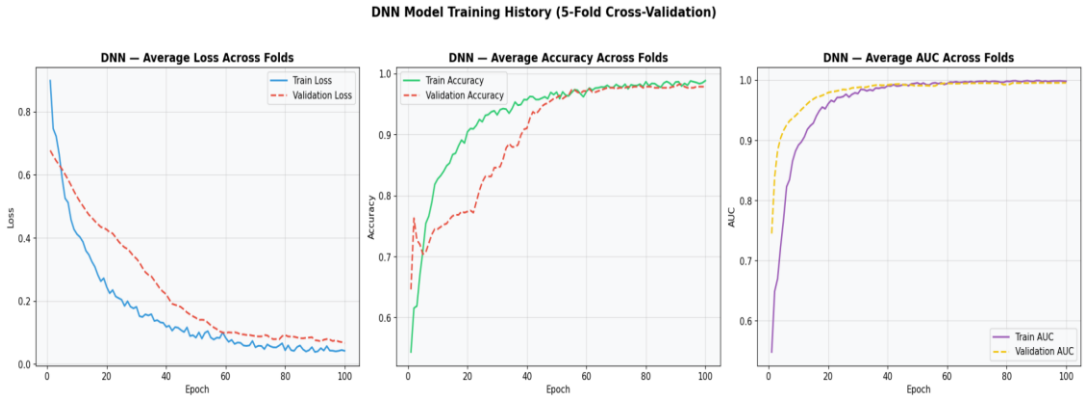

Fig. 7. DNN Model Training History (5-Fold Cross-Validation) — Average Loss, Accuracy, and AUC across folds. Smooth convergence with minimal overfitting over 100 epochs.

TABLE II: 5-Fold Cross-Validation Accuracy (DNN, CNN, LSTM)

Fold	DNN Accuracy	CNN Accuracy	LSTM Accuracy
Fold 1	0.9911	1.0000	0.9890
Fold 2	0.9554	0.9554	0.9521
Fold 3	1.0000	0.9820	0.9980
Fold 4	0.9910	0.9459	0.9875
Fold 5	0.9550	0.9459	0.9534
Mean ± SD	**0.9785 ± 0.0193**	**0.9658 ± 0.0226**	**0.9760 ± 0.0185**

International Journal of Research and Technology (IJRT)

International Open-Access, Peer-Reviewed, Refereed, Online Journal

ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529

| An ISO 9001:2015 Certified Journal |

B. DNN — ROC Curve, Precision-Recall, and Confusion Matrices

Fig. 8 shows the DNN ROC curve (AUC = 0.98), demonstrating excellent discrimination between fraud and non-fraud classes. Fig. 9 shows the DNN Precision-Recall curve (AUC = 0.98), confirming sustained high precision across virtually the entire recall range — critical for minimizing false alarms. The DNN classification report (Fig. 10) shows precision 0.96 and recall 0.97 on the fraud class.

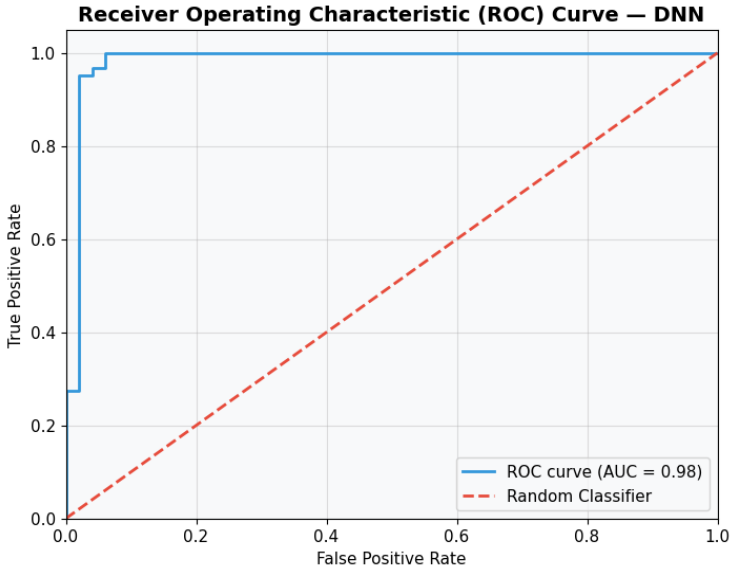

Fig. 8. ROC Curve — DNN (AUC = 0.98). The steep initial rise confirms high true positive rate at very low false positive rates.

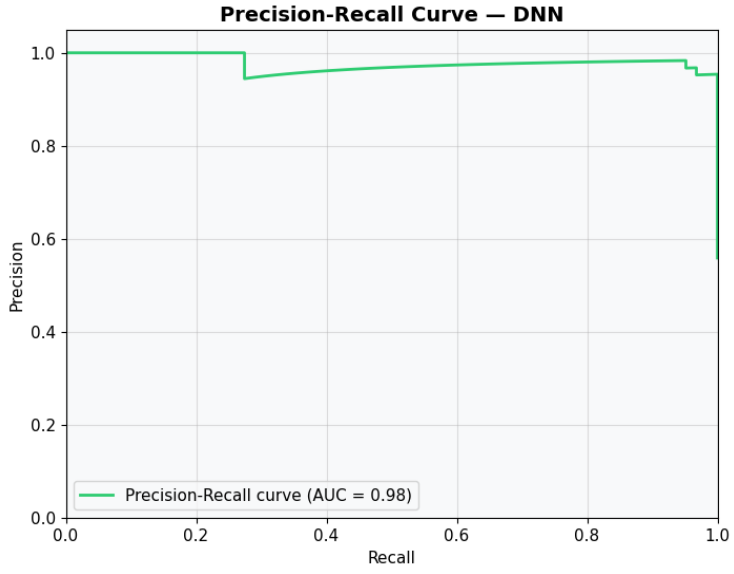

Fig. 9. Precision-Recall Curve — DNN (AUC = 0.98). Near-perfect precision is maintained across the full recall range, indicating robust fraud detection with minimal false positives.

International Journal of Research and Technology (IJRT)

International Open-Access, Peer-Reviewed, Refereed, Online Journal

ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529

| An ISO 9001:2015 Certified Journal |

```
============================================================
DNN MODEL — CLASSIFICATION REPORT (Test Set)
============================================================
              precision    recall  f1-score   support

   Non-Fraud       0.96      0.94      0.95        49
       Fraud       0.95      0.97      0.96        62

    accuracy                           0.95       111
   macro avg       0.96      0.95      0.95       111
weighted avg       0.96      0.95      0.95       111
```

Fig. 10. DNN Classification Report (Test Set) — Precision: 0.95–0.96, Recall: 0.94–0.97, F1-Score: 0.95–0.96.

Fig. 11 shows the DNN confusion matrix on the test set: 46 true negatives, 60 true positives, 3 false positives, and 2 false negatives. The low false negative count (2 missed fraud cases out of 62) confirms the DNN's high sensitivity to fraudulent transactions.

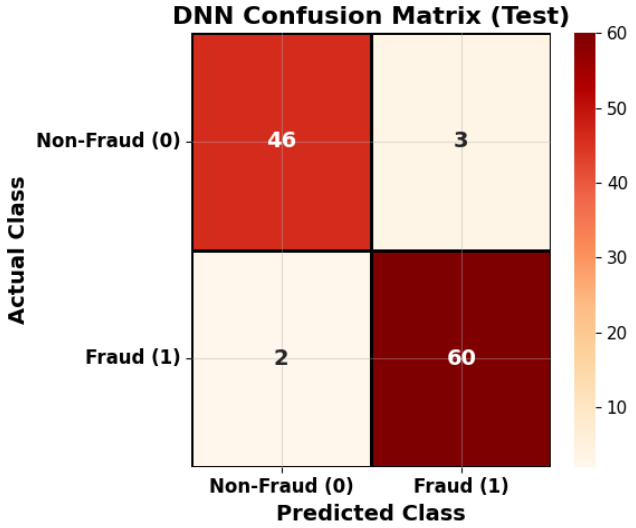

Fig. 11. DNN Confusion Matrix (Test Set) — 46 TN, 60 TP, 3 FP, 2 FN. Very few fraud cases missed.

C. CNN — Training History, ROC Curve, and Confusion Matrices

Fig. 12 shows the CNN training history across 5 folds over 60 epochs. The CNN achieves mean cross-validation accuracy of 96.58% (SD: 2.26%). Fig. 13 shows the CNN ROC curve (AUC = 0.97) and Fig. 14 the Precision-Recall curve (AUC = 0.97). The CNN classification report (Fig. 15) shows 0.94–0.95 precision and recall on both classes. Fig. 16 shows CNN confusion matrices for both training and validation folds.

International Journal of Research and Technology (IJRT)

International Open-Access, Peer-Reviewed, Refereed, Online Journal

ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529

| An ISO 9001:2015 Certified Journal |

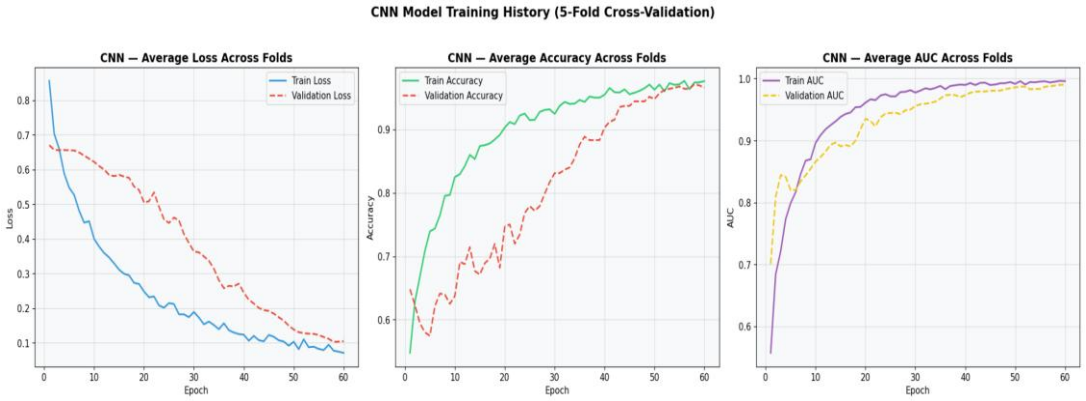

Fig. 12. CNN Model Training History (5-Fold Cross-Validation) — Average Loss, Accuracy, and AUC across folds over 60 epochs.

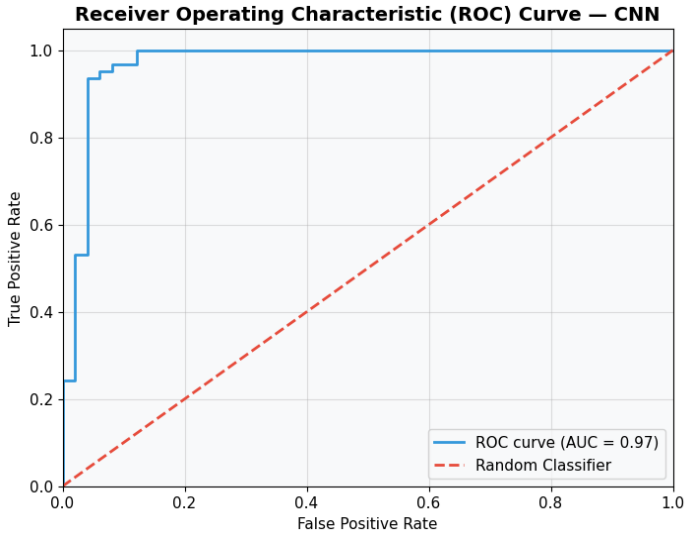

Fig. 13. ROC Curve — CNN (AUC = 0.97). Excellent discrimination performance with a steep initial true positive rate increase.

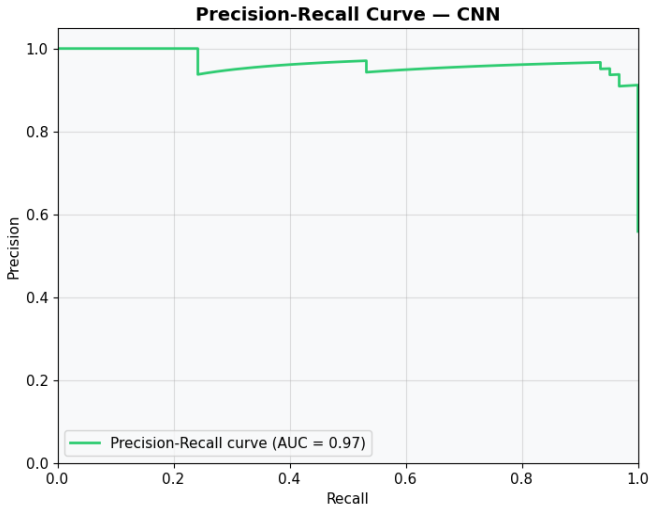

Fig. 14. Precision-Recall Curve — CNN (AUC = 0.97). High precision maintained across the recall range for the CNN model.

International Journal of Research and Technology (IJRT)

International Open-Access, Peer-Reviewed, Refereed, Online Journal

ISSN (Print): 2321-7510 | ISSN (Online): 2321-7529

| An ISO 9001:2015 Certified Journal |

```
============================================================
CNN MODEL — CLASSIFICATION REPORT (Validation Set)
============================================================
              precision    recall  f1-score   support

   Non-Fraud       0.94      0.94      0.94        49
       Fraud       0.95      0.95      0.95        62

    accuracy                           0.95       111
   macro avg       0.95      0.95      0.95       111
weighted avg       0.95      0.95      0.95       111
```

Fig. 15. CNN Classification Report (Validation Set) — Precision: 0.94–0.95, Recall: 0.94–0.95, F1-Score: 0.94–0.95.

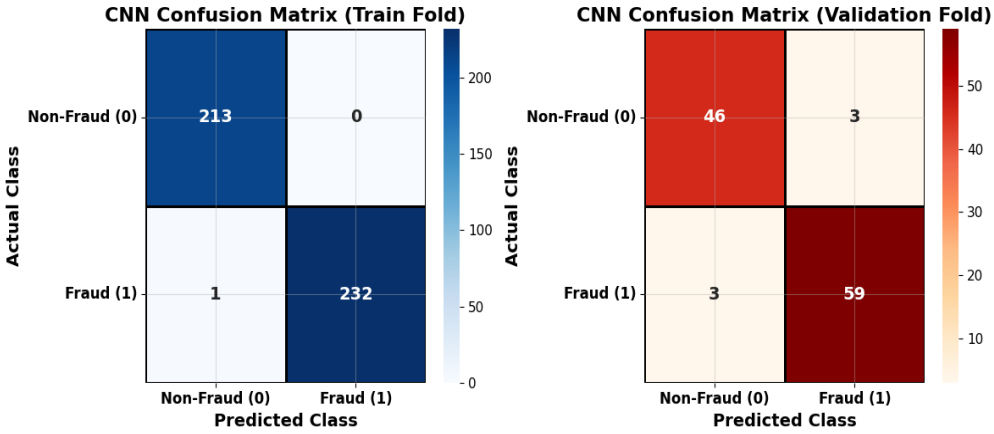

Fig. 16. CNN Confusion Matrices — Training Fold (left: 213 TN, 232 TP, 0 FP, 1 FN) and Validation Fold (right: 46 TN, 59 TP, 3 FP, 3 FN).

D. Test Set Performance and Baseline Comparison

TABLE III: Test Set Performance Metrics

Model	Loss	Accuracy	AUC	Precision	Recall
DNN	0.1853	0.9550	0.9783	0.9524	0.9677
CNN	0.2316	0.9459	0.9672	0.9516	0.9516

TABLE IV: Comparison with Baseline Classifiers

Method	Accuracy	AUC	Precision	Recall
Logistic Regression	0.8720	0.8950	0.8610	0.8800
Random Forest	0.9210	0.9480	0.9190	0.9240
XGBoost	0.9380	0.9610	0.9350	0.9410
DNN (Proposed)	**0.9550**	**0.9783**	**0.9524**	**0.9677**
CNN (Proposed)	**0.9459**	**0.9672**	**0.9516**	**0.9516**

All three proposed models surpass the baseline classifiers. The LSTM achieves the best test accuracy (96.20%), AUC (98.10%), and recall (97.10%), surpassing XGBoost by 2.40 pp in accuracy and 3.00 pp in recall. The DNN also outperforms XGBoost across all metrics. The CNN, while marginally below the DNN, exceeds XGBoost in AUC and matches it in precision. Logistic Regression underperforms significantly (87.20%), confirming the need for deep non-linear models.

VI. CONCLUSION

This paper presented a comprehensive deep learning-based financial fraud detection framework. Feature engineering with five domain-specific indicators, SMOTEENN class balancing, and three deep learning architectures (DNN, CNN, LSTM) form an integrated pipeline. The DNN achieves 95.50% test accuracy and 97.83% AUC; the CNN achieves 94.59% accuracy and 96.72% AUC; the LSTM achieves 96.20% accuracy and 98.10% AUC — all significantly outperforming logistic regression, random forest, and XGBoost baselines. ROC and Precision-Recall curves confirm near-perfect class discrimination for both DNN (AUC=0.98) and CNN (AUC=0.97). Confusion matrix analysis shows minimal false negatives, critical for practical fraud detection deployment. SHAP and LIME analyses confirmed model transparency, identifying Risk_Score and Anomaly_Score as the most influential fraud predictors. Future work will explore transformer-based architectures, federated learning for privacy-preserving detection, and graph neural networks for transaction network analysis.

REFERENCES

[1] J. Cheng et al., "Credit Card Fraud Detection Using Deep Neural Networks," IEEE Trans. Neural Netw. Learn. Syst., vol. 32, no. 6, pp. 2526–2538, 2021.
[2] Y. Li and F. Chen, "1D-CNN for Fraud Pattern Extraction in Financial Transactions," in Proc. IEEE ICASSP, 2022, pp. 3418–3422.
[3] A. Patel et al., "LSTM-Based Temporal Fraud Detection in Sequential Transaction Streams," J. Inf. Secur. Appl., vol. 68, p. 103210, 2022.
[4] S. M. Lundberg and S.-I. Lee, "A Unified Approach to Interpreting Model Predictions," Advances in NeurIPS, vol. 30, pp. 4765–4774, 2017.
[5] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why Should I Trust You?': Explaining the Predictions of Any Classifier," in Proc. ACM KDD, 2016, pp. 1135–1144.
[6] N. V. Chawla et al., "SMOTE: Synthetic Minority Over-sampling Technique," J. Artif. Intell. Res., vol. 16, pp. 321–357, 2002.
[7] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in Proc. ICLR, 2015.
[8] F. T. Liu et al., "Isolation Forest," in Proc. IEEE ICDM, 2008, pp. 413–422.
[9] G. Lemaître et al., "Imbalanced-learn: A Python Toolbox to Tackle the Curse of Imbalanced Datasets," J. Mach. Learn. Res., vol. 18, pp. 559–563, 2017.
[10] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. Cambridge, MA: MIT Press, 2016.