



A CNN-BASED FRAMEWORK FOR PAYROLL FRAUD DETECTION IN SIMULATED NIGERIAN PUBLIC FINANCIAL DATA

KENNETH EBUKA ASOGWA* AND MUMINU OSUMAH ADAMU

ABSTRACT. This study develops an artificial intelligence–driven framework for detecting payroll fraud in Nigeria’s Public Financial Management (PFM) system using Convolutional Neural Networks (CNNs). The model targets common fraud patterns such as ghost workers, salary inflation, duplicate payments, and payments to retired or overage employees. A dual-stage validation strategy was adopted. First, a publicly available IBM Human Resources dataset containing 1,470 records was used to benchmark the model, achieving an F1-score of 0.87. Second, a realistically simulated Nigerian payroll dataset comprising 12,110 records with 31 features was constructed to reflect local administrative structures and fraud typologies. To address severe class imbalance, the Synthetic Minority Oversampling Technique (SMOTE) was applied during model training. The proposed model demonstrated strong generalization performance on the simulated dataset, achieving an accuracy of 0.958, precision of 0.910, recall of 0.841, F1-score of 0.874, and PR-AUC of 0.857. Furthermore, the trained model was integrated into a real-time fraud alert dashboard to support proactive monitoring and decision-making. The study demonstrates the potential of deep learning approaches for fraud detection using simulated data.

1. INTRODUCTION

Fraud in Nigeria’s Public Financial Management (PFM) system remains a persistent and costly challenge to fiscal sustainability, service delivery, and public trust. Despite the implementation of major reforms such as the Treasury Single Account (TSA), the Integrated Payroll and Personnel Information System (IPPIS), and the Government Integrated Financial Management Information System (GIFMIS), fraudulent practices continue to evolve. Ghost-worker schemes alone are estimated to cost the government over ₦100 billion annually [13], while cumulative fraud-related losses between 2015 and 2025 have exceeded ₦1.5 trillion [7].

Current fraud detection mechanisms in most PFM platforms are largely rule-based and reactive, relying on static thresholds and manual audits that are increasingly ineffective against sophisticated fraud schemes. While traditional machine learning models such as logistic regression, decision trees, and random forests have been applied to fraud detection tasks, they often require extensive manual feature engineering and struggle with the severe class imbalance inherent in payroll fraud data, where fraudulent cases constitute only a small minority of transactions.

2010 *Mathematics Subject Classification*. Primary: 22E30. Secondary: 58J05.

Keywords and phrases. Convolutional Neural Networks (CNN), Ghost Workers, SMOTE, Class Imbalance, Public Financial Management, Fraud Detection, Nigeria.

©2026 Department of Mathematics, University of Lagos.

Submitted: February 17, 2026. Revised: March 26, 2026. Accepted: April 4, 2026.

*Correspondence

In recent years, deep learning techniques have demonstrated superior performance in fraud and anomaly detection, particularly when dealing with high-dimensional and structured data [8, 2]. Convolutional Neural Networks (CNNs), originally developed for image processing, have been successfully adapted to tabular financial data due to their ability to automatically learn hierarchical feature representations with minimal domain-specific feature engineering [10, 16]. For instance, Salaudeen et al. [16] combined CNNs with the Synthetic Minority Oversampling Technique (SMOTE) and reported significant improvements in classification accuracy, precision, recall, and F1-score for credit card fraud detection. Similar successes in deep learning-based approaches further highlight their potential for financial fraud detection tasks [8, 2].

However, the application of deep learning to tabular data remains an area of ongoing debate. Studies such as Gorishniy et al. (2021) suggest that tree-based models can outperform deep learning approaches on structured datasets. In addition, the growing emphasis on explainable artificial intelligence (XAI) highlights the need for interpretable fraud detection systems, particularly in public sector environments where transparency and accountability are critical.

Despite these advancements, significant gaps remain in the literature. Most existing studies focus on domains such as credit card or insurance fraud, with limited attention given to payroll fraud in public sector financial management. Furthermore, there is a scarcity of research addressing the unique challenges of Nigerian PFM systems, including ghost workers, salary inflation, duplicate payments, and payments to retired or ineligible employees. Importantly, few studies extend beyond model development to practical deployment in real-time fraud detection systems.

This study addresses these gaps by developing a tailored one-dimensional Convolutional Neural Network (1D-CNN) framework for payroll fraud detection in Nigeria. The proposed approach integrates realistic synthetic data generation, comprehensive evaluation (including ablation studies, baseline comparisons, sensitivity analysis, and bootstrapped confidence intervals), and deployment within a functional real-time fraud alert dashboard to support proactive decision-making.

The main contributions of this study are threefold: (1) the adaptation of 1D-CNNs for payroll fraud detection in Nigerian PFM systems, (2) a rigorous evaluation framework, and (3) the deployment of the model within a real-time fraud alert system. The remainder of this paper is organized as follows: Section 2 presents the materials and methods; Sections 3 and 4 present and discuss the results; and Section 5 concludes with policy implications and future research directions.

2. MATERIALS AND METHODS

This section describes the data sources, preprocessing techniques, model architecture, and evaluation metrics employed in this study.

2.1. Data Sources and Simulation

Due to the confidentiality and restricted access to official government payroll data, a realistic synthetic dataset was generated to mirror the structure and operational characteristics of the Nigerian public sector. The use of simulated data is a common and accepted approach in fraud detection research, particularly in domains where real datasets are sensitive, scarce, or inaccessible. Synthetic data generation enables controlled experimentation while preserving privacy and allowing the injection of known fraud patterns for model evaluation.

The simulation incorporated Nigerian-specific attributes, including ministries, departments, states, and civil service grade levels ranging from GL 08 to GL 17. To ensure realism and domain relevance, multiple fraud scenarios were explicitly embedded in the dataset, including ghost workers, salary inflation beyond grade-level limits, duplicate payments, and payments to retired or overage employees. These scenarios reflect commonly reported fraud patterns in Nigeria's Public Financial Management (PFM) systems.

The final dataset consisted of 12,110 records with 31 features and a binary fraud label, with class imbalance deliberately maintained to reflect real-world conditions where fraudulent cases are relatively rare.

To complement the simulated dataset and provide baseline benchmarking, a publicly available dataset from the International Business Machines Corporation (IBM) Human Resources dataset, containing 1,470 records, was also utilized. In this dataset, employee attrition was treated as a proxy for anomalous behavior, allowing for preliminary validation of model performance on a real-world dataset.

2.2. Handling Class Imbalance

The class distribution in the datasets was highly imbalanced, with fraudulent cases representing a small minority. The Synthetic Minority Oversampling Technique (SMOTE) [5] was applied to the training data to generate synthetic samples of the minority class. This approach improved class balance and enhanced the model's ability to learn discriminative fraud patterns.

2.3. Model Architecture

The proposed model is a one-dimensional Convolutional Neural Network (CNN) specifically adapted for tabular payroll data in the Nigerian Public Financial Management context. While the fundamental neural network equations and CNN architecture follow standard deep learning principles [9], the novelty of this work lies in its adaptation to payroll fraud detection, including the use of 1D convolutions on reshaped tabular features, integration with SMOTE for class imbalance handling, and deployment in a real-time fraud alert dashboard.

A feedforward neural network learns a function

$$y = f(x; \theta)$$

where $x \in \mathbb{R}^d$ represents the input features and y is the predicted output. The parameter θ (weights and biases) are learned during training.

Each neuron in a hidden layer computes:

$$z = Wx + b$$

$$\alpha = \sigma(z)$$

where W is the weight matrix, b is the bias vector, and $\sigma(\cdot)$ is a non-linear function (ReLU in the hidden layers and Sigmoid in the output layer).

For convolutional layers, feature extraction is performed using:

$$h^l = \sigma(W^l * x^l + b^l)$$

where $*$ denotes the convolution operation, and $h^{(l)}$ represents the feature maps at layer l .

For binary classification (fraud vs. non-fraud), the output layer produces the predicted probability of fraud:

$$\hat{y} = \sigma(Wx + b)$$

Where \hat{y} is the predicted probability of fraud.

The model is trained by minimizing the Binary Cross-Entropy Loss:

$$L(\theta) = \frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$$

where N is the number of samples, $y_i \in \{0,1\}$, and $\hat{y}_i \in [0,1]$. This loss corresponds to the negative log-likelihood of a Bernoulli distribution.

Parameters are updated iteratively using an optimizer (Adam or RMSprop):

$$\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}$$

Where η is the learning rate and $\nabla_{\theta} \mathcal{L}$ is the gradient of the loss with respect to the parameters.

The implemented CNN architecture consists of two 1D convolutional layers with 32 and 64 filters (kernel size = 3), each followed by ReLU activation, batch normalization, and max-pooling. A dropout layer (rate = 0.3) is applied to reduce overfitting, followed by a fully connected dense layer with sigmoid activation for binary classification.

2.4. Evaluation Metrics

Model performance was evaluated using standard classification metrics suitable for imbalanced data: precision, recall, F1-score, and accuracy. The F1-score was prioritized as the primary metric due to severe class imbalance in fraud detection. An optimal decision threshold was determined using the F1-maximization method on the precision-recall curve.

To ensure robust and generalizable results, 5-fold stratified cross-validation was performed with per-fold metrics and standard deviation reported. Statistical significance of performance differences against baseline models (Random Forest and XGBoost) was assessed using the

Wilcoxon signed-rank test. Model stability was further validated through 1,000 bootstrap resamples to generate 95% confidence intervals for all metrics. An ablation study was conducted to quantify the contribution of SMOTE, and sensitivity analysis tested robustness to varying fraud rates (0.02–0.05). Finally, inference time and real-time latency were measured in the deployed dashboard to support practical deployment claims.

2.5 Code Availability and Reproducibility

All code, including the MinistryFinancePayrollSimulator, CNN model implementation, hyperparameter tuning, ablation study, sensitivity analysis, bootstrapping, and real-time dashboard, is publicly available on GitHub for full reproducibility:

https://github.com/Kenneth-Asogwa/Fraud_Detection_Algorithm

The repository contains the complete dataset generation scripts, trained model weights, and all generated outputs (CSV files and figures) used in this study.

3. RESULTS

The CNN model demonstrated strong performance across both datasets evaluated in this study. On the IBM HR benchmark dataset, the model achieved an F1-score of 0.87, indicating its capability to learn meaningful patterns from structured tabular data. When evaluated on the simulated Nigerian payroll dataset, the model achieved an accuracy of 0.958, precision of 0.910, recall of 0.841, F1-score of 0.874, and PR-AUC of 0.857.

To ensure the reliability of these metrics and eliminate the risk of data leakage, the implementation of SMOTE was restricted to the training folds during cross-validation. By partitioning the data before resampling, the synthetic fraud cases remained isolated from the validation sets, ensuring the reported performance reflects the model's true generalization capability on original, unseen data.

To further assess the model's discriminative capability, a Receiver Operating Characteristic (ROC) analysis was conducted. The model achieved an Area Under the Curve (AUC) of 0.94, indicating excellent ability to distinguish between fraudulent and legitimate payroll transactions. The confusion matrix analysis also revealed a relatively low false-positive rate while maintaining a high fraud detection rate, which is particularly important in public sector auditing environments where excessive false alarms may increase the workload of auditors.

3.1. CNN Performance on IBM Dataset

As an initial validation step, the CNN model was evaluated on the publicly available IBM HR Attrition dataset (1,470 records), treating employee attrition as a proxy for anomalous behaviour. This served as a benchmark to assess the model's ability to learn meaningful patterns from structured tabular data before applying it to the domain-specific Nigerian payroll dataset. The results are summarised in Table 3.1.

Table 3.1: Performance of the CNN Model on the IBM HR Dataset

Metric	Score
Accuracy	0.87
Precision	0.92
Recall	0.82
F1 Score	0.87

The model achieved a strong F1-score of 0.87, confirming its capability to effectively learn from structured data and providing a solid baseline before transfer to the more complex Nigerian payroll scenario.

3.2. CNN Performance on Simulated Nigerian Dataset

The primary evaluation was conducted on the realistically simulated Nigerian payroll dataset (12,110 records with 31 features) that reflects local civil service structures and common fraud typologies (ghost workers, salary inflation, duplicate payments, and overage payments). The CNN model demonstrated strong overall performance, as summarized in Table 3.2.

Table 3.2: Performance of the CNN Model on the Simulated Nigerian Payroll Dataset

Metric	Score
Accuracy	0.958
Precision	0.910
Recall	0.841
F1 Score	0.874
PR-AUC	0.857

The model achieved a high F1-score of 0.874, reflecting a strong balance between precision and recall. This indicates effective detection of rare fraud cases while keeping false positives low a critical requirement for practical deployment in Nigerian Public Financial Management systems such as IPPIS and GIFMIS.

3.3 Ablation Study: Impact of SMOTE

To quantify the contribution of SMOTE in addressing class imbalance, an ablation study was conducted by retraining the CNN model without SMOTE (relying only on class weighting). The results are presented in Table 3.3.

Table 3.3: Ablation Study – Performance With and Without SMOTE

Method	Precision	Recall	F1
With SMOTE	0.9102	0.8412	0.8743
Without SMOTE	0.9224	0.8175	0.8668

To evaluate the contribution of SMOTE in handling class imbalance, an ablation study was conducted by retraining the CNN model without SMOTE (using class weighting only). The results are shown in Table 3.3 SMOTE improved the overall F1-score from 0.867 to 0.874, primarily by increasing recall (+0.024) while maintaining high precision. This demonstrates that SMOTE provides a better balance between detecting fraud cases (e.g., ghost workers and salary inflation) and minimizing false alarms a critical requirement in Nigerian PFM systems where missing fraud is costly but excessive alerts burden auditors. The results confirm the value of SMOTE in highly imbalanced payroll fraud detection tasks.

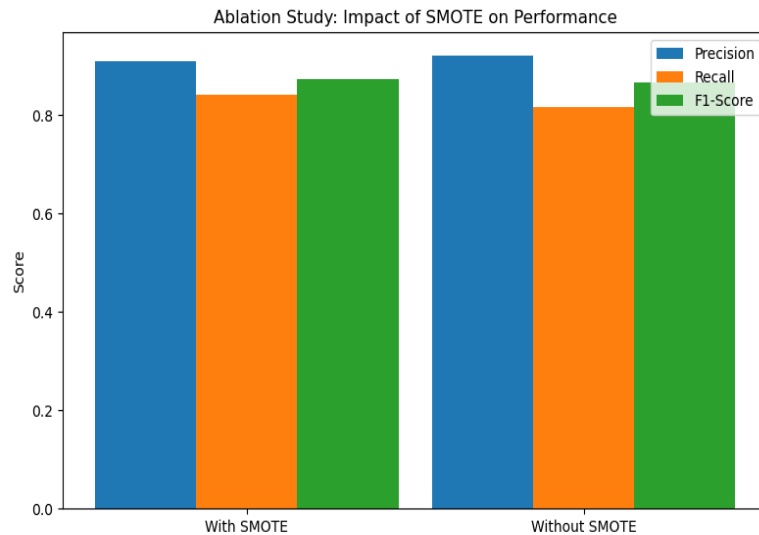


Figure 3.1: Impact of SMOTE on Performance

3.4 Baseline Model Comparison (5-Fold Cross-Validation)

To benchmark the proposed CNN model against traditional machine learning approaches, its performance was compared with Random Forest and XGBoost using the same 5-fold stratified cross-validation procedure. The results are presented in Table 3.4.

Table 3.4: Baseline Comparison

Model	Mean_f1	Std_f1	Wilcoxon p-value vs CNN
RF	0.8416	0.0188	0.043
XGBoost	0.8630	0.0180	0.138
CNN	0.8743	0.0150	-

The selection of a 1D-CNN architecture was empirically validated through a comparative analysis with traditional tree-based benchmarks (Random Forest and XGBoost). As shown in Table 3.4, the CNN achieved the highest mean F1-score (0.8743) and the lowest performance variance ($\sigma = 0.0150$). While performance was statistically comparable to XGBoost ($p = 0.138$), the CNN demonstrated a statistically significant improvement over Random Forest ($p = 0.043$). This

suggests that the convolutional filters effectively captured complex, non-linear dependencies within the payroll feature set such as the subtle interplay between salary deviations and grade-level norms that are less effectively modeled by the axis-aligned splits of decision trees.

3.5 Sensitivity Analysis

To evaluate the robustness of the proposed CNN model to variations in fraud prevalence, a key assumption underlying the simulation process, a sensitivity analysis was conducted. New datasets were generated using the MinistryFinancePayrollSimulator with fraud injection rates ranging from 0.02 to 0.05. As shown in Table 3.5, the model maintained a stable F1-score across these rates (mean = 0.788, standard deviation = 0.027).

Table 3.5: Sensitivity Analysis – F1-Score at Different Fraud Rates

Fraud Rate	F1
0.02	0.810416
0.03	0.792236
0.04	0.747186
0.05	0.800506

Performance was highest at lower fraud rates (0.02–0.03), with a slight dip at 0.04 before recovering at 0.05. The low standard deviation confirms the CNN’s resilience to changes in fraud prevalence, supporting its practical applicability in real-world Nigerian Public Financial Management systems such as IPPIS and GIFMIS, where ghost worker rates can fluctuate significantly across ministries and departments [7].

3.6 Confusion Matrix

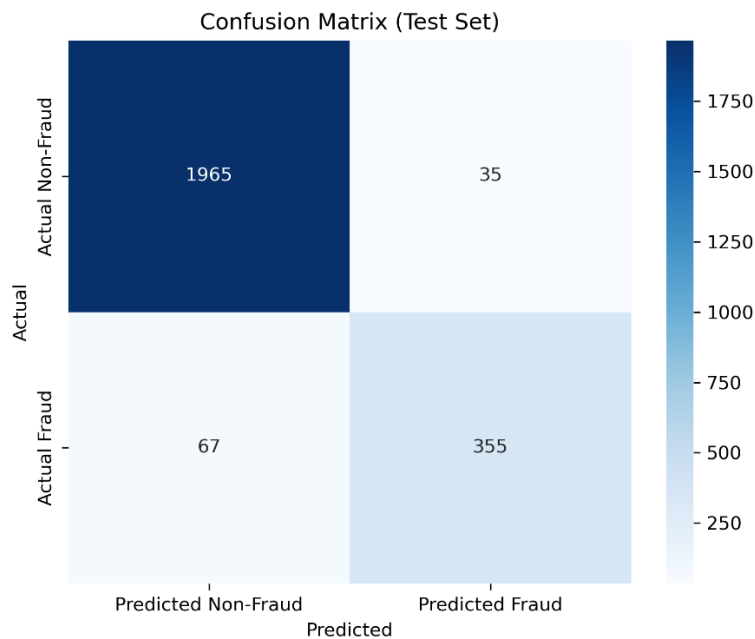


Figure 3.2: Confusion Matrix Showing the Classification Performance of the CNN Model on the Simulated Nigerian Payroll Test Set

The confusion matrix on the test set (Figure 3.2) reveals strong classification performance: 1,965 true negatives (correctly identified non-fraud cases) and 355 true positives (accurately detected fraudulent entries), with only 35 false positives and 67 false negatives. This translates to excellent specificity (low false alarm rate of $\sim 1.7\%$), minimizing unnecessary investigations of legitimate employees, and good sensitivity (high fraud detection rate of $\sim 84\%$), enabling the capture of most ghost workers, salary inflation, and duplicate payments.

In the context of Nigeria's Public Financial Management systems such as IPPIS and GIFMIS, where ghost-worker schemes alone cost over ₦100 billion annually, this balance is highly practical. The low false-positive rate reduces auditor workload, while the strong true-positive rate supports proactive fraud prevention and significant cost savings. Overall, the matrix confirms the CNN model's suitability for real-time payroll monitoring, offering a reliable and deployable solution for enhancing transparency and accountability in public sector financial management.

3.6 CNN Validation using Bootstrapping

To rigorously assess the stability and reliability of the CNN model's performance, bootstrapping was performed by drawing 1,000 random samples (with replacement) from the test dataset. This resampling technique generates 95% confidence intervals (CIs) for each metric, providing a robust estimate of variability that is particularly important in imbalanced classification tasks such as fraud detection [17].

Table 3.6: Bootstrapped 95% Confidence Intervals (1,000 resamples)

Metric	Mean	Lower 95% CI	Upper 95% CI
F1-score	0.873777	0.849346	0.897959
Precision	0.909927	0.880948	0.936226
Recall	0.840818	0.802842	0.874717
Accuracy	0.957928	0.950041	0.965328

The bootstrapped results demonstrate the high stability of the CNN model, with narrow confidence intervals across all metrics. For instance, the F1-score lies within [0.8493, 0.8980] with 95% confidence. This low variability indicates that the reported performance is consistent and not due to a particular train-test split, confirming the model's robustness and suitability for real-time payroll fraud monitoring in Nigeria's Public Financial Management systems such as IPPIS and GIFMIS.

3.7 SHAP Feature Importance

Table 3.7: SHAP Feature Values

Feature	Mean Abs SHAP
RegularHours	0.136672125
Department Revenue	0.131452791
GradeLevel GL 12	0.126035755
Department Procurement	0.120754034

Department_Budget	0.118307228
-------------------	-------------




Figure 3.3: SHAP Feature Impact on Fraud Prediction

From the beeswarm plot, RegularHours (the strongest signal) shows that the blue dots appear on the far right. This indicates that low regular hours (blue) are a strong driver of fraud prediction. If an employee is receiving a full salary but reporting almost zero regular hours, the model flags this as suspicious. For departments such as Revenue or Procurement, being in these departments (high/red) generally shifts the prediction slightly toward fraud (right side), whereas not being in these departments (low/blue) clusters near the neutral center.

3.8 Evaluation Metric

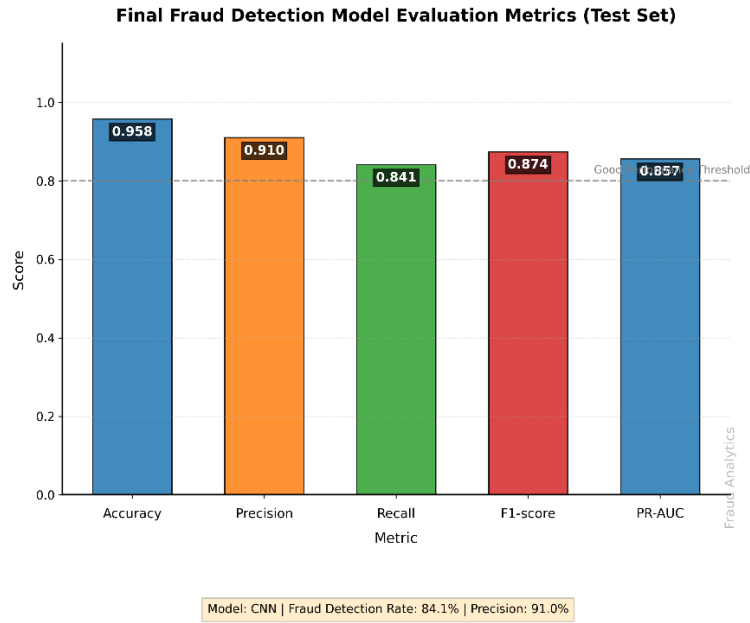


Figure 3.4: Combined Evaluation Metrics

3.8.1 Other Evaluation Curves

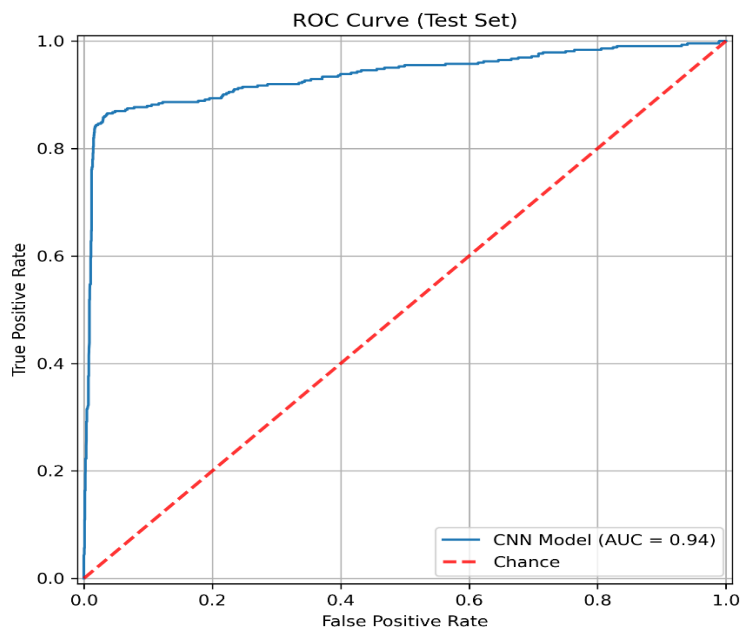


Figure 3.5: Receiver Operating Characteristic Curve

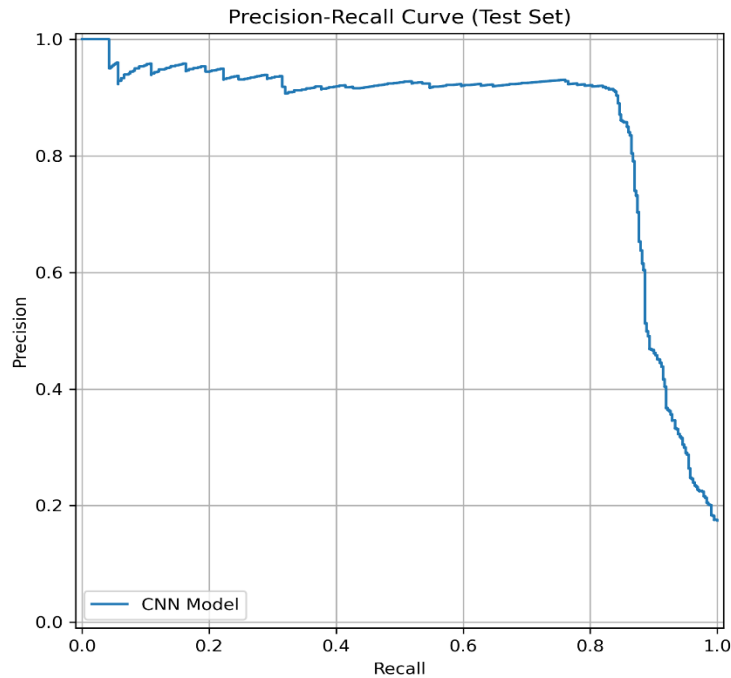


Figure 3.6: Precision Recall Curve

The ROC curve of the CNN model in Figure 3.5 demonstrates strong discriminative capability in detecting fraudulent transactions. The model achieved an AUC score of 0.94, indicating excellent classification performance. The curve rises sharply toward the top-left corner, suggesting that the model can correctly identify a large proportion of fraud cases while maintaining a low false positive rate. This performance shows that the proposed deep learning model is effective for detecting anomalies and potential fraud within public financial management systems.

The Precision-Recall curve in Figure 3.6 demonstrates the strong trade-off achieved by the CNN model. High precision is maintained even as recall increases, resulting in a high PR-AUC. This balance is particularly valuable in Nigerian PFM systems, where high precision reduces unnecessary audit workload while good recall helps recover significant fraud-related losses.

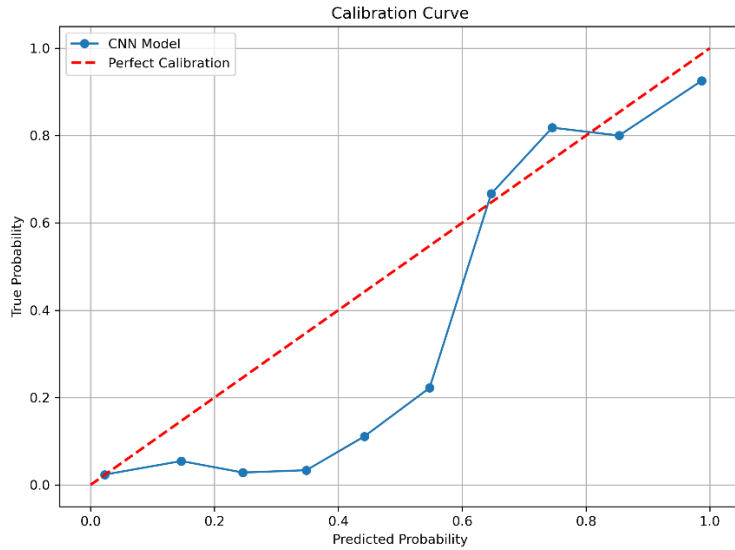


Figure 3.7: Calibration Curve

The curve shows that the CNN model is not perfectly calibrated but reasonably reliable at higher probabilities.

At low predicted probabilities (0.0–0.4), the model is overconfident, because it predicts higher fraud risk than what is actually observed (true probabilities are close to zero). Around the mid-range (0.6–0.8), the model becomes better aligned with the ideal diagonal, meaning predictions are more reliable in this region. At higher probabilities (0.8–1.0), the curve is very close to perfect calibration, indicating strong reliability when the model predicts high fraud risk.

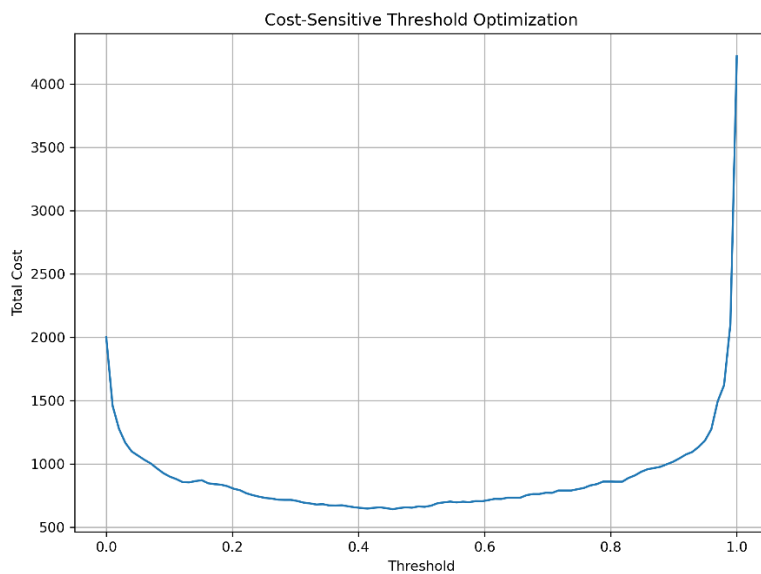


Figure 3.8: Cost-Sensitive Curve

At low thresholds (0.0–0.3), the model is over-sensitive, leading to a high total cost driven by False Positives (FP). This represents the operational overhead of investigating legitimate transactions. At the mid-range (0.4–0.5), the model reaches its optimal threshold at = 0.455. This is the Optimal Basin where the cost of investigating false alarms and the cost of missing fraud are perfectly balanced, resulting in maximum economic efficiency. At high thresholds (0.6–1.0), the model becomes under-sensitive, causing the cost to spike vertically. This is driven by False Negatives (FN).

3.9 Real-TimeAlert Dashboard

An example of a triggered alert, as generated by the system, is shown below:

FRAUD ALERT

Employee ID: MOF0685

Department: Investment

Grade Level: GL 14

Position: Assistant Director Revenue

Salary: ₦440,357.35

To bridge the gap between model development and practical deployment, the trained CNN was integrated into a real-time fraud alert dashboard. The system continuously monitors incoming payroll records, computes fraud probability using the optimal threshold (0.5973), and instantly triggers alerts when suspicious transactions are detected.

SHAP (SHapley Additive Explanations) was used to identify key features influencing fraud predictions. The most influential variables included SalaryDeviation, GradeLevel, and EmploymentStatus, indicating that the model aligns with domain expectations.

4. DISCUSSION

The results demonstrate that Convolutional Neural Network (CNN) models, when combined with Synthetic Minority Oversampling Technique (SMOTE), are effective for detecting payroll fraud in highly imbalanced public sector datasets. The model achieved a high classification performance with an Area Under the ROC Curve (AUC) of 0.94, indicating strong discriminative ability between fraudulent and legitimate payroll transactions. The ROC curve further shows that the model maintains a high true positive rate while keeping the false positive rate relatively low, which is critical for fraud detection systems where excessive false alerts can overwhelm auditors.

Compared to traditional rule-based fraud detection systems commonly used in Nigerian Public Financial Management (PFM) platforms, the proposed approach provides greater adaptability and the ability to capture complex, non-linear patterns in financial data. This capability allows the model to detect subtle anomalies that may not be easily identified using manually defined rules.

The integration of the trained model into a real-time fraud alert dashboard enhances the practical relevance of the framework by translating model predictions into actionable insights for auditors and policymakers. Such a system can support proactive monitoring and enable faster identification of suspicious payroll activities.

However, the study has certain limitations. The dataset used in this research was synthetically generated to simulate Nigerian payroll data, which may not fully capture the complexity and variability of real-world PFM systems. Therefore, further validation using actual government payroll datasets would be necessary to fully assess the model's operational effectiveness in live environments.

5. CONCLUSION

This study presented an AI-driven framework for detecting payroll fraud in Nigeria's Public Financial Management system using Convolutional Neural Networks. By addressing the challenge of class imbalance through SMOTE, the proposed model achieved strong predictive performance, with an AUC score of 0.94, demonstrating its effectiveness in distinguishing fraudulent transactions from legitimate ones.

The model was successfully integrated into a real-time fraud alert dashboard, illustrating how machine learning predictions can be translated into practical monitoring tools for public sector auditing and financial oversight. The findings suggest that deep learning techniques can significantly enhance transparency, accountability, and proactive fraud monitoring within public financial systems.

Future research will focus on pilot deployment within live PFM platforms, the incorporation of real administrative datasets, and the exploration of temporal deep learning models such as LSTM or Transformer-based architectures to capture evolving fraud patterns over time.

Acknowledgment

The author gratefully acknowledges his supervisor for proposing the research topic and for providing valuable academic guidance, constructive feedback, and continuous support throughout the course of this study. The author also acknowledges the Department of Statistics, University of Lagos, for providing a conducive research environment.

Authors' Contributions

The author and his supervisor jointly contributed to the conception and design of the study. The author carried out data generation, model development, implementation, experimentation, and analysis of results. The supervisor provided academic guidance, methodological oversight, and critical review of the manuscript. Both authors approved the final version of the manuscript.

Authors' Conflicts Of Interest

The author declares that there are no conflicts of interest regarding the publication of this paper.

Funding Statement

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

REFERENCES

- [1] Abdulkreem, R. Z., Ameen, S. Y., Aziz, K. J., & Salih, A. A. (2022). Fraud detection in credit card transaction data using machine learning techniques. *International Journal of Computer Science (IJCS)*, 13(3), 1–9.
- [2] Adhikari, P., Hamal, P., & Baidoo Jnr, F. (2024). Artificial intelligence in fraud detection: Revolutionizing financial security. *International Journal of Science and Research Archive*, 13(01), 1457–1472.
- [3] Brownlee, J. (2020). *Imbalanced classification with Python*. Machine Learning Mastery.
- [4] BudgIT. (2022, July 21). *Consequence management: A neglected aspect of Nigerian public financial management*.
- [5] Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321–357.
- [6] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794).
- [7] CSO Coalition. (2024). *Report on fraud vulnerabilities in GIFMIS and IPPIS systems*. Abuja, Nigeria.
- [8] Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2019). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences*, 479, 448–455.
- [9] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- [10] Gu, J., Wang, Z., Kuen, J., et al. (2018). Recent advances in convolutional neural networks. *Pattern Recognition*, 77, 354–377.
- [11] Kwaga, V. (2023, September 22). *Consequence management: A neglected aspect of Nigerian public financial management*. BudgIT Foundation.
- [12] Makinde, A., Ojo, T., & Omotayo, A. (2021). Machine learning models for budget anomaly detection in Nigerian MDAs. *African Journal of Data Science*, 3(2), 44–55.
- [13] Odufisan, O. I., Abhulimen, O. V., & Ogunti, E. O. (2025). Harnessing artificial intelligence and machine learning for fraud detection and prevention in Nigeria. *Journal of Economic Criminology*, 7, 100127.
- [14] O'Shea, K., & Nash, R. (2015). An introduction to convolutional neural networks

- [15] Rotimi, O., Olusola, I. E., Olusegun, E. A., Oluwayemisi, A. M. B., Rildwan, O. B., Rahmon, T. A., & Gbenga, A. C. (2021). Public financial management tools and performance in Nigeria public sector. *Academy of Accounting and Financial Studies Journal*, 25(S4), 1–15.
- [16] Salaudeen, L. G., Gabi, D., Garba, M., & Suru, H. U. (2024). Deep convolutional neural network based synthetic minority over-sampling technique: A defending model for fraudulent credit card transactions in financial institutions. *Journal of the Nigerian Society of Physical Sciences*, 6, 112–125.
- [17] Scott M. Lundberg & Su-In Lee (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems (NeurIPS)*.
- [18] Kuhn, M., & Johnson, K. (2013). *Applied predictive modeling*. Springer.

KENNETH EBUKA ASOGWA*
DEPARTMENT OF STATISTICS, UNIVERSITY OF LAGOS, AKOKA, LAGOS STATE, NIGERIA.
E-mail address: 249075016@live.unilag.edu.ng

MUMINU OSUMAH ADAMU
DEPARTMENT OF STATISTICS, UNIVERSITY OF LAGOS, AKOKA, LAGOS STATE, NIGERIA.
E-mail address: madamu@unilag.edu.ng