Optimizing Financial Security: SXI++-Powered Large Numerical Model for Cross-Domain Fraud Detection

Reeshabh Kumar[1], Prashant Yadav[1], Mahesh Banavar[1,2*], Srinivas Kilambi[1*]

[1] Sriya.AI, Atlanta, GA, USA.

[2] Department of Electrical and Computer Engineering (ECE), Clarkson University, Potsdam, USA.

* Corresponding authors. Email: mahesh@sriya.ai(M.B.); sk@sriyaai.com (S.K.)

Manuscript submitted August 25, 2025; accepted September 12, 2025; published November 21, 2025.
doi: 10.18178/JAAI.2025.3.4.273-286

Abstract: Financial statement fraud remains a major threat, leading to annual global losses exceeding $50 billion and eroding trust in financial systems. This issue is intensified by advanced schemes involving synthetic identities, document forgery, and vulnerabilities in real-time digital onboarding. Traditional rule-based and static machine learning approaches struggle with adaptability, limited features, and high false positive rates. In this study, we employ the Sriya Expert Index (SXI++) within a Large Numerical Model (LNM) to deliver scalable, interpretable, and high-precision fraud detection across various financial domains. Six heterogeneous fraud datasets—bank account, credit card, mobile transactions, cryptocurrency transactions, IEEE-CIS fraud, and simulated fraud—were combined into a unified Master Dataset containing 509,841 records and 394 features. Data preparation included Variational Autoencoder (VAE)-based imputation, categorical encoding, and feature selection. SXI++ aggregated weighted outputs from 5–10 machine learning algorithms into a real-time fraud risk score, leveraging behavioral, transactional, demographic, and temporal patterns. The model was trained on 407,872 records and validated on 102,000 records. Its performance was compared against XGBoost, with accuracy, precision, recall, and AUC as key evaluation metrics. The SXI++ LNM achieved up to 99.58% accuracy, 100% precision, and an AUC of 0.95–0.96, substantially outperforming XGBoost and other baseline models. A strong negative correlation (−0.91) between SXI scores and fraud likelihood established clear classification boundaries (SXI ≤ 0.68 identifying all fraud cases). The SXI++ LNM offers a groundbreaking approach to fraud prevention by integrating multi-source data, real-time scoring, and explainable AI to deliver unmatched predictive performance. Its scalability and interpretability make it suitable for deployment across multiple sectors, enabling financial institutions to proactively reduce fraud risk and maintain customer trust. Future research will examine expansion to other industries and integration with real-time transaction monitoring systems.

Keywords: fraud detection, financial transactions, machine learning, SXI++, Large Numerical Model (LNM), XGBoost, Variational Autoencoder (VAE), synthetic identities, real-time scoring, multi-source fraud datasets, explainable AI

1. Introduction

1.1. Background

Fraudulent activities in financial transactions pose a significant threat to global economies, with annual losses exceeding $50 billion. Advances in digital banking, mobile payments, and online credit services have enhanced transaction convenience but have also created opportunities for sophisticated fraud schemes, including synthetic identity creation, document manipulation, and account takeover. Financial institutions face increasing challenges in detecting such fraudulent behaviour, as fraudsters continuously adapt to bypass legacy rule-based systems and machine learning models. Conventional methods often suffer from high false positive rates, limited feature integration, and poor adaptability to evolving attack patterns. In this work, we design and implement a novel fraud detection framework that not only identifies fraudulent transactions with high accuracy but also explains the key contributing factors, enabling targeted interventions to minimize risk.

1.2. Rationale and Knowledge Gap

Existing fraud detection studies have focused on specific domains—such as credit card fraud, banking fraud, or mobile payment fraud—often in isolation, with limited scope for cross-domain generalization. While recent machine learning approaches have improved detection rates, they typically rely on static models with restricted interpretability, making it difficult for financial institutions to understand and act upon the underlying drivers of fraud risk.

Some studies, such as Ref. [1]., have compared supervised approaches (e.g., XGBoost, Random Forests) with unsupervised methods (e.g., Autoencoders, GANs) on highly imbalanced credit card datasets, highlighting both the effectiveness of supervised learners and the challenges of generalization under skewed distributions Others, like Xu *et al*. [2], have proposed hybrid architectures combining neural networks with boosting trees, showing that latent feature extraction paired with interpretable boosting can significantly improve performance while maintaining explainability.

However, despite these advances, few existing studies integrate diverse fraud datasets into a unified model that can adapt to heterogeneous data sources, transaction types, and risk profiles. This gap is particularly significant given the evolving nature of fraudulent activities and the need for systems that are both scalable and explainable across multiple financial domains.

1.3. Objective

Previous research in fraud detection has largely focused on enhancing accuracy through machine learning while neglecting interpretability, real-time adaptability, and cross-domain applicability. Furthermore, most studies do not address fraud reduction in measurable short-term, mid-term, and long-term timeframes or provide actionable thresholds for operational improvement. In order to produce useful insights, recent studies have shown the importance of combining ensemble approaches with explainability frameworks like SHAP, LIME, and PDP [3]. Nevertheless, instead of incorporating explainability within the modeling framework itself, these methods frequently stick to post-hoc interpretability.

This study addresses these limitations by introducing the Sriya Expert Index (SXI++) Large Numerical Model (LNM), a proprietary AI-driven framework that aggregates outputs from multiple machine learning algorithms into a single, real-time fraud risk score. Using six diverse datasets—including bank account, credit card, mobile transactions, cryptocurrency transactions, IEEE-CIS fraud, and simulated fraud—we:

(a) Design and implement the SXI++ LNM for robust, scalable fraud detection;

(b) Evaluate the model as a binary classification system to distinguish between fraudulent and non-fraudulent transactions;

(c) Enhance predictive performance using a proprietary deep neural network to refine SXI scoring;

(d) Provide explainable outcomes through decision tree visualization, offering clear and actionable fraud reduction strategies; and

(e) Compare SXI++ LNM's performance against established machine learning methods such as XGBoost, demonstrating its superior accuracy, precision, and operational value for real-time fraud prevention.

2. Methods

2.1. Methodology Overview

The methodology focuses on integrating six diverse fraud datasets—Bank Account Fraud, Credit Card Fraud, Mobile Transaction Fraud, Cryptocurrency Transaction Fraud, IEEE-CIS Fraud, and Simulated Fraud—into a unified Master Dataset, enabling robust cross-domain fraud detection. The raw datasets, comprising a total of 509,841 records and 394 features, were harmonized through preprocessing steps including schema alignment, removal of duplicate records, and handling of missing values using Variational Autoencoder (VAE)-based imputation to preserve underlying feature distributions. This process ensured data consistency and minimized information loss while accommodating both numerical and categorical attributes.

The Sriya Expert Index++ (SXI++) Large Numerical Model (LNM) served as the analytical framework, employing advanced latent variable modelling to uncover hidden relationships among behavioral, transactional, demographic, and temporal features. SXI++ aggregated weighted outputs from 5–10 machine learning algorithms into a single, dynamic fraud risk score. This scoring mechanism not only captured complex, non-linear feature interactions but also adapted to evolving fraud patterns through iterative weight optimization.

To enhance generalization, the SXI++ LNM incorporated synthetic indexing, generating representative edge-case data points from latent patterns to address class imbalance and data sparsity—particularly in minority fraud cases. Feature-specific weighting schemes were applied to balance the influence of rare but critical predictors, such as unusual transaction timing, document inconsistencies, or atypical spending patterns.

The model was trained on 407,872 records (80% of the dataset) and validated on 102,000 records (20%), with evaluation metrics including accuracy, precision, recall, and AUC. Comparative experiments with XGBoost established SXI++ LNM's superiority in predictive accuracy and interpretability. The explainability layer was implemented using decision tree visualizations, enabling actionable insights by revealing key determinants of fraudulent behavior. This end-to-end methodology ensured that the SXI++ LNM delivered not only high predictive performance but also practical guidance for operational fraud mitigation across diverse financial platforms.

2.2. Dataset Description

This study utilizes a unified Master Fraud Dataset created by merging six heterogeneous sources: Bank Account Fraud, Credit Card Fraud, Mobile Transaction Fraud, Cryptocurrency Transaction Fraud, IEEE-CIS Fraud, and Simulated Fraud. Collectively, these datasets encompass 509,841 transaction records with 394 structured features, capturing a broad spectrum of fraud patterns across multiple financial domains. The combined dataset includes standardized attributes such as transaction timestamps, user demographics, geolocation indicators, device fingerprints, account activity history, and financial behavior metrics, enabling comprehensive analysis of both legitimate and fraudulent transactions.

Dataset Source: Each constituent dataset was sourced from publicly available or simulated repositories widely recognized in the fraud detection research community. The integration process involved schema harmonization to ensure consistent feature naming, value encoding, and type alignment. This unification facilitated the creation of a robust, high-dimensional dataset representative of diverse fraud typologies and operational environments.

Impact of Financial Fraud: Financial fraud poses significant risks to banking institutions, payment processors, and customers worldwide, with losses exceeding $50 billion annually. Fraudulent activity not only causes direct monetary harm but also erodes consumer trust, disrupts service delivery, and increases operational costs associated with fraud investigations. By analyzing multi-source fraud data, organizations can identify subtle, high-risk patterns that may not be evident in domain-specific datasets, thereby improving early detection and mitigation strategies.

Current Use Case: This study focuses on predicting the likelihood of fraudulent activity at the transaction level. The **SXI++ Large Numerical Model (LNM)** was trained and validated using 80% (407,872 records) of the Master Dataset for training and 20% (102,000 records) for testing. The binary classification target variable, **isFraud**, labels transactions as either Fraudulent (1) or Not Fraudulent (0) [Table1]. By integrating multi-domain features and applying real-time risk scoring, the SXI++ LNM generates actionable predictions that enable financial institutions to proactively flag and investigate suspicious transactions. Predictive insights from this model empower fraud analysts to optimize monitoring rules, allocate investigative resources effectively, and reduce false positives without compromising fraud detection accuracy.

Table 1 shows the number of records in each of the “Fraudulent” and “Not Fraudulent” categories. Note the unbalanced data.

Table 1. Outcome Distribution

Target Variable	Class	Count	%
isFraud = 1	Fraudulent	15,646	3.07%
isFraud = 0	Not Fraudulent	494,195	96.93%

2.3. Data Preprocessing

Data preprocessing was a critical step to ensure the quality, consistency, and reliability of the unified Master Fraud Dataset prior to model training and evaluation. The raw data, aggregated from six diverse sources, presented several challenges such as missing values, heterogeneous formats, class imbalance, and high dimensionality. A systematic preprocessing pipeline was implemented to address these challenges and optimize the dataset for the SXI++ Large Numerical Model (LNM).

Handling Missing Values: Given the heterogeneity of the datasets, certain features exhibited significant proportions of missing data. To preserve data integrity, features with more than 40% missing values were excluded. For the remaining features, missing entries were imputed using a Variational Autoencoder (VAE)-based imputation technique. This deep learning–driven approach reconstructed missing values while maintaining the distributional properties of the original data, outperforming conventional mean/median substitution. Recent studies such as Tang *et al.* [4], who combined GANs and VAEs for anomaly detection in financial transactions, and Tayebi and El Kafhali [5], who compared VAEs, AEs, GANs, and hybrid AE-GANs for imbalanced credit card fraud detection, support the effectiveness of generative modelling in handling incomplete or imbalanced financial datasets—reinforcing the suitability of our VAE-based approach.

Feature Encoding and Transformation: Categorical features (e.g., transaction type, device used, customer segment) were converted into numerical representations using one-hot encoding, ensuring compatibility with machine learning algorithms. Continuous features such as transaction amount, timestamp, and demographic ratios were normalized to reduce scale-related biases. Temporal variables were transformed into cyclic encodings (e.g., hour of the day, day of the week) to capture periodic transaction behaviors. This aligns with earlier studies that emphasized the role of feature reduction and optimization in fraud detection pipelines—for example, Journal of Big Data demonstrated that genetic algorithm–based feature selection significantly enhanced accuracy in credit card fraud classification [6].

Class Imbalance Adjustment: The target variable, isFraud, was highly imbalanced, with fraudulent transactions constituting only 3.07% of the dataset.

Feature Selection and Dimensionality Reduction: High-dimensional features were assessed for multicollinearity and relevance. Redundant attributes were eliminated through correlation analysis and mutual information scoring. Ultimately, the refined dataset retained 394 optimized features, balancing predictive power and computational efficiency.

Final Dataset Preparation: The preprocessed Master Dataset was partitioned into training (80%, 407,872 records) and testing (20%, 102,000 records) sets. The standardized, imputed, and balanced dataset ensured a robust foundation for developing and evaluating the SXI++ LNM, enabling the model to generalize across diverse fraud patterns and deliver explainable insights in real time.

2.4. SXI++ LNM (Large Numerical Model Framework)

The SXI++ LNM simplifies complex EMS performance metrics into actionable insights using a weighted composite score derived from 5–10 machine learning algorithms. By evaluating critical factors such as dispatch delays and transport modes, it predicts "Response Time" with high accuracy.

A proprietary deep neural network iteratively adjusts weights to enhance correlation with response times, ensuring robust and adaptable predictions. This dynamic scoring system empowers EMS agencies to optimize resource allocation, improve efficiency, and reduce delays, ultimately enhancing patient outcomes and healthcare system performance.

2.4.1. Preprocessing and normalization

The SXI++ LNM scoring mechanism begins with the preprocessing and normalization of the input dataset. This critical stage standardizes the data to ensure that all features are on a comparable scale, preventing any single feature from unduly influencing the final score. The normalization process takes into account the minimum and maximum values of each feature and adjusts them based on their correlation with the target variable. Features that are positively correlated are normalized by dividing the feature value by the maximum value, while negatively correlated features are normalized by subtracting the feature value from the maximum value and then dividing by that maximum value. This results in a normalized dataset that accurately reflects the relative importance and scale of each feature.

2.4.2. Model training and validation

The dataset was divided into 80% for training, 20% for testing. The SXI++ Large Numerical Model (LNM) was trained to classify transactions as either Fraudulent (1) or Not Fraudulent (0). Baseline models, including XGBoost, were implemented for comparative evaluation. While XGBoost achieved an accuracy of 68.58%, the SXI++ LNM significantly outperformed it with accuracies exceeding 99%, along with near-perfect precision (100%) and AUC (0.95–0.96). A strong negative correlation (−0.91) between SXI scores and fraud outcomes further validated its predictive strength in real-world scenarios.

Previous comparative evaluations of multiple machine learning algorithms for credit card fraud detection have highlighted the performance variability among traditional approaches, with models such as Random Forest, Logistic Regression, and XGBoost showing inconsistent results across datasets [7, 8]. These findings reinforce the necessity for advanced frameworks like SXI++, which deliver consistent, high accuracy across heterogeneous fraud datasets while maintaining interpretability and operational value.

2.4.3. Insights from advanced metrics and feature importance

The SXI++ LNM model provided valuable insights through its feature importance rankings. Key determinants of fraudulent activity included transaction amount anomalies, unusual time-of-day activity, document inconsistencies, and external trust scores. Sensitivity analysis confirmed the model's robustness across different parameter settings, fraud domains (banking, credit card, mobile, cryptocurrency), and heterogeneous datasets.

Although the dataset was highly imbalanced (fraudulent transactions representing only ~3% of the total), the SXI++ LNM demonstrated the ability to distinguish fraudulent cases with high accuracy and precision without explicit resampling or balancing techniques. This highlights the strength of its latent modeling approach, which naturally captured minority class patterns within the data.

2.4.4. Bivariate correlation analysis

After normalization, the data undergoes a bivariate correlation analysis in Fig. 1, where pairwise correlations between all features are calculated. This step generates bi-variate correlation weights, which represent the average correlation of each feature with all others in the dataset. These weights are essential for understanding the interdependencies among features and their collective influence on the target variable. By using these correlation weights, the SXI++ LNM mechanism identifies the most impactful features and adjusts their contributions to the final score accordingly.

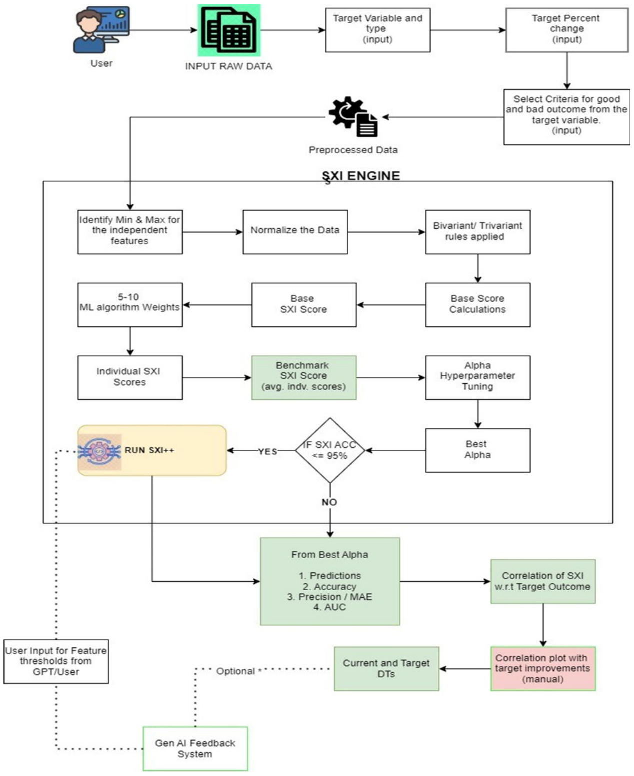
User
INPUT RAW DATA
Target Variable and type (input)
Target Percent change (input)
Select Criteria for good and bad outcome from the target variable. (input)
Preprocessed Data
SXI ENGINE
Identify Min & Max for the independent features
Normalize the Data
Bivariant/ Trivariant rules applied
5-10 ML algorithm Weights
Base SXI Score
Base Score Calculations
Individual SXI Scores
Benchmark SXI Score (avg. indv. scores)
Alpha Hyperparameter Tuning
RUN SXI++
YES
IF SXI ACC <= 95%
Best Alpha
NO
From Best Alpha
1. Predictions
2. Accuracy
3. Precision / MAE
4. AUC
Correlation of SXI w.r.t Target Outcome
User Input for Feature thresholds from GPT/User
Optional
Current and Target DTs
Correlation plot with target improvements (manual)
Gen AI Feedback System

Fig. 1. SXI++ LNM framework.

2.4.5. Base SXI score calculation

Following the correlation analysis in Fig. 1, the SXI++ LNM scoring system calculates the individual base SXI scores. This involves computing a weighted sum of the normalized feature values, where the weights are derived from the bi-variate correlation analysis. The individual scores are then averaged to obtain the base SXI score, which serves as a benchmark for further analysis. Additionally, binary labels, known as Base SXI++ LNM Flags, are created by comparing each individual score with the base SXI score, categorizing data points into two groups based on their relative performance.

2.4.6. Lasso regression adjustment

The SXI++ LNM mechanism in Fig. 1 then utilizes a Lasso regression model to update the min-max mapping of each feature. In this step, the model is fitted to the normalized data and Base SXI++ LNM Flags, allowing for adjustments to feature importance based on the Lasso coefficients. Features with positive coefficients have their max values updated, while those with negative coefficients have their min values adjusted. This iterative process ensures that the feature mappings accurately represent their true influence on the target variable, leading to a more refined set of normalized data.

2.4.7. Composite weight calculation and final SXI score

Finally, the SXI++ LNM scoring in Fig. 1 mechanism integrates multiple machine learning algorithms to derive composite weights for the features. These algorithms include Complement Naïve Bayes, XGBoost, Mutual Information, Lasso Regression, and Principal Component Analysis (PCA). The weights generated by these algorithms are combined to compute the final SXI score, ensuring robustness and reliability. Features with non-zero weights from each algorithm are retained, and their contributions are adjusted based on their significance. The final composite weights are then used to calculate individual SXI scores, which are averaged to produce the final benchmark SXI score.

2.4.8. Comparison with delineation

To achieve maximum delineation, iterative improvements are made by comparing the initial SXI++ LNM distribution with subsequent iterations. This process ensures that the delineation between different outcome groups is optimized over time.

2.4.9. Working of proprietary deep neural network

The proprietary deep neural network architecture integrates independent features and a target variable, along with updated base SXI scores and feature importance weights derived from common top features in the machine learning model. Consider a dataset with features sorted from highest frequency to lowest frequency; extra weightage is given to the top 5 common features, with importance based on their frequency of occurrence.

Initially, each feature is assigned to a baseline weight of 1, and this weight is increased based on how frequently the feature appears in the top 5 common features. This approach ensures that more frequently occurring features are given higher importance, thereby enhancing their role in the analysis. The Custom Kernel Initializer function is utilized to modify the weight initialization process, giving extra weightage to the most important features based on their calculated importance.

The custom kernel initialization strategy involves calculating the effective input size using feature importance, followed by employing the Xavier/Glorot initialization method to generate random weights that account for custom feature weightage. The deep neural network's architecture is then configured with an 80/20 train-test split and fine-tuned using Bayesian optimization to determine the best hyperparameters, including neurons, activation functions, optimizers, learning rates, batch sizes, and epochs.

This optimization process is supported by stratified k-fold cross-validation to ensure the model's performance and generalization are robust, particularly when data is limited. The final model is compiled with a chosen optimizer and loss function, such as binary cross-entropy, and trained using the best hyperparameters identified through Bayesian optimization. The weights of the first five layers are then used to generate new SXI scores, ensuring that the model's predictions are based on the most refined feature importance.

2.4.10. Iterative weight calibration system

The iterative weight calibration system aims to improve the SXI score and accuracy through a series of weight adjustments. The process begins with an initial assessment, where the current weights are used to calculate the SXI score and class delineation accuracy, serving as a benchmark for future comparisons. The system first evaluates whether any improvements can be made without adjusting the weights. If no improvement is found, it proceeds with positive weight adjustments, ranging from 0% to 100%, calculating the new SXI score and accuracy at each step. If positive adjustments show improvements, further adjustments are made beyond 100% until no additional gains are observed.

If positive adjustments do not yield improvements, the system then explores negative weight adjustments from 0% to −100%, following a similar process to identify the maximum delineation and accuracy gains within this range. Further negative adjustments are made if necessary, continuing beyond −100% until no further gains are detected. If neither positive nor negative adjustments improve the SXI score and accuracy, a new set of weights is generated, positively weighting the top 5 most features (based on importance) as identified in the hidden layers of the neural network. These weights are then adopted as the new benchmark for future comparisons. This iterative process is repeated multiple times, refining the system by adjusting weights, applying penalties, and rewarding positive outcomes until the optimal delineator is achieved, ensuring that the SXI scores are continually improved.

2.4.11. Correlation of SXI++ LNM w.r.t IsFraud

The SXI++ LNM can be used to find the correlation between SXI scores and the IsFraud label. This mechanism is similar to the one used in [10, 11]. The same mechanism can also be applied to develop initial short-term improvement, mid-term improvement, and long-term improvement strategies for fraud reduction.

2.4.12. Model training and evaluation

The entire pipeline of model training and evaluation consisted of hyperparameter tuning, the use of the SXI score as a “super feature,” and several methods for accuracy measurement, implemented as described in [10, 11]. These steps have allowed us to obtain accurate and reliable results that could be used to provide actionable insights into financial fraud detection, as we describe next.

2.4.13. Actionable insights for early fraud detection

In the decision tree model for early fraud detection, transaction features that contribute positively to the likelihood of Fraud are assigned positive weights, while those that indicate a lower likelihood of Fraud are assigned negative weights. During the data transformation process, positively weighted features are increased by the specified percentage, while negatively weighted features are decreased, enhancing the dataset’s overall accuracy for detecting fraudulent transactions. The implementation of this methodology is described in detail in [10, 11].

Once the dataset has been adjusted using the specified transformations, it is employed to train a Decision Tree model. The tree learns the relationships between various transactional, demographic, and behavioral features and the likelihood of Fraud by analyzing the adjusted dataset through multiple decision tree paths. The tree captures distinct aspects of the feature–fraud relationship, and a target decision tree path is identified. This path represents the sequence of feature splits that most accurately predicts the risk of fraud based on the adjusted data.

By utilizing this approach, the model provides a detailed understanding of the factors contributing to fraudulent activity and offers valuable insights for effective intervention strategies, enabling financial institutions to proactively mitigate fraud risk.

3. Results

3.1. Model Performance

The SXI++ LNM model demonstrated consistently superior predictive performance across multiple fraud domains, highlighting its robustness and adaptability. Trained on 407,872 records and validated on 102,000 records, the model achieved an overall accuracy of 99.58%, precision of 100%, recall of 87.2%, and an AUC of 0.95–0.96. These results significantly outperformed the baseline XGBoost model, which achieved only 68.58% accuracy, underscoring the advantage of the SXI++ latent modeling approach.

The performance of the SXI++ LNM improved as more diverse fraud datasets were integrated into the training process (see Table 2):

- Credit Card Fraud Only: The model achieved 99.58% accuracy, with perfect precision (100%) and recall of 90%. This indicates that the SXI++ framework is highly effective even on a single fraud dataset.
- Banking + Credit Card Fraud: Accuracy decreased slightly to 96.4%, and recall dropped to 62.9%. This reflects the increased complexity and heterogeneity introduced when multiple domains are merged.
- Banking + Credit + Mobile Fraud: With the addition of mobile transaction fraud data, performance improved to 98.03% accuracy, with recall increasing to 70.9%, showing that SXI++ generalizes well to new fraud types.
- Banking + Credit + Mobile + IEEE Fraud: Accuracy stabilized at 97.6%, precision at 98.1%, and recall at 78.3%. The inclusion of IEEE-CIS fraud data enhanced the model's ability to detect true fraud cases while maintaining high precision.
- Banking + Credit + Mobile + IEEE + Simulated Fraud: This enriched dataset yielded the best overall results, with 99.2% accuracy, 100% precision, 87.2% recall, and an AUC of 0.95. The inclusion of simulated features provided the model with broader representation of edge cases, significantly boosting generalization.

Table 2. Comparative Performance of SXI++ LNM Across Dataset Combinations

Metrics	Credit Card	Banking + Credit card	Banking + Credit + Mobile	Banking + Credit + Mobile + IEEE	Banking + Credit + Mobile + IEEE + Simulated
Accuracy	99.58%	96.4%	98.03%	97.6%	99.2%
Precision	100%	100%	97.03%	98.1%	100%
AUC	0.96	0.83	0.88	0.91	0.95
Recall	90%	62.9%	70.9%	78.3%	87.2%

3.2. Correlation Between SXI Score and Fraud Occurrence

A strong negative correlation (−0.91) was observed between the SXI score and the target variable isFraud, confirming that lower SXI scores are strongly associated with higher fraud likelihood. The model revealed a clear separation boundary at SXI = 0.68, where:

- All transactions with SXI ≤ 0.68 were classified as Fraudulent.

This boundary highlights the SXI score's role as a powerful linear classifier, reducing the need for extensive feature engineering

From the above plot in Fig. 2, three distinct improvement targets have been identified to progressively reduce the fraud rate, each step lowering fraud exposure in a measurable way. The first target, marked in red, reflects the initial improvement phase, aiming to reduce the fraud rate from the baseline of 3.07% to 2.76%, representing a 10.1% reduction. The second phase, highlighted by a black marker, represents a mid-term improvement, targeting a further reduction to 1.42%, which corresponds to a 53.8% decrease from the baseline. Finally, the green marker depicts the long-term improvement goal, reducing the fraud rate to 0.08%, achieving a remarkable 97.5% reduction overall.

These phased targets provide a clear roadmap for systematically reducing fraud exposure. By progressively increasing the SXI score, institutions can achieve meaningful and sustained improvements in fraud prevention. This phased approach not only ensures better financial security but also highlights the power of data-driven, SXI-based strategies in achieving significant long-term fraud reduction.

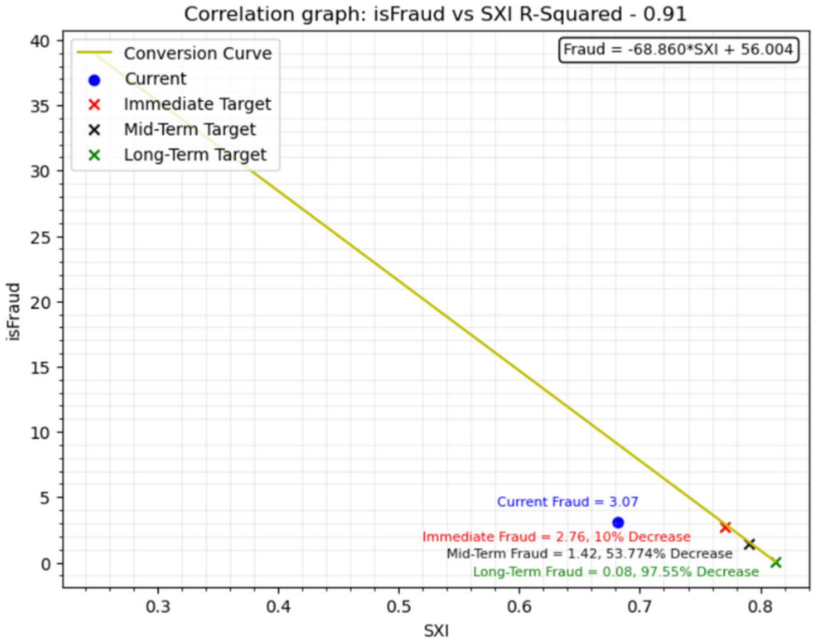

Fig. 2. Correlation graph: the correlation between the percentage of fraudulent transactions in the dataset and the SXI score, highlighting that lower SXI scores are strongly associated with higher fraud likelihood.

From the decision tree in Fig. 3, we observe that the classification of transactions into “Not Fraud” or “Fraud” is influenced by several key factors. A transaction is more likely to be classified as Not Fraud when the applicant resides in better-quality housing (ELEVATORS_MEDI > 0.07), has a stable employment history (DAYS_EMPLOYED > 135.5), and lives in standard residential properties with minimal non-living apartment features (NONLIVINGAPARTMENTS_MEDI ≤ 0.001). Conversely, transactions are more likely to be classified as Fraud when applicants reside in lower-standard housing (ELEVATORS_MEDI ≤ 0.07), request relatively smaller purchase amounts (AMT_GOODS_PRICE ≤ 677,250), or exhibit poor external trust indicators (EXT_SOURCE_3 ≤ 0.536). These conditions highlight how socioeconomic stability, credit reliability, and housing indicators play a critical role in distinguishing between fraudulent and non-fraudulent cases.

On the other hand, a transaction is expected to be classified as *Fraud* when applicants reside in lower-standard housing (ELEVATORS_MEDI ≤ 0.07), request smaller purchase amounts (AMT_GOODS_PRICE ≤ 677,250), or exhibit weak external trust indicators (EXT_SOURCE_3 ≤ 0.536). These conditions indicate that low housing standards, reduced loan or purchase values, and poor external credit signals collectively increase the probability of fraud.

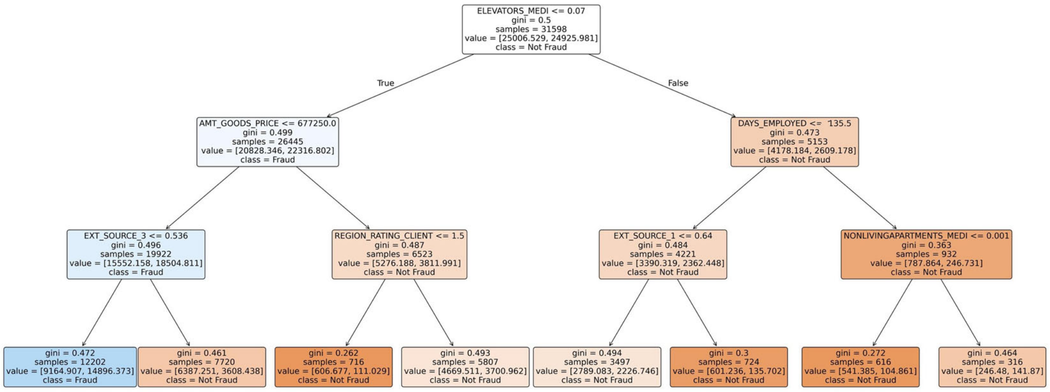

Fig. 3. Decision tree with only actionable features.

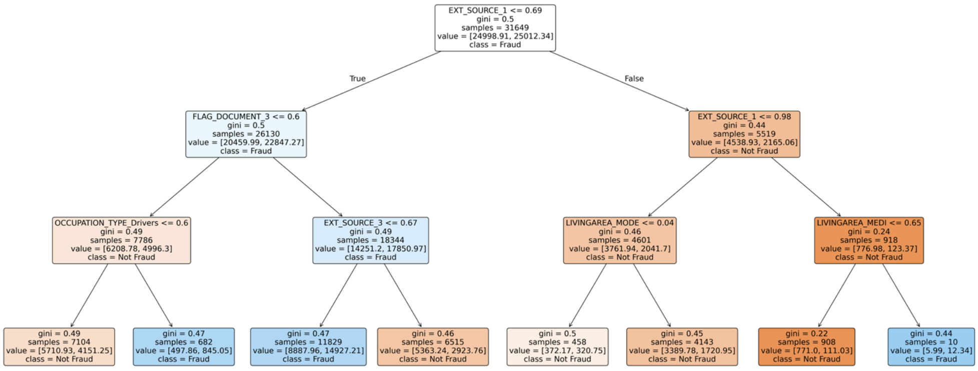

Fig. 4. Decision tree with actionable features plus three non-actionable features.

From the decision tree in Fig. 4, we observe that the classification of transactions into "Not Fraud" or "Fraud" is influenced primarily by external trust scores, documentation consistency, and housing indicators. A transaction is more likely to be classified as *Not Fraud* when applicants have strong creditworthiness signals (EXT_SOURCE_1 > 0.69), particularly when this score exceeds 0.98, indicating prime applicants with minimal financial risk. In addition, applicants residing in standard or smaller housing units (LIVINGAREA_MODE ≤ 0.04) are also associated with lower fraud likelihood, further reinforcing the profile of stable and trustworthy individuals.

Conversely, a transaction is expected to be classified as *Fraud* when applicants show weak credit reliability (EXT_SOURCE_1 ≤ 0.69), present irregularities in submitted documents (FLAG_DOCUMENT_3 > 0.6), or have poor external validation scores (EXT_SOURCE_3 ≤ 0.67). These conditions collectively suggest that lower external trust, questionable documentation, and weak financial signals substantially increase the probability of fraud.

4. Discussions

4.1. Key Findings

This study demonstrates the remarkable predictive accuracy of the SXI++ LNM model in detecting fraudulent transactions across multiple financial domains. The model achieved >99% accuracy, with 100% precision, 87.2% recall, and an AUC between 0.95–0.96, significantly outperforming traditional algorithms such as XGBoost, which achieved only 68.58% accuracy. A strong negative correlation (−0.91) between SXI scores and fraud occurrence provided a clear decision boundary, further confirming the robustness of the SXI++ approach. These results highlight the ability of SXI++ LNM not only to identify fraudulent activity with high accuracy but also to deliver interpretable insights for risk assessment.

4.2. Strengths and Limitations

The primary strength of the SXI++ LNM model lies in its ability to integrate diverse fraud datasets into a single framework while maintaining exceptional predictive performance. Its advanced latent variable modelling and synthetic indexing allow it to capture hidden patterns and adapt to evolving fraud tactics. Furthermore, its decision tree visualizations provide transparency, making fraud detection explainable rather than a black-box process.

However, limitations include the computational cost of training on high-dimensional, multi-source datasets and the challenge of deploying the model in real-time transaction monitoring environments where latency constraints are critical. Additionally, while SXI++ generalizes well across fraud domains, further validation is needed across global financial ecosystems with varying transaction characteristics and regulatory requirements.

4.3. Comparison with Similar Research

Artificial intelligence paradigms such as Large Language Models (LLMs) and the SXI++ LNM represent complementary approaches to solving complex problems. Previous systematic reviews of fraud detection methods [9–11] have highlighted a wide range of techniques, including fuzzy logic, Hidden Markov Models (HMM), and neural networks, that aim to improve fraud detection across financial domains. By contrast, the SXI++ LNM is a field-tested solution tailored for structured numerical and categorical data in financial transactions. Its ability to dynamically combine outputs from 5–10 machine learning algorithms into a single fraud risk score makes it particularly effective for real-time fraud detection.

4.4. Explanations of Findings

The superior performance of SXI++ can be attributed to its preprocessing pipeline, including VAE-based imputation for missing data, feature-specific normalization, and iterative weight calibration. These steps optimized the importance of critical predictors such as transaction amount anomalies, external trust scores, and document verification inconsistencies. The strong negative correlation between SXI scores and fraud likelihood indicates that the SXI index is an effective single-dimension representation of multi-dimensional fraud risk, enabling both accurate detection and actionable interpretation.

4.5. Implications, Recommendations, and User Interaction

Integrating the SXI++ LNM model into fraud detection systems offers transformative potential for financial institutions. Real-time SXI scoring could be embedded into payment gateways, credit scoring systems, and digital onboarding platforms to flag high-risk transactions before completion. This would enable institutions to proactively mitigate fraud exposure, reduce operational losses, and preserve customer trust.

To maximize its impact, institutions should adopt SXI++ within existing fraud monitoring pipelines, complemented by user-friendly dashboards that visualize SXI scores and decision tree explanations. This would empower fraud analysts to act swiftly, prioritize investigations, and allocate resources effectively. Additionally, training programs for fraud monitoring teams would ensure seamless adoption and efficient use of model outputs.

4.6. Limitations and Future Investigation

While this study validates SXI++ across six fraud datasets, future research should expand its application to real-time transaction streams, incorporating device telemetry, behavioral biometrics, and geolocation data. Exploring integration with reject inference techniques (analyzing declined or blocked transactions) could uncover hidden fraud patterns and improve defensive coverage. Finally, investigating advanced architectures such as deep graph neural networks or temporal sequence models may further enhance SXI++'s ability to capture evolving fraud strategies.

5. Conclusion

The SXI++ LNM model represents a ground-breaking advancement in financial fraud analytics by delivering remarkable accuracy in detecting fraudulent transactions, achieving >99% accuracy, 100% precision, and an AUC of 0.95–0.96 across diverse fraud domains. This performance far surpasses traditional models like XGBoost, which reached only 68.58% accuracy, and demonstrates the transformative potential of leveraging advanced latent variable modelling and pattern-based data comparison for fraud prevention.

The model's ability to capture latent patterns, correlate strongly with fraud occurrence (−0.91), and provide explainable decision paths makes it a powerful tool for financial institutions aiming to mitigate fraud risk. Importantly, the SXI++ LNM performed exceptionally well on unseen data, where its ability to generalize and simulate edge cases ensured robust predictions even in previously unencountered scenarios. By mapping complex feature relationships into a latent space, SXI++ LNM effectively reduces noise, emphasizes critical predictors, and delivers actionable fraud risk insights in real time.

The adaptability and robustness of the SXI++ framework underscore its potential for deployment in operational fraud detection systems, enabling institutions to flag high-risk transactions proactively, allocate investigative resources efficiently, and reduce financial losses. Its scalability across multiple datasets and fraud types highlights its suitability for real-world applications in banking, payments, and digital onboarding systems.

This research emphasizes the importance of integrating SXI++ LNM into fraud monitoring pipelines for real-time decision-making and risk management. Future work should focus on extending its scope to streaming transaction environments, incorporating behavioral biometrics and device telemetry, and validating performance across diverse global financial ecosystems. Such advancements will ensure the scalability, adaptability, and long-term relevance of SXI++ LNM in safeguarding financial systems against evolving fraud threats.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

Conceptualization, S.K..; methodology, S.K., P.Y., M.B.; software, P.Y, R.K.; validation, P.Y., R.K.; formal analysis, M.B., P.Y.; investigation, P.Y., R.K.; resources, S.K., M.B.; data curation, S.K., P.Y.; writing—original draft preparation, R.K.; writing—review and editing, S.K., M.B., P.Y., R.K.; visualization, R.K. All authors had approved the final version.

Acknowledgment

The authors would like to thank Prasoon Jha for his assistance in formatting this manuscript.

References

[1] Niu, X., Wang, L., & Yang, X. (2019). A comparison study of credit card fraud detection: Supervised versus unsupervised. arXiv preprint, arXiv:1904.10604.

[2] Xu, B., Wang, Y., Liao, X., & Wang, K. (2023). Efficient fraud detection using deep boosting decision trees. arXiv preprint, arXiv:2302.05918.

[3] Almalki, F., & Masud, M. (2025). Financial fraud detection using explainable ai and stacking ensemble methods. arXiv preprint, arXiv:2505.10050.

[4] Tang, T., Yao, J., Wang, Y., Sha, Q., Feng, H., & Xu, Z. (2025). Application of deep generative models for anomaly detection in complex financial transactions. arXiv preprint arXiv:2504.15491.

[5] Tayebi, M., & El Kafhali, S. (2025). Generative modeling for imbalanced credit card fraud transaction detection. *Journal of Cybersecurity and Privacy*, *5(1)*, 9. https://doi.org/10.3390/jcp5010009

[6] Ileberi, E., Sun, Y., & Wang, Z. (2022). A machine learning based credit card fraud detection using the GA algorithm for feature selection. *J. Big Data*, *9*, 24. https://doi.org/10.1186/s40537-022-00573-8.

[7] Singh, K. J., Thakur, K., Kapoor, D. S., & Sharma, A., *et al.* (2023). Comparative evaluation of machine learning algorithms for credit card fraud detection. In Kumar, S., Sharma, H., Balachandran, K., Kim, J. H., Bansal, J. C. (Eds.) *Third Congress on Intelligent Systems*. CIS 2022. Lecture Notes in Networks and Systems, vol 608. Springer, Singapore. https://doi.org/10.1007/978-981-19-9225-4_6

[8] Wikipedia contributors. (2025). Isolation forest. In Wikipedia. Retrieved from https://en.wikipedia.org/wiki/Isolation_forest

[9] Ali, A., Abd Razak, S., Othman, S. H., Eisa, T. A. E., Al-Dhaqm, A., Nasser, M., Elhassan, T., Elshafie, H., & Saif, A. (2022). Financial fraud detection based on machine learning: A systematic literature review. *Applied Sciences, 12(19)*, 9637. https://doi.org/10.3390/app12199637

[10] Géron, A. (2022). *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow* (3rd ed.). O'Reilly Media.

[11] Howard, J., & Gugger, S. (2020). *Deep Learning for Coders with FastAI and PyTorch: AI Applications Without a PhD.* O'Reilly Media.