

Genetic Algorithm Optimized LSTM-Light GBM Framework for Explainable Financial Fraud Detection

Ms. Arya M. Palve¹, Prof. Jayaprabha M. Kanase², Prof. Dr. B. D. Phulpagar³

¹*Pursuing MTech in Computer Engineering from Progressive Education Society's Modern College of Engineering, Pune*

²*Assistant Professor at Progressive Education Society's Modern College of Engineering, Pune*

³*Professor at Progressive Education Society's Modern College of Engineering, Pune*

Abstract—Financial fraud detection presents a continuous challenge for contemporary banking institutions due to high transactional velocity, extreme class imbalance, and the sophisticated evolving tactics of fraudsters. Traditional machine learning techniques often fail to capture complex temporal patterns across sequential transactions, while standalone deep learning architectures frequently suffer from opaque decision-making processes. To address these limitations, this paper proposes a novel Genetic Algorithm (GA) optimized hybrid architecture integrating Long Short-Term Memory (LSTM) networks with Light Gradient Boosting Machine (LightGBM) for explainable financial fraud detection. In the proposed framework, the LSTM network functions as an automated temporal feature extractor, generating low-dimensional deep sequential embeddings from transaction histories. These embeddings are subsequently concatenated with domain-engineered tabular features and fed into a LightGBM classifier. To optimize performance without manual fine-tuning, a multi-objective Genetic Algorithm automatically searches the hyperparameter space of both the LSTM and LightGBM components to maximize classification metrics while minimizing architectural complexity. Furthermore, post-hoc model interpretability is established by integrating SHapley Additive exPlanations (SHAP), providing global feature importance and local instance-level justifications for automated fraud alerts. Experimental evaluations conducted on benchmark financial datasets demonstrate that the proposed GA-LSTM-LightGBM framework significantly outperforms baseline models, achieving superior performance in terms of AUC-ROC, Precision-Recall AUC, and F1-Score, while maintaining minimal computational latency. The integration of high-performance sequential feature extraction with explicit post-hoc explainability offers a robust, regulatory-compliant solution for real-time financial fraud monitoring systems.

Index Terms—Financial Fraud Detection, Genetic Algorithm, Long Short-Term Memory (LSTM), LightGBM, Hybrid Deep Learning, Explainable AI (XAI), SHAP.

I. INTRODUCTION

In today's hyper-connected digital economy, electronic financial transactions—ranging from online banking and credit card payments to peer-to-peer transfers—have experienced unprecedented growth. While this digital shift has significantly enhanced transactional convenience and operational efficiency, it has simultaneously exposed financial institutions to increasingly sophisticated forms of fraud. Modern financial fraud operations are characterized by high volume, dynamic behavior, and temporal complexity, resulting in multibillion-dollar annual losses globally. Traditional fraud detection systems relying on hand-crafted static rules or conventional statistical models are increasingly inadequate; they fail to adapt to evolving fraudulent behavior and suffer from elevated false-positive rates that disrupt legitimate consumer activities.

The rapid expansion of digital financial services, online banking, electronic payments, and mobile-based transactions has transformed the global financial ecosystem. Although these developments have improved transaction efficiency and accessibility, they have also created new opportunities for fraudulent activities. Financial fraud encompasses a wide range of deceptive practices, including unauthorized transactions, payment manipulation, identity misuse, account takeover, and other forms of financial misconduct. The increasing sophistication of fraudulent behavior presents substantial challenges for

financial institutions, which must identify suspicious activities promptly while maintaining the reliability and security of legitimate transactions. Consequently, the development of intelligent, scalable, and adaptive fraud detection systems has become an important area of research in financial technology and applied machine learning.

Traditional financial fraud detection methods generally rely on predefined rules, statistical analysis, and manually engineered detection criteria. Although rule-based systems can identify known fraudulent patterns, they may be less effective against previously unseen fraud strategies and rapidly changing transaction behaviors. Statistical approaches can provide valuable insights into transaction irregularities; however, their performance may be constrained by complex nonlinear relationships, high-dimensional datasets, and severe class imbalance. In practical financial datasets, legitimate transactions typically outnumber fraudulent transactions by a substantial margin, making it difficult to achieve high fraud detection sensitivity without increasing false-positive rates. These limitations have encouraged researchers to investigate machine learning and deep learning techniques capable of learning complex patterns directly from transaction data.

Machine learning algorithms, particularly gradient boosting methods, have demonstrated considerable potential for financial fraud classification. Light Gradient Boosting Machine (LightGBM) is an efficient gradient boosting framework that can process large-scale structured datasets while capturing nonlinear feature interactions. Its computational efficiency and ability to model complex decision boundaries make it suitable for transaction-level fraud detection. Nevertheless, the effectiveness of LightGBM depends on appropriate hyperparameter selection, including learning rate, number of leaves, tree depth, and regularization settings. Improper parameter configuration may result in overfitting, reduced generalization, or suboptimal classification performance. Therefore, systematic optimization strategies are required to identify suitable model configurations for fraud detection applications.

Financial transactions often exhibit temporal dependencies, sequential patterns, and behavioral relationships across multiple transactions. For example, unusual transaction frequency, changes in

spending behavior, and consecutive activities across accounts may provide important indications of fraudulent behavior. Long Short-Term Memory (LSTM) networks, a specialized recurrent neural network architecture, are capable of learning dependencies across sequential data through memory cells and gating mechanisms. This capability makes LSTM a promising technique for extracting temporal representations from financial transaction sequences. However, LSTM models may require substantial computational resources and careful hyperparameter tuning. Furthermore, their predictions can be difficult to interpret, which creates challenges for financial institutions that require transparent and auditable fraud-related decisions.

To address optimization challenges, Genetic Algorithms (GAs) offer an evolutionary search-based approach for identifying effective model configurations. Inspired by natural selection, a Genetic Algorithm iteratively evaluates candidate solutions through a fitness function and improves them using selection, crossover, and mutation operations. In the proposed research, the Genetic Algorithm is utilized to optimize the hyperparameters of the LSTM–LightGBM framework. This optimization strategy aims to improve the integration of sequential feature learning and gradient boosting classification while reducing reliance on manually selected parameters. The resulting framework is intended to provide a systematic approach to enhancing predictive performance and computational efficiency in financial fraud detection.

Despite the predictive capabilities of advanced machine learning models, their limited interpretability remains a significant concern in financial applications. Fraud detection decisions may affect transaction approvals, customer access to financial services, and institutional risk management. Consequently, stakeholders require explanations regarding why a particular transaction has been classified as fraudulent or legitimate. Explainable Artificial Intelligence (XAI) techniques can help address this requirement by providing interpretable information about model predictions. SHapley Additive exPlanations (SHAP) is a widely used model explanation method that estimates the contribution of individual features to a prediction. By integrating SHAP analysis, the proposed framework can identify influential transaction characteristics, support model

transparency, and assist analysts in investigating suspicious activities. The explanations should be interpreted as model-based feature contributions rather than direct proof of fraudulent intent.

This research proposes a Genetic Algorithm Optimized LSTM–LightGBM framework for explainable financial fraud detection. The framework combines LSTM-based sequential representation learning, LightGBM-based classification, Genetic Algorithm-driven hyperparameter optimization, and SHAP-based interpretability. The proposed methodology is designed to address key challenges associated with financial fraud detection, including temporal dependencies, class imbalance, nonlinear transaction relationships, model optimization, and decision transparency. Its effectiveness can be assessed through established evaluation measures such as precision, recall, F1-score, ROC-AUC, and Precision–Recall AUC, with particular emphasis on detecting fraudulent transactions while controlling false-positive errors.

The principal contributions of this research are the development of an integrated optimization and classification framework, the incorporation of sequential transaction information into gradient boosting-based fraud detection, and the introduction of an explainability layer for interpreting model decisions. The study further aims to establish a comprehensive evaluation methodology for examining predictive performance, optimization effectiveness, and interpretability. By integrating these components, the research seeks to contribute to the development of reliable and transparent artificial intelligence solutions for modern financial fraud detection systems.

II. LITERATURE REVIEW

Evolution of Financial Fraud Detection

Financial fraud detection has evolved from traditional rule-based and statistical approaches to sophisticated machine learning and deep learning systems. Early detection mechanisms primarily relied on predefined transaction thresholds, expert-designed rules, and statistical deviations from established customer behavior. Although these approaches remain useful for identifying known patterns, their dependence on manually specified conditions limits adaptability when fraudulent strategies change. Machine learning

methods have subsequently enabled automated pattern recognition by learning relationships between transaction characteristics and fraud labels. However, the effectiveness of these methods depends on data quality, feature representation, class distribution, and the ability to generalize to evolving fraud patterns.

Machine Learning Techniques for Fraud Classification

Supervised machine learning algorithms, including logistic regression, decision trees, random forests, support vector machines, and ensemble classifiers, have been extensively investigated for financial fraud detection. These methods can classify transactions using features such as transaction amount, payment method, transaction frequency, and account-related attributes. Ensemble learning techniques are particularly valuable because they combine multiple decision-making mechanisms to improve predictive stability. Nevertheless, financial fraud datasets commonly exhibit substantial class imbalance, in which fraudulent transactions represent only a small proportion of total observations. Consequently, conventional accuracy may provide an incomplete assessment of model performance. Precision, recall, F1-score, ROC-AUC, and Precision–Recall AUC are more informative for evaluating fraud detection systems, especially when identifying fraudulent transactions is the primary objective.

LightGBM for Financial Fraud Detection

Light Gradient Boosting Machine (LightGBM) is a gradient boosting framework designed for efficient learning from structured datasets. It constructs an ensemble of decision trees sequentially, with each subsequent tree reducing the errors of the preceding model. Its computational efficiency, flexible feature handling, and ability to represent nonlinear feature interactions make it a promising approach for financial transaction classification. LightGBM can also support feature importance analysis, enabling researchers to investigate which transaction attributes influence predictive outcomes.

Despite these advantages, LightGBM performance is influenced by hyperparameters such as learning rate, number of leaves, maximum tree depth, feature sampling, and regularization. Excessive model complexity may lead to overfitting, whereas overly restrictive configurations can reduce the model's

ability to identify complex fraudulent patterns. Furthermore, a conventional LightGBM model may not fully capture temporal dependencies when transaction records are treated as independent observations. These considerations motivate the integration of sequential learning and optimization techniques into the fraud detection pipeline.

LSTM Networks for Sequential Transaction Analysis

Long Short-Term Memory networks are recurrent neural network architectures developed to learn dependencies across sequential observations. Through memory cells and gating mechanisms, LSTM networks can retain relevant information over multiple time steps and reduce the limitations associated with conventional recurrent neural networks. In financial fraud detection, transaction sequences may reveal behavioral characteristics that are not adequately represented by isolated transaction records. Examples include sudden changes in transaction frequency, repeated high-value payments, unusual geographic movement, and deviations from historical spending behavior.

LSTM-based models can extract temporal representations from ordered transaction sequences and provide learned features for downstream classification. However, their performance depends on sequence construction, input feature selection, sequence length, network architecture, and training configuration. In addition, LSTM models can be computationally demanding and may be difficult to interpret directly. Therefore, combining LSTM-based feature extraction with a structured-data classifier such as LightGBM represents a potential strategy for leveraging both temporal and nonlinear feature-learning capabilities.

Hybrid LSTM–LightGBM Frameworks

Hybrid deep learning and machine learning architectures have attracted interest because different algorithms can address complementary aspects of complex classification problems. In an LSTM–LightGBM framework, the LSTM network may learn sequential representations from transaction histories, while LightGBM uses the resulting representations together with relevant transaction features for classification. This architecture can potentially combine temporal pattern recognition with efficient gradient boosting-based decision boundaries.

The effectiveness of such a hybrid framework depends on the design of information transfer between the two models. Important considerations include whether LSTM outputs are used as additional features, whether the models are trained sequentially or jointly, and whether the classification stage uses raw transaction attributes alongside learned representations. Without careful validation, a hybrid system may introduce redundant features, computational overhead, or information leakage. Consequently, a well-defined experimental protocol is necessary to establish whether the combined framework provides measurable improvements over individual LSTM and LightGBM models.

Genetic Algorithms for Hyperparameter Optimization

Genetic Algorithms are population-based optimization techniques inspired by evolutionary processes. They search for effective solutions through the iterative application of selection, crossover, and mutation. In machine learning, GAs can optimize hyperparameters that significantly influence model performance. For an LSTM–LightGBM fraud detection framework, potential optimization variables include LSTM hidden units, number of recurrent layers, dropout rate, sequence length, learning rate, LightGBM number of leaves, tree depth, boosting iterations, and regularization parameters.

The application of Genetic Algorithms may help explore complex hyperparameter spaces that are difficult to optimize using manual experimentation. However, evolutionary optimization can require substantial computational resources because numerous candidate configurations must be trained and evaluated. The choice of fitness function is also critical. Optimizing only accuracy may produce models that perform well on legitimate transactions but fail to detect fraudulent cases. A fraud-sensitive fitness function should therefore consider suitable metrics such as recall, F1-score, or Precision–Recall AUC, depending on the research objective. The use of validation data that is strictly separated from the final test set is essential to prevent optimistic performance estimates.

III. RESEARCH GAP:

Despite significant advancements in artificial intelligence and machine learning-based financial

fraud detection, several research challenges remain unresolved. Existing studies have investigated conventional machine learning algorithms, deep learning architectures, ensemble classifiers, and explainable artificial intelligence techniques. However, the integration of sequential transaction learning, evolutionary hyperparameter optimization, and interpretable gradient boosting classification remains an important area for further investigation. The following research gaps motivate the development of the proposed Genetic Algorithm Optimized LSTM–LightGBM framework.

Limited Integration of Temporal and Structured Transaction Features

Many financial fraud detection approaches analyze transactions as independent records, which may limit their ability to capture behavioral dependencies across consecutive transactions. Although LSTM networks can learn temporal patterns, conventional gradient boosting models such as LightGBM are primarily designed for structured feature-based learning. There is a need to investigate hybrid architectures that effectively combine sequential representations from LSTM networks with the nonlinear classification capabilities of LightGBM.

Insufficient Optimization of Hybrid Model Hyperparameters

The performance of deep learning and gradient boosting models is strongly influenced by their hyperparameter configurations. Existing implementations may rely on manual tuning, grid search, or default parameter settings, which can be inefficient for complex hybrid architectures. The effectiveness of Genetic Algorithm-based optimization for simultaneously identifying suitable LSTM and LightGBM parameters in financial fraud detection requires systematic evaluation.

Challenges Associated with Class Imbalance

Financial fraud datasets are frequently characterized by a significant imbalance between legitimate and fraudulent transactions. Models optimized primarily for overall accuracy may fail to identify a meaningful proportion of fraudulent cases. Although resampling and cost-sensitive learning approaches have been explored, further research is required to investigate how class imbalance handling and fraud-sensitive

optimization objectives influence the performance of GA-optimized hybrid models.

Limited Explainability in Sequential Hybrid Architectures

Deep learning-based fraud detection systems can generate accurate predictions while providing limited transparency regarding the underlying decision process. SHAP and other XAI methods offer mechanisms for interpreting model outputs, but the explanation of hybrid architectures involving learned sequential features presents additional challenges. It is necessary to investigate how SHAP-based explanations can be used to identify influential transaction characteristics while distinguishing between original financial features and latent representations generated by LSTM networks.

Inadequate Comparative Evaluation of Optimization and Model Integration

The individual contributions of LSTM, LightGBM, and Genetic Algorithm optimization are not necessarily equivalent. A hybrid model may demonstrate improved performance due to feature representation, classifier selection, or parameter optimization. Without appropriate ablation studies and baseline comparisons, it is difficult to determine which component contributes most significantly to the final results. Therefore, a systematic evaluation comparing standalone models, non-optimized hybrid models, and GA-optimized configurations is required.

Need for Robust Evaluation and Practical Applicability

Fraud detection systems must balance the identification of fraudulent transactions with the reduction of false alarms. In addition to predictive accuracy, practical evaluation should consider recall, precision, F1-score, ROC-AUC, Precision–Recall AUC, computational cost, and model stability. Furthermore, temporal validation and leakage prevention are important when transaction sequences are used. Research is needed to establish whether the proposed framework can achieve reliable fraud detection performance under realistic experimental conditions.

Consolidated Research Gap

The central research gap lies in the need for a systematically optimized and explainable framework

that combines LSTM-based sequential feature learning with LightGBM classification for financial fraud detection. The proposed research addresses this gap by employing a Genetic Algorithm to optimize selected model hyperparameters and incorporating SHAP analysis to interpret classification outcomes. The framework will be evaluated using appropriate baseline models, fraud-sensitive performance metrics, and ablation experiments to determine its effectiveness, limitations, and practical relevance.

IV. PROPOSED METHODOLOGY

The proposed study develops a Genetic Algorithm Optimized LSTM–LightGBM framework for explainable financial fraud detection. The framework combines the sequential learning capability of Long Short-Term Memory (LSTM), the classification efficiency of Light Gradient Boosting Machine (LightGBM), the optimization capability of Genetic Algorithm (GA), and the interpretability of SHapley Additive exPlanations (SHAP). The methodology is designed to identify complex fraudulent transaction patterns while providing an explanation of the factors influencing the prediction.

The overall methodology consists of the following stages:

- Data Collection → Data Preprocessing → Feature Engineering → Transaction Sequencing → LSTM Feature Learning → Genetic Algorithm Optimization → LightGBM Classification → SHAP Explainability → Performance Evaluation

V. DATA COLLECTION

A publicly available financial transaction dataset is selected for evaluating the proposed framework. The dataset contains transaction-related attributes and a target variable indicating whether a transaction is legitimate or fraudulent. Depending on the selected dataset, the available attributes may include transaction amount, transaction type, transaction time, customer or account information, merchant-related characteristics, and other behavioral indicators.

The dataset is examined to understand its size, feature distribution, class distribution, missing values, and data characteristics. The selected dataset is divided into training, validation, and testing subsets. Wherever

timestamp information is available, chronological splitting is preferred because it provides a more realistic representation of fraud detection in real financial environments.

Data Preprocessing

The raw dataset is preprocessed before being supplied to the proposed model. Missing values, duplicate observations, inconsistent records, and irrelevant attributes are identified and appropriately treated. Categorical variables are converted into numerical representations using suitable encoding techniques.

Numerical variables are scaled when required, particularly for the LSTM component. Outliers are investigated carefully because extreme transaction values may themselves represent meaningful fraud indicators. Therefore, suspicious observations are not automatically removed without considering their potential relevance to fraud detection.

Special attention is given to the highly imbalanced nature of financial fraud datasets. Appropriate class-balancing strategies are applied only to the training data to avoid data leakage. The validation and testing datasets are retained in their original distribution to provide a realistic assessment of the model.

Feature Engineering

Feature engineering is performed to improve the representation of financial transaction behavior. In addition to the original transaction attributes, relevant behavioral characteristics can be generated from the available data.

Examples include transaction frequency, average transaction amount, changes in transaction value, time intervals between consecutive transactions, and deviations from historical customer behavior. The final feature set is selected based on data characteristics and experimental requirements.

Feature engineering is particularly important because fraud may not always be identifiable from a single transaction. A transaction that appears normal independently may become suspicious when considered in relation to previous transaction behavior.

Transaction Sequencing

A major component of the proposed methodology is the transformation of individual transaction records into sequential observations. Transactions belonging

to the same customer, account, card, or other appropriate entity are arranged chronologically.

The resulting sequences are provided to the LSTM network. This enables the model to learn changes in transaction behavior over time rather than treating every transaction as an isolated observation.

For example, a sudden increase in transaction frequency or a sequence of unusually large transactions may provide stronger evidence of potential fraud than the characteristics of a single transaction alone.

LSTM-Based Sequential Feature Learning

The LSTM network is used to learn temporal dependencies within the transaction sequences. LSTM is particularly suitable for this task because its memory mechanism allows relevant information from previous transactions to influence the representation of subsequent transactions.

The network receives the transaction sequence and generates a learned representation containing information about historical transaction behavior. Instead of using the LSTM directly as the final fraud classifier, the proposed methodology uses the learned representation as an input to the subsequent LightGBM classification stage.

This approach allows the LSTM component to focus primarily on temporal and behavioral feature learning, while LightGBM focuses on nonlinear classification and feature interactions.

Genetic Algorithm-Based Optimization

The performance of both LSTM and LightGBM depends strongly on their hyperparameter configuration. Selecting these parameters manually may not produce the most effective model configuration.

Therefore, a Genetic Algorithm is incorporated into the proposed framework to automatically search for an effective combination of hyperparameters.

The optimization process includes:

- Generation of an initial population of candidate configurations.
- Evaluation of each configuration using the validation dataset.
- Selection of better-performing configurations.
- Application of crossover to create new candidate solutions.

- Application of mutation to maintain population diversity.
- Re-evaluation of the new configurations.
- Repetition of the evolutionary process until the stopping criterion is reached.

Parameters such as LSTM hidden units, dropout rate, learning rate, sequence length, number of LightGBM leaves, tree depth, number of estimators, and regularization parameters can be included in the optimization process.

The objective is to identify a configuration that provides strong fraud detection performance while avoiding unnecessary model complexity.

VI. FITNESS EVALUATION

The Genetic Algorithm requires a fitness criterion to determine which candidate model configurations perform better.

Because financial fraud datasets are generally imbalanced, accuracy alone is not considered an appropriate optimization objective. The proposed methodology emphasizes fraud-sensitive measures such as precision, recall, F1-score, and Precision–Recall AUC.

Particular importance is given to recall because failure to detect an actual fraudulent transaction can have significant financial consequences. At the same time, precision is considered to control excessive false alarms that could unnecessarily affect legitimate customers.

Integration of LSTM and LightGBM

After the Genetic Algorithm identifies the optimized configuration, the LSTM model is trained using the selected parameters. The learned sequential features are extracted from the LSTM.

These learned features are then combined with selected original transaction features and supplied to LightGBM. The resulting LightGBM model performs the final classification of transactions as legitimate or fraudulent.

The integration provides complementary capabilities:

- LSTM: captures temporal and sequential transaction behavior.
- LightGBM: learns nonlinear relationships and feature interactions.

- Genetic Algorithm: searches for an effective model configuration.
- SHAP: explains the factors contributing to model predictions.

This integrated architecture forms the core of the proposed fraud detection framework.

LightGBM-Based Fraud Classification

LightGBM is selected as the final classification algorithm because of its ability to efficiently process structured data and model complex nonlinear relationships.

The classifier receives the combined representation containing transaction-level characteristics and LSTM-derived sequential information. It then produces a fraud probability for each transaction.

An appropriate classification threshold is selected using the validation data. This is important because the default classification threshold may not necessarily provide the best balance between fraud detection and false-positive reduction in an imbalanced financial dataset.

Explainability Using SHAP

The proposed framework incorporates SHAP to address the interpretability limitations of complex machine learning models.

SHAP analysis identifies the contribution of individual features to a particular prediction. Two forms of analysis are considered.

Global explanation is used to identify the most influential features across the complete dataset. This helps determine which transaction characteristics are generally important for fraud detection.

Local explanation is used to examine individual transactions. It can show why a particular transaction received a high or low fraud probability.

For example, the explanation may indicate that an unusual transaction amount, transaction frequency, or behavioral deviation contributed strongly to a fraud prediction. Such explanations can assist financial analysts in reviewing suspicious transactions.

VII. EXPERIMENTAL DESIGN

The proposed framework is evaluated against several baseline models to establish whether the integration of

the different components provides meaningful improvement.

The comparative models may include:

- Logistic Regression
- Random Forest
- LSTM
- LightGBM
- LSTM–LightGBM
- GA-optimized LightGBM
- Proposed GA-optimized LSTM–LightGBM

All models are evaluated using the same dataset and comparable experimental conditions wherever applicable.

Performance Evaluation

The proposed model is evaluated using multiple performance indicators, including:

- Accuracy
- Precision
- Recall
- F1-score
- ROC-AUC
- Precision–Recall AUC
- False-positive rate
- False-negative rate
- Computational time

Since the dataset is expected to be imbalanced, greater emphasis is placed on precision, recall, F1-score, and Precision–Recall AUC rather than accuracy alone.


The confusion matrix is also analyzed to understand the number of correctly and incorrectly classified fraudulent and legitimate transactions.

Ablation Study

An ablation study is conducted to determine the individual contribution of the major components of the proposed framework.

The following configurations can be investigated:

- LightGBM without optimization.
- LSTM without LightGBM.
- LSTM–LightGBM without Genetic Algorithm.
- GA-optimized LightGBM.
- GA-optimized LSTM.
- Complete GA-optimized LSTM–LightGBM framework.



The results of the ablation study help determine whether improvements are primarily associated with sequential feature learning, hybrid classification, or Genetic Algorithm-based optimization.

VIII. EXPLAINABILITY ANALYSIS

Following performance evaluation, SHAP results are analyzed to identify the most influential features associated with fraud predictions. Global feature importance is used to identify common fraud-related patterns, while local explanations are generated for representative fraudulent and legitimate transactions. The explainability analysis provides an additional dimension to the research by moving beyond the question of "How accurately can the model detect fraud?" to also address "Why did the model classify this transaction as fraudulent?"

LSTM Module:

The Long Short-Term Memory (LSTM) module is the sequential learning component of the proposed Genetic Algorithm Optimized LSTM-LightGBM framework. Its primary purpose is to identify temporal and behavioral patterns within financial transaction sequences that may not be captured when transactions are analyzed independently.

Financial fraud is often associated with changes in transaction behavior over time rather than with a single transaction. For example, an unusual increase in transaction frequency, repeated high-value transactions, or a sudden deviation from previous spending behavior may indicate suspicious activity. The LSTM module is therefore used to learn these sequential dependencies.

Input Sequence Formation

Before applying LSTM, transactions are arranged chronologically according to an appropriate entity such as customer, account, card, or merchant. A sequence consists of a predefined number of consecutive transactions.

Each transaction is represented by its selected features, such as:

- Transaction amount
- Transaction type
- Transaction time
- Transaction frequency
- Previous transaction interval
- Customer/account-related attributes
- Behavioral features

The resulting sequence is provided to the LSTM as a three-dimensional input consisting of samples, time steps, and features.

LSTM Architecture

The LSTM network consists of memory cells and three primary gates:

- Input gate: determines which new information should be retained.
- Forget gate: determines which previously stored information should be discarded.
- Output gate: determines which information should be passed to the next stage.

These mechanisms allow the network to retain important historical transaction information while reducing the influence of irrelevant information.

A typical architecture for the proposed study can consist of:

Sequential Transaction Input → LSTM Layer → Dropout Layer → Dense/Feature Representation Layer → LSTM Feature Output

The exact number of LSTM layers, hidden units, dropout rate, learning rate, and sequence length are not fixed manually. They are treated as candidate parameters for optimization using the Genetic Algorithm.

Temporal Feature Extraction

The principal output of the LSTM module is a learned representation of transaction behavior. Instead of directly using this representation as the final fraud prediction, it is transferred to the LightGBM module.

The LSTM-generated representation can capture information about the relationship between current and previous transactions, providing additional information to the LightGBM classifier.

Genetic Algorithm Optimization of LSTM


The Genetic Algorithm is used to identify an effective LSTM configuration. Candidate parameters may include:

Parameter	Optimization Role
Number of LSTM Layers	Control Network Depth
Hidden Units	Control representation capacity
Sequence Length	Determines historical information considered
Learning Rate	Control training updates
Dropout Rate	Help control overfitting
Batch Size	Control training batches
Number of epochs	Control training duration

Each candidate configuration is evaluated using validation data. Configurations producing better fraud-detection performance are retained for subsequent generations through selection, crossover, and mutation.

LSTM Output for LightGBM

After the optimal LSTM configuration is identified, the network is trained and its learned representations are extracted. These features are combined with selected original transaction features before being supplied to LightGBM.



This creates a complementary relationship between the two models:

LSTM → learns temporal and behavioral patterns

LightGBM → performs nonlinear fraud classification

The information flow is therefore:

Transaction Sequences
↓
Preprocessing & Feature Engineering
↓
LSTM Network
↓
Temporal/Behavioral Feature Representation
↓
Feature Fusion with Original Transaction Features
↓
LightGBM
↓
Fraud Probability
↓
SHAP Explanation

This design enables the proposed framework to exploit both temporal dependencies through LSTM and nonlinear feature interactions through LightGBM, while the Genetic Algorithm searches for an effective configuration of the combined architecture.

IX. RESEARCH OBJECTIVES:

The primary objective of this research is to develop a Genetic Algorithm Optimized LSTM–LightGBM framework for accurate, robust, and explainable financial fraud detection. The specific objectives of the study are as follows:

Sequential & Tabular Feature Extraction (LSTM-to-LightGBM Integration)

- To extract high-level, temporal embeddings from transaction sequences using an LSTM network.
- To pass these extracted sequence embeddings directly as feature inputs into LightGBM along with engineered tabular features, enhancing gradient boosting performance on sequential patterns.

Automated Hyperparameter Optimization via Genetic Algorithms (GA)

- To build a GA optimization loop that automatically tunes hyperparameter search spaces (e.g., LSTM layer depth, learning rate, sequence length, LightGBM tree depth, num_leaves, feature fraction).
- To maximize multi-objective fitness functions balancing precision, recall, AUC-ROC, and model complexity.

Hybrid Model Ensembling & Performance Benchmarking

- To construct a robust hybrid decision framework combining deep sequential representations and tree-based gradient boosting.
- To evaluate and benchmark the proposed GA-LSTM-LightGBM framework against baseline models (standalone LightGBM, standalone LSTM, XGBoost, Random Forest, and unoptimized variants) on standard financial fraud datasets.

Model Interpretability and Explainability (XAI)

- To integrate SHAP (SHapley Additive exPlanations) to provide both global model explainability

(identifying key risk drivers across all transactions) and local explainability (explaining individual flagged transactions).

- To bridge the "black-box" nature of LSTM sequence models with actionable insights for fraud analysts and regulatory compliance.

Deployment Viability & Real-Time Decision Making

- To assess the latency, computational overhead, and throughput of the hybrid model under real-time transaction processing constraints.
- To define a dual-path decision pipeline (automated approval/blocking for high-confidence predictions vs. manual review for edge cases).
- To develop an integrated LSTM–LightGBM framework capable of identifying fraudulent financial transactions by combining temporal behavioral learning with nonlinear machine learning-based classification.
- To utilize LSTM networks for extracting sequential transaction patterns and identifying temporal dependencies that may indicate abnormal or fraudulent financial behavior.
- To employ a Genetic Algorithm for hyperparameter optimization of the LSTM–LightGBM framework and reduce dependence on manual parameter selection.
- To address the challenges associated with imbalanced financial transaction datasets by applying appropriate data preprocessing and class-balancing strategies while maintaining reliable evaluation conditions.
- To investigate the effectiveness of combining LSTM-derived sequential features with LightGBM-based classification for improving fraud detection performance.
- To incorporate SHAP-based Explainable Artificial Intelligence (XAI) to identify and interpret the major transaction features contributing to fraudulent and legitimate predictions.
- To evaluate the proposed framework using appropriate performance metrics, including precision, recall, F1-score, ROC-AUC, and Precision–Recall AUC, with particular emphasis on fraudulent transaction detection.
- To compare the proposed framework with conventional and individual machine learning/deep

learning models such as LSTM, LightGBM, and other suitable baseline classifiers.

- To conduct ablation analysis to determine the individual contribution of LSTM-based feature learning, Genetic Algorithm optimization, LightGBM classification, and SHAP-based analysis.
- To establish an interpretable and practically relevant fraud detection framework that can support financial analysts and institutions in identifying suspicious transactions and understanding the factors influencing automated fraud predictions.

X. CONCLUSION AND FUTURE SCOPE

This research proposed a Genetic Algorithm Optimized LSTM–LightGBM framework for explainable financial fraud detection. The proposed approach integrates the temporal learning capability of Long Short-Term Memory networks with the nonlinear classification and efficient feature handling capabilities of LightGBM. The Genetic Algorithm is incorporated to systematically optimize important model hyperparameters, reducing dependence on manual parameter selection. In addition, SHAP-based Explainable Artificial Intelligence is introduced to improve the transparency of the fraud detection process.

The LSTM component enables the framework to capture sequential and behavioral patterns within financial transaction histories, while LightGBM provides an effective classification mechanism for distinguishing fraudulent and legitimate transactions. The integration of these models is intended to address the limitations of approaches that consider transactions independently or rely on a single learning architecture. The optimization stage further provides a systematic mechanism for identifying suitable model configurations according to fraud-sensitive performance objectives.

The incorporation of SHAP adds an important interpretability layer to the proposed framework. Instead of providing only a fraud or legitimate prediction, the system can identify the transaction characteristics that contribute to a particular prediction. This can support financial analysts in investigating suspicious activities and improve confidence in AI-assisted fraud detection.

The proposed framework is evaluated using multiple performance measures, including precision, recall, F1-score, ROC-AUC, and Precision–Recall AUC. The comparative and ablation analyses are designed to determine the contribution of sequential feature learning, Genetic Algorithm optimization, and LightGBM classification. Overall, the proposed methodology provides a unified approach toward developing a fraud detection system that considers predictive performance, optimization, temporal behavior, and explainability.

The research demonstrates the potential of combining evolutionary optimization, deep learning, gradient boosting, and explainable AI within a single financial fraud detection framework. However, the final effectiveness of the framework should be established through experiments on representative datasets and comparison with appropriate state-of-the-art methods.

Future Scope

Several opportunities exist for extending the proposed research. First, the framework can be evaluated using larger and more diverse real-world financial datasets containing multiple transaction channels, customer behaviors, and geographical characteristics. Such evaluation can provide stronger evidence regarding its scalability and generalization capability.

Second, future research can investigate real-time fraud detection, in which the optimized LSTM–LightGBM model processes incoming transactions continuously and generates fraud-risk scores with minimal latency. This would increase the practical applicability of the framework in banking, digital payment, and financial technology environments.

Third, the framework can be extended using online and incremental learning techniques to address concept drift. Fraud patterns continuously evolve, and a model trained on historical transactions may gradually lose effectiveness when new fraud strategies emerge. Adaptive learning mechanisms could allow the system to update its knowledge without requiring complete retraining.

Fourth, future studies can investigate advanced sequence architectures such as Bidirectional LSTM, GRU, Transformer, and attention-based models and compare them with the proposed LSTM architecture.

Graph-based learning can also be explored to model relationships between customers, accounts, merchants, devices, and transactions.

Fifth, future work can improve explainability by combining SHAP with other XAI approaches such as LIME, counterfactual explanations, and feature interaction analysis. This could provide both global and transaction-level explanations that are easier for financial investigators to interpret.

Finally, future research can investigate cost-sensitive optimization and multi-objective Genetic Algorithms that simultaneously optimize fraud recall, precision, computational efficiency, and false-positive reduction. Integration with banking security systems, fraud investigation platforms, and privacy-preserving learning technologies could further enhance the practical deployment of the proposed framework.

Thus, future development can transform the proposed framework from an experimental fraud classification model into an adaptive, real-time, explainable, and scalable financial fraud intelligence system.

REFERENCES

- [1] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [2] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, San Francisco, CA, USA, 2016, pp. 785–794.
- [3] G. Ke et al., "LightGBM: A highly efficient gradient boosting decision tree," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.
- [4] S. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.
- [5] S. M. Lundberg et al., "From local explanations to global understanding with explainable AI for trees," *Nature Machine Intelligence*, vol. 2, pp. 56–67, 2020.
- [6] A. Géron, *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow*, 2nd ed. Sebastopol, CA, USA: O'Reilly Media, 2019.
- [7] C. C. Aggarwal, *Neural Networks and Deep Learning: A Textbook*. Cham, Switzerland: Springer, 2018.
- [8] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
- [9] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002.
- [10] H. He and E. A. Garcia, "Learning from imbalanced data," *IEEE Transactions on Knowledge and Data Engineering*, vol. 21, no. 9, pp. 1263–1284, 2009.
- [11] V. Vapnik, *The Nature of Statistical Learning Theory*, 2nd ed. New York, NY, USA: Springer, 2000.
- [12] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.
- [13] J. H. Holland, *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI, USA: University of Michigan Press, 1975.
- [14] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA, USA: Addison-Wesley, 1989.
- [15] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
- [16] C. Elkan, "The foundations of cost-sensitive learning," in *Proc. 17th Int. Joint Conf. Artificial Intelligence (IJCAI)*, 2001, pp. 973–978.
- [17] A. Dal Pozzolo, O. Caelen, R. A. Johnson, and G. Bontempi, "Calibrating probability with undersampling for unbalanced classification," in *2015 IEEE Symposium Series on Computational Intelligence*, Cape Town, South Africa, 2015, pp. 159–166.
- [18] A. Dal Pozzolo, G. Boracchi, O. Caelen, C. Alippi, and G. Bontempi, "Credit card fraud detection: A realistic modeling and a novel learning strategy," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 8, pp. 3784–3797, 2018.
- [19] A. Carcillo, Y.-A. Le Borgne, O. Caelen, and G. Bontempi, "Streaming active learning strategies for real-life credit card fraud detection," in *Machine Learning and Knowledge Discovery in Databases*, Springer, 2018, pp. 286–300.

- [20] A. Carcillo, Y.-A. Le Borgne, O. Caelen, Y. Kessaci, F. Oblé, and G. Bontempi, “Combining unsupervised and supervised learning in credit card fraud detection,” *Information Sciences*, vol. 557, pp. 317–331, 2021.
- [21] Y.-A. Le Borgne, W. Siblini, B. Lebichot, and G. Bontempi, *Reproducible Machine Learning for Credit Card Fraud Detection—Practical Handbook*. Université Libre de Bruxelles, 2022.
- [22] A. Pumsirirat and L. Yan, “Credit card fraud detection using deep learning based on auto-encoder and restricted Boltzmann machine,” *International Journal of Advanced Computer Science and Applications*, vol. 9, no. 1, pp. 30–36, 2018.
- [23] A. Fiore, A. De Santis, F. Perla, P. Zanetti, and F. Palmieri, “Using generative adversarial networks for improving classification effectiveness in credit card fraud detection,” *Information Sciences*, vol. 479, pp. 448–455, 2019.
- [24] A. Jurgovsky et al., “Sequence classification for credit-card fraud detection,” *Expert Systems with Applications*, vol. 100, pp. 234–245, 2018.
- [25] E. A. López-Rojas and S. Axelsson, “Money laundering detection with a deep learning approach,” in *2016 European Intelligence and Security Informatics Conference (EISIC)*, 2016, pp. 1–8.
- [26] M. Weber et al., “Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics,” *arXiv preprint arXiv:1908.02591*, 2019.
- [27] M. Ribeiro, S. Singh, and C. Guestrin, ““Why should I trust you?”: Explaining the predictions of any classifier,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 1135–1144.
- [28] A. Adadi and M. Berrada, “Peeking inside the black-box: A survey on explainable artificial intelligence (XAI),” *IEEE Access*, vol. 6, pp. 52138–52160, 2018.
- [29] R. Guidotti et al., “A survey of methods for explaining black box models,” *ACM Computing Surveys*, vol. 51, no. 5, pp. 1–42, 2018.
- [30] F. Pedregosa et al., “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.