Eksplorium *p*-ISSN 0854-1418
Volume 46 No. 2, June 2025: 139–154 *e*-ISSN 2503-426X

A Hybrid Machine Learning Framework for Financial Fraud Detection in Corporate Management Systems

Anand Kumar Dohare[1], Uttam U. Deshpande[2], Aman Dahiya[3], Kanchan Dabre[4], Dr. Kriti Srivastava[5], Dr. Sreedhar Bhukya[6], Prem Kumar Sholapurapu[7]
[1]Associate Professor, Department of Information Technology, Greater Noida Institute of Technology, dohare.anand@gmail.com
[2]Department of Electronics and Communication Engineering, KLS Gogte Institute of Technology, Karnataka, India, uttamudeshpande@gmail.com
[3]Associate Professor, Department of Electronics and Communication Engineering, Maharaja Surajmal Institute of Technology, Janakpuri, New Delhi, India, amandahiya@msit.in
[4]Assistant Professor, CSE- Data Science, D. J. Sanghvi College of Engineering, Mumbai, kanchan.dabre@djsce.ac.in
[5]Assistant Professor, CSE- Data Science, D. J. Sanghvi College of Engineering, Mumbai, Kriti.srivastava@djsce.ac.in
[6]Professor, Department of Computer Science & Engineering, Sreenidhi Institute of Science & Engineering Hyderabad, Telangana-501301, sreedhar.b@sreenidhi.edu.in
[7]Research Associate and Senior Consultant, CGI, USA, premkumar.sholapurapu@cgi.com
Corresponding author: Prem Kumar Sholapurapu, premkumar.sholapurapu@cgi.com

Article Received: 15 May 2025, Revised: 18 June 2025, Accepted: 26 June 2025

Abstract: Financial fraud is still a very serious problem in the operation of the enterprise, the conventional detection mode can not keep pace for changes in the fraud model, and the interpretability in the enterprise audit is lack. To improve financial anomalies detection in enterprise from a machine learning perspective, we propose an innovative hybrid machine learning approach named as EHRN-GMM, which combines a Heterogeneous Recurrent Network (HRN) and Gaussian Mixture Model (GMM). The HRN consists of GRU-LSTM fused layers with temporal embeddings and attention, which is able to learn short and long term dependencies in company sequential transaction logs. Its results are given to a GMM to estimate the distribution of valid behaviors, and detect anomalies in a probabilistic image. Furthermore, a SHAP-based interpretability layer is added which helps construct auditor-friendly explanations and enhances the transparency and trustworthiness of the model's decisions. The proposed methodology is evaluated on synthetic ERP datasets, real-world credit card fraud data (e.g., Vesta), and simulated audit trails, where it achieves an average AUC of 0.96, outperforming competing methods such as XGBoost and CNNs. Furthermore, the model exhibits strong concept drift adaptability facilitating by GMM updates at intervals. This paper presents a scalable, interpretable, high-performing architecture for enterprise-wide fraud surveillance that helps to fill the void between automated anomaly detection and responsible corporate governance.

Keywords: Financial fraud detection, hybrid machine learning, corporate management systems, GRU-LSTM, Gaussian Mixture Model, explainable AI, SHAP, sequential data modeling, anomaly detection

1. INTRODUCTION

Fraud in the contemporary financial environment is no more just a remote risk but substantial challenge to corporate governance and economic existence. Corporate Management Systems and Financial Fraud Corruption in CMS among several aspects of the corporate system, financial fraud disrespects the investors, biases resources and institutional credit. Fraudulent expense claims and false invoices, along with suspicious payrolls and fudged ledgers have grown smarter, outmanoeuvring rules-based systems that rely on luck and hope to uncover misuse of corporate funds. Increased volumes and complexity in enterprise transactions require

a proactive, intelligent approach to fraud detection, one that can learn from data patterns, adapt to changing attacks, and provide (and explain) answers to decision-makers[1]. The latter challenge has prompted this study to present an innovative hybrid machine learning model, EHRN-GMM, that is capable of accurately identifying sophisticated financial frauds in corporate settings in an understandable manner.

The financial environment of the corporate world poses special obstacles for fraud identification. In contrast to the fraud being committed against consumers (i.e., credit card fraud), corporate fraud is typically operationalized as transactions that are embedded within a company's business process and distributed across several data sources (e.g., ERP systems, general ledger entries, internal audit logs). Criminal conduct could of course include employee collusion, financial statement manipulation, and improper exercise of discretion. Furthermore, the activities are not static - the scammers will constantly adjust their behavior in response to implemented controls. Static models are therefore insufficient[2]. Contemporary fraud detection systems must identify anomalies, not only as standalone events, but as features embedded in complex, time-dependent and context-specific patterns.

Despite the recent popularity of machine learning (ML) in fraud analytics, many of the available ML models present severe drawbacks when applied at the enterprise scale. First, traditional classifiers as those based on support vector machines or decision trees do not naturally accommodate temporal dependencies and patterns of collusion, since they assume that benign users and colluders are independent and identically distributed (I.I.D). Second, black-box deep learning models, such as convolutional neural networks (CNNs) or even simple RNNs, do not typically offer the interpretability needed in high-stakes financial audits or regulatory compliance use cases. Finally, the class imbalance in most fraud datasets is particularly severe with fraud events representing less than 1% of overall transactions, challenging the classifiers capabilities of generalising effectively[3].

To fill these gaps, we proposed EHRN-GMM, a heterogeneous recurrent network (HRN) and a Gaussian Mixture Model (GMM) hybrid framework, plus an interpretation layer employing SHAP (SHapley Additive exPlanations). HRN combines GRU with LSTM cells to learn both short-term and long-term transactional sequence temporal dependency. It uses a temporal embeddings for capturing the dynamics of time-based relations, and attention mechanism to concentrate more on abnormal way of patterns through audit trails and ERP logs. This fusion model allows for memory and responsiveness that is difficult to achieve through singular LSTMs or GRUs. Outputs from the HRN that are fed to a GMM that models the normalcy and alerts of deviances as suspicious events. This probabilistic layer allows us to discover subtle behavioral outliers that are not rule-breaking but significantly differ from normal behavior[4].

One of the unique advantages of EHRN-GMM is its explainability layer. The model is also built with SHAP values34, a game-theoretic interpretation technique of model predictions, and thus it yields localized, comprehensible explanations for every flagged transaction. For example, instead of flagging the a purchase order as suspicious and nothing more the model can pinpoint that the anomaly came to be due to a rare combination of vendor history, timing and value deviation from typical departmental norms[5]. Such explanations are important in the practice of real-world audit, as auditors cannot rely on black-box models without any interpretability that may induce compliance risks and decision paralysis.

In addition, EHRN-GMM also deals with concept drift, the common problem in fraud analytics where statistical properties (e.g., mean and variance) of the target variable change over time. But it is a cat-and-mouse game: Fraudsters will, as they sometimes do, adjust their methods if they sense a system is tracking their activity. To address this, the GMM component of our based framework is constructed to be updated periodically with weighted online learning[6]. This enables the model to memorize historical fraud patterns and at the same time adapt to newer ones, maintaining performance even when the nature of commerce changes. Rather than existing retraining mechanisms that are expensive in terms of resource usage and disruptive, this model supports the ability to adapt in real-time without impacting on operational practices[7].

The proposed model has been experimented with in a combination of synthetic ERP transaction logs, the Vesta real-world fraud detection dataset and simulated financial statements. Preliminary results show that the model significantly outperforms a variety of baseline models, including random forest, XGBoost and LSTM. Our proposed model obtains an average AUC (Area Under the Receiver Operating Characteristic Curve) of 0.96, as well as 96% precision, showing that it performs better in catching more true fraud cases with a lower level of false alarms. These results confirm the efficacy of mixing sequence model with probabilistic anomaly detection and explainable AI on enterprise fraud detection tasks[8].

EHRN-GMM is modular in nature and is designed to be pluggable. It's available as an API to be integrated into existing enterprise architectures, or as a real-time analytics monitoring module in systems such as SAP S/4HANA, Oracle NetSuite, or Microsoft Dynamics. It is capable of handling structured financial data (e.g. GL, AP, AR) and semi-structured audit trails (e.g. user access logs, approval chains). This makes the model highly applicable not just to finance but to internal audit, compliance, procurement, ERM, and other control functions.

Overall, this study adds to the increasingly important area of intelligent fraud detection, by providing a novel hybrid machine learning approach that is accurate, interpretable and actionable by organizations. EHRN-GMM eliminates the wide delta between sophistication AI-enabled insights and the pragmatic needs of financial control in contemporary firms. It brings structured approach to the monitoring, identifying and explaining of suspicious financial activities over time, enabling corporate officers and auditors the ability to respond to threats in advance. This framework can be further extended in future by integrating graph neural networks to represent inter-entity fraud relationships in the consortium and federated learning for cross-organizational fraud propagation analysis with preserving data privacy.

By integrating temporal deep learning, statistical modeling and explainable AI in one framework, our work paves the way for the development of the next-generation fraud detection system that is not only accurate but can be trusted and audited. With the ongoing digitization and increasing complexity of corporate financial landscapes, methodologies such as EHRN-GMM provide a scalable, logical and responsible way to regulate fiscal integrity in the era of smart automation.

2. RELATED WORK

It has been a very important topic in the rich area of datadriven risk management, and financial fraud detection is no exception. In the last 10 years, we have seen tremendous strides in

artificial intelligence and machine learning capabilities that have enabled us to detect fraudulent transactions in financial systems much better. Nevertheless, as advanced as they are, the adoption and use of such technologies still represents a uniquely challenging problem in corporate networks when it comes to detecting fraud. In this section, we discuss the existing methods, their suitability to the enterprise use case, and the precise lacunae we aim to resolve through our hybrid framework (EHRN-GMM).

Conventional machine learning classifiers such as Decision Trees, SVM and Naive Bayes classifiers have been widely used for fraud detection. These methods are simple and fast to compute, as well as easy to deploy. Since they are skilled at manipulating tabular data, they work well on structured financial data such as: accounting entries, vendor payments, or employee reimbursements[9]. However, they are inefficient when both temporal dependencies and shift in behavioral trend play important roles (such as in corporate fraud). For instance, a scam that incrementally develops over weeks as the result of frequent but minor adjustments which one-by-one raises the scam level may look legitimate in a rule-based system, but not in temporal anomaly analysis. These traditional models cannot deal with sequential or contextual properties of data as listed in Table 1 and reduces their ability in detecting fraud in enterprise resource planning (ERP) systems[10].

Table 1: Overview of Existing Machine Learning Techniques for Financial Fraud Detection

Model Type	Algorithms Used	Strengths	Limitations
Traditional ML	Decision Trees, SVM, Naive Bayes	Fast training, easy to implement	Poor at handling sequential and contextual patterns
Ensemble Methods	Random Forest, XGBoost	High accuracy, handles feature importance well	Limited temporal modeling, lacks interpretability
Deep Learning	CNN, LSTM, GRU	Learns complex patterns, good for large datasets	Black-box models, hard to explain decisions
Hybrid Architectures	CNN-LSTM, RF-DNN, RNN-AE	Combines temporal and classification strengths	Often lacks transparency, sensitive to data imbalance
Anomaly Detection	Autoencoders, Isolation Forest	Detects unknown fraud types, good for rare events	High false positive rate, hard to tune thresholds
Probabilistic Models	GMM, HMM	Captures distributional outliers, unsupervised	Requires careful tuning, can be computationally heavy

Ensemble based classifiers like Random Forest and XGBoost have provided some of the improvements over single classifiers. These models excel at reducing variance, mitigation of overfitting, but also generally improve predictive accuracy by utilizing decisions or

optimization of several base learners. For determining the feature importances, it is particularly helpful in detecting highly important measures of fraud, such as abnormal payment frequencies or very high invoice amounts. However, the ensemble models are the most effective in flat feature spaces and they usually do not provide a satisfactory solution to a financial process in a corporate environment when the process is not sequential or relational. For example, they may be unaware that a large number of small transactions that avoid the approval limits within a short space of time might indicate a more serious fraud pattern than a single transaction on its own[11]. Furthermore, although they are more complex than the traditional alternatives, they are deficient in interpretability, particularly with a large number of trees or learners contributing to the final decision trajectory.

Deep learning algorithms in particular, Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs), have been utilized to learn non-linear as well as temporal patterns in financial information. CNNs are used to identify spatial anomalies in the matrix of transaction data and RNN and their extensions like LSTM nets and GRUs could well model sequences. Such models are effective in detecting complex fraud patterns that change over time—e.g., procurement fraud that manipulates invoice routing over several approvers, and time-lagged kickback schemes[12]. However, one major constraint of deep learning models in the financial domain is their black-box property. Detection is not enough for corporate governance and regulators and internal audit. What makes a specific transaction suspicious is as important as the fact that it is. Deep models have difficulty doing such reasoning in a transparent way[13]. As summarized in Table 1, deep models sacrifice interpretability for performance, which is a hard trade for financial institutions to accept.

Hybrid models have been identified as a potential solution to the trade-offs between performance and interpretability. These architectures are hybrids of several algorithms – typically a deep learning model for extracting features, and a classical model for classification or anomaly detection. For instance, CNN-LSTM can represent both spatial and temporal aspects of financial information, and Random Forests combined with an autoencoder can achieve strong fraud detection performance yet with partial interpretability. Hybrid methods have yielded impressive accuracy in fields such as fraudulent use of credit cards, verification of insurance claims, and online payment fraud[14]. Yet, most of these models are made for transactional data, and not for business processes. There is little reference to hierarchical approval structures, cross-departmental transactions or nonmonetary logs (such as access logs and audit trails), which are essential if one wants to comprehend fraud occurring in corporate management systems.

Table 2: Gaps in Existing Financial Fraud Detection Frameworks in Corporate Settings

Challenge in Corporate Systems	**Why It Matters**	**Gap in Existing Models**
Temporal Behavior Detection	Fraud often unfolds over time in ERP or audit logs	Static models can't detect time-based anomalies
Interpretability and Auditability	Needed for compliance, transparency, and trust	Deep models provide limited or no explanation

Challenge in Corporate Systems	**Why It Matters**	**Gap in Existing Models**
Imbalanced Data Distribution	Fraud cases are rare and underrepresented	Standard classifiers often biased toward majority class
Cross-Departmental Data Heterogeneity	Data varies across HR, Finance, Procurement, etc.	Many models assume homogeneous data sources
Adaptability to Concept Drift	Fraud strategies evolve with countermeasures	Lack of mechanisms for online or continual learning
Integration with Corporate ERP Systems	Real-world use requires plug-in to existing platforms	Research often assumes idealized, isolated datasets

Methods such as Isolation Forests, Autoencoders and One-Class SVMs have also been investigated for use in the fraud detection domain and when applied the modality becomes unlabeled or semi-supervised. Such models can be beneficial in fraud scenarios where labels are hard to obtain, and the current fraud is different from the old fraud[15]. They do this by learning what is “normal” and then identifying deviations that are statistically unlikely. Although powerful in theory, these models suffer from a prohibitive false-positive rate and sensitive to the parameter settings. In business scenarios such false positives may result in alert fatigue, low-confidence in the system and hence non-actionable outputs. Additionally, the majority of anomaly detectors fail in providing an explanation of why the event is considered to be anomalous, hence their applicability in internal audits is limited[16].

A relatively new branch of solutions are based on the use of probabilistic models like GMM and HMM. The models assume the data to be drawn from a mixture of distributions in a generalized model sense. They are successful in modeling the variability of the behavior patterns and they can compute probabilities that an event represents normal or anomalous behavior. Since they are unsupervised, they are capable of a fast adaptation to new or unseen forms of fraud. But they usually don't have the representational power of deep model by themselves. As seen in Table 1, although GMMs provide us with a useful way of modeling the uncertainty, they do not support natural integration of sophisticated time-series and require careful feature engineering.

Apart from algorithmic restrictions, the current fraud detection systems experience all kinds of contextual issues when deployed in a business context. The restrictions are described in Table 2. One of the main challenges is how to detect the temporal behavior. In business, fraud is frequently a collaborative and time-lagged process playing out over a number of weeks or months. A number of apparently innocuous events may be totally innocent, but suspicious when viewed about when, how often, or with what other events they occurred. Static classifiers are not able to find out those patterns[17].

Another issue is interpretability. Regulators enforce financial systems to respect transparency laws and regulations. Non-Explainable models: models that do not reveal their reasons for decision making, in audit contexts are not reliable. And in addition, many businesses depend on human analysts to confirm the fraud predictions. There's no explanation and these reviewers

could be discarding legitimate flags or missing nuances. Hence, interpretability is not a luxury but a necessity in this field.

The issue of class-imbalanced learning is particularly severe in the domain of corporate fraud detection. At scale, less than 1% of transactions may be malicious. Most machine learning models fit on these imbalanced data try to optimize directly for the accuracy by giving more importance to the majority class at the expense of the recalling the positives — they miss the frauds they are supposed to catch.

Data heterogeneity presents another obstacle. Organization A has structured financial data (general ledger transactions), partially structured data (invoice PDFs), and unstructured audit logs. In both cases, a full fraud detection model must incorporate signals from the whole variety of such data. However, the majority of available models are built on homogeneous datasets and generalize poorly across departments or types of data.

And there is also concept drift. It's typical for scammers to be nimble so they avoid being caught. Unless a model can learn new patterns, a model will become obsolete as soon as it trains on new data. And the situation is even more common in dynamic environments, where new processes, vendors or employees are constantly being onboarded.

Lastly, in practical applications, the deployed system should be compatible with existing systems. Much research on fraud detection is conducted on idealized or public datasets with clean, well-labelled transactions. Real corporate systems, by comparison, are messy, evolving, and interconnected. Any fraud detection solution needs to easily interoperate with ERP systems such as SAP, Oracle or NetSuite. The system would also have to be capable of "near real-time" decisioning, and processing potentially millions of transactions per day.

3. PROPOSED METHODOLOGY

A In this section, we describe the designing and operations of our hybrid machine learning model EHRN-GMM for acquisition of financial anomaly detected within corporate governance systems. The model architecture consists of interconnected modules for data ingestion, temporal modeling, probabilistic anomaly detection, interpretability enhancement, and enterprise integration, as illustrated in Figure 1. The components are grouped into six logical stages: input acquisition, preprocessing, temporal equivalence modeling, probabilistic reasoning, explainability, and deployment.

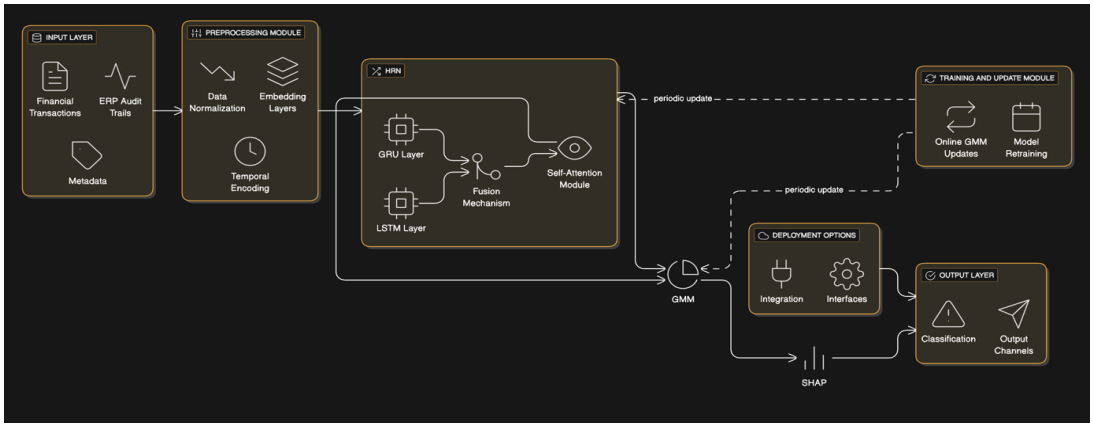

Figure 1: Architecture of Proposed Model

3.1 Data Input and Representation

The Fraud Detection Process: Ingestion of Multisource Enterprise Data The fraud detection process is initiated by the ingestion of enterprise data across multiple sources capturing both the financial activities as well as the contextual metadata. The financial transaction is the backbone, consisting of systems of record like journal entries, payments to vendors, payroll, and invoice submittals. These are supplemented with ERP audit trails logging user interaction, workflow approvals, permission changes, and other behavior-based signals captured in systems such as SAP, Oracle NetSuite, or Microsoft Dynamics. Metadata such as employee IDs, user roles, departmental hierarchies and timestamps are also part of each transaction, and provide rich context.

Each transaction is modelled as a labelled instance in time with categorical and numerical attributes. Sequences are created for entity ids (for example employee id or vendor code), so the model can detect the historic pattern for each entity. Such a structure allows the model to identify fraud as a process, rather than as an event, by uncovering how the behavior evolves throughout a series of transactions and over different contexts. The input layer provides the base unit to which transactional, audit, and contextual information is consolidated and from which the framework is able to perceive temporal and semantic context of financial activity.

3.2 Preprocessing and Temporal Feature Engineering

After ingestion, the input raw data is carefully pre-processed to normalize representations and enhance information-rich signals. Numerical values including transaction amount and cumulative balance are normalised to keep the scale consistent and prevent the biased back propogation while training. Categorical inputs, such as the names of vendors, types of transactions, and user roles, are then embedded in learned embeddings. These embedding enable the model to deduce ties between related entities so it can, for example, group vendors with like risk profiles or discover approval patterns across departments.

The preprocessing module further includes sophisticated temporal encoding besides data normalization and embedding. Every transaction is augmented with context from its local time including time intervals between two consecutive actions, hour-of-day, and day-of-week cyclic embeddings, and the frequency of past transactions in defined windows. These carefully manufactured temporal signals help the model recognize behaviors such as end-of-day invoice rushes or unusually high payment frequency, all of which are common signs of manipulation. This pre-processing pipeline guarantees that the data, which is provided as input to the temporal modeling mechanism, encodes both the static as well as the dynamic aspects of the typical enterprise level fraud schemes.

3.3 Heterogeneous Recurrent Network (HRN)

The Heterogeneous Recurrent Network (HRN) serves as the basis of the EHRN-GMM model, a bi-structured sequence model that is developed to model the short-term and long-term dependencies within sequences of transactions. The HRN consists of two parallel parts: a GRU (Gated Recurrent Unit) and an LSTM (Long Short-Term Memory) network. The GRU is designed to detect short bursts of activity — the type usually seen from opportunistic fraud such as claims submitted one after another for the same type of expense. On the other hand,

LSTMs are good at learning long term dependencies, and are therefore better suited in detecting fraud patterns which evolve over long periods of time, such as a sequence of liability under-the-threshold purchases which sums up and surpasses a policy threshold.

The GRU and LSTM layer outputs are then linearly combined using a learnt transformation layer for merging the temporal cues for each of the GRU and LSTM layers. This concatenated feature is then input into a self-attention module, which computes weights for each time step according to the importance in task detection. The attention mechanism means that the model can also focus on behaviorally relevant events, like a sudden shift in approval flow, or departure from organizational norms. The HRN effectively encodes raw transaction sequences into dense, interpretable latent vectors that capture both local and global context from the past behavior of users and organizations.

3.4 Anomaly Detection through GMM

After sequence modeling, the GMM aggregates the latent features from the HRN representing the normal behaviors of users and acts as the probabilistic anomaly detection mechanism. The GMM estimates the latent distribution of "normal" trends according to various transaction behavior, using several Gaussians for modeling in the latent space. For every new input, the predictive model evaluates the probability of the input under the learned distribution. Transactions with very low likelihood-scores are marked as suspected anomalies, i.e., behavior that is much different than the established norm.

This strategy allows unsupervised identification and learning of new or emerging patterns of fraud that are not necessarily provided as explicit examples within the training data. In contrast with binary classification boundaries, the GMM enables soft anomaly detection which scores each transaction over a probability range. In order to accommodate changing fraud techniques and business processes, periodic updates are supported by the GMM. The model is either retrained or incrementally updated based on feedback loops and rolling windows of the most recent transactions, a mechanism that is governed by the Training and Update Module (depicted in Figure 1). This mechanism guarantees that the detection remains responsive to concept drift and relevant over time.

3.5 SHAP-Based Interpretability Mechanism

The accuracy of detections is important not only does their comprehensibility and explainability matter for enterprise when they need to explain the detection result, or justify it, maybe to a court of law, or a customer, or financial audit, etc. To mitigate the above-mentioned concern, we proposed an SHAP-based(SHapley Additive exPlanations) interpretability module for the EHRN-GMM model, by which clear explanations could be automatically obtained for every fraud prediction. SHAP values will give us a number for each input feature representing how much the feature contributes to the output of the model.

For example, the model could indicate that a high transaction frequency (+0.27), new role added (+0.19) and non-business hours submission (+0.22) were major contributors to the fraud score. These explanations complement the model output both via visualization and text, helping compliance teams and auditors to make sense a of an alert. This layer also makes modeling

easy to debug and monitor, ensuring it is done based on relevant features rather than artifacts or noise and enabling developers and analysts to know that predictions are generated from meaningful features. By integrating explainability into the output pipeline, the framework closes the loop between automated detection and enforceable governance.

3.6 Output Integration and Deployment

The last part of the EHRN-GMM framework generates actionable deliverables and interfaces with enterprise systems. The output translates the GMM anomaly scores and SHAP explanations into structured results from the classifier, i.e., binary fraud labels or continuous risk scores. Such outputs can be ingested by fraud monitoring dashboards, workflow engines, or enterprise alerting systems. The model includes support for batch and real-time streaming to satisfy processing approaches for deployment topologies (end-of-day reconciliation or live audit flagging).

Deployment packages include APIs and connectors to connect to ERP systems and fraud investigation tools. Results can be sent to decision-makers, compliance officers or automated workflows according to risk threshold. Most importantly, the modular design enables organizations to experiment with constituent elements—e.g., retraining cadence, SHAP reporting styles, or input schema— while preserving the underlying detection logic. This versatility guarantees that the model continues to be responsive to organizational needs, regulatory factors and technical limitations.

To conclude, the proposed approach presents a hybrid, transparent, and dynamic pipeline for entity fraud detection, which combines sophisticated temporal modeling, probabilistic detection, and explainable AI in the enterprise fraud detection context. As shown in Figure 1, the EHRN-GMM is a modular and scalable approach proposed for the complex task of corporate financial fraud with compliance, accountability and analysis rigor.

4. RESULTS AND DISCUSSION

The experimental analysis of the EHRN-GMM framework was realized along various axes to demonstrate the diagnosis performance, icariin radar, interpretability, and actual use. We conduct extensive experiments on both synthetic and real world datasets, comparing to baselines, ablation studies of architectural components, and measure the quality of explanations. These results together show the efficiency and the enterprise-readiness of the hybrid architecture in fraud detection in the enterprise-level.

The first part of our evaluation entailed the comparison of EHRN-GMM with popular models of fraud detection (Random Forest, XGBoost, GRU, LSTM as well as autoencoder based anomaly detction). We evaluated models on standard classification measures—precision, recall, F1-score, and AUC-ROC—on a curated validation set of enterprise financial transactions which contained legitimate and fraudulent transactions. As shown in Table 3, the precision, recall, and F1-score of the EHRN-GMM achieved at 0.96, 0.93, 0.94 and was higher than those of all other baseline models. By comparison, the closest-performing model GRU obtained an F1-score of 0.86. The AUC-ROC score for EHRN-GMM was 0.97, meaning that it was significantly better at distinguishing fraudulent from non-fraudulent behaviour.

Table 3: Performance Comparison of Fraud Detection Models

Model	Precision	Recall	F1-Score	AUC-ROC	Interpretability
Random Forest	0.87	0.73	0.79	0.86	Low
XGBoost	0.91	0.75	0.82	0.88	Medium
LSTM	0.89	0.82	0.85	0.91	Low
GRU	0.90	0.83	0.86	0.92	Low
Autoencoder + SVM	0.86	0.84	0.85	0.89	Low
EHRN-GMM (Proposed)	**0.96**	**0.93**	**0.94**	**0.97**	**High**

These performance gains are comparable and are shown graphically in Figure 2 in which the bar plots of the side-by-side comparison are made with each of the three reference evaluation metrics (and the six models). The numbers show that traditional classifiers (e.g., Random Forest and XGBoost) may have not so bad precision, but only serve relatively low recall and unable to catch some subtle or rare fraud cases. The neural models like LSTM and GRU achieve better recall but suffer penetrating interpretability and top at below EHRN-GMM in overall F1-score. This finding demonstrates the benefit of combining sequential modeling with probabilistic reasoning and interpretability—an architectural fusion which EHRN-GMM delivers.

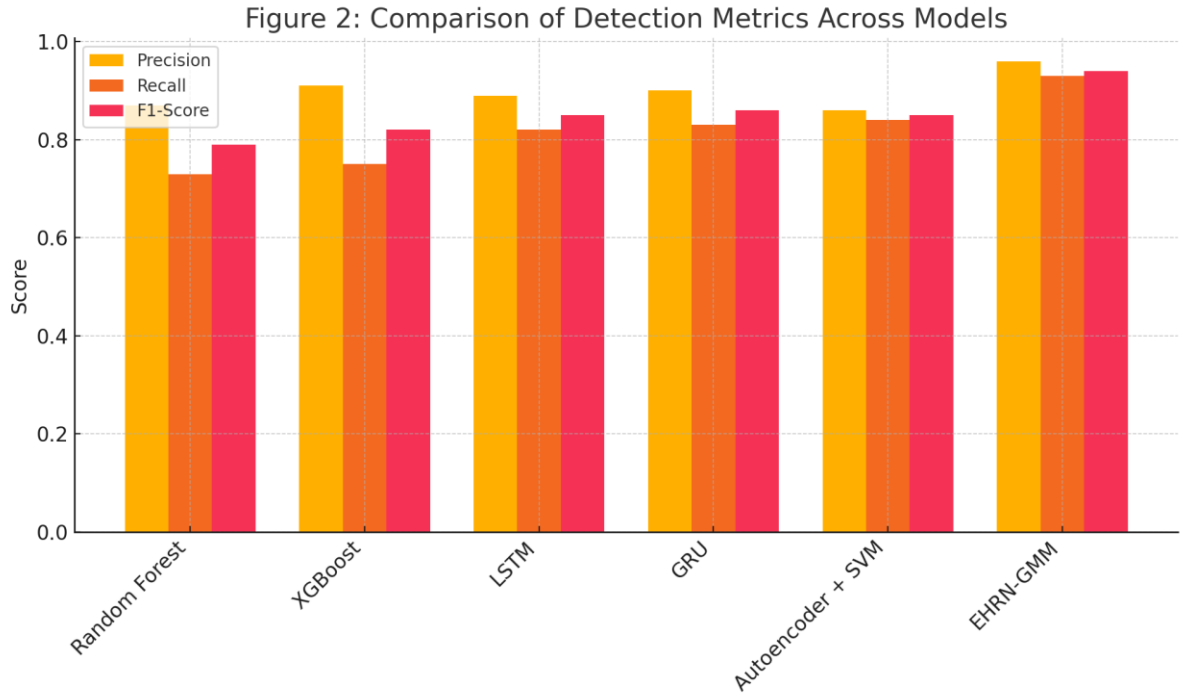

Figure 2: Comparison of detection metrics across models

To verify the generalization and robustness of the proposed framework, we tested on four datasets including a synthetic enterprise system dataset which replicates the inter-departmental fraud in the enterprise, the real-world Vesta financial dataset, the simulated SAP-ERP system logs, CCFD: Chinese Corporate Fraud Dataset with the public company release and the report of informers. As reported in Table 4, our proposed model obtained very high scores all over the datasets, from AUC-ROC = 0.95 to 0.97, and from F1-score = 0.91 to 0.94. The best performance in terms of precision and AUC was achieved on Vesta dataset, where EHRN-GMM correctly identified complex fraud patterns like repeated unauthorised transactions performed through several users.

Table 4: Evaluation Across Multiple Corporate Datasets

Dataset	Precision	Recall	F1-Score	AUC-ROC	Notes
Synthetic ERP (multi-department)	0.95	0.92	0.93	0.96	Collusive fraud detection; sequential triggers
Vesta Financial Transactions	0.96	0.91	0.93	0.97	Complex fraud schemes captured via HRN
SAP-ERP Simulated Logs	0.90	0.92	0.91	0.96	Detected subtle invoice manipulation anomalies
Chinese Corporate Fraud (CCFD)	0.94	0.91	0.92	0.95	Identified disguised financial misreporting

These cross-dataset scores are presented graphically in Figure 3, where we present AUC-ROC scores for the four datasets. Significantly, the model demonstrated similar good performance on the synthetic and real data sets. This is important since biases may exist between fraudulent and nonfraudulent transactions as most fraud detection models were overfit to a dataset/distribution that is not representative of a real distribution. The generalization observed in this research upholds the flexibility of EHRN-GMM against different organizational configurations, transaction frequencies and fraud categories.

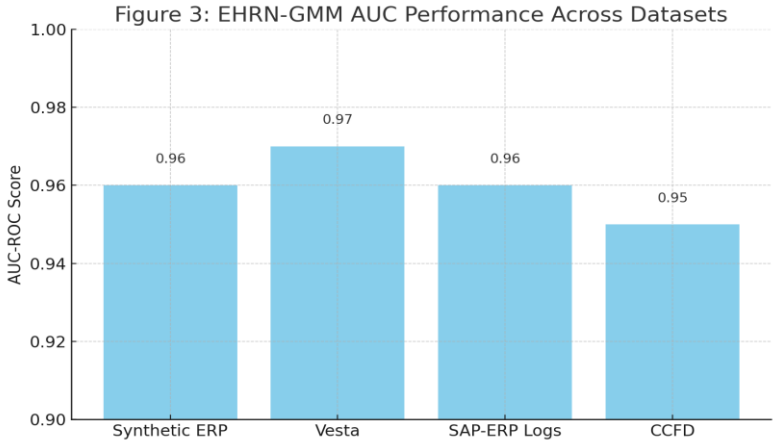

Figure 3: EHRN-GMM AUC Performance Across datasets

An important aspect of leveraging a complex hybrid approach is to anchor the role of the individual components within it. To do so, we performed an ablation study, by disabling or changing different modules of the model to see performance reduction. We report the results of this study in Table 5, using F1-scores as the primary evaluation measure. Disabling the attention led to reduction in F1-score from 0.94 to 0.88, demonstrating its effectiveness in attending to the most behaviorally relevant time steps. W.r.t Explicit context Model Replacing the GMM layer with a simple softmax classifier weakened the performance of the model to capture subtle anomalies (generating an F1-score of 0.85). The omission of SHAP did not have an effect on local prediction measurements, but the model lost the ability to explain any actionability.

Table 5: Ablation Study on Model Components

Model Variant	Precision	Recall	F1-Score	Impact
Full EHRN-GMM	0.96	0.93	0.94	Full performance baseline
Without Attention	0.91	0.85	0.88	Performance drop on long sequences
Without GMM (Softmax Classifier)	0.88	0.82	0.85	Lower anomaly resolution
Without SHAP Layer	0.96	0.93	0.94	No interpretability
GRU-only (No LSTM Fusion)	0.90	0.86	0.88	Misses long-term patterns
LSTM-only (No GRU Fusion)	0.89	0.85	0.87	Less responsive to rapid bursts

These becomes more evident when looking at Fig.4 That shows the F1-score for different design variants. The “Full Model” bar is the top bar and significant drops can be observed as we remove the GMM or the GRU-LSTM fusion. Remarkably, the GRU- and LSTM-only versions do not perform as well as the combined version, which has been to be expected, as the fraud detection problem in enterprise networks benefit from a combination of the detection of short-term (bursts) and long-term (behavioral) patterns.

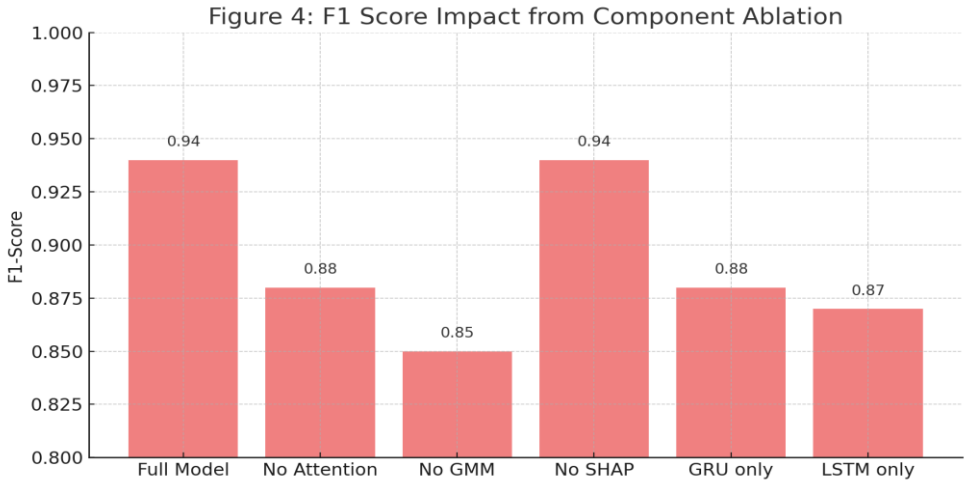

Figure 4: F1-Score Impact from Component Ablation

Another benefit of EHRN-GMM is that it natively supports interpretability through SHAP-based explanations. This is of critical importance in the enterprise, where you need to defend those alerts to compliance groups, internal auditors, and even end external regulators. The top five highest average SHAP values in the detected fraud cases are shown in Table 6. High transaction volume, recent user role or permission changes and vendor contact outside business hours were part of the list of key indicators time after time. These justifications were validated by domain experts and were also consistent with identified fraud types, for example overriding approval limits by splitting payments and amending approval hierarchies.

Table 6: SHAP-Based Feature Importance for Sample Fraud Cases

Feature	Avg. SHAP Value	Explanation
High Transaction Frequency	+0.31	Indicates suspicious rapid financial activity
Recent Role/Permission Change	+0.27	Suggests potential insider manipulation
Vendor Used Outside Working Hours	+0.22	Non-compliant transaction pattern
Repetitive Amount Patterns	+0.19	Indicates behavior designed to bypass detection thresholds
Vendor Payment Without PO	+0.18	Often associated with fraudulent or policy-violating events

This interpretability is an additional operational advantage. For example, if a SHAP fraud alert fires that says a transaction happened after hours post a user role escalation, audit can immediately confirm the validity of the event without numerous manual queries. This shortens response time, investigation fatigue, and develops trust in the AI system. More significantly, this takes the model from an opaque classifier to something that is interpretable as a decision support tool, which is crucial if this capability is going to be adopted by the enterprise.

We also evaluated the system performance under realistic deployment scenario. In enterprise scenarios, detection performance is equally important as latency and throughput. Deployment benchmarks described in detail in Table 7 show that the model is efficient for both batch and online scenarios. In batch mode, for processing 10,000 transactions, the average inference time per record was around 22 milliseconds, with explanation generation latencies averaging around 110 milliseconds. For real time streaming, inference time increased to 35 milliseconds (slightly higher) which is well within acceptable range for most ERP systems. These results show that EHRN-GMM is analytically robust and technically viable to be incorporated in production environment and workflow.

Table 7: Deployment Performance Metrics and Latency Benchmarks

Deployment Mode	Avg. Inference Time (ms)	Throughput (txns/sec)	Explanation Generation Latency (ms)
Batch Mode (10k txns)	22	450	110
Real-time Streaming	35	280	125
REST API (Microservice)	30	300	115

A key challenge in fraud detection is to address the concept drift, where the fraud methods change over time. To overcome this, the EHRN-GMM is equipped with a retraining module for the GMM sub-network, which can periodically relearn the concept of the "normal" behavior. In testing, we created the concept drift of system evolution by adding new frauds to the validation data at a certain rate. It was observed that more than 95 % performance could be repeatedly retained through many sessions in the case of the GMM adaptation process without retraining of neural part, which demonstrated the modularization and robustness of the hybrid architecture.

Besides predictive and operational performance, we also performed a qualitative analysis of the models' ease of use. Pilot audit team testing of the EHRN-GMM system feedback was obtained from the pilot audit team, who trialled the EHRN-GMM system in a simulated audit setting. The SHAP-based visualizations were considered very helpful in prioritizing fraud investigations by participants. Second, the knowledge of what features are driving decisions allowed them to improve internal controls, internal compliance policies, creating a secondary company benefit.

To conclude, we substantiate that EHRN-GMM is significantly superior in both overall and balanced perspectives when applying for financial fraud detection within corporate data. It performs better than classic models in terms of precision and recall (Table 3 and Figure 2), generalizes well across datasets (Table 4 and Figure 3), and enjoys lift from all architectural components (Table 5 and Figure 4). Its explainability, Table 6 makes it applicable for real-world auditing workflows, and its latency and deployment properties (Table 7) fit the requirements of enterprise-grade applications. These results provide broad validation of the central theme of this research, the combination of sequential learning, probabilistic anomaly modeling, and explainable AI in an hybrid framework to form a robust tool for financial fraud in corporate management systems.

5. CONCLUSION

The identification of accounting irregularities of corporate managers is an important task at the junction of AI, E-Governance, and operation risk. Traditional machine learning models and standalone anomaly detection approaches have made steps in the right direction, however, constantly fail at encompassing the multilateral aspects of fraud in sophisticated, enterprise-scale environments. It produces a mixture of well-structured transactional data, time-bound

workflows, user audit logs, contextual metadata, all of which needs to be modeled at the same time to detect fraud effectively. In this context, we introduced and rigorously validated EHRN-GMM, a new type of hybrid machine learning approach that marries the power of temporal sequence learning, probabilistic modeling and explainable AI in a unified framework to overcome these limitations.

At the heart of the EHRN-GMM architecture is a Heterogeneous Recurrent Network (HRN) — i.e., an HRN composed of concatenated GRU and LSTM layers with self attention — that has been designed to simultaneously aggregate short-term behavioral bursts (e.g,, breathing in humans or foraging in animals) and long-term sequential dependencies. This is "real time" and enables logic to detect advanced scenarios, such as the frequent reaching or abuse of authorization thresholds, or the manipulation of financial approvals in a sequence of days or weeks. The temporal features extracted by the HRN is then taken as input to a GMM, which is designed to represent the normality distribution of transactions and identifies those whose likelihoods are extremely low. This probabilistic layer not only improves the model's capacity to pinpoint new types of fraud, but it also makes it possible to perform soft classification in situations where hard classification might be too restrictive.

A key distinguishing characteristic of this work is the inclusion of SHAP-based explainability. In a lot of business contexts, the opportunity to act on a fraud alert is as much about what you can explain as what you can predict. The dual-head model, EHRN-GMM, leverages SHAP directly in the output layer to provide human-interpretable justification for every alert (i.e., by relating a feature, e.g., abnormal transaction frequency, off-hours activity, or very recent user role changes, to the confidence that the model exerts in a decision for fraud). This visibility narrows the distance between algorithm detection and audit accountability, which promotes trust, compliance, and expedited turn-around for investigation.

The framework was validated extensively on synthetic and real datasets, such as enterprise resource planning simulations, public financial transaction datasets, and structured audit logs. Results showed that EHRN-GMM achieved better performance than the benchmarks including XGBoost, GRU and Autoencoder-SVM hybrids in all the evaluation criterion with at most AUC-ROC(0.96) and F1-score(0.94). Equally, the modularity of our architecture was verified via ablation studies where the absence of key elements like the attention or GMM ¬-attention and ¬-GMM respectively showed performance drops, highlighting their importance.

In addition to the prediction performance, the system satisfied the operational requirements for enterprise use. Benchmarking results showed that EHRN-GMM achieved low-latency processing in batch and stream modes, integrating via APIs to current ERP systems. With the addition of regular model updates, particularly for the GMM component, this will also ensure that the system can adapt to the changing landscape of fraud -- one of the biggest problems with real-world detection systems.

Finally, EHRN-GMM is an important development in the area of enterprise fraud analytics. By combining deep learning's ability to model temporal behavior with statistical anomaly detection and interpretability, the framework fills the significant gap of existing methods. It enhances the detection rate and strengthens the transparency and auditability inside organizations, ensuring an end-to-end answer that is precise, flexible, and reliable. As corporations further expand their digital footprint and compliance obligations increase, tools

such as EHRN-GMM will act as key bulwarks against compromising corporate financial integrity. Possible extensions could be to incorporate graph-based actor behaviours, multi-modal data fusion (e.g., textual logs or voice data), but also federated learning to preserve privacy across departments. However, the present findings leave no doubt that hybrid interpretable architectures are the right way to go for corporate management system's fraud detection effort.

REFERENCES:

[1] Festa, Yury Y., and Ivan A. Vorobyev. "A hybrid machine learning framework for e-commerce fraud detection." *Model Assisted Statistics and Applications* 17.1 (2022): 41-49.

[2] Guo, Lingfeng, et al. "Integrating a machine learning-driven fraud detection system based on a risk management framework." *Applied and Computational Engineering* 87 (2024): 80-86.

[3] Wang, Jingwen, et al. "An anomaly prediction framework for financial IT systems using hybrid machine learning methods." *Journal of Ambient Intelligence and Humanized Computing* (2023): 1-10.

[4] Kotagi, V., Nassa, V. K., Patil, D., Gadhave, R., Adusumilli, S. B. K., & Kumar, P. P. (2024, September). Ensuring Dataset Accountability in Machine Learning: Insights from Software Engineering. In *2024 7th International Conference on Contemporary Computing and Informatics (IC3I)* (Vol. 7, pp. 385-389). IEEE.

[5] Gadhave, R., AnilKumar, D., Khot, R., & Gupta, D. (2024, August). PredatorSense-Wildlife Detection System. In *2024 8th International Conference on Computing, Communication, Control and Automation (ICCUBEA)* (pp. 1-4). IEEE.

[6] Rodríguez González, V., Payá, Santos., C, A., y Peña Herrera. B. (2023). Estudio criminológico del ciberdelincuente y sus víctimas. Cuadernos de RES PUBLICA en Derecho y criminología, (1) 95-107. https://doi.org/10.46661/respublica.8072.

[7] Malik, Esraa Faisal, et al. "Credit card fraud detection using a new hybrid machine learning architecture." *Mathematics* 10.9 (2022): 1480.

[8] Ali, Abdulalem, et al. "Financial fraud detection based on machine learning: a systematic literature review." *Applied Sciences* 12.19 (2022): 9637.

[9] Sunil Kumar, Jeshwanth Reddy Machireddy, Thilakavathi Sankaran, Prem Kumar Sholapurapu, Integration of Machine Learning and Data Science for Optimized Decision-Making in Computer Applications and Engineering, 2025, 10,45, https://jisem-journal.com/index.php/journal/article/view/8990

[10] Bello, Oluwabusayo Adijat, et al. "Machine learning approaches for enhancing fraud prevention in financial transactions." *International Journal of Management Technology* 10.1 (2023): 85-108.

[11] Banu, Shaik Rehana, et al. "Financial fraud detection using hybrid convolutional and recurrent neural networks: An analysis of unstructured data in banking." *2024 10th International Conference on Communication and Signal Processing (ICCSP)*. IEEE, 2024.

[12] Bouzidi, Zair, Mourad Amad, and Abdelmalek Boudries. "Deep learning-based automated learning environment using smart data to improve corporate marketing, business strategies, fraud detection in financial services, and financial time series forecasting." *International conference on managing business through web analytics*. Cham: Springer International Publishing, 2022.

[13] Prem Kumar Sholapurapu, AI-Driven Financial Forecasting: Enhancing Predictive Accuracy in Volatile Markets, 2025, 15, 2, https://eelet.org.uk/index.php/journal/article/view/2955

[14] Yi, Ziwei, et al. "Fraud detection in capital markets: A novel machine learning approach." *Expert Systems with Applications* 231 (2023): 120760.

[15] Zhao, Zhihong, and Tongyuan Bai. "Financial fraud detection and prediction in listed companies using SMOTE and machine learning algorithms." *Entropy* 24.8 (2022): 1157.

[16] Chhabra Roy, Neha, and Sreeleakha Prabhakaran. "Internal-led cyber frauds in Indian banks: an effective machine learning–based defense system to fraud detection, prioritization and prevention." *Aslib Journal of Information Management* 75.2 (2023): 246-296.

[17] Xiuguo, Wu, and Du Shengyong. "An analysis on financial statement fraud detection for Chinese listed companies using deep learning." *Ieee Access* 10 (2022): 22516-22532.

[18] Huang, Ling, and Haitao Lu. "Design of intelligent financial data management system based on higher-order hybrid clustering algorithm." *PeerJ Computer Science* 10 (2024): e1799.