Cureus Journal of Computer Science
https://doi.org/10.7759/s44389-026-00109-8

Original Article

A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning

Ankita N. Patil [1], Mritunjay K. Ranjan Mr. [1], Tushar Thorat [1], Vaidik Gharat [1], Samruddhi Sonawane [1]

1. *Department of Computer Sciences and Application, Sandip University, Nashik, IND*

Received: April 10, 2026 | Review began: April 29, 2026 | Review ended: May 28, 2026 | Published: June 08, 2026

Abstract

This paper presents a hybrid forensic auditing framework that integrates Large Language Models (LLMs), Graph Neural Networks (GNNs), Retrieval-Augmented Generation (RAG), and a rule-based Finite State Machine (FSM) inspired by Automata Theory to detect corporate fraud hidden within unstructured financial data. Traditional fraud detection systems primarily rely on structured data and predefined rules, which often fail to identify narrative inconsistencies and behavioral patterns present in audit trails, financial statements, and corporate communications. The proposed framework utilizes LLMs to identify semantic and linguistic indicators of deception, while the FSM models fraud progression through interpretable state transitions to support audit traceability and explainability. GNNs are employed to analyze relationships among entities and identify collusive activities within organizational networks, including suspicious communication and transaction patterns. In addition, the RAG module retrieves relevant financial regulations, audit precedents, and compliance guidelines to enhance contextual understanding and decision-making. By combining linguistic intelligence, relational reasoning, and rule-based symbolic analysis, the framework provides an explainable and scalable approach to corporate fraud detection. The framework was evaluated using real-world datasets, including U.S. Security and Exchange Commission filings and the Enron email corpus, and achieved an accuracy of 96%, precision of 95%, recall of 94%, F1-score of 94%, and area under the receiver operating characteristic curve of 0.94, demonstrating improved fraud detection performance, auditing transparency, and regulatory compliance.

Categories: Forensic Analysis, Security and Privacy, Machine Learning (ML)

Keywords: corporate fraud, large language models, graph neural networks (gnn), retrieval-augmented generation, automata theory

Introduction

Corporate fraud is one of the most enduring and harmful risks to financial robustness, organizational integrity, and stakeholder confidence [1]. Cooking the books, overstating liabilities, and collusion network are common aspects of fraudulent dealings, and its nature is such that only little tangible evidence may be left. Conventional systems of fraud detection are largely rule-based and structured data-driven, ill-suited to detect fraud in unstructured financial text, like in the audit report, other pieces of corporate correspondence, fraud-related legal filings, and so forth [2]. This shortcoming is especially important when a large part of the data about modern corporations involves free-text data, records of correspondence, and abnormal transaction records, which are likely the first places where problematic trends and language peculiarities can be spotted [3]. Recent developments in the area of artificial intelligence have led to the development of new possibilities in the field of forensic auditing since now machines can process and reason jointly across natural language, relational constructs and formal logic. The former have demonstrated amazing power over

semantic and linguistic cues to intentional deception, whereas Graph Neural Networks (GNNs), with their focus on discovering indirect connections and consorting relationships in organizations networks, have become essential [4]. In complement, the Automata Theory provides a formal structure which can be used to encode fraudulent event series onto auditable state translations allowing auditors to be aware of the narrative process that could have taken place during the development of possible frauds. Besides, Retrieval-Augmented Generation (RAG) methods enable the systems to factor into the process domain-specific financial regulations, case precedents, and compliance needs and greatly enhance contextual accuracy. In this paper, we introduce a hybrid neuro-symbolic "framework specifically designed for forensic auditing that integrates Large Language Models (LLMs), GNNs, RAG, and a rule-based Finite State Machine (FSM) inspired by Automata Theory, where each component contributes distinct functionalities for detecting corporate fraud in unstructured financial data. The proposed approach aligns the linguistics intelligence, the structured inference and the relational reasoning so as to address the gap between the statistical detection and the explainable audit reports. This framework is then tested on actual data sets, such as U.S. Securities and Exchange Commission (SEC) filings and the Enron email corpus and proves to be meaningfully better at detecting fraud, the transparency of the decision-making process and regulatory compliance as shown in Figure *1*.

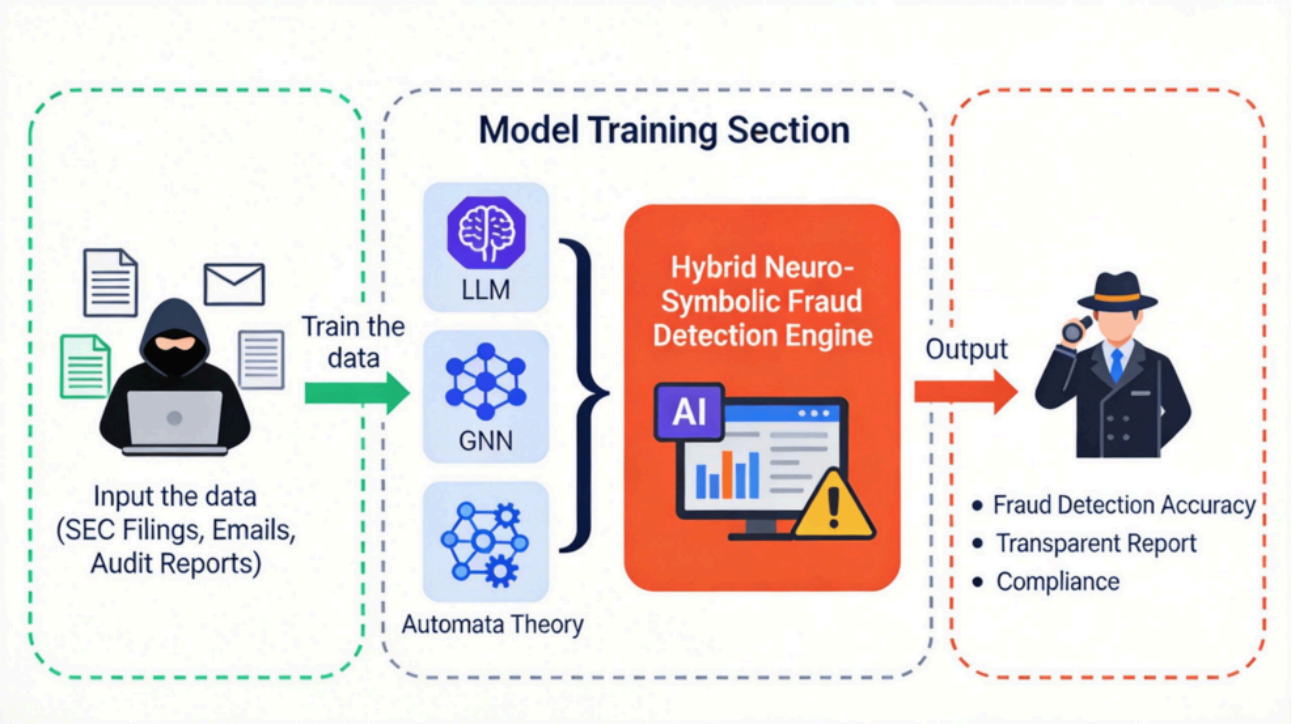

FIGURE 1: Conceptual framework of hybrid neuro-symbolic fraud detection system.

LLM, Large Language Model; GNN, Graph Neural Network, SEC, U.S. Securities and Exchange Commission; AI, Artificial Intelligence

Objective

· To design a hybrid neuro-symbolic forensic auditing system where LLMs, Automata Theory, GNNs, and RAG are incorporated into it to discover the corporate fraud residing in the unstructured financial data.

·To make the fraud detection more explainable and transparent by transferring linguistic and relational indicators of fraud into formal state-transition models in order to allow auditors to re-construct the development of the fraudulent activities.

·To confirm the validity of the scheme on real-life data sets like SEC filings and the Enron email collection to test the accuracy, scalability, and consistency with regulations vis-a-vis the other models of fraud detection like the rule based and the machine learning only schemes.

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

Related Study

Citation	Domain	Key contribution	Results	Advantages	Limitations	Research relevance
[5]	Supply-chain finance fraud detection	Proposes MultiFraud, a heterogeneous GNN framework with multitask learning and explainability mechanisms for fraud detection.	Demonstrated improved detection performance on heterogeneous relational datasets with interpretable outputs.	Captures complex entity relationships and supports explainable graph-based analysis.	Requires graph construction and sufficient labeled training data.	Highlights the effectiveness of GNN-based relational learning for fraud-oriented network analysis.
[6]	Graph-based transaction fraud detection	Introduces label-aware message passing for semi-supervised fraud detection in graph structures.	Achieved higher AUC and F1-score compared with conventional GNN baselines.	Improves learning under limited labeled data and class imbalance conditions.	Primarily evaluated on transaction datasets rather than corporate forensic datasets.	Demonstrates the applicability of adaptive GNN architectures for fraud detection tasks.
[7]	Retrieval-Augmented Generation in financial documents	Evaluates RAG-based pipelines for improving factual consistency in financial document question-answering systems.	Improved answer reliability and contextual grounding compared with standalone LLM approaches.	Reduces hallucination and improves regulatory contextualizatio n.	Retrieval quality and knowledge base maintenance significantly affect performance.	Supports the integration of contextual retrieval mechanisms in financial auditing systems.
[8]	Machine learning for financial statement fraud detection	Reviews recent ML-based approaches for financial fraud detection and identifies limitations in explainability and unstructured data analysis.	Summarizes methodological trends, dataset limitations, and evaluation challenges across multiple studies.	Provides broad analytical coverage of financial fraud detection methods.	Does not propose a unified explainable or hybrid analytical framework.	Identifies research gaps related to explainability and multi-modal fraud analysis.
[9]	LLMs for financial text	Investigates methods for	Demonstrated improved	Enhances reliability of	Focused primarily on	Demonstrates the usefulness

	analysis	improving factual consistency and information extraction reliability in financial text processing.	precision and reduced hallucination in earnings-report analysis tasks.	financial text interpretation using LLM-based techniques.	information extraction rather than end-to-end fraud detection.	of LLM-based semantic analysis for financial document interpretation.

TABLE 1: Comparative analysis of existing research studies in fraud detection, graph-based learning, retrieval-augmented generation, and explainable AI approaches.

GNN, Graph Neural Network; RAG, Retrieval-Augmented Generation; LLM, Large Language Models; ML, Machine Learning; AUC, Area under the Curve

Problem statement

Corporate fraud still poses great economic, reputational and legal harm on organizations but most of current forensic auditing systems rely heavily on structured data, fixed rules, and limited anomaly detection algorithms [10]. These approaches do not detect deceptive indicators that are embedded within unstructured financial narratives, inter-entity relationships and regulatory environments where the precursors of fraudulent acts are likely to manifest themselves first. Moreover, traditional machine learning models are primarily focused on pattern recognition and often lack interpretability, transparency, and effective integration across heterogeneous data sources [11]. Such a gap translates to misses in detection, poor explainability and diminished confidence on the part of auditors. A scalable and interpretable fraud detection framework is required to address the multi-modal and complex nature of corporate fraud in modern financial environments. Therefore, this study proposes a hybrid forensic auditing framework that integrates linguistic analysis, relational reasoning, contextual regulatory retrieval, and interpretable fraud progression modeling to improve fraud detection accuracy, transparency, and regulatory compliance. The effectiveness of the proposed framework is evaluated using measurable performance metrics, including accuracy, precision, recall, F1-score, and area under the receiver operating characteristic curve (AUC-ROC), on real-world datasets such as SEC filings and the Enron email corpus in Figure *2*.

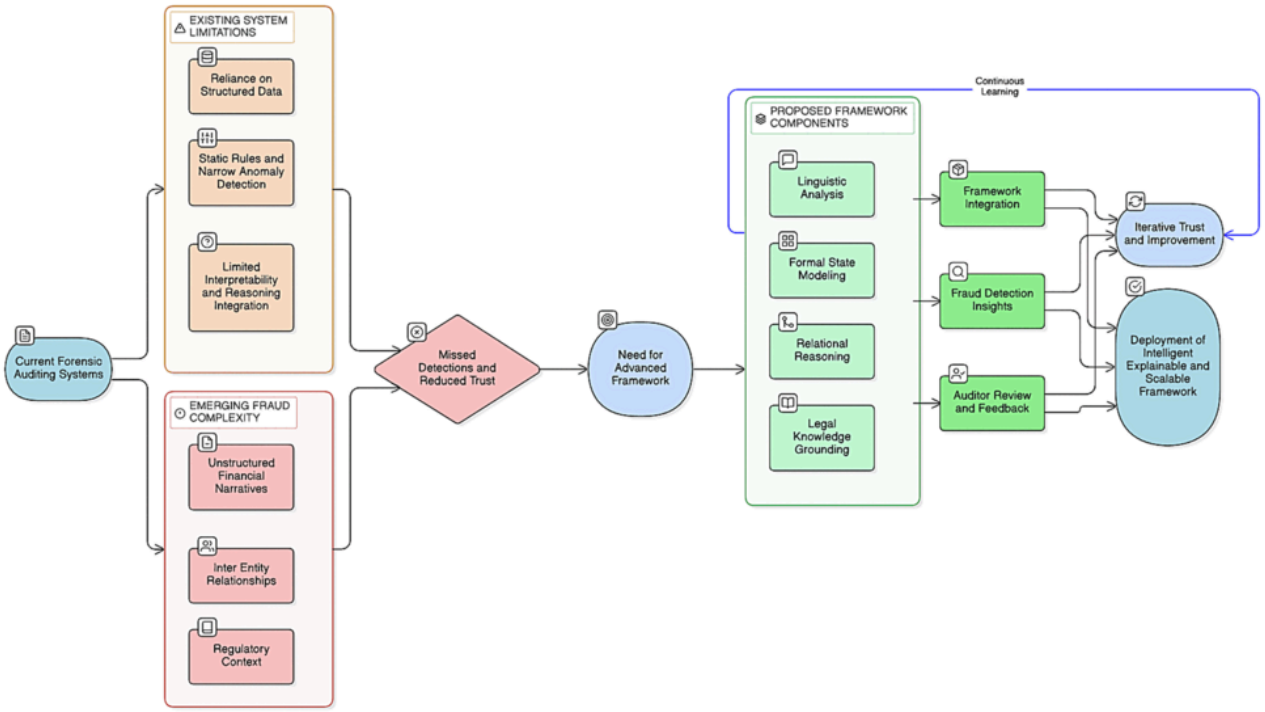

FIGURE 2: Proposed neuro-symbolic framework for scalable and explainable fraud detection.

S1: Data Collection def collect_data(): r eturn load_corporate_datasets() # SEC filings, emails, audit reports, transactions S2: Preprocessing def preprocess_data(raw_text, structured_data): text = text_cleaning_pipeline(raw_text) # Tokenization, masking return extract_entities_events(text, structured_data) S3: Dataset Split def split_dataset(features, labels): return stratified_split(features, labels, ratio= (0.7, 0.15, 0.15)) S4: Initialize Modules def initialize_hybrid_modules(): return define_fraud_fsm(), build_regulatory_index(), init_gnn_architecture(), load_finetuned_llm() S5: Train Core Models def train_core_models(train_X, train_Y):	**gnn = train_gnn(train_X['graph'], train_Y) llm = finetune_llm(train_X['text'], train_Y) weights = learn_fusion_weights(train_X, train_Y) return gnn, llm, weights S6: Integrated Training def train_hybrid_system(fsm, rag, gnn, llm, weights, train_X, train_Y, val_X, val_Y): return hybrid_training_loop(fsm, rag, gnn, llm, weights, train_X, train_Y, val_X, val_Y, epochs=50) S7: Evaluation def evaluate_system(fsm, rag, gnn, llm, weights, test_X, test_Y): preds = hybrid_inference(fsm, rag, gnn, llm, weights, test_X) return compute_metrics(test_Y, preds) S8: Deployment def deploy_system(fsm, rag, gnn, llm, weights): save_components(fsm, rag, gnn, llm, weights, "HybridFraudDetection") export_for_cloud_or_onprem()**

TABLE 2: Methodological workflow of hybrid fraud detection system.

SEC, U.S. Securities and Exchange Commission

Materials And Methods

The FSM was designed as an interpretable rule-based component to model the progression of fraudulent activities across multiple analytical stages commonly observed in forensic auditing and financial fraud investigations. The selected states, namely, Clean, Narrative Inconsistency, Obfuscation, Collusion Detected, Material Misstatement, Fraud Confirmed, Cleared, represent sequential behavioral and analytical indicators associated with potential corporate fraud progression. The "Clean" state denotes normal operational behavior without detectable anomalies, while "Narrative Inconsistency" captures semantic irregularities identified through LLM-based textual analysis. The "Obfuscation" state represents concealment-related activities, including misleading disclosures or hidden transactional behavior. "Collusion Detected" corresponds to anomalous relational patterns identified through GNN-based entity interaction analysis, and "Material Misstatement" reflects significant financial reporting inconsistencies or regulatory violations. The terminal states, "Fraud Confirmed" and "Cleared," represent final investigative outcomes based on integrated module outputs. These states were defined using domain knowledge derived from forensic auditing workflows, fraud investigation practices, and financial compliance analysis to improve interpretability, audit traceability, and explainable fraud progression modeling within the proposed framework. in Table *2* [12].

In this data collection phase, heterogeneous sources are pooled together so that the framework is able to reveal fraud patterns in the multi-modalities [13] as shown in Table *1*. The dataset combines raw text, more structured numerical information, and regulation expertise to offer language and relational background to fraud detection. Such multimodal set empowers the system to measure implicit narrative disparities, transactions abnormalities, and violation deviation as shown in Figure *2* depicts in Table *3* [14].

Data type	Source	Purpose in framework
Unstructured text	SEC filings, Enron emails, audit reports	LLM extracts deception cues & inconsistencies
Structured tabular data	Ledger entries, payment records, vendor databases	GNN detects anomalous transactions & relationships
Regulatory knowledge base	GAAP, IFRS, SOX regulations, SEC rules, and financial case records	RAG retrieves relevant regulations & precedents
Metadata & Context	Audit metadata, time windows, org hierarchies	FSM models fraud evolution via state transitions

TABLE 3: Mapping of data sources to functional roles in fraud detection.

LLM, Large Language Models; GNN, Graph Neural Network; RAG, Retrieval-Augmented Generation; FSM, Finite State Machine; GAAP, Generally Accepted Accounting Principles; SOX, Sarbanes-Oxley Act; SEC, U.S. Securities and Exchange Commission; IFRS, International Financial Reporting Standards

The bar graph in Figure *3* framework integrates multiple data modalities, including unstructured financial text, structured transactional records, regulatory documents, and contextual metadata, to support comprehensive fraud analysis. Unstructured textual data contributes semantic and linguistic indicators of deceptive behavior, while structured transactional data supports the identification of anomalous financial activities. Regulatory documents provide contextual

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

compliance knowledge for retrieval-augmented analysis, and metadata contributes additional temporal and organizational context. Together, these complementary modalities enhance the interpretability and analytical capability of the proposed hybrid fraud detection framework.

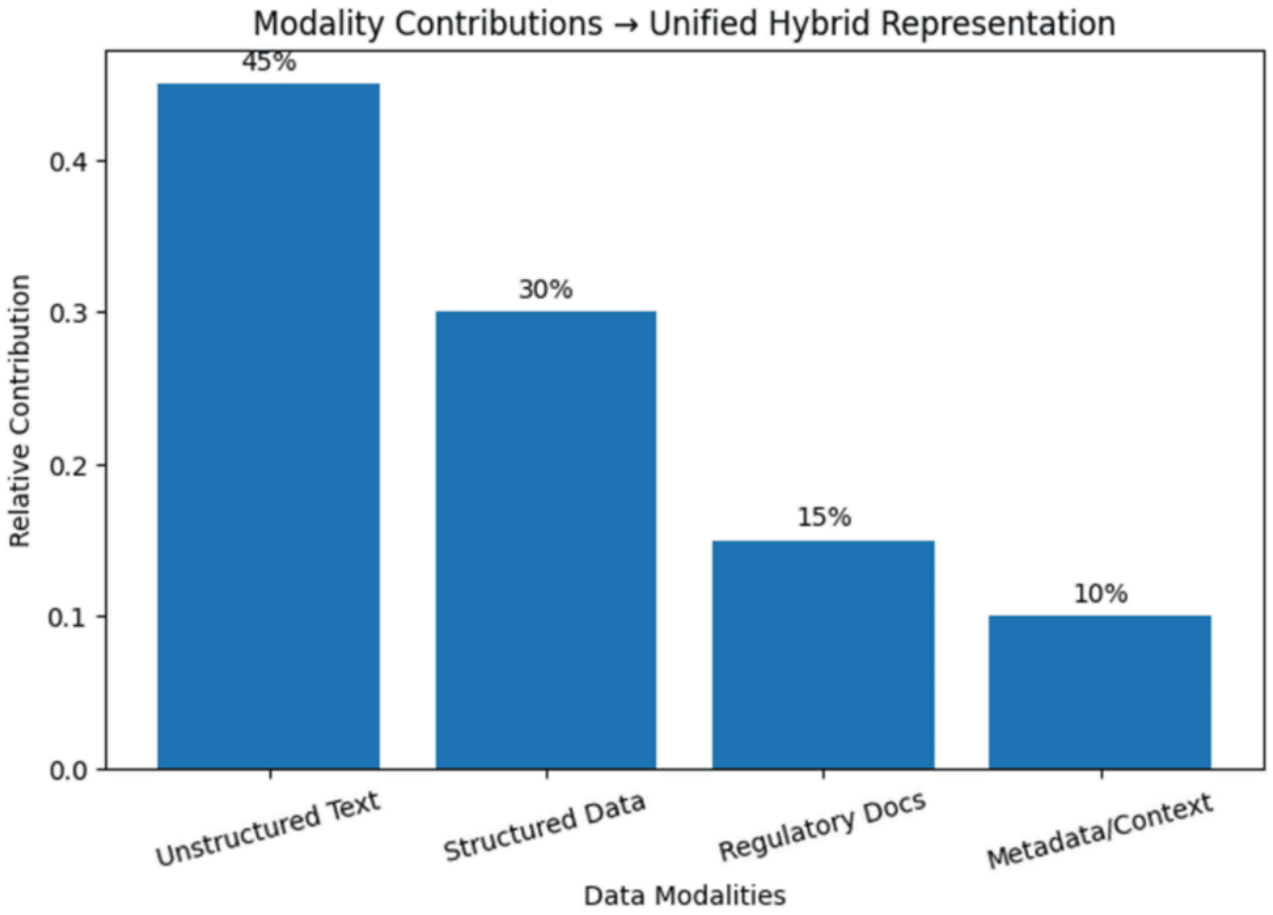

FIGURE 3: Distribution of modality importance in unified representation.

The preprocessing of the data so that the heterogeneous corporate data size and composition, including unstructured narratives, structured financial transactions, and domain knowledge are transformed into a unified and feature-rich representation adequate to the hybrid LLM-FSM-GNN-RAG framework [15]. This is to eliminate noise, normalize format, standardize modalities, and sanitize the data to extract multimodal features and reason at the graph level. The cleaned dataset Dclean is derived from the raw dataset drawn through sequential transformations as shown in Equation 1:

$$(D_{\text{clean}} = f_{\text{text}}(D) \oplus f_{\text{num}}(D) \oplus f_{\text{graph}}(D)) \quad (1)$$

Where,

- D: Raw input dataset (unprocessed data)

- $f_{\text{text}}(\cdot)$: Text preprocessing function (e.g., tokenization, cleaning, NLP feature extraction)

- $f_{\text{num}}(\cdot)$: Numerical preprocessing function (e.g., normalization, scaling, missing value handling)

- $f_{\text{graph}}(\cdot)$: Graph-based transformation function (e.g., relationship modeling, feature aggregation)

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

- D_{clean}: Final cleaned and structured dataset ready for modeling

Draw represents the original multi-modal dataset containing unstructured text, structured transactions, and metadata, while D_{clean} denotes the processed dataset used for training and inference. The functions $f_{text}(\cdot)$, $f_{num}(\cdot)$, and $f_{graph}(\cdot)$ correspond to text preprocessing (tokenization, embedding), numerical normalization, and graph construction, respectively. The operator $\oplus$ denotes the concatenation of features from different modalities into a unified representation. Text Processing & Embeddings Unstructured corporate documents are tokenized, lemmatized, and stripped of stop words and personally identifiable information (PII). Each token w_i is mapped into a semantic vector via the LLM encoder ϕ_{LLM} and pooled into a document embedding as shown in Equation 2:

$$(e_i = \phi_{\mathbf{LLM}}(w_i), \qquad z_{\mathbf{tx}} = \mathbf{Pool}(e_{1:T}) \in \mathbb{R}^{d_t}) \tag{2}$$

Here,

- w_i: Input token (word/subword at position i)

- $\phi_{LLM}(\cdot)$: Embedding function of the Large Language Model (LLM)

- e_i:Token embedding vector for w_i

- $e_{1:T}$: Sequence of embeddings from token 1 to T

- T: Total number of tokens in the sequence

- Pool($\cdot$): Pooling operation (e.g., mean, max, CLS token)

- z_{tx}: Aggregated text representation (sentence/document embedding)

- R^{d_t}: d_t-dimensional vector space

- d_t: Dimension of text embedding

This representation captures deception cues, semantic inconsistencies, and fraud-relevant context. Numerical Normalization & Encoding Structured data, such as transaction amounts or time intervals, is scaled to a consistent range using Min-Max normalization as shown in Equation 3:

$$\tilde{x} = \frac{x - \min(x)}{\max(x) - \min(x)} \tag{3}$$

Here,

- x: Original data value

- min(x): Minimum value in the dataset

- max(x): Maximum value in the dataset

- $\tilde{x}$: Normalized value (scaled between 0 and 1)

Figure 4 illustrates the correlation between module output scores generated by the LLM, GNN, FSM, and RAG components across evaluation samples within the proposed hybrid fraud detection framework. Moderate correlations were observed between FSM and RAG (0.75), GNN and FSM (0.70), and LLM and GNN (0.72), indicating partial overlap in the information captured by different analytical modules. These correlations suggest that certain fraud-related patterns may be simultaneously reflected through linguistic, relational, contextual, and rule-based analyses. However, each module still contributes distinct analytical capabilities within the framework. The combined integration of semantic

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

analysis (LLM), relational pattern detection (GNN), interpretable fraud progression tracking (FSM), and contextual regulatory retrieval (RAG) supports a more comprehensive and explainable fraud detection process, which is reflected in the improved overall performance of the hybrid model.

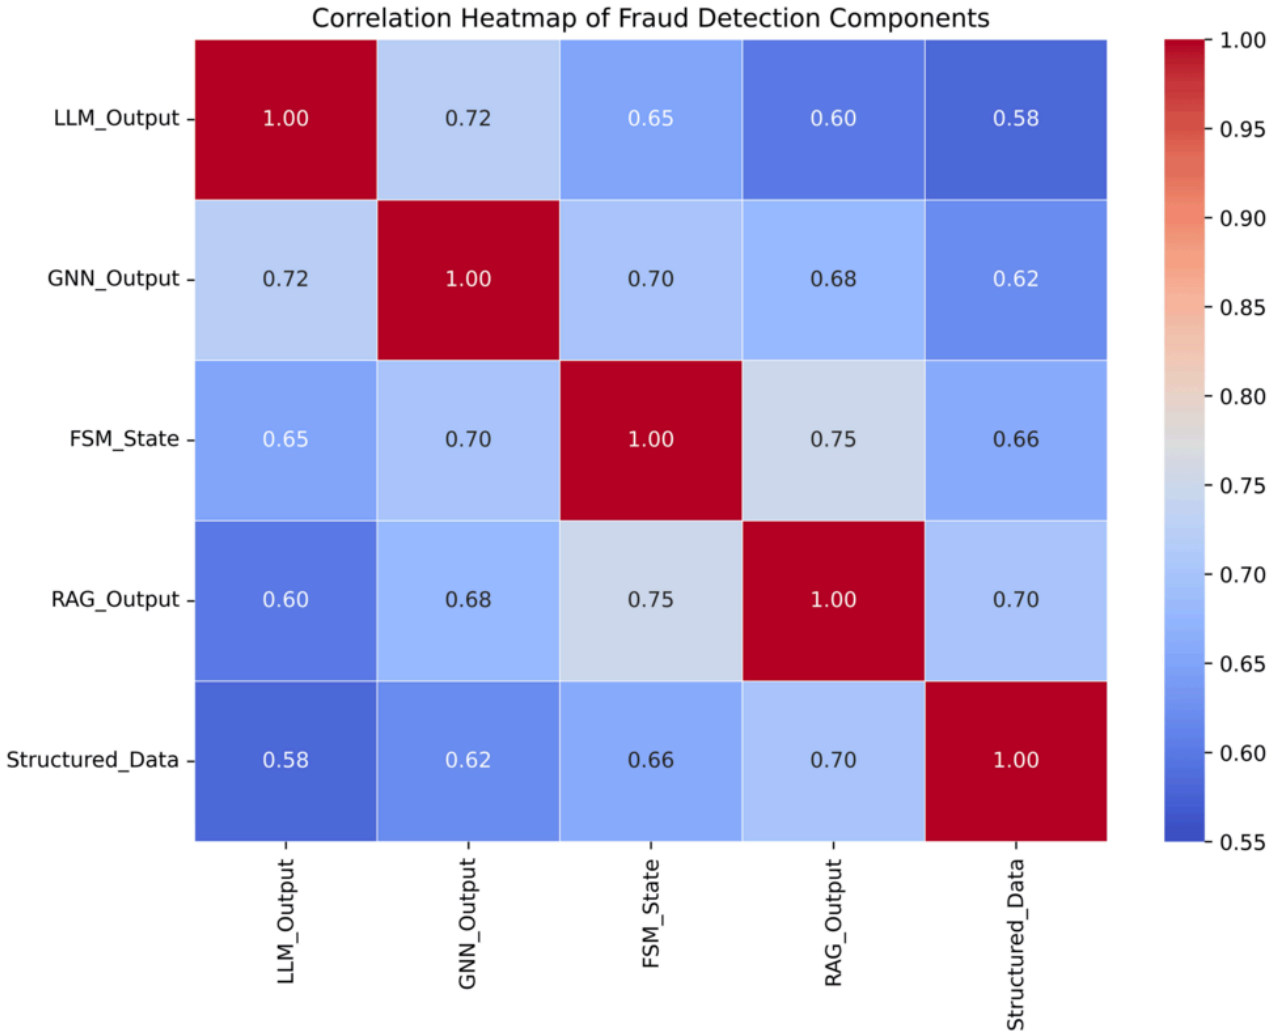

FIGURE 4: Correlation heatmap of hybrid fraud detection components.

GNN, Graph Neural Network; RAG, Retrieval-Augmented Generation; LLM, Large Language Models; FSM, Finite State Machine

Event Linking in events from text (e.g., contract signing) and transactions (e.g., payment execution) are aligned temporally and relationally to form cross-modal pairs as shown in Equation 4:

$$L = (e_i^{txn}, e_j^{comm}) \mid |t_i - t_j| \leq \Delta t \ \wedge \ \mathrm{Match}(e_i^{txn}, e_j^{comm}) = 1 \tag{4}$$

where,

- L: Set of linked event pairs

- $(e_i^{txn}$: Transaction event i

- e_j^{comm}: Communication event j

- $t_i - t_j$: Timestamps of events

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

- Δt Time threshold for linking events

- $|\ t_i - t_j\ | \leq \Delta t$: Temporal proximity condition

- Match(e_i, e_j): Matching function (semantic/behavioral similarity condition)

This linkage allows the automata layer to reconstruct fraud progression sequences over time. Graph Construction Entities and their relationships are modeled as a heterogeneous graph as shown in Equation 5:

$$G = \left(\mathcal{V}, \mathcal{E}, \mathbf{X}^V, \mathbf{X}^E\right), \quad \mathbf{h}_v = \left[\mathbf{z}_{\text{text}}(v) \,\|\, \mathbf{z}_{\text{num}}(v) \,\|\, \mathbf{z}_{\text{cat}}(v)\right] \tag{5}$$

Where,

- G: Graph representation

- V: Set of nodes (vertices)

- $\mathcal{E}$: Set of edges

- X^V: Node feature matrix

- X^E: Edge feature matrix

- *v*: A node in the graph

- h_v: Final feature vector of node *v*

- $z_{text}(v)$: Text-based embedding of node *v*

- $z_{num}(v)$: Numerical feature embedding

- $z_{cat}(v)$: Categorical feature embedding

- $\|$: Concatenation operator

The GNN uses this structure to detect collusion patterns and abnormal inter-entity interactions. Feature Fusion The final feature vector for each case is obtained by concatenating text, numerical, and categorical embeddings as shown in Equation 6:

$$\boldsymbol{f} = \left[\mathbf{z}_{\text{text}} \,\|\, \mathbf{z}_{\text{num}} \,\|\, \mathbf{z}_{\text{cat}}\right] \in \mathbb{R}^{d_t + d_n + d_c} \tag{6}$$

Where,

- *f*: Final fused feature vector

-z_{text}: Text feature embedding

- z_{num}: Numerical feature embedding

- z_{cat}: Categorical feature embedding

- $\|$: Concatenation operator

- d_t: Dimension of text features

- *dn*: Dimension of numerical features

- d_c: Dimension of categorical features

- $R^{d_t+d_n+d_c}$: Combined feature space

To ensure balanced representation of fraud and non-fraud cases, the dataset is partitioned into training, validation, and testing subsets using stratified sampling. This method maintains the proportion of positive (fraud) and negative (non-fraud) samples across all subsets, preventing bias and improving generalization as shown in Table 4 and as shown in Equation 7.

$$\mathbf{D_{train}}, \mathbf{D_{val}}, \mathbf{D_{test}} \xleftarrow{\text{stratify}} \frac{|\mathcal{D}_k^+|}{|\mathcal{D}_k|} \approx \frac{|\mathcal{D}^+|}{|\mathcal{D}|}, \quad k \in \{\mathbf{train}, \mathbf{val}, \mathbf{test}\} \tag{7}$$

where (D) denotes the complete dataset, $\boldsymbol{D^+}$ denotes the set of positive fraud samples, $\boldsymbol{D_k}$ denotes each dataset subset, and $\boldsymbol{D_k^+}$ denotes the positive fraud samples within subset (k). This equation ensures that the class distribution is approximately preserved across the training, validation, and testing subsets.

Subset	Proportion	Purpose
Training (D_{train})	70%	Model training and weight optimization
Validation (D_{val})	15%	Hyperparameter tuning and early stopping
Testing (D_{test})	15%	Final performance evaluation and generalization check

TABLE 4: Dataset partitioning strategy for model development.

At this step, the four main constituents of the analysis core of the put forward hybrid forensic auditing framework are launched and set up to collaborate with each other: the FSM, the RAG index, the GNN, and the LLM. The modules have specialized functions, but they all work together and lead to explainable and context-sensitive fraud detection outcomes. Let the hybrid module set be as shown in Equation 8. The hybrid analytical framework consists of four primary modules represented as:

$$H = \{F, R, G, L\} \tag{8}$$

where (F) denotes the rule-based FSM, (R) represents the RAG module, (G) denotes the GNN, and (L) represents the LLM. These modules operate collaboratively to perform linguistic analysis, relational reasoning, contextual knowledge retrieval, and interpretable fraud progression tracking within the proposed hybrid fraud detection framework.

F is the FSM that gets initialized with a set of states S, transition function δ, set of start states 0, and a set of accept states A_0.The feature importance analysis Figure 5 illustrates the relative contribution of different attributes in fraud detection, where transaction-related variables such as transaction amount, transaction frequency, and geographic distance emerged as dominant predictors. Behavioral indicators like failed login attempts and contextual metadata also demonstrated strong influence, highlighting their role in distinguishing fraudulent from legitimate activities. By ranking features based on their importance scores, the model not only enhances interpretability but also justifies the selection of key attributes, ensuring that subsequent training focuses on the most discriminative signals for accurate and robust fraud detection.

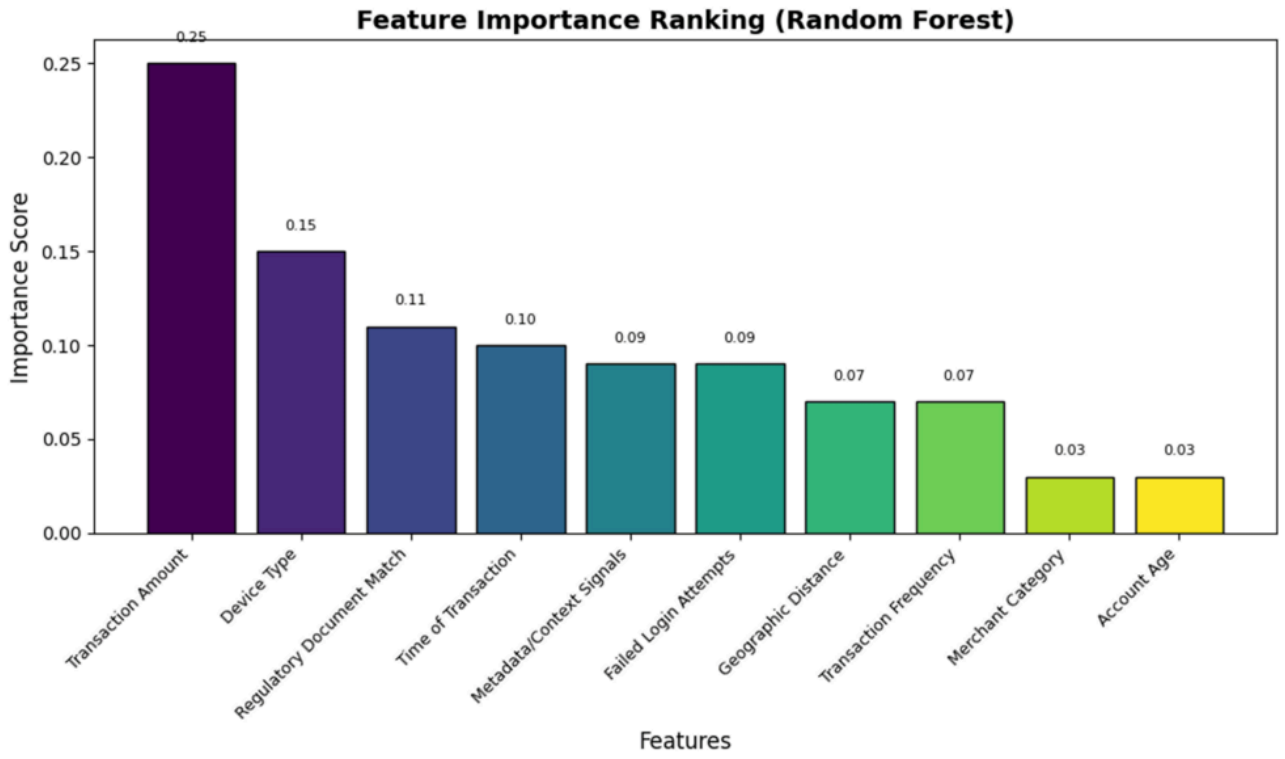

FIGURE 5: Feature importance distribution in fraud detection (Random Forest).

R is the RAG index created based on financial regulations, compliance documents, and case precedents; G is the heterogeneous GNN trained with architecture Ag, layer depth l, and feature dimensions df; and L is the fine-tuned LLM that takes the responsibility to find linguistic cues related to fraud.

Hybrid module initialization

In this step, the four foundational elements of analysis within proposed hybrid forensic auditing framework are initialized and optimized to be able to operate in tandem: the FSM based on Automata Theory, the RAG index, the GNN, and the LLM [16]. Automata Theory Integration of the FSM is modeled as a Deterministic Finite Automaton (DFA) as shown in Equation 9:

$$F = (Q, \Sigma, \delta, q_0, F) \tag{9}$$

Where,

- Q: Set of states

- Σ: Input alphabet

- δ: Transition function

- q_0: Initial state

- F: Set of final (accepting) states

Q is the finite set of fraud states; Σ the alphabet of observable events from LLM, GNN, and RAG; δ: Q×Σ→Q the transition function; q_0 the "Clean/No Fraud Detected" start state; and F the accept states: "Fraud Confirmed" or "Case Cleared." Q= {S0 : Clean, S1: Narrative Inconsistency, S2: Obfuscation, S3: Collusion Detected, S4: Material Misstatement, SA:

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

Fraud Confirmed, SR: Cleared}. Fraud Detection Transition Diagram The Figure *6* is depicted as a stepwise sequence in which a transaction in its individual state progresses through other states depending on prior set conditions and in carrying out the fraud analysis in Table *5*. The Normal state commences a transaction and during the transaction, when the transaction value is more than a threshold value (1) the transaction enters Credential Abuse Detected. Positive verification shows it back to Normal, whereas high fraud indicators detected by an AI move it to the state of High-Risk Purchase [17]. The decision is then made here: a low score (<0) means that the transaction is considered Cleared and logged; a high score (0≥) results in Fraud Confirmed, and to blocking and the escalation of the case. This formal flow guarantees constant surveillance, risk-judging decision making, and success in anti-fraud management.

Current state	Input condition/Analytical trigger	Associated module	Action performed	Next state
Clean	No anomaly detected	LLM/GNN/RAG	Continue monitoring	Clean
Clean	Linguistic inconsistency detected in financial text	LLM	Flag semantic anomaly	Narrative Inconsistency
Narrative Inconsistency	Concealment-related disclosure patterns identified	LLM + RAG	Increase fraud suspicion level	Obfuscation
Obfuscation	Suspicious entity interaction or collusive relationship detected	GNN	Trigger relational anomaly analysis	Collusion Detected
Collusion Detected	Significant reporting irregularity or compliance violation identified	RAG + FSM	Escalate fraud investigation	Material Misstatement
Material Misstatement	Fraud risk score exceeds decision threshold	Fusion Layer	Confirm fraudulent activity	Fraud Confirmed
Material Misstatement	Insufficient evidence of fraud	Fusion Layer	Mark case as cleared	Cleared

TABLE 5: Finite state transition table for fraud investigation.

RAG, Retrieval-Augmented Generation; GNN, Graphic Neural Network; LLM, Large Language Models; FSM, Finite State Machine

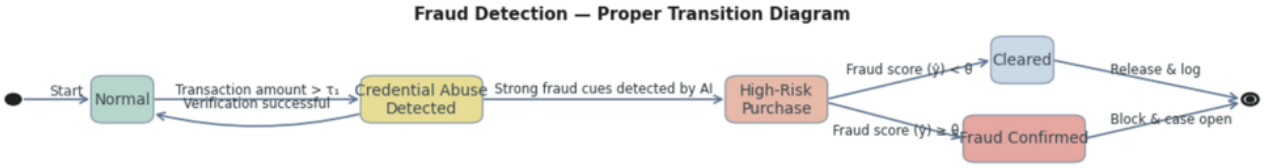

FIGURE 6: Finite state transition table for fraud investigation.

AI, Artificial Intelligence

Train core models

This stage focuses on training the analytical components GNN, LLM, and the Fusion Layer while aligning their outputs with the FSM from Automata Theory and the RAG system. Each model is optimized independently but designed to interoperate in the hybrid inference process as shown in Figure *7* [18].

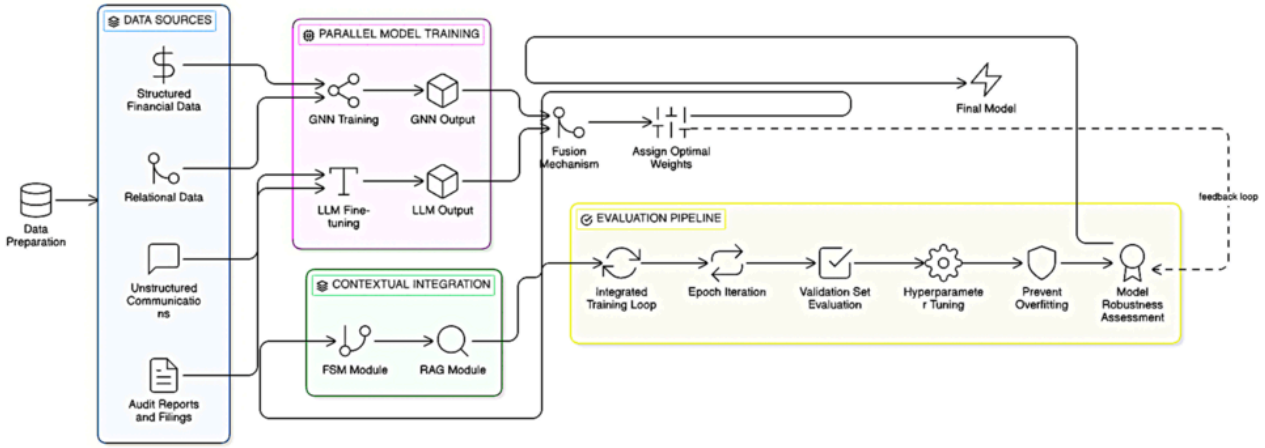

FIGURE 7: System workflow from data preparation to final Fraud Detection Model.

GNN, Graph Neural Network; RAG, Retrieval-Augmented Generation; LLM, Large Language Models; FSM, Finite State Machine

GNN training

The GNN is also trained on the heterogeneous entity event graph G constructed during preprocessing and presented in Eq 10:

$$G = (\mathcal{V}, \mathcal{E}, X^V, X^E) \tag{10}$$

Where,

-$\mathcal{V}, \mathcal{E}$: Nodes and edges of graph

- X^V: Node features

- X^E: Edge features

The model learns a node classification function as shown in Equation 11:

$$f_\theta : \left(\mathcal{V}, \mathcal{E}, X^V, X^E\right) \rightarrow y_v \in \{\text{fraud}, \text{non-fraud}\} \tag{11}$$

Where,

- f_θ: Model (parameterized by θ\thetaθ)

- $\mathcal{V}, \mathcal{E}$: Nodes and edges of graph

- X^V,X^E: Node and edge features

- y_v: Predicted label for node *v*

- {fraud,non-fraud}: Output classes

Parameters θ are optimized by minimizing the binary cross-entropy loss as shown in Eq 12:

$$\mathrm{L}_{\mathcal{GNN}} = -\frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} \left[y_v \log\left(\hat{y}_v\right) + (1 - y_v) \log\left(1 - \hat{y}_v\right) \right] \tag{12}$$

Where,

- $\mathbf{L}_{\mathcal{GNN}}$: Loss function (Binary Cross-Entropy for GNN)

- V: Set of nodes

- $|\,V\,|$: Total number of nodes

- *v*: A node in the graph

- y_v: True label (0 or 1)

- $\hat{y}_v$: Predicted probability for node v

The GNN outputs relationship anomaly scores, which are discretized into FSM alphabet symbols.

LLM fine-tuning

The LLM is fine-tuned on labeled corporate narratives to detect linguistic deception cues [19]. Given token embeddings e1: T for a document d as shown in Equation 13:

$$\hat{y}_d = \sigma\left(W_{\text{cls}} \cdot \text{Pool}\left(e_{1:T}\right) + b_{\text{cls}}\right) \tag{13}$$

Where,

- $\hat{y}_d$: Predicted output (probability)

- $\sigma(\cdot)$: Activation function (e.g., sigmoid)

- W_{cls}: Classification weight matrix

- b_{cls}: Bias term

- $\text{Pool}\left(e_{1:T}\right)$: Pooled embedding of sequence

- $e_{1:T}$ Token embeddings from position 1 to T

The observation symbols from LLM, GNN, and RAG, δ the transition function, q_0 the "Clean" start state, and F the accept states: "Fraud Confirmed" and "Cleared." The transition function delta is defined using expert-driven rule-based transitions derived from observed fraud analysis patterns and outputs generated by the LLM, GNN, and RAG modules [20,21]. The proposed neuro-symbolic fusion combines the outputs of all modules shown in Equation 14:

$$\hat{y} = \sigma\left(w_1 s_{\text{LLM}} + w_2 s_{\text{FSM}} + w_3 s_{\text{GNN}} + w_4 s_{\text{RAG}} + b\right) \tag{14}$$

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

Here,

- $\hat{y}$: Final predicted output

- σ(•): Activation function (e.g., sigmoid)

- w_1, w_2, w_3, w_4: Learnable weights

- s_{LLM}: Score from Large Language Model

- s_{FSM}: Score from Finite State Machine

- s_{GNN}: Score from Graph Neural Network

- s_{RAG}: Score from Retrieval-Augmented Generation

- b: Bias term

Once GNN, Linear Layer Memory, and the Fusion Layer (S5) are trained independently, the framework is integrated so as to train all of FSM, RAG, GNN, and LLM. The stage is used to match outputs, cal further thresholds, and optimize end-to-end detection. In every iteration, the LLM will give a linguistic score numerical measure called an anomaly score (sLLM), the GNN a relational score (sGNN), and the RAG a compliance score (sRAG) [22]. They are converted into FSM symbols yielding a state-based score (sFSM) that is combined with sLLM, sGNN, and sRAG by the Fusion Layer in totaling up to a single fraud risk score ().

Joint loss function

The integrated objective combines module-level losses and the fusion classification loss as shown in Equation 15:

$$\mathcal{L}_{\text{total}} = \alpha \mathcal{L}_{\text{LLM}} + \beta \mathcal{L}_{\text{GNN}} + \gamma \mathcal{L}_{\text{fuse}} + \lambda \mathcal{R}_{\theta} \quad (15)$$

Where,

- $\mathcal{L}_{\text{total}}$: Total loss

- $\mathcal{L}_{\text{LLM}}$: Loss from LLM module

- $\mathcal{L}_{\text{GNN}}$: Loss from GNN module

- $\mathcal{L}_{\text{fuse}}$: Fusion/ensemble loss

- $\mathcal{R}_{\theta}$: Regularization term

- $\alpha, \beta, \gamma, \lambda$: Weighting coefficients

The joint loss combines for linguistic cue detection, LGNN for graph-based fraud detection, and for final classification, with $R(\theta)$ as a regularization term. Weighting coefficients α, β, γ, λ are tuned on the validation set to balance each component's contribution.

FSM integration in training

The FSM operates as a rule-based interpretability and fraud progression tracking component within the proposed framework. The FSM does not dynamically update its transition function during model training; instead, predefined state transitions are triggered using outputs generated by the LLM, GNN, and RAG modules.

Calibration and threshold optimization

After module-level processing, the output scores ((s_{LLM}, s_{GNN}, s_{RAG}, s_{FSM})) are combined using a logistic regression-based fusion layer. Decision thresholds and temperature scaling parameters are optimized on the validation dataset to improve calibration performance and maximize the F1-score of the final fraud detection system.

To gauge the predictive performance as well as interpretability, the trained hybrid within the parameters of a held-out test set is done. The outputs of each of the modules (sLLM, sGNN, sRAG, sFSM) and the overall output of the whole system, which is the score of the fraud risk are evaluated with standard classification measures accuracy, precision, recall, F1-score, and AUC-ROC in order to measure quality of detection [23].

System architecture

The envisioned architecture of the system is a modular, microservice-based pipeline that parses multimodal corporate data through three sequential stages of preprocessing (structured text, numerical, and graph representation) and subsequently pipes that data through four major intelligence modules: (1) a super-fine-tuned LLM to detect linguistic signs of deception; (2) a GNN that recognizes anomalous relations between any entities; (3) an RAG engine, which grounds the ends of these modules (sLLM, sGNN, sRAG, sFSM) are combined into a freely interpretable fraud risk score *y*), with explanatory artifacts, such as an FSM state history, highlighted text strings by the LLM, GNN subgraphs, and regulatory passages it has quoted. It is architected to enable real-time and batch processing, is implemented as secure containerized services, and integrates continuous monitoring and drift detection and human-in-the-loop correction feedback to preserve accuracy, interpretability, and regulatory compliance over the long-term as shown in Figure *8*.

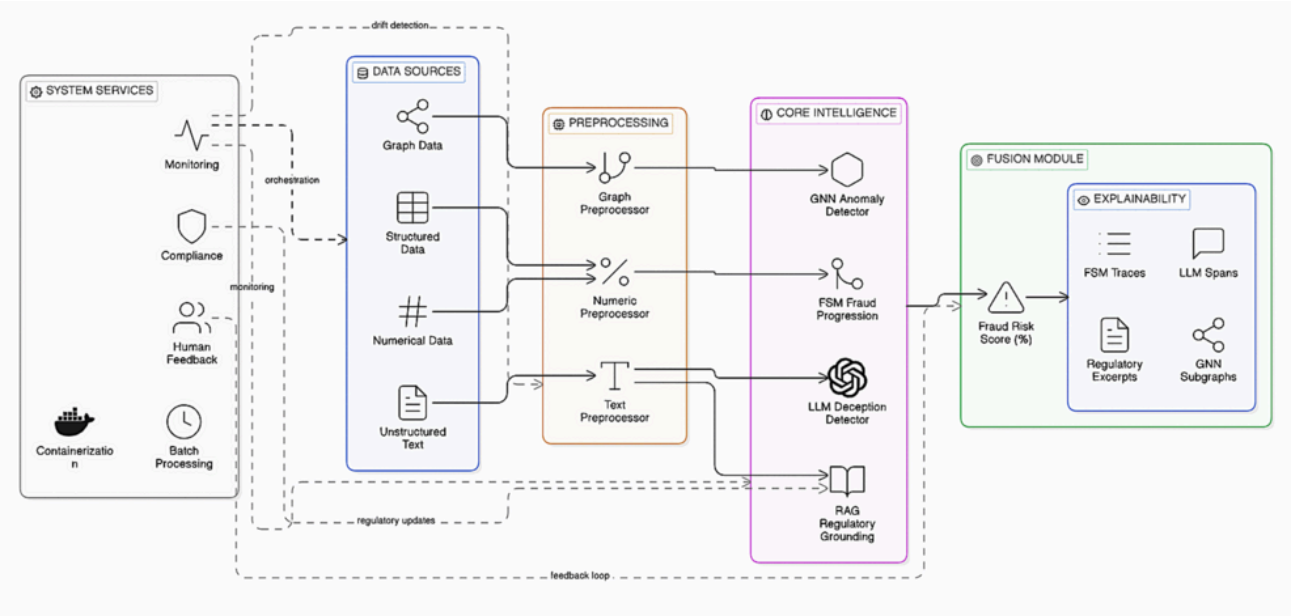

FIGURE 8: System architecture of hybrid neuro-symbolic fraud detection framework.

LLM, Large Language Models; GNN, Graph Neural Network; RAG, Retrieval-Augmented Generation; FSM, Finite State Machine

Results And Discussion

Experimental setup

To rigorously evaluate the effectiveness of the proposed hybrid neuro-symbolic fraud detection framework, a comprehensive experimental setup was designed encompassing dataset selection, preprocessing strategies, model configuration, training protocols, and evaluation metrics.

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

Dataset Description

The experimental evaluation was conducted using a combination of real-world and benchmark datasets to ensure robustness and diversity. The SEC Financial Filings dataset provides rich unstructured textual data such as corporate reports and disclosures, while the Enron Email Corpus offers communication data useful for detecting anomalous patterns and collusion. The experimental evaluation was conducted using real-world datasets, including SEC filings, audit reports, corporate emails, and financial transaction records containing structured attributes such as transaction amounts, timestamps, and account metadata. Collectively, the dataset represents a multi-modal structure comprising unstructured textual data, structured financial records, and contextual metadata, enabling comprehensive fraud analysis within the proposed framework.

Data Preprocessing

The preprocessing pipeline was designed to ensure consistency and feature richness across heterogeneous data modalities. Textual data was processed through tokenization, stop-word removal, and lemmatization, followed by Named Entity Recognition (NER) to extract relevant entities such as persons and organizations. Semantic embeddings were generated using a fine-tuned LLM encoder. Structured data was normalized using Min-Max scaling, and categorical variables were transformed using one-hot encoding. Furthermore, entities and events were mapped into a heterogeneous graph structure, where nodes represent entities (e.g., accounts, employees) and edges capture interactions such as transactions and communications.

Model Configuration

The proposed hybrid framework integrates four key components to enable comprehensive fraud detection. A fine-tuned LLM is employed to extract semantic and linguistic fraud cues from unstructured text. A GNN, specifically a multi-layer Graph Convolutional Network (GCN), is used to capture relational dependencies and detect anomalous interactions among entities. A FSM, modeled as a Deterministic Finite Automaton (DFA), represents the progression of fraudulent activities through state transitions. Additionally, a RAG module retrieves relevant regulatory knowledge (e.g., GAAP, SEC rules) to provide contextual grounding and improve decision reliability.

Training Strategy

The dataset was partitioned using stratified sampling into training (70%), validation (15%), and testing (15%) sets to maintain class balance. The models were trained using the Adam optimizer with a learning rate of 0.001, batch size of 32, and 50 epochs. Binary cross-entropy loss was used for training the LLM and GNN modules, while a joint loss function was applied for the integrated hybrid framework. Outputs from all modules (LLM, GNN, FSM, and RAG) were combined using a weighted fusion mechanism, where optimal weights were learned using the validation set to balance performance and interpretability in Table *6*.

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

Parameter	Value
Learning rate	0.001
Batch size	32
Epochs	50
Optimizer	Adam
Hidden dimension	128
GNN layers	2
Dropout rate	0.3
Loss function	Binary Cross-Entropy
Fusion threshold	0.5

TABLE 6: Hyperparameter configuration of the proposed hybrid fraud detection framework.

GNN, Graph Neural Network

Evaluation Metrics

The performance of the proposed framework was evaluated using standard classification metrics, including accuracy, precision, recall, F1-score, and AUC-ROC. In addition, confusion matrices were used to analyze class-wise performance, ROC curves were employed to assess discrimination capability, and loss curves were monitored to evaluate training stability and convergence behavior.

Baseline Models for Comparison

To validate the effectiveness of the proposed hybrid model, it was compared against several baseline approaches, including a traditional rule-based fraud detection system, an LLM-only model, a GNN-only model, and a combined LLM+GNN model. These comparisons provide a clear benchmark to demonstrate the advantages of integrating neuro-symbolic components for improved fraud detection performance.

Implementation Environment

All experiments were implemented in Python using PyTorch for deep learning model development, Scikit-learn for evaluation, and NetworkX for graph processing. The models were executed on a GPU-enabled system with NVIDIA CUDA support. The experiments were conducted on a workstation equipped with an NVIDIA RTX 3080 GPU with 10GB VRAM, 32GB system RAM, and an Intel Core i7 processor., ensuring efficient training and scalability for large-scale datasets.

The binary confusion matrix was evaluated using 5,000 test samples to measure the fraud versus non-fraud classification performance of the proposed framework. The results demonstrated strong diagonal dominance between correctly predicted fraud and non-fraud instances, indicating high classification accuracy and effective integration of the LLM, GNN, FSM, and RAG modules within the hybrid fraud detection framework. There is some overlap, especially between Unstructured Text and Structured Data, and Metadata/Context and LLM Output indicating the natural semantic and

contextual relationships-existing between financial documents and audit trails. On the whole, the outcomes demonstrate the resilience of the system to process multi-modal inputs and the high accuracy of its most important aspects, such as regulatory retrieval, anomaly detection in graphs, and fraud state models, which support the conclusion regarding its scalability and reliability to be utilized in the real-world forensic audit as shown in Figure *9* as shown in Equation 16.

$$\text{Accuracy} = \frac{TP + TN}{TP + FP + FN + TN} \tag{16}$$

Here,

- TP: True Positives

- TN: True Negatives

- FP: False Positives

- FN: False Negatives

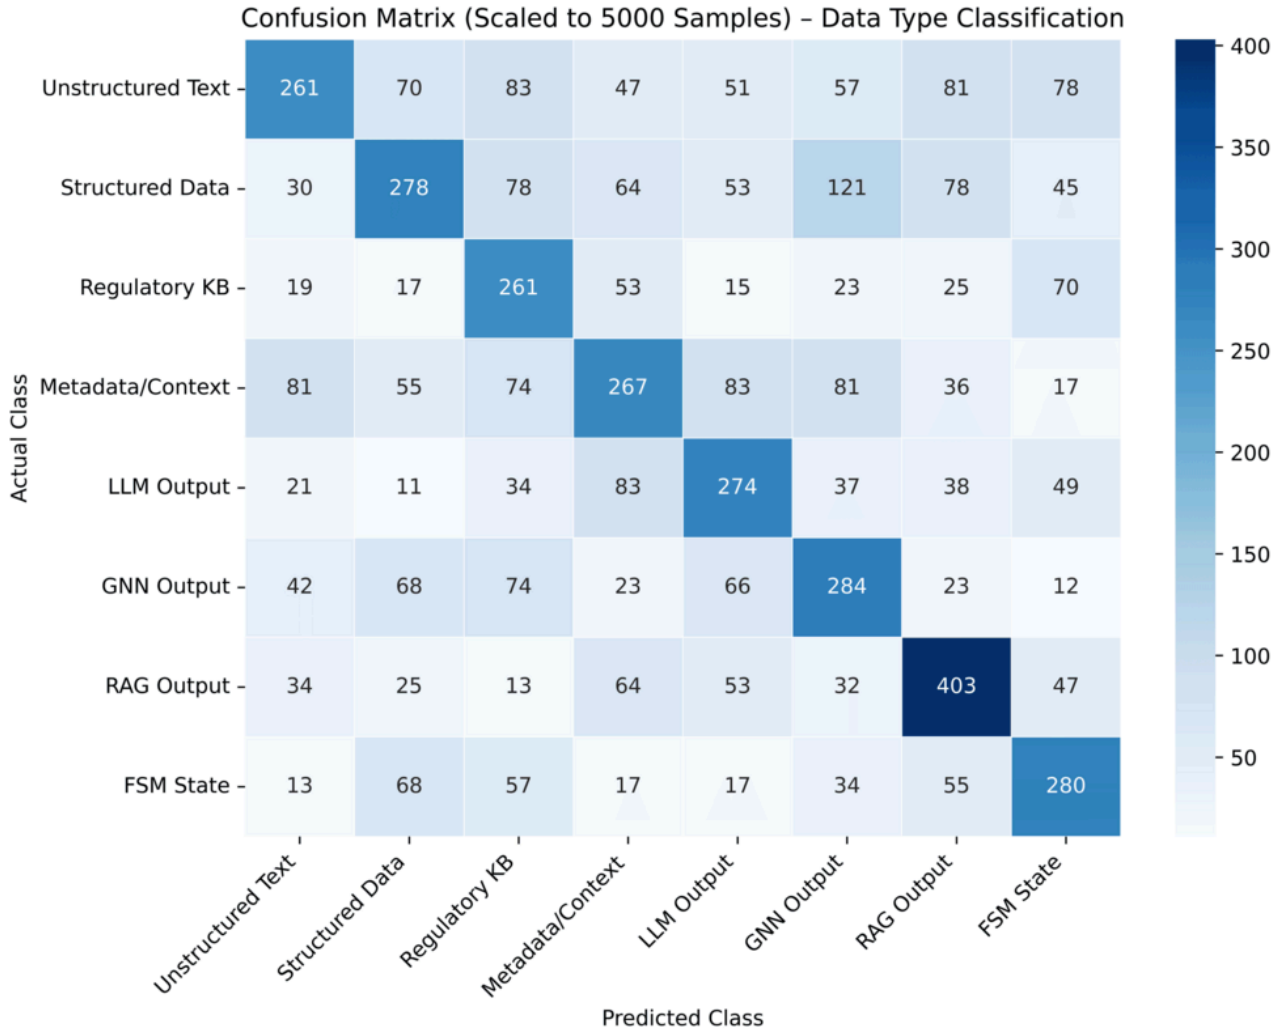

FIGURE 9: Multi-class confusion matrix for fraud detection data modalities.

GNN, Graph Neural Network; LLM, Large Language Model; RAG, Retrieval-Augmented Generation; FSM, Finite State Machine

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

The comparative performance analysis evaluated five model configurations, namely the Rule-based Baseline, LLM, GNN, LLM+GNN, and the Proposed Hybrid Framework (LLM+GNN+FSM+RAG), using Accuracy, Precision, Recall, and F1-score as evaluation metrics. The results demonstrated progressive performance improvement from the baseline model to the proposed hybrid framework. Among all evaluated approaches, the Proposed Hybrid Framework achieved the highest performance with an Accuracy of 0.96, Precision of 0.95, Recall of 0.94, and F1-score of 0.94. These findings indicate that the integration of semantic analysis through LLMs, relational reasoning using GNNs, rule-based fraud progression tracking via FSM, and contextual regulatory retrieval through RAG contributes to improved fraud detection capability and balanced classification performance. The comparative results presented in Table *7* and Figure *10* demonstrate the effectiveness of the proposed hybrid framework over individual and partially integrated models.

Model	Accuracy	Precision	Recall	F1-score
Rule-based baseline	0.91	0.90	0.89	0.90
LLM	0.92	0.91	0.91	0.91
GNN	0.93	0.92	0.92	0.92
LLM+GNN	0.94	0.93	0.93	0.93
Proposed Hybrid Model	0.96	0.95	0.94	0.94

TABLE 7: Comparative performance of baseline, single-module, and hybrid fraud detection models.

LLM, Large Language Models; GNN, Graphic Neural Network

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

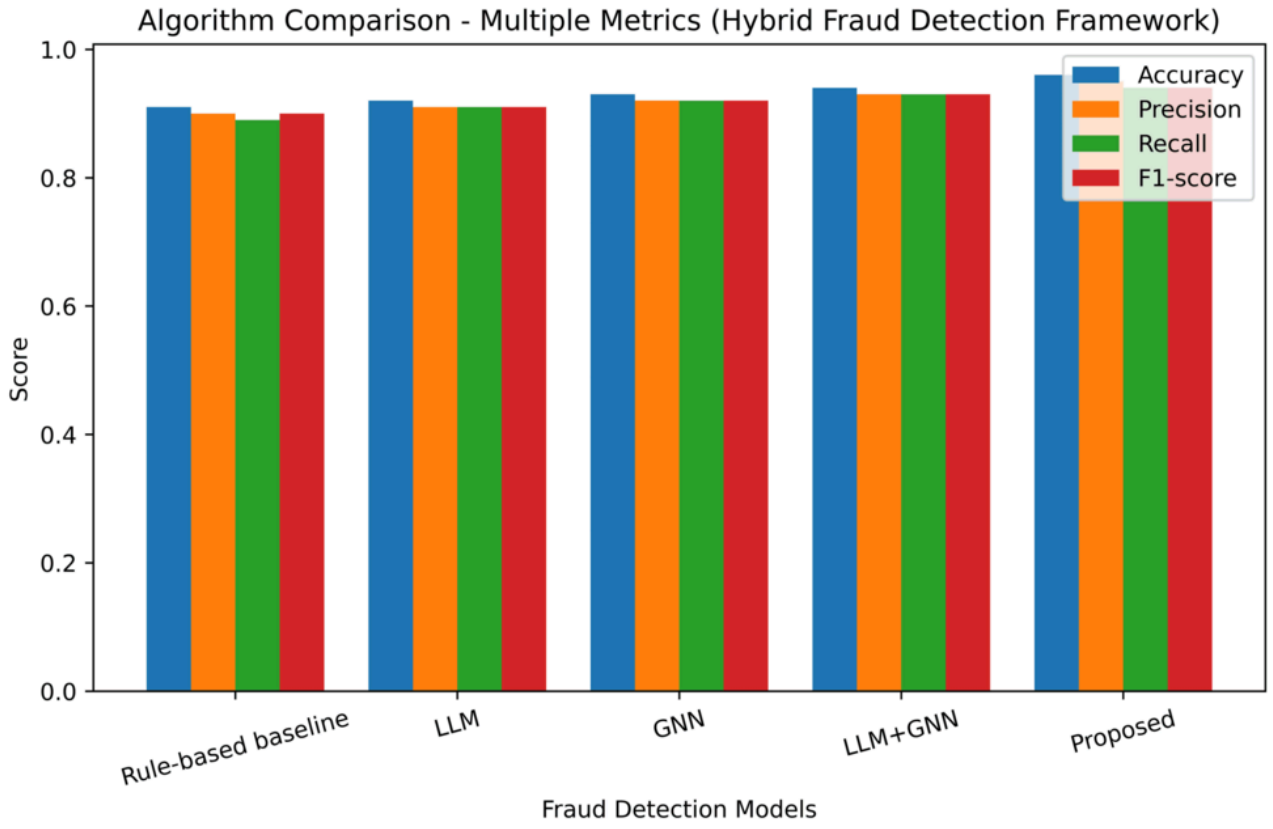

FIGURE 10: Algorithm comparison across multiple metrics in hybrid fraud detection framework.

GNN, Graph Neural Network; LLM, Large Language Model

The Figure *11* displays the Multi-Class ROC Curve of the Hybrid Fraud Detection Model with all classes displaying good discriminatory properties of 0.91-0.97 AUC and the AUC being 0.94 on overall micro-average scales. Both curves exist well above the random bases, and this shows that the true positive rates are high, with a low false positive. It is interesting to see that Class 4 has the highest AUC 0.97 and the lowest-performing Class 0 still robustly reports the AUC 0.91. The findings bring out how useful the model has been in terms of adequately recognizing fraudulent practices on a variety of different categories, which makes the model reliable and highly predictive on the fraud detection part.

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

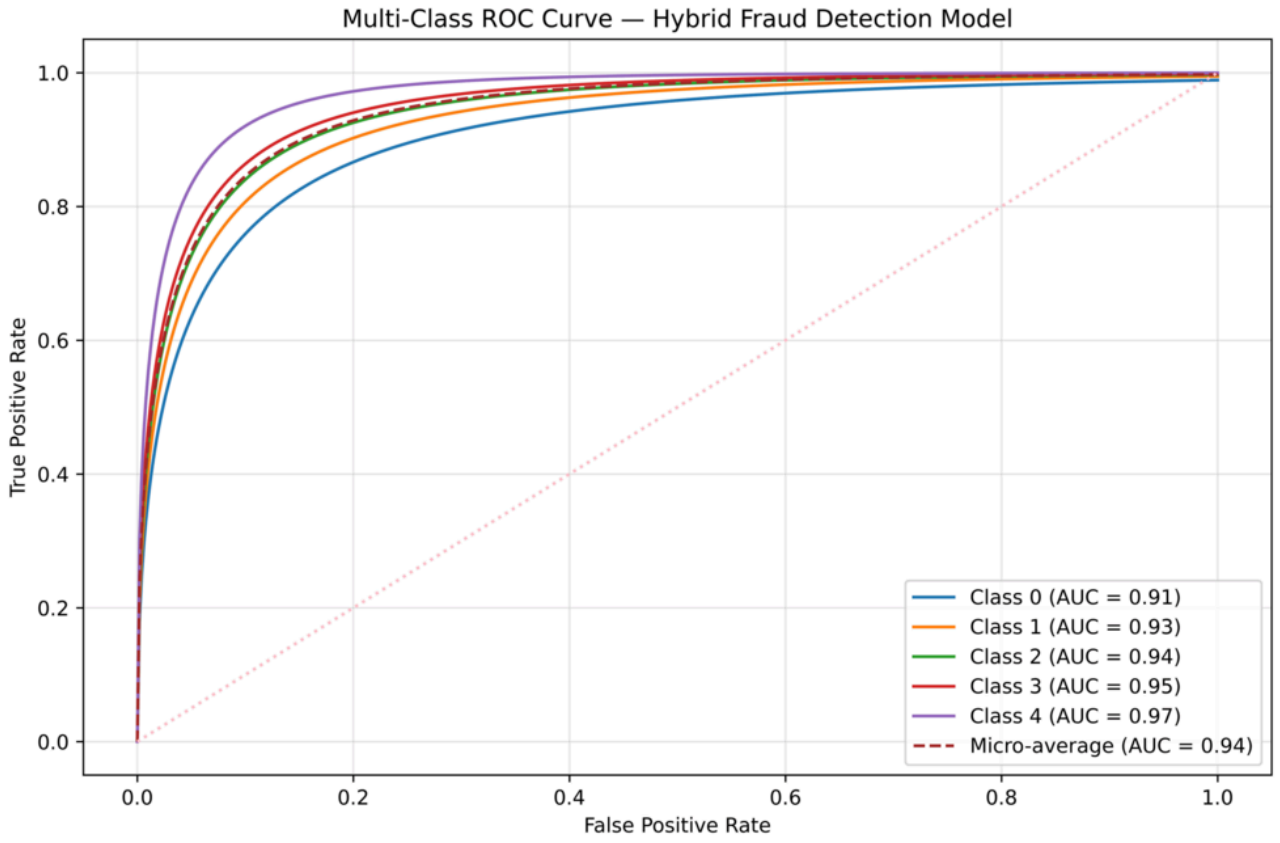

FIGURE 11: Multi-Class ROC curve for hybrid fraud detection model.

ROC, Receiver Operating Characteristic curve

Figure *12* illustrates the training and validation loss curves of the proposed hybrid fraud detection framework over 50 training epochs using binary cross-entropy as the optimization objective. Both curves demonstrate a gradual reduction in loss throughout the training process, with the training loss decreasing rapidly during the initial epochs before stabilizing after approximately epoch 35. The validation loss follows a similar trend, indicating stable learning behavior and effective generalization on unseen data. The relatively small gap between the training and validation loss curves suggests limited overfitting and supports the robustness of the proposed framework during model optimization.

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

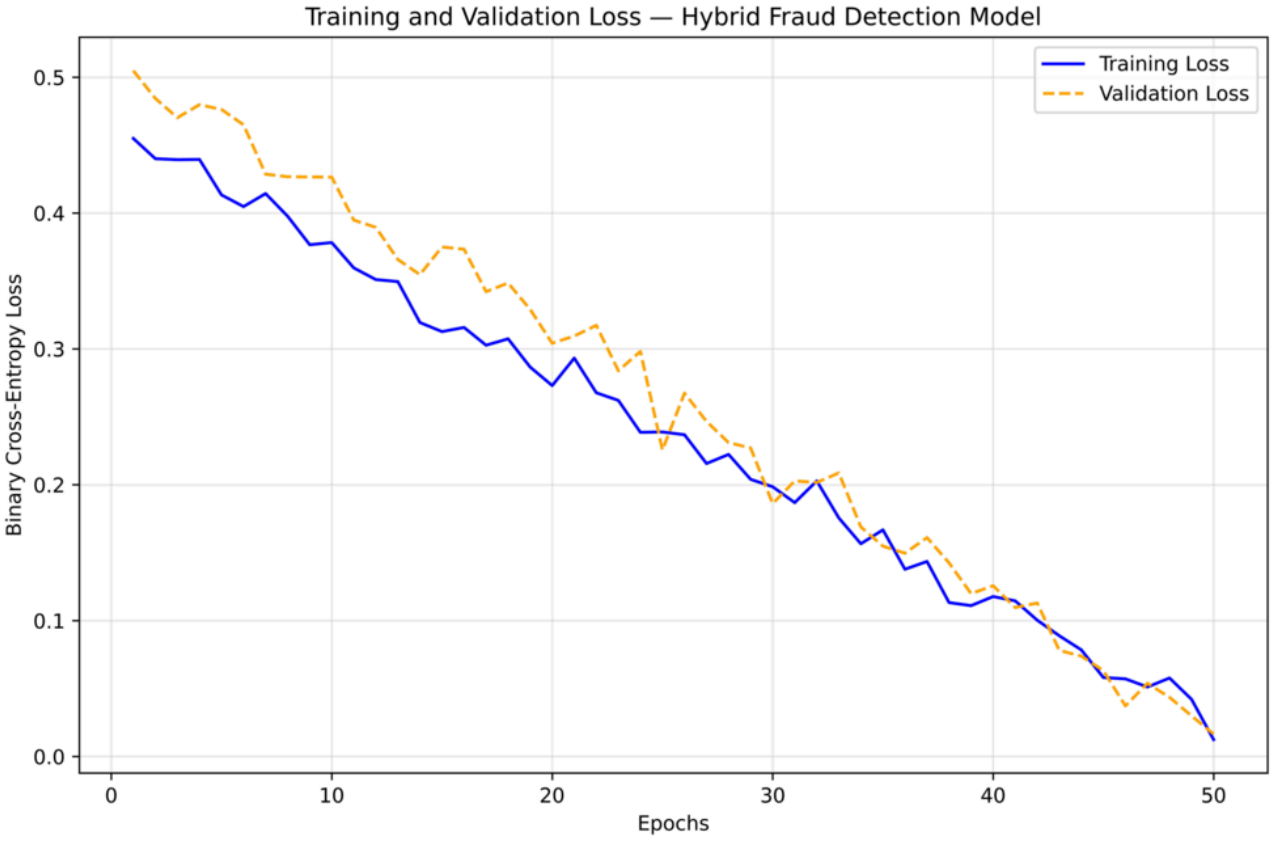

FIGURE 12: Training and validation loss curve of hybrid fraud detection model.

Conclusions

The proposed research presents a hybrid forensic auditing framework that integrates LLMs, GNNs, RAG, and a rule-based Finite State Machine (FSM) inspired by Automata Theory for analyzing heterogeneous financial data and supporting corporate fraud detection. The framework combines semantic analysis, relational pattern analysis, contextual regulatory retrieval, and interpretable fraud progression tracking to assist in identifying potentially suspicious financial activities within unstructured and structured financial records. Experimental evaluation conducted using SEC filings and the Enron email corpus demonstrated encouraging fraud detection performance across multiple evaluation metrics, including accuracy, precision, recall, and F1-score. The study additionally highlights the potential of integrating multiple analytical modules to improve interpretability and audit traceability within forensic auditing workflows. However, the current implementation includes several limitations, including the use of predefined rule-based FSM transitions, limited ablation analysis, absence of multi-seed statistical validation, and challenges associated with dataset labeling and class imbalance in fraud detection tasks. Therefore, the reported findings should be interpreted as preliminary experimental results rather than definitive evidence of generalized real-world performance. Future work will focus on adaptive automata learning, advanced explainability analysis, statistical robustness evaluation, and broader real-world validation to further strengthen the reliability and scalability of the proposed framework.

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Ankita N. Patil, Mritunjay K. Ranjan Mr., Tushar Thorat, Vaidik Gharat, Samruddhi Sonawane

Acquisition, analysis, or interpretation of data: Ankita N. Patil, Mritunjay K. Ranjan Mr.

Drafting of the manuscript: Ankita N. Patil, Tushar Thorat, Vaidik Gharat

Critical review of the manuscript for important intellectual content: Ankita N. Patil, Mritunjay K. Ranjan Mr., Samruddhi Sonawane

Supervision: Mritunjay K. Ranjan Mr.

Disclosures

Human subjects: All authors have confirmed that this study did not involve human participants or tissue. **Animal subjects:** All authors have confirmed that this study did not involve animal subjects or tissue. **Conflicts of interest:** In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

Data Availability Statements

The datasets (and/or code) supporting this study are available from the corresponding author upon reasonable request.

Acknowledgements

The authors would like to acknowledge the use of publicly available datasets, including the Enron Email Dataset (https://www.kaggle.com/datasets/wcukierski/enron-email-dataset?resource=download) and the Enron Network Dataset (https://www.cs.cmu.edu/~enron/) , which were used for experimental evaluation. These datasets significantly contributed to validating the proposed hybrid fraud detection framework.

References

1. Yian W, Fang H: Research on the relationship between financial ecological environment and enterprise financial risk based on multiple regression model. 2nd International Conference on Economic Management and Model Engineering (ICEMME). 2020, 247-259. 10.1109/icemme51517.2020.00055
2. Scur G, Kunimura AA, Chang J: Aligning corporate strategy with supplier selection. IEEE Engineering Management Review. 2022, 50:40-53. 10.1109/emr.2022.3167558
3. Wang L: Abnormal transaction detection and fraud prediction of bank accounts based on big data association analysis algorithm. 4th International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE). 2025, 1-6. 10.1109/icdcece65353.2025.11035453
4. Lou C, Wang Y, Li J, Qian Y, Li X: Graph neural network for fraud detection via context encoding and adaptive aggregation. Expert Systems with Applications. 2025, 261:125473. 10.1016/j.eswa.2024.125473
5. Wu B, Chao KM, Li Y: Heterogeneous graph neural networks for fraud detection and explanation in supply chain finance. Information Systems. 2024, 121:102335. 10.1016/j.is.2023.102335
6. Iaroshev I, Pillai R, Vaglietti L, Hanne T: Evaluating retrieval-augmented generation models for financial report question and answering. Applied Sciences. 2024, 14:9318. 10.3390/app14209318

7. LH Aros, Molano LXB, Gutierrez-Portela F, Hernandez JJM, Barrero MSR: Financial fraud detection through the application of machine learning techniques: a literature review. Humanities and Social Sciences Communications. 2024, 11:1-22. 10.1057/s41599-024-03606-0

8. Sarmah B, Mehta, D, Pasquali S, Zhu T: Towards reducing hallucination in extracting information from financial reports using large language models. AIMLSystems '23: Proceedings of the Third International Conference on AI-ML Systems. 2023, 1-5. 10.1145/3639856.3639895

9. Deng, G, Wu H, He W, Ding S, Pan Q, Ma Z : Anomaly detection algorithm for multivariate complex data based on deep neural network. 4th International Conference on Artificial Intelligence, Internet and Digital Economy (ICAID). 2025, 229-232. 10.1109/icaid65275.2025.11034428

10. Wang, L, Song C, Liu C : A method of heterogeneous network information data integration based on XML technology. 6th International Conference on Wireless Communications and Applications (ICWCAPP). 2022, 92-96. 10.1109/icwcapp57292.2022.00030

11. Ammann, L, Ott S, Landolt CR, Lehmann MP : Securing rag: a risk assessment and mitigation framework. IEEE Swiss Conference on Data Science. 2025, 127-134. 10.1109/sds66131.2025.00024

12. Tural B, Örpek Z, Destan Z: Retrieval-augmented generation (RAG) and LLM integration. 8th International Symposium on Innovative Approaches in Smart Technologies (ISAS). 2024, 1-5. 10.1109/isas64331.2024.10845308

13. Ahmad A, Saad M, Al Ghamdi M, Nyang D, Mohaisen D: BlockTrail: a service for secure and transparent blockchain-driven audit trails. IEEE Systems Journal. 2021, 16:1367-1378. 10.1109/jsyst.2021.3097744

14. van Erp B, de Vries B : Hybrid inference with invertible neural networks in factor graphs. 30th European Signal Processing Conference (EUSIPCO). 2022, 1397-1401. 10.23919/eusipco55093.2022.9909873

15. Wu Y, Dong Y, Qin J, Pedrycz W: Linguistic distribution and priority-based approximation to linguistic preference relations with flexible linguistic expressions in decision making. IEEE Transactions on Cybernetics. 2021, 51:649-659. 10.1109/tcyb.2019.2953307

16. Shadangi AR, Das SS, Chakrabarti I, Pilla PK: A novel FSM-based hardware-efficient, power, and delay optimized transmitter architecture for OTFS. IEEE 21st India Council International Conference (INDICON). 2024, 1-6. 10.1109/indicon63790.2024.10958498

17. Chaudhary P, Singh P, Boral S, Gupta M, Vedantan K, Nayaka RJ: Pattern recognition using QML algorithm for fraud detection in banking. International Conference on Networks and Cryptology (NETCRYPT). 2025, 1864-1868. 10.1109/netcrypt65877.2025.11102584

18. Herman MBR, Ng P, Low MYH, Remy D: Integrating RAG with face recognition for personalized guest services for hospitality industries. TENCON 2024 - 2024 IEEE Region. 2024, 1394-1397. 10.1109/tencon61640.2024.10902965

19. Chen J, Quintano J, Qiang Y: FinFraud-LLM: exploring large language models for financial fraud detection. IEEE International Conference on Data Mining Workshops (ICDMW). 2025, 3087-3091. 10.1109/icdmw69685.2025.00407

20. Rajendran S: RAG-enhanced LLM framework for real-time fraud detection in financial systems. International Journal of Advances in Signal and Image Sciences. 2025, 11:868-889. 10.29284/2zbf4r73

21. Bahaweres RB, Nuraini IK: Cost-sensitive approach for improving AUC-ROC curve of software defect prediction. International Seminar on Intelligent Technology and Its Applications (ISITIA). 2024, 178-183. 10.1109/isitia63062.2024.10668184

22. Sharma A: FraudRAG: knowledge graph-augmented retrieval for real-time financial statement fraud detection using a domain-driven Medallion architecture. 2026, 1-9. 10.2139/ssrn.6714178

23. Zhang Z, Zhang H: An online transfer learning model for intrusion detection using FT-transformer and KSWIN-driven concept drift detection mechanism. 5th International Conference on Computer Engineering and Application (ICCEA). 2024, 128-131. 10.1109/iccea62105.2024.10603489

How to cite this article:
Patil A N, Ranjan M K, Thorat T, et al. (June 08, 2026) A Hybrid Neuro-Symbolic Framework for Corporate Fraud Detection Using Large Language Models, Graph Neural Networks, and Automata-Based Reasoning. Cureus J Comput Sci 3 : es44389-026-00109-8. DOI https://doi.org/10.7759/s44389-026-00109-8