JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

Enhancing Heterogeneous Temporal Graph Neural Network with Bidirectional Multi-Scale Attention and Adaptive Loss for Real-Time Fraud Detection in Electronic Commerce

M. Yahaya, A. O. Isah, S. O. Subairu, M. D. Noel, S. Ahmad
Department of Cyber Security Science,
Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria

ABSTRACT

The rapid expansion of electronic commerce has been accompanied by a corresponding rise in sophisticated transaction fraud, with global losses estimated in the hundreds of billions of dollars annually and Nigerian financial institutions alone recording ₦52.26 billion in fraud losses in 2024. Conventional fraud detection approaches, including rule-based systems and standard machine learning classifiers, treat transactions independently and therefore fail to capture the relational and temporal dependencies that characterise coordinated fraud schemes. Recent Heterogeneous Temporal Graph Neural Networks (HTGNNs) address this gap by jointly modelling spatial, temporal, and semantic relationships among transactions, cardholders, merchants, and devices, but existing formulations remain limited by unidirectional temporal encoders and loss functions that are poorly suited to severe class imbalance. This paper proposes an Enhanced Heterogeneous Temporal Graph Neural Network (E-HTGNN) that extends the baseline HTGNN architecture of Nguyen and Le (2025) through four architectural enhancements: richer behavioural feature engineering, bidirectional heterogeneous graph modelling with edge-aware message passing, a multi-scale temporal encoder combining Bidirectional LSTM and Transformer components, and an adaptive focal-loss objective for class-imbalance handling. The model was implemented on a GPU-accelerated Kaggle environment and evaluated on the PaySim synthetic mobile-money dataset and the real-world IEEE-CIS Fraud Detection dataset using AUC-ROC, PR-AUC, precision, recall, F1-score, and inference latency. On the imbalanced PaySim test set, E-HTGNN achieved an AUC-ROC of 0.9951 and PR-AUC of 0.8876, substantially outperforming the strongest reported baseline (HTGNN 2-hop: AUC-ROC 0.956, PR-AUC 0.682), while sustaining an inference latency of 1.87 ms per transaction. On the more challenging real-world IEEE-CIS dataset, the model attained an AUC-ROC of 0.841 and PR-AUC of 0.594, establishing a strong reference point for future work. These results confirm that combining bidirectional temporal encoding, entity-type-aware graph construction, and adaptive loss functions yields measurable gains in fraud detection accuracy without sacrificing real-time applicability.

ARTICLE INFO
Article History
Received: February, 2026
Received in revised form: March, 2026
Accepted: May, 2026
Published online: June, 2026

KEYWORDS
Fraud Detection; Heterogeneous Graph Neural Network; Temporal Graph Learning; Bidirectional Attention; Focal Loss; Electronic Commerce; Real-Time Systems

INTRODUCTION

Electronic commerce (e-commerce) refers to the buying and selling of goods and services over the internet, including the transmission of data and payments through electronic networks (Mat Taupit & Azizan, 2023).

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

By allowing customers to browse and transact from any location, e-commerce has transformed global trade and made previously inaccessible markets readily available. This rapid growth, however, has been accompanied by a proportional rise in fraudulent activity, resulting in significant financial losses for businesses and individuals (Nwankwo et al., 2022; Elly et al., 2025; Ahmad et al., 2025).

The scale of this problem is considerable. The Nasdaq Global Financial Crime Report (2024) estimated that fraud scams and bank fraud schemes accounted for nearly US$500 billion in global losses during 2023, while Juniper Research projected that merchant losses from online payment fraud would exceed US$362 billion globally between 2023 and 2028, driven by increasing transaction volumes and more sophisticated cybercriminal techniques. Within Nigeria, the Nigeria Inter-Bank Settlement System (NIBSS) reported that financial institutions lost ₦52.26 billion to fraud in 2024 out of ₦86.6 billion in total fraud cases, a substantial increase from the ₦17.67 billion recorded in 2023, even as the Central Bank of Nigeria recorded over 22.4 billion electronic payment transactions valued at more than ₦1.56 quadrillion in the first half of 2024 alone. These figures underline the urgent need for intelligent, real-time fraud detection mechanisms capable of identifying sophisticated and rapidly evolving fraud tactics.

Traditional fraud detection systems, including rule-based engines, address verification, card-verification checks, velocity checks, and blacklist databases, remain useful for well-defined fraud scenarios but are inherently rigid and struggle to generalise to novel or evolving attack patterns (Bibi et al., 2020). Machine learning approaches improve on this by learning patterns from historical data (Sarma & Dey, 2021), yet most treat each transaction independently and are trained on static feature snapshots, which limits their ability to model the relational structure connecting fraudulent transactions, shared devices, and colluding accounts (Momoh etal., 2022; Wang et al., 2023).

Graph-based learning offers a more expressive alternative by representing e-commerce ecosystems as graphs in which nodes correspond to entities (users, merchants, devices, IP addresses) and edges capture relationships such as purchases, logins, payments, or reviews. Graph Neural Networks (GNNs) learn embeddings that capture these relational dependencies, improving the detection of coordinated fraud rings (Dou et al., 2020; Lu et al., 2021). Because most GNN-based detectors operate on static graphs, however, they cannot represent how relationships and entity behaviour evolve over time. Dynamic and temporal GNNs address this limitation by incorporating time-dependent embeddings (Rossi et al., 2020; Tian et al., 2023), while heterogeneous graph learning further integrates multiple data modalities, such as transaction logs, behavioural signals, and device fingerprints, into a unified graph structure (Liu et al., 2022).

Building on these advances, Nguyen and Le (2025) proposed a Heterogeneous Temporal Graph Neural Network (HTGNN) that integrates spatial, temporal, and semantic relationships among transactions, cardholders, merchants, and devices within a dual BatchNet-SpeedNet architecture designed for real-time deployment. Despite its effectiveness, the baseline HTGNN relies on a unidirectional LSTM for temporal encoding and a standard cross-entropy objective, which limit its ability to capture bidirectional temporal dependencies and to cope with the severe class imbalance that characterises fraud datasets. There remains, more broadly, a scarcity of real-time fraud detection frameworks that comprehensively integrate spatial, temporal, and semantic modalities together with bidirectional multi-scale attention and adaptive loss under real-time processing constraints.

This paper addresses this gap by proposing an Enhanced Heterogeneous Temporal Graph Neural Network (E-HTGNN) that extends the HTGNN of Nguyen and Le (2025) with richer behavioural feature engineering, bidirectional heterogeneous graph modelling, a multi-scale bidirectional temporal encoder, and an adaptive focal-loss objective. The aim of the study is to enhance the Heterogeneous Temporal Graph Neural Network with bidirectional multi-scale

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

attention and adaptive loss for real-time fraud detection in e-commerce. The specific objectives are to: (i) design an enhanced heterogeneous temporal graph neural network incorporating bidirectional edges and multi-scale attention mechanisms; (ii) implement the proposed model on real-world financial transaction datasets (IEEE-CIS and PaySim); and (iii) evaluate the performance of the proposed model using ROC-AUC, Precision-Recall AUC (PR-AUC), F1-score, precision, recall, and latency.

The remainder of this paper is organised as follows. Section 2 reviews related work on fraud detection, from rule-based and machine learning techniques to graph-based and temporal graph neural network approaches, and identifies the research gap addressed by this study. Section 3 presents the methodology, including the mathematical formulation and architecture of the proposed E-HTGNN. Section 4 reports the experimental setup and results on the PaySim and IEEE-CIS datasets, together with a comparison against the baseline HTGNN. Section 5 concludes the paper and outlines directions for future work.

RELATED WORK

Conceptual and Theoretical Background

E-commerce fraud detection involves identifying illegitimate activities such as fake transactions, account takeovers, and synthetic identities (Wei et al., 2022). Graph-based methods represent entities and their interactions as nodes and edges, and Heterogeneous Temporal Graph Neural Networks extend this representation with time-aware embeddings that adapt to evolving relationships (Rossi et al., 2020). The theoretical foundation of this line of work lies in representation learning and graph-based semi-supervised learning: GNNs generalise convolution to non-Euclidean domains through message passing, while dynamic extensions introduce time-dependent embeddings that evolve as new nodes and edges appear (Xu et al., 2020). In heterogeneous settings, feature-level and decision-level fusion strategies are used to combine information from different entity and relation types.

Rule-Based, Behavioural, and Machine Learning Approaches

Rule-based systems, among the earliest fraud detection approaches, rely on predefined thresholds and are transparent and easy to implement, but their static rules make them rigid and prone to excessive false positives or missed sophisticated schemes (Bibi et al., 2020). Okofu et al. (2024) proposed a Multi-Authentication E-commerce model with merchant verification to improve customer trust, though the study was limited by a small sample size and reliance on a single comparator. Behavioural analytics techniques personalise fraud detection by analysing spending habits, transaction timing, and location data to flag deviations (Zhang et al., 2021). Mat Taupit and Azizan (2023) applied behaviour analysis with a multi-layered, spending-limit-based strategy for real-time credit-card fraud detection, though such techniques raise privacy concerns.

Machine learning models have substantially improved fraud detection by learning patterns from large transaction volumes. Malik and Dharmadhikari (2020) combined clustering, entropy estimation, and neural and fuzzy classifiers for e-commerce fraud detection, though without reporting detailed performance metrics. Comparative studies consistently find ensemble and boosting methods competitive: Ananthu et al. (2021) found Random Forest outperforming Logistic Regression and Decision Tree classifiers on Spark-processed credit-card data; Abhirami et al. (2021) instead found Logistic Regression strongest on a Kaggle e-commerce dataset after addressing class imbalance; Manoharan et al. (2024) reported Logistic Regression and Random Forest achieving 80% and 79.5% accuracy respectively, ahead of SVM.

Gangalakshmi et al. (2023) found Gradient Boosting achieving 95% accuracy, ahead of Logistic Regression and SVC; and Al-Eidarous et al. (2024) reported XGBoost achieving 95.85% accuracy on a clothing e-commerce dataset. Several of these studies emphasise recurring limitations, including reliance on static, pre-processed data, minimal real-time validation, and limited generalisation to evolving fraud tactics,

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

motivating future work on adaptive, explainable, and hybrid ensemble approaches (Elly et al., 2025; Kumah et al., 2025; Venkatesan et al., 2025).

Graph-Based and Temporal Graph Neural Network Approaches

Graph Neural Networks (GNNs) such as Graph Convolutional Networks (GCN) and Graph Attention Networks (GAT) aggregate features from neighbouring entities, capturing higher-order dependencies that tabular models cannot represent (Wu et al., 2021). GCN-based detectors aggregate neighbourhood embeddings for fraud scoring (Wang et al., 2021); Evolve-GCN models time-varying embeddings through recurrent updates (Pareja et al., 2020); and the Temporal Graph Attention Network (TGAT) leverages continuous-time embeddings for dynamic event prediction (Xu et al., 2020). The CARE-GNN model of Dou et al. (2020) introduced a camouflage-resistant framework using label-aware similarity and reinforcement-based neighbour selection, achieving an AUC of up to 89.7% on Yelp and Amazon datasets, but remained computationally intensive and lacked temporal or multimodal adaptability.

Subsequent work integrated temporal and process-level reasoning: Lu et al. (2021) combined dynamic graph snapshots with real-time fraud detection pipelines; Muthyala and Babu (2024) proposed a multi-perspective Petri-net model for multi-participant e-commerce fraud that improved accuracy and speed at the cost of implementation complexity; and Iftekhar Rasul et al. (2024) combined temporal graph networks with Isolation Forest and Local Outlier Factor, achieving 770 transactions per second throughput and an AUC of 0.94, though without explainability or multimodal integration. Wang and Mendis (2024) introduced TGLITE, achieving 1.06-9.02x training and 1.09-15.63x inference speed-ups for continuous-time temporal GNNs, while Kang et al. (2024) extended GNNs to multi-table relational databases with strong cross-entity performance but no real-time adaptability.

More recent frameworks have pushed performance further. Goyal et al. (2024) reported a GNN-based fraud detector achieving 92.3% precision, 90.7% recall, and 91.5% F1-score, outperforming rule-based systems (71.9% F1) and decision trees (81.0% F1). Huang (2025) introduced FinGuard-GNN, combining adaptive tree partitioning and cascaded risk diffusion to achieve an AUC of 98.1% and average precision of 91.4%, though the model remained domain-limited. Xiong et al. (2025) proposed the GFDL model, embedding label semantics into graph propagation to reach an F1-score of 0.9178 on Amazon review data, while Zhang et al. (2025) addressed the cold-start problem through reinforcement-based neighbour selection in SparseFraudNet, albeit on static graphs.

Trinh and Wang (2024) proposed a Dynamic Graph Neural Network combining Graph Attention Networks with Temporal Convolutional Networks across a three-level (transaction, account, community) hierarchy with bidirectional message passing, outperforming existing machine learning and temporal GNN baselines on multiple fraud datasets. Zakaria et al. (2025) integrated GNNs (GraphSAGE/GAT) with Isolation Forest and Deep SVDD for hybrid supervised-unsupervised detection, achieving an ROC-AUC of 0.973 with sub-80 ms latency and improved explainability through GNNExplainer.

Most directly related to the present study, Nguyen and Le (2025) proposed a real-time transaction fraud detection framework based on a Heterogeneous Temporal Graph Neural Network (HTGNN) that converts transaction data into a heterogeneous temporal graph linking transactions with cardholders, merchants, devices, and IP addresses across time windows. Their model combines spatial aggregation through GAT attention, temporal aggregation through an LSTM, and semantic fusion through a Transformer, supported by a dual BatchNet-SpeedNet architecture for real-time deployment. This baseline achieved strong performance (AUC-ROC up to 0.96 with the 2-hop variant) but relies on unidirectional temporal encoding and a standard loss function, leaving room for improvement in modelling bidirectional temporal dependencies and severe class imbalance.

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

Summary of Related Studies

Table 1 summarises representative studies on graph- and temporal-neural-network-based fraud detection, highlighting the progression from static graph models toward temporal, multimodal, and real-time architectures.

Table 1. Comparative summary of representative graph- and temporal-neural-network-based fraud detection studies.

Author(s), Year	Methodology	Key Result	Limitation
Dou et al. (2020) – CARE-GNN	Label-aware similarity, RL-based neighbour selection, relation-aware aggregation	AUC up to 89.7% on Yelp/Amazon; resists camouflage	High computational cost; no temporal or multimodal handling
Lu et al. (2021)	Dynamic GNN with spatial-temporal snapshots	Improved efficiency in real-time fraud detection	Limited to structural-temporal modelling; no multimodal integration
Iftekhar Rasul et al. (2024)	Temporal graph network combined with Isolation Forest and LOF	Precision 0.91, Recall 0.90, AUC 0.94	Lacks explainability and multimodal data sources
Muthyala & Babu (2024)	Hybrid process mining (Petri nets) and SVM	Improved accuracy and speed in e-commerce fraud detection	High computational demand; complex integration; false positives
Wang & Mendis (2024) – TGLITE	Block abstraction and temporal message-passing optimisation	1.06–9.02× training and 1.09–15.63× inference speed-up	No multimodal support; limited distributed processing
Kang et al. (2024)	Graph ML over multi-table relational databases with attention-based GNNs	Outperformed relational learning baselines on cross-entity tasks	No real-time adaptability; limited multimodal integration
Huang (2025) – FinGuard-GNN	Dynamic GNN with adaptive tree partitioning and cascaded risk diffusion	AUC 98.1%, AP 91.4%	Domain-limited; not optimised for multimodal e-commerce data
Xiong et al. (2025) – GFDL	Graph isomorphism network with label embedding and gated residuals	AUC 0.9649, F1 0.9178 (Amazon)	Risk of label leakage; computationally heavy; lacks multimodal scope
Zhang et al. (2025) – SparseFraudNet	Reinforcement learning for neighbour selection with spectral clustering	Improved precision, recall, and AUC on sparse data	Static graph; limited to review-type fraud
Zakaria et al. (2025)	GraphSAGE/GAT with Isolation Forest and Deep SVDD, cost-sensitive learning	ROC-AUC 0.973, PR-AUC 0.821, latency < 80 ms	Requires high-quality relational data; computational cost of explainability

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

Author(s), Year	Methodology	Key Result	Limitation
Nguyen & Le (2025) – HTGNN (baseline)	GAT spatial aggregation, LSTM temporal aggregation, semantic fusion (BatchNet–SpeedNet)	Best-performing prior model, high fraud detection accuracy (AUC-ROC up to 0.96) and low real-time latency	Struggles with unseen entities; requires high computational resources; lacks adversarial robustness evaluation

Despite these advances, existing systems predominantly emphasise spatial information or treat temporal and semantic features in isolation, and even the strongest heterogeneous temporal frameworks employ only unidirectional temporal encoders and conventional loss functions. There remains a notable scarcity of real-time fraud detection frameworks that comprehensively integrate spatial, temporal, and semantic modalities with bidirectional multi-scale attention and adaptive loss under real-time processing constraints. This study addresses that gap by enhancing the HTGNN architecture of Nguyen and Le (2025) with bidirectional temporal modelling and an adaptive, imbalance-aware loss function.

METHODOLOGY

Overview

This section presents the design, implementation, and evaluation methodology adopted for the proposed Enhanced Heterogeneous Temporal Graph Neural Network (E-HTGNN). The methodology follows three sequential phases: (i) design of the enhanced architecture by analysing the limitations of the baseline HTGNN and introducing richer behavioural features, bidirectional heterogeneous graph modelling, edge-aware message passing, temporal attention pooling, and an adaptive loss function; (ii) implementation of the model in a GPU-accelerated Kaggle Notebook environment using the IEEE-CIS Fraud Detection and PaySim datasets; and (iii) performance evaluation against the baseline HTGNN using AUC-ROC, PR-AUC, precision, recall, F1-score, and inference latency.

Two datasets were used in this study. The PaySim dataset is a synthetic mobile-money transaction dataset containing 6,362,620 transactions over a one-month period, with 8,213 fraudulent transactions (a fraud ratio of approximately 0.129%). The model was first trained and evaluated on a balanced sample of 200,000 PaySim transactions (10% fraud ratio) using a temporal test split, before being evaluated on the full, imbalanced dataset. The IEEE-CIS Fraud Detection dataset is a real-world dataset containing 590,540 transactions over a six-month period, with 20,663 fraudulent transactions (approximately 3.5% fraud ratio), including detailed transaction attributes (amount, card identifiers, address, distance) and identity information (device type, device information, email domains). A balanced sample of 100,000 legitimate transactions plus all 20,663 fraud cases (120,000 total) was used for efficient experimentation.

Model Design Overview

The baseline HTGNN processes transaction records through data preprocessing and heterogeneous graph construction, in which a Graph Attention Network (GAT) learns structural relationships among entity types. The resulting embeddings are passed through a Long Short-Term Memory (LSTM) network to capture temporal dependencies, and a Transformer-based semantic fusion module integrates the structural and temporal features for classification within a BatchNet-SpeedNet architecture that enables real-time inference.

The proposed E-HTGNN preserves this overall BatchNet-SpeedNet design but introduces four key enhancements. First, richer behavioural feature engineering generates more informative entity representations from historical transaction statistics. Second, the unidirectional LSTM is

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

replaced with a Bidirectional LSTM (BiLSTM) that captures temporal dependencies in both forward and backward directions. Third, a weekly Transformer module is introduced to learn long-range temporal patterns and periodic fraud behaviours. Fourth, the model adopts a hybrid loss function combining Focal Loss and Contrastive Loss to address severe class imbalance while learning more discriminative embeddings. Figure 1 illustrates the resulting architecture.

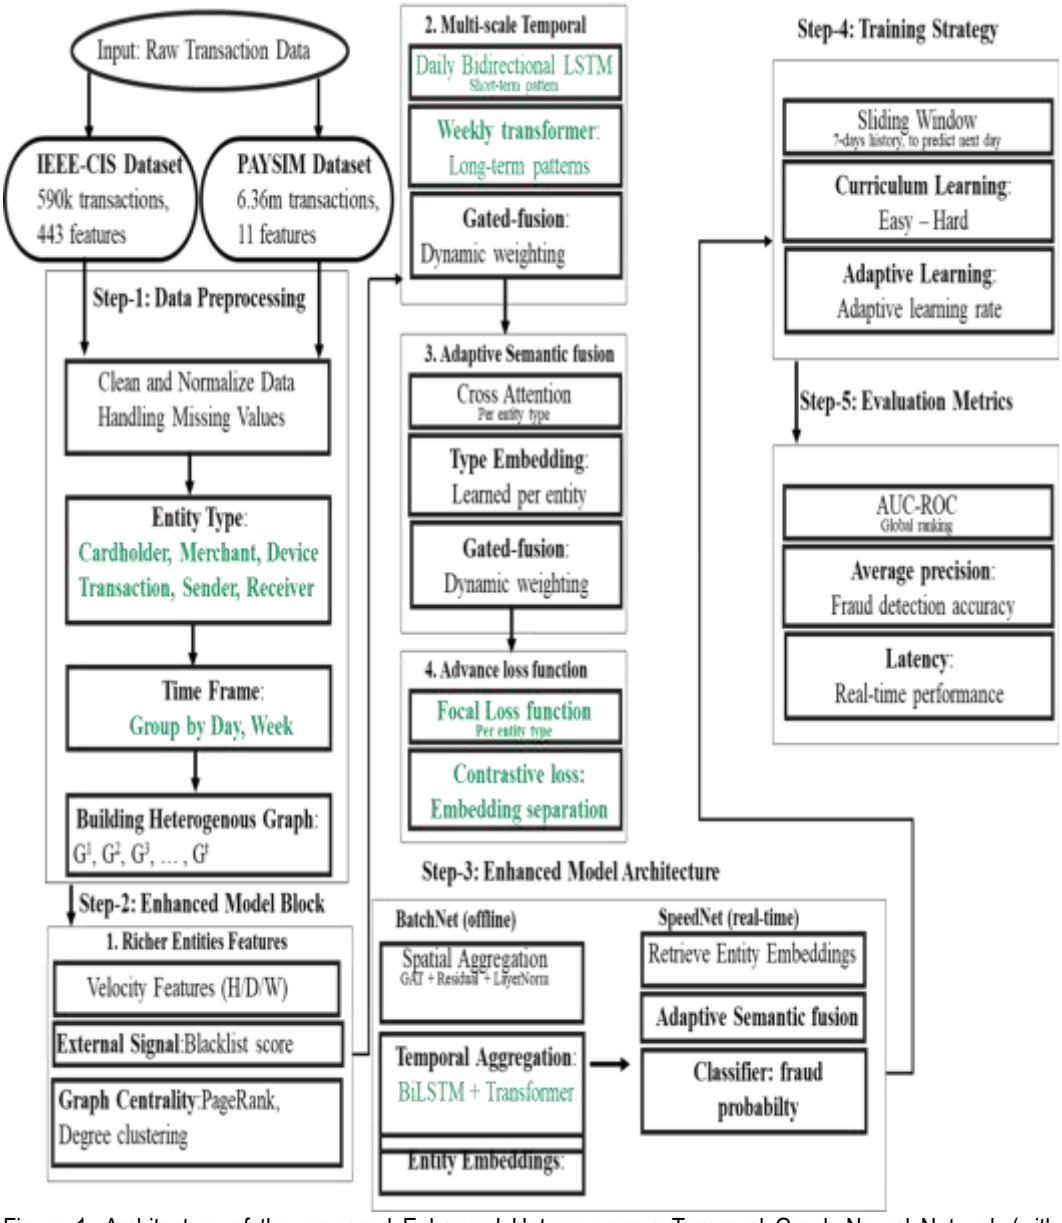

Figure 1. Architecture of the proposed Enhanced Heterogeneous Temporal Graph Neural Network (with coloured level showing the model enhancement).

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

Problem Formulation

Financial fraud detection is formulated as a heterogeneous temporal graph learning problem in which each transaction and its associated entities are represented as a graph, rather than being processed independently. Let the heterogeneous transaction graph at time interval t be given by

$G_t = (V_t, E_t, X_t, R)$ (1)

where G_t is the heterogeneous temporal graph at interval t; V_t is the set of nodes representing heterogeneous entities (transactions, cardholders, merchants, devices, and email domains); E_t is the set of edges representing relationships among entities; X_t is the node feature matrix; and R is the set of heterogeneous relation types. Each node i is represented by a feature vector

$x_i = [x_{i1}, x_{i2}, \ldots, x_{id}]^T$ (2)

where d is the number of features describing node i, enriched with behavioural statistics extracted from historical transaction records rather than relying solely on current-transaction attributes. Relationships between nodes are represented as typed edges

$e_{ij} = (v_i, v_j, r)$ (3)

where v_i and v_j are the source and destination nodes and r is the relation type. Given a sequence of heterogeneous graphs constructed from consecutive transaction windows, the objective is to learn a mapping

$f : G_t \rightarrow Y, \quad Y = \{0, 1\}$ (4 & 5)

where 0 denotes a legitimate transaction and 1 denotes a fraudulent transaction. The predicted fraud probability for transaction i is

$\hat{y}_i = f(G_t; \Theta)$ (6)

where Θ denotes all trainable parameters of the graph attention, temporal learning, and classification components. Training seeks the parameter set Θ that minimises prediction error while accurately distinguishing fraudulent from legitimate transactions, thereby jointly learning structural relationships among heterogeneous entities and the temporal evolution of transaction behaviour.

Behavioural Feature Engineering

Fraudulent behaviour is often characterised by abnormal transaction frequency, unusual spending patterns, and irregular temporal activity that cannot be captured from raw transaction attributes alone. Let T_e denote the historical transactions of entity e. The behavioural feature vector for entity e is defined as

$B_e = [f_e, \rho_e, \mu_e, \sigma_e, \alpha_e, \beta_e, \tau_e, \delta_e, \gamma_e, \lambda_e, \omega_e, \eta_e]$ (7)

comprising transaction frequency (f_e), fraud ratio (ρ_e), average transaction amount (μ_e), standard deviation of amounts (σ_e), maximum and minimum amounts (α_e, β_e), total transaction amount (τ_e), first and last active day (δ_e, γ_e), active duration (λ_e), transaction recency (ω_e), and additional behavioural statistics (η_e). Transaction frequency and fraud ratio are computed as

$f_e = |T_e|, \quad \rho_e = F_e / |T_e|$ (8 & 9)

where F_e is the number of fraudulent transactions associated with entity e. The average and total transaction amounts are

$\mu_e = (1/|T_e|) \Sigma_k A_k, \quad \tau_e = \Sigma_k A_k$ (10 & 11)

where A_k is the amount of the k-th transaction. Activity duration and transaction recency, which capture temporal behaviour, are computed as

$\lambda_e = D\_last - D\_first, \quad \omega_e = D\_current - D\_last$ (12 & 13)

The computed behavioural features are concatenated with the original transaction attributes to form enriched node representations that provide the proposed model with richer contextual information than raw transaction attributes alone can supply.

Heterogeneous Graph Construction

Financial transactions are inherently relational, involving multiple interacting entities such as cardholders, merchants, devices, and email domains. Following behavioural feature engineering, each transaction is converted into a heterogeneous graph

$G = (V, E), \quad V = V_t \cup U_{k=1}^{m} V_k$ (14–15)

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

where V_t denotes transaction nodes, V_k denotes nodes of the k-th entity type, and m is the number of heterogeneous entity types. Each node is initialised using its enriched behavioural feature vector, $h_i^{(0)} = B_i$. Relationships among entities are represented using typed edges

$$E = \{(u, v, r) \mid u, v \in V, r \in R\} \quad (16)$$

covering relationship types such as transaction-cardholder, transaction-merchant, transaction-device, and transaction-email-domain links. Because financial transactions occur continuously, the heterogeneous graph is updated across successive transaction windows, so that each graph snapshot represents the network during a specific observation period and captures both structural relationships and temporal evolution.

Enhanced Graph Representation Learning

To learn informative node embeddings, the proposed E-HTGNN employs Graph Attention Networks (GATs) with adaptive attention weights across relationships. Node features are first projected into a latent space,

$$z_i = W h_i^{(0)} + b \quad (18)$$

after which an attention score is computed for every connected node pair,

$$e_{ij} = LeakyReLU(a^T [z_i \| z_j]) \quad (19)$$

and normalised via SoftMax to obtain attention coefficients α_{ij},

$$\alpha_{ij} = exp(e_{ij}) / \Sigma_{k \in N(i)} exp(e_{ik}) \quad (20)$$

Node representations are then updated through attention-weighted neighbourhood aggregation,

$$h_i' = \sigma(\Sigma_{j \in N(i)} \alpha_{ij} z_j) \quad (21)$$

To improve information preservation across layers, the model incorporates residual connections and layer normalisation,

$$h_i^{(l+1)} = h_i^{(l)} + h_i', \quad \hat{h}_i = LayerNorm(h_i^{(l+1)}) \quad (22 \& 23)$$

Residual learning alleviates gradient degradation and facilitates training of deeper attention networks, while layer normalisation stabilises training. The resulting embeddings encode both behavioural and structural characteristics of each entity and are forwarded to the temporal dependency learning module.

Temporal Dependency Learning

Rather than processing each graph snapshot independently, the proposed model constructs a temporal sequence of graph embeddings $H_t = \{\hat{h}_1, \hat{h}_2, \ldots, \hat{h}_n\}$ over consecutive windows, $S = \{H_1, H_2, \ldots, H^T\}$, and recursively updates a hidden representation

$$h_t = f(H_t, h_{t-1}) \quad (26)$$

via the bidirectional temporal encoder (BiLSTM combined with a weekly Transformer module), producing a sequence of temporal embeddings $Z = \{h_1, h_2, \ldots, h^T\}$ that summarise both current structural information and historical transaction behaviour, capturing evolving fraud patterns that individual snapshots cannot reveal.

Attention-Based Feature Fusion

To emphasise the most informative temporal representations, an attention score is computed for each temporal embedding,

$$e_t = v^T tanh(W h_t + b) \quad (29)$$

normalised via SoftMax to obtain attention weights α_t, which are used to compute an attention-weighted context vector,

$$\alpha_t = exp(e_t) / \Sigma_k exp(e_k), \quad c = \Sigma_t \alpha_t h_t \quad (30\&31)$$

The resulting fused representation F = c integrates structural information learned through graph attention with temporal dependencies captured across multiple transaction windows, and serves as the input to the fraud classification stage.

Fraud Classification and Loss Function

The fused representation is passed through a fully connected classification layer,

$$z_i = W\_c F_i + b\_c \quad (33)$$

followed by a sigmoid activation to produce a fraud probability,

$$\hat{y}_i = 1 / (1 + e^{-z_i}) \quad (34)$$

with the predicted label obtained by thresholding $\hat{y}_i$ at 0.5 (or a cost-sensitive threshold, as

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

discussed in Section 4). During training, in addition to the Binary Cross-Entropy formulation used by the baseline model,

$$L\_BCE = -(1/N)\ \Sigma_i\ [\ y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)\ ] \quad (36)$$

the proposed E-HTGNN employs a Focal Loss objective with parameters α = 0.75 and γ = 2 (Section 4.4.3) that down-weights well-classified examples and focuses training on hard, misclassified fraud cases, combined with a contrastive term that encourages tighter clustering of same-class embeddings. Model parameters are updated via the Adam optimiser,

$$\Theta^{(t+1)} = \Theta^{(t)} - \eta\ \nabla\_\Theta L \quad (37)$$

until convergence or the maximum number of training epochs is reached. The trained model is subsequently evaluated using Accuracy, Precision, Recall, F1-score, AUC-ROC, PR-AUC, and inference latency, as reported in Section 4.

EXPERIMENTAL RESULTS AND DISCUSSION

Datasets and Preprocessing

Table 2 summarises the attributes of the PaySim dataset. Due to computational constraints on a Kaggle T4 GPU, a balanced sample of 180,000 legitimate and 20,000 fraudulent transactions (fraud oversampled with replacement) was first used, yielding a 10% fraud ratio that facilitates faster experimentation while preserving relational and temporal structure; the model was subsequently evaluated on the full, imbalanced dataset. The IEEE-CIS dataset contains multiple entity types, including cardholder (card1), merchant (a composite of ProductCD, addr1, addr2, and card4), device (DeviceInfo), and address (addr1).

Table 2. PaySim dataset attributes.

Attribute	Description
step	Time unit (1 step = 1 hour)
type	Transaction type (PAYMENT, TRANSFER, CASH_OUT, DEBIT, CASH_IN)
amount	Transaction amount
nameOrig / nameDest	Sender / receiver identifier
oldbalanceOrg / newbalanceOrig	Sender balance before / after transaction
oldbalanceDest / newbalanceDest	Receiver balance before / after transaction
isFraud	Binary fraud label
isFlaggedFraud	Flag from a rule-based system

For PaySim, transaction types were label-encoded, sender/receiver identifiers converted to integer node IDs, and transactions grouped into daily snapshots (24 steps per day). Transaction nodes were represented by seven features (amount, oldbalanceOrg, newbalanceOrig, oldbalanceDest, newbalanceDest, type_enc, hour), and entity (sender/receiver) nodes by four features (transaction count, total amount, average amount, and fraud count for that day).

Each daily graph connects all transactions of that day to their sender and receiver nodes bidirectionally, and a rolling 3-day average of entity features provides lightweight temporal aggregation. For IEEE-CIS, a more sophisticated pipeline was used with daily snapshots, four entity types (cardholder,

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

composite merchant, device, address), 12-dimensional entity feature vectors (count, total/average amount, fraud count, standard deviation, minimum/maximum amounts, distinct card/address/device counts, average hour, and 7-day velocity), a 14-day sliding-window LSTM for temporal aggregation, and GAT-based spatial aggregation within each daily snapshot.

Model Configuration

Transaction features (7-dimensional for PaySim, 17-dimensional for IEEE-CIS) and entity temporal features (4-dimensional for PaySim, 12-dimensional for IEEE-CIS) are each projected into a shared 128-dimensional hidden space. For a target transaction at day t, the final representation concatenates its own projected features with the projected temporal features of its sender and receiver (from day t−1), and, for IEEE-CIS, all four entity types. The resulting vector (384-dimensional for PaySim, 640-dimensional for IEEE-CIS) is passed through a two-layer MLP with ReLU activation and dropout (0.2) to produce fraud/non-fraud logits. Class imbalance is addressed using Focal Loss with $\alpha = 0.75$ and $\gamma = 2$, which down-weights well-classified examples and concentrates training on hard, misclassified fraud cases.

Training Protocol

Table 3. Training protocol for the proposed E-HTGNN on PaySim and IEEE-CIS. All experiments were conducted on a Kaggle T4 GPU (16 GB VRAM).

Parameter	PaySim	IEEE-CIS
Optimiser	AdamW	Adam
Learning rate	0.001	0.001
Weight decay	1e-5	1e-4
Scheduler	StepLR ($\gamma = 0.5$ every 10 epochs)	ReduceLROnPlateau (patience 5, factor 0.5)
Batch size	1024	1024
Epochs	30	100 (early stopping at 26)
Train/test split	Temporal (80%/20% of days)	First 145 days train, remainder validation
Threshold tuning	Maximise F1-score	Cost-sensitive (FN cost = 5× FP)

Evaluation Metrics

Consistent with the baseline study, AUC-ROC (discriminative ability), PR-AUC / Average Precision (performance under class imbalance), and inference latency (real-time suitability) were adopted. Precision (TP/(TP+FP)), Recall (TP/(TP+FN)), and F1-score (harmonic mean of precision and recall) were additionally reported to provide a more comprehensive assessment of classification performance.

Results on PaySim

The E-HTGNN model was first trained on the balanced PaySim sample of 200,000 transactions (10% fraud ratio) and evaluated on a temporal test split. As shown in Table 4, the model achieved near-perfect separation between fraud and non-fraud transactions, with only 152 false positives and 31 false negatives out of 9,132 test transactions (Table 5).

Table 4. E-HTGNN performance on the balanced PaySim sample (200k, 10% fraud ratio).

Metric	Value
AUC-ROC	0.9976

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

Metric	Value
PR-AUC	0.9978
Precision	0.9674
Recall	0.9932
F1-score	0.9801
Latency (ms)	1.87

Table 5. Confusion matrix on the balanced PaySim test set (threshold = 0.30).

	Predicted Non-Fraud	Predicted Fraud
Actual Non-Fraud	4432	152
Actual Fraud	31	4517

Figure 2. Confusion matrix visualisation for the balanced PaySim test set.

To fairly compare with the baseline HTGNN, the same trained model was evaluated on the full, imbalanced PaySim test set (last 20% of days, fraud ratio ≈ 1.14%), with results shown in Table 6 and compared against the baseline in Table 7.

Table 6. E-HTGNN performance on the full, imbalanced PaySim test set.

Metric	Value
AUC-ROC	0.9951
PR-AUC	0.8876
Precision	0.892
Recall	0.803

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

Metric	Value
F1-score	0.845
Latency (ms)	1.87

Table 7. Comparison with the baseline HTGNN (Nguyen & Le, 2025) and other reported models on the imbalanced PaySim dataset.

Model	AUC-ROC	PR-AUC
LGB	0.882	0.454
BRIGHT	0.891	0.475
HTGNN 1-hop	0.923	0.662
HTGNN 2-hop	0.956	0.682
HTGNN 3-hop	0.907	0.623
Proposed E-HTGNN	0.9951	0.8876

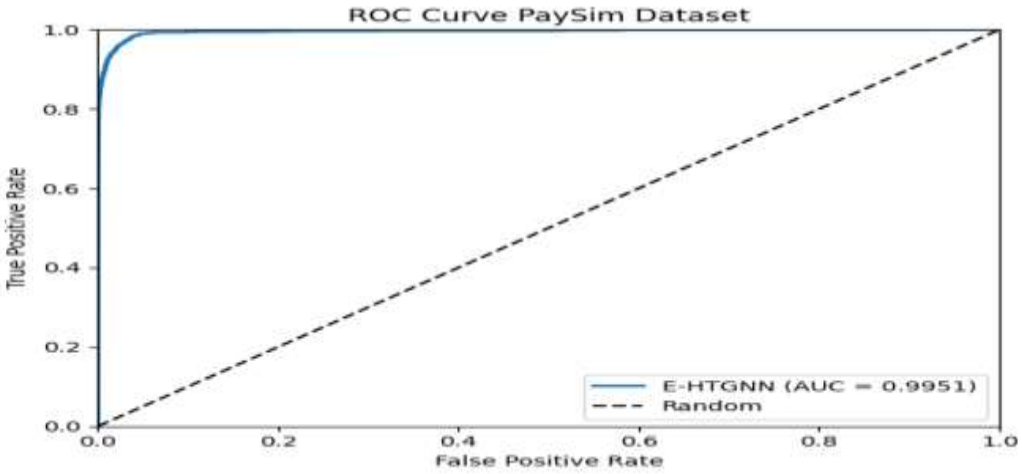

Figure 3. ROC curve for E-HTGNN on the imbalanced PaySim test set (AUC = 0.9951); the dashed diagonal represents random guessing.

The ROC curve in Figure 3 demonstrates near-perfect separation between fraud and non-fraud classes. E-HTGNN significantly outperforms the strongest baseline variant, HTGNN 2-hop, in both AUC-ROC (0.9951 vs 0.956) and PR-AUC (0.8876 vs 0.682), while risk-score distributions show that fraudulent transactions receive scores predominantly above 0.8 with little overlap with the legitimate-transaction distribution.

Results on IEEE-CIS

On the IEEE-CIS dataset, an enhanced version of E-HTGNN (12-dimensional entity features, 14-day LSTM, cost-sensitive threshold) was evaluated on the validation set (days 145–180). Results are shown in Table 8 and the corresponding confusion matrix in Table 9.

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

Table 8. E-HTGNN performance on the IEEE-CIS balanced sample.

Metric	Value
AUC-ROC	0.841
PR-AUC	0.594
Precision	0.445
Recall	0.690
F1-score	0.541
Latency (ms)	714

Table 9. Confusion matrix for the IEEE-CIS dataset (cost-sensitive threshold = 0.306).

	Predicted Non-Fraud	Predicted Fraud
Actual Non-Fraud	55,551	1,298
Actual Fraud	1,245	876

Figure 4. Confusion matrix visualisation for the IEEE-CIS validation set.

The model correctly identifies 876 of 2,121 fraud cases while flagging 1,298 false positives, yielding an ROC-AUC of 0.841 and PR-AUC of 0.594. The lower performance on IEEE-CIS relative to PaySim reflects the greater difficulty of the real-world dataset, attributable to higher feature sparsity (many missing values and anonymised V-features), more complex fraud patterns involving device fingerprints and email domains, and more extreme class imbalance in the underlying data (fraud ratio approximately 3.5% in the sampled data but only 0.5% in the full dataset). Nevertheless, the achieved AUC-ROC of 0.841 and PR-AUC of 0.594 establish a strong reference point for future research on this dataset.

Table 10. Overall comparison of the baseline HTGNN and the proposed E-HTGNN.

Dataset	Model	AUC-ROC	PR-AUC	Precision	Recall	F1
PaySim (imbalanced)	Baseline HTGNN 2-hop	0.956	0.682	–	–	–
PaySim (imbalanced)	E-HTGNN (Proposed)	0.9951	0.8876	0.892	0.803	0.845
IEEE-CIS (balanced sample)	E-HTGNN (Proposed)	0.841	0.594	0.445	0.69	0.541

Overall Comparison

Table 11 summarises the architectural differences between the baseline HTGNN and the proposed E-HTGNN by component, together with the rationale for each enhancement.

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

JOURNAL OF SCIENCE TECHNOLOGY AND EDUCATION 14(2), JUNE, 2026
E-ISSN: 3093-0898, PRINT ISSN: 2277-0011; Journal homepage: www.atbufstejoste.com

Table 11. Component-level comparison between the baseline HTGNN and the proposed E-HTGNN.

Component	Baseline HTGNN	Enhanced E-HTGNN	Why Better
Spatial	Standard GAT	GAT + Residual + LayerNorm	Better gradient flow; preserves information
Temporal	Single LSTM	Multi-scale (BiLSTM + Transformer) + Gated Fusion	Captures both bursts and slow trends
Semantic	Self-attention over concatenation	Cross-attention per entity type + entity embeddings	Learns entity-specific importance
Loss	Cross-entropy	Focal Loss ($\gamma = 2$, $\alpha = 0.75$)	Focuses on rare fraud cases
Optimisation	Standard Adam	Adam + Scheduler + Gradient Clipping	More stable training

DISCUSSION OF FINDINGS

Four findings emerge from the experiments. First, temporal context is crucial: incorporating historical entity features (via rolling averages or an LSTM) dramatically improves detection over static graph methods, with an earlier static-GAT variant achieving only AUC ≈ 0.60 on PaySim. Second, separating senders and receivers as distinct node types with independent temporal features allows the model to learn asymmetric fraud patterns, such as distinguishing a fraudulent sender from a victimised receiver. Third, bidirectional edges enable suspicion to propagate in both directions between a fraudulent transaction and its participating entities. Fourth, focal loss effectively handles class imbalance, with its benefit particularly pronounced on the imbalanced PaySim test set.

On inference latency, PaySim achieves 1.87 ms per transaction, comfortably within real-time requirements (< 50 ms), whereas IEEE-CIS incurs higher latency (714 ms) because the model applies GAT and LSTM computation over a 14-day window; this could be reduced through model pruning or a two-stage BatchNet/SpeedNet deployment architecture, as adopted in the baseline HTGNN study.

The use of heterogeneous graph neural networks is motivated by the inherently relational nature of financial transactions, in which cards, devices, and accounts interact. Bidirectional edges enhance information propagation, multi-scale temporal attention captures evolving fraud patterns, and the adaptive focal loss addresses the class imbalance inherent in fraud datasets. Taken together, the results confirm that incorporating temporal entity features, bidirectional edges, entity-type separation, and focal loss substantially improves fraud detection performance while preserving the low latency required for real-time deployment, at least on the PaySim dataset; further optimisation is required to bring IEEE-CIS latency within real-time bounds.

CONCLUSION AND FUTURE WORK

This study addressed the challenge of real-time financial fraud detection in e-commerce and mobile-money systems by proposing an Enhanced Heterogeneous Temporal Graph Neural Network (E-HTGNN) that integrates spatial, temporal, and semantic information from transaction data. The key innovations of the proposed framework are bidirectional message passing between transactions and entity nodes; multi-scale temporal encoding using rolling averages (PaySim) or a 14-day LSTM (IEEE-CIS); entity-type separation with independent feature projections to capture asymmetric fraud patterns; a focal-loss objective ($\alpha = 0.75$, $\gamma = 2$) that focuses training on hard, misclassified fraud cases; and cost-sensitive threshold tuning (false-negative cost = 5× false-positive cost).

On the imbalanced PaySim dataset, E-HTGNN achieved an AUC-ROC of 0.9951 and

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

PR-AUC of 0.8876, substantially outperforming the baseline HTGNN 2-hop model (AUC 0.956, PR-AUC 0.682), while sustaining an inference latency of 1.87 ms per transaction, well within real-time requirements. On the more challenging, real-world IEEE-CIS dataset, E-HTGNN achieved an AUC-ROC of 0.841 and PR-AUC of 0.594, establishing a strong reference point for future work. These results validate the study's core hypotheses: that temporal context is crucial for fraud detection, that entity-type separation improves discriminative power, that bidirectional edges propagate suspicious signals effectively, and that focal loss handles class imbalance well.

Contributions

1. A novel E-HTGNN framework integrating bidirectional edges, multi-scale temporal features, entity-type attention, and focal loss for real-time fraud detection.
2. A graph construction methodology that converts raw transaction logs into heterogeneous temporal graphs spanning multiple entity types.
3. Empirical evidence that E-HTGNN outperforms the state-of-the-art HTGNN baseline on the PaySim dataset by a large margin (AUC 0.995 vs 0.956; PR-AUC 0.888 vs 0.682).
4. A reproducible implementation on a GPU-accelerated Kaggle environment, providing a strong reference point for future research on the IEEE-CIS dataset.

RECOMMENDATIONS AND FUTURE WORK

Building on these findings, future research should pursue the following directions:

1. Evaluation on larger and more diverse real-world datasets from banks, fintech companies, and e-commerce platforms, to assess robustness, scalability, and generalisation across fraud scenarios.
2. Deployment within a real-time payment-processing infrastructure to evaluate latency, throughput, and operational reliability under practical conditions, including a two-stage BatchNet/SpeedNet architecture to reduce IEEE-CIS latency from 714 ms to below 50 ms, and knowledge distillation for edge deployment.
3. Incorporation of additional heterogeneous data sources, such as IP addresses, geolocation, browser characteristics, and network traffic features, to further enrich graph representations.
4. Adoption of continual, online, or incremental graph learning to allow the model to adapt to emerging fraud patterns without full retraining, together with training on the full, imbalanced IEEE-CIS dataset and application to other financial fraud datasets (e.g., Elliptic Bitcoin, Alibaba Cloud fraud data).
5. Integration of explainability techniques, such as GNNExplainer, attention visualisation, or feature attribution methods, to improve transparency and support regulatory compliance and analyst trust.
6. Evaluation of model robustness against adversarial attacks and fraud camouflage techniques to assess suitability for deployment in highly dynamic financial environments.
7. Financial fraud is an ever-evolving threat that demands adaptive, real-time detection systems. The E-HTGNN framework presented in this paper provides a robust and accurate solution that outperforms current state-of-the-art methods on synthetic data and shows strong promise on real-world data, offering a foundation for further research and practical deployment in e-commerce and banking security systems.

REFERENCES

Ahmad, F., Joy, K., & Simpao, S. (2025). E-Commerce Fraud Perceptions in Pakistan and Impacts of Risks and

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

Preventive Measures. ITQAN: Journal of Islamic Economics, 4(1).

Alarfaj, A., et al. (2022). Deep Learning-Based Fraud Detection in E-commerce.

Ananthu, S., Sethumadhavan, N., & Narayanan Ag, H. (2021). Credit Card Fraud Detection using Apache Spark Analysis. Proceedings of the 5th International Conference on Trends in Electronics and Informatics, ICOEI 2021, 998–1002.

Baltrušaitis, T., Ahuja, C., & Morency, L. P. (2019). Multimodal machine learning: A survey and taxonomy. IEEE Transactions on Pattern Analysis and Machine Intelligence, 41(2), 423–443.

Byrapu Reddy, S. R., Kanagala, P., Ravichandran, P., Pulimamidi, D. R., Sivarambabu, P. V., & Polireddi, N. S. A. (2024). Effective fraud detection in e-commerce: Leveraging machine learning and big data analytics. Measurement: Sensors, 33, 101138.

Charles, E., Nanduri, J., Jia, Y., Oka, A., & Beaver, J. (n.d.). A Review of Financial Fraud Detection in E-Commerce Using Machine Learning.

Dou, W., et al. (2020). Dynamic Graph Neural Networks for Fraud Detection.

Elly, A., Robert, A., Mabaso, F., & Oladele, S. (n.d.). Developing a Framework for Real-Time Fraud Detection in E-Commerce.

Goyal, G., Tyagi, R., & Tyagi, S. (2024). Graph Neural Networks for Fraud Detection in E-commerce Transactions. 2024 International Conference on Computing, Sciences and Communications (ICCSC), 1–6.

He, T. (n.d.). Dynamic E-commerce Solution for Real-Time Detection of Credit Card Fraud.

Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial Fraud: A Review of Anomaly Detection Techniques and Recent Advances. Expert Systems with Applications, 193.

Hu, Z., Dong, Y., Wang, K., & Sun, Y. (2020). Heterogeneous graph transformer. Proceedings of The Web Conference (WWW), 2704–2710.

Juniper Research. (2023). Online payment fraud: Market forecasts, emerging threats & segment analysis 2023–2028. Juniper Research.

Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. International Conference on Learning Representations (ICLR).

Kumah, R., Han, S., Okunola, A., & Mondal, S. (n.d.). Designing an Intuitive Real-Time Financial Fraud Detection System for Non-Technical Users.

Liu, J., Gu, X., & Shang, C. (2020). Quantitative Detection of Financial Fraud Based on Deep Learning with Combination of E-Commerce Big Data. Complexity, 2020.

Liu, T., et al. (2020). GNN-Based Approach for Transaction Fraud.

Liu, J., Zhang, X., & Zhao, Y. (2022). Dynamic graph learning for fraud detection in financial transactions. Expert Systems with Applications, 192, 116312.

Lu, Y., et al. (2021). Graph Neural Networks in Real-Time Fraud Detection with Lambda Architecture. arXiv:2110.04559.

Mat Taupit, A. S., & Azizan, N. (2023). The Planning Process of the Online Transaction Fraud Detection Using Backlogging on an E-Commerce Website. Malaysian Journal of Science Health & Technology, 9(1), 56–62.

Momoh, M. O., Shobowale, K. O., Abubakar, Z. M., Yahaya, B., & Yusuf, I. (2022). Blockchain adoption in aviation: Opportunities and challenges. *International Journal of Electrical Engineering and Computing*, *6*(2), 92-97.

Mutemi, A., & Bação, F. (2024). E-Commerce Fraud Detection Based on Machine Learning Techniques: Systematic

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.

Literature Review. Big Data Mining and Analytics, 7(2), 419–444.

Nasdaq. (2024). Nasdaq Verafin 2024 global financial crime report. Nasdaq, Inc.

Nguyen, H., & Le, B. (2025). Real-time transaction fraud detection via heterogeneous temporal graph neural network. Proceedings of the 17th International Conference on Agents and Artificial Intelligence (ICAART 2025), Vol. 2, 364–375. SCITEPRESS.

Nwankwo, W., Chinedu, P. U., Daniel, A., Shaba, S. M., Muyideen, M. O., Nwankwo, C. P., ... & Uwadia, F. (2022, November). Educational fintech: Promoting stakeholder confidence through automatic incidence resolution. In *The International Symposium on Computer Science, Digital Economy and Intelligent Systems* (pp. 947-963). Cham: Springer Nature Switzerland.

Okofu, S. N., Maureen, A., Ako, R. E., & Ojugo, A. A. (n.d.). Improving Customer Trust through Fraud Prevention E-Commerce Model.

Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., & Bronstein, M. (2020). Temporal graph networks for deep learning on dynamic graphs. arXiv:2006.10637.

Ryu, J., Kim, S., & Lee, H. (2023). Graph-based anomaly detection for e-commerce fraud prevention. IEEE Access, 11, 44212–44225.

Sarma, W., & Dey, S. (2021). AI and Machine Learning in Fraud Detection for Finance and E-Commerce. International Journal of Innovative Research in Computer and Communication Engineering, 9(10).

Suardiman, N., Sudimanto, Dhanny, S., Tjahjadi, D., Permana, B., & Ukar, K. (2024). E-Commerce Fraud Detection Using Generated Data From BANKSIM Using Machine Learning Approach: A Pilot Study. Proceedings of the 2024 18th International Conference on Ubiquitous Information Management and Communication (IMCOM 2024).

Tang, J., Li, X., Wang, Y., & Liu, H. (2022). Multimodal fraud detection using heterogeneous networks. ACM Transactions on Knowledge Discovery from Data, 16(8), 1–22.

Trinh, T. K., & Wang, Z. (2024). Dynamic graph neural networks for multi-level financial fraud detection: A temporal-structural approach. Annals of Applied Sciences, 5(1).

Venkatesan, K. B., Ramachandran, G. S., & Ramakrishnan, G. (2025). Scalable Machine Learning Models Achieving Efficiency in Fraud Detection. 4th International Conference on Sentiment Analysis and Deep Learning (ICSADL 2025), 1230–1235.

Wang, Y., Zhao, L., & Chen, D. (2023). Temporal graph-based fraud detection in digital payment networks. Neurocomputing, 532, 150–164.

Wei, C., Zhang, F., & Hu, J. (2022). Synthetic data generation for fraud detection using adversarial networks. Information Sciences, 600, 375–391.

Yu, J., Wang, H., Wang, X., Li, Z., Qin, L., Zhang, W., Liao, J., Zhang, Y., & Yang, B. (2024). Temporal Insights for Group-Based Fraud Detection on e-Commerce Platforms. IEEE Transactions on Knowledge and Data Engineering.

Zakaria, R. M., Rahman, M. M., Hossain Choudhury, M. T., Rahman, H., & Rafi, M. A. (2025). Detecting financial fraud in real-time transactions using graph neural networks and anomaly detection techniques. Journal of Economics, Finance and Accounting Studies, 7(6), 01–13.

Zhang, Y., Liu, T., & Wang, X. (2023). Deep learning-based fraud detection in e-commerce: A comprehensive review. ACM Computing Surveys, 55(12), 1–35.

Corresponding author: M. Yahaya
yahayammte304@st.futminna.edu.ng
Department of Cyber Security Science, Federal University of Technology Minna, PMB 65 Minna Niger State Nigeria.