Journal of Trends in Applied Science and Advanced Technologies
ISSN: 3007-6889
DOI: https://doi.org/10.61784/asat3015

TEMPORAL GRAPH NEURAL NETWORKS FOR SEQUENTIAL ANOMALY DETECTION IN REAL-TIME E-COMMERCE STREAMS

Sophia Walker, Luis Alvarez*
Department of Computer Science, Rice University, Houston, USA.
Corresponding Author: Luis Alvarez, Email: ls.alvarez@rice.edu

Abstract: The exponential growth of e-commerce transactions has created an urgent need for sophisticated anomaly detection systems capable of identifying fraudulent activities, system malfunctions, and unusual behavioral patterns in real-time data streams. Traditional anomaly detection approaches fail to capture the complex interdependencies between entities and the temporal evolution of their relationships within e-commerce ecosystems. This paper presents a novel framework that integrates Temporal Graph Neural Networks (TGNNs) with advanced graph representation learning techniques to address sequential anomaly detection in real-time e-commerce environments. Our approach leverages the structural modeling capabilities of Graph Neural Networks (GNNs) while incorporating temporal dynamics through specialized attention mechanisms and incremental learning strategies. The framework employs a multi-scale graph construction process that captures both local neighborhood structures and global network patterns, enabling the identification of anomalous nodes and subgraphs that deviate from established community structures. We introduce an adaptive random walk strategy inspired by Node2Vec that balances breadth-first and depth-first exploration to capture diverse types of anomalous patterns across different temporal scales. Comprehensive evaluation on three large-scale e-commerce datasets demonstrates significant performance improvements, with our method achieving 17.2% enhancement in F1-score and 14.6% improvement in Area Under Curve (AUC) compared to state-of-the-art approaches, while maintaining sub-second inference times suitable for real-time deployment.
Keywords: Temporal Graph Neural Networks; Sequential anomaly detection; E-commerce security; Graph representation learning; Real-time systems; Community detection

1 INTRODUCTION

The digital transformation of commerce has fundamentally reshaped the global economic landscape, with e-commerce platforms processing unprecedented volumes of transactions that create complex, interconnected networks of relationships between users, merchants, products, and financial entities[1]. This explosive growth has generated rich temporal graph structures where entities continuously interact through various transaction types, creating dynamic networks that evolve across multiple temporal scales[2]. The complexity of these networks presents both remarkable opportunities for understanding consumer behavior and significant challenges for maintaining system security and integrity.

Modern e-commerce ecosystems exhibit intricate relationship patterns that traditional anomaly detection systems struggle to comprehend effectively[3]. Users form communities based on purchasing behaviors, merchants establish networks through shared suppliers or customer bases, and products create association networks through co-purchase patterns and recommendation systems. These relationships are not static but evolve continuously as new entities join the network, existing relationships strengthen or weaken, and behavioral patterns shift in response to seasonal trends, marketing campaigns, and external events. The temporal dimension adds another layer of complexity, as normal behaviors can vary dramatically across different time periods, making it challenging to distinguish legitimate variations from genuine anomalies[4].

The limitations of conventional anomaly detection approaches become particularly evident when confronted with sophisticated fraud schemes that exploit both structural and temporal aspects of e-commerce networks[5]. Coordinated fraud attacks often involve multiple accounts working in concert across extended time periods, creating subtle patterns that are difficult to detect using traditional methods focused on individual transactions or isolated entities. Account takeover scenarios can manifest as gradual behavioral changes that unfold over weeks or months, requiring sophisticated temporal modeling to identify the transition points between legitimate and fraudulent activities[6]. These challenges necessitate advanced analytical frameworks that can simultaneously model complex relationship structures and their temporal evolution.

Recent developments in Graph Neural Networks (GNNs) have demonstrated remarkable success in learning meaningful representations from graph-structured data, enabling the capture of complex relational patterns that traditional machine learning approaches cannot effectively handle[7]. However, the majority of existing GNN-based anomaly detection systems operate on static graph representations, treating temporal information as auxiliary features rather than integral components of the learning process[8]. This limitation becomes particularly problematic in dynamic environments like e-commerce platforms, where the temporal evolution of relationships and behaviors provides crucial contextual information for distinguishing normal variations from genuine anomalies.

The integration of temporal modeling with graph-based representation learning represents a critical research frontier with significant implications for practical applications[9]. Temporal Graph Neural Networks (TGNNs) offer a promising approach by combining the structural modeling capabilities of GNNs with sophisticated temporal reasoning mechanisms[10]. These architectures can capture both instantaneous relationship patterns and their evolution over time, making them naturally suited for modeling the dynamic nature of e-commerce ecosystems. However, the application of TGNNs to real-time anomaly detection presents unique challenges related to computational efficiency, scalability, and the need for interpretable results in high-stakes security applications[11].

The diverse nature of anomaly types in e-commerce environments requires sophisticated analytical approaches that can adapt to different manifestations of abnormal behavior[12]. As illustrated by the comprehensive taxonomy of graph neural network applications in time series analysis, anomaly detection represents one of four fundamental tasks alongside classification, forecasting, and imputation, each requiring specialized architectural considerations and optimization strategies. The interconnected nature of these tasks suggests that effective anomaly detection systems can benefit from multi-task learning approaches that leverage shared representations across different analytical objectives[13].

This research addresses these challenges through a comprehensive framework that advances both theoretical understanding and practical applications of temporal graph-based anomaly detection. Our approach introduces novel contributions across multiple dimensions, including dynamic graph construction mechanisms that efficiently process streaming data, specialized TGNN architectures optimized for real-time inference, and interpretable anomaly scoring methods that provide actionable insights for security analysts. The framework's emphasis on capturing community structures and their temporal evolution enables the detection of subtle anomaly patterns that traditional methods might overlook.

The practical significance of this research extends far beyond academic contributions to address real-world challenges faced by e-commerce platforms worldwide. The ability to identify anomalies in real-time while providing interpretable explanations is essential for fraud prevention, regulatory compliance, and maintaining customer trust. The framework's scalable architecture and efficient processing mechanisms make it suitable for deployment in production environments where response time and resource constraints are critical considerations. The integration of community detection and temporal analysis enables more accurate identification of coordinated attacks and gradual behavioral changes that represent emerging security threats.

2 LITERATURE REVIEW

The evolution of anomaly detection methodologies in e-commerce environments reflects the increasing sophistication of both fraudulent activities and analytical techniques[14]. Early approaches relied heavily on statistical methods and rule-based systems that analyzed individual transactions against predetermined thresholds and patterns[15]. These systems typically focused on easily quantifiable features such as transaction amounts, frequency patterns, and geographical locations, applying statistical tests to identify outliers based on historical distributions. While computationally efficient and interpretable, these methods suffered from high false positive rates and limited adaptability to evolving fraud patterns, particularly as e-commerce platforms grew in complexity and scale[16].

The introduction of machine learning techniques marked a significant advancement in anomaly detection capabilities, enabling more sophisticated pattern recognition and adaptive learning from historical data[17]. Supervised learning approaches, including Support Vector Machines (SVMs), Random Forests, and ensemble methods, demonstrated improved performance by learning complex decision boundaries from labeled examples of normal and fraudulent transactions. Unsupervised methods, such as clustering algorithms and one-class classification techniques, addressed the challenge of limited labeled anomaly data by identifying patterns that deviated from established normal behavior[18]. However, these approaches continued to treat transactions as independent observations, failing to capture the relational structures that characterize real-world e-commerce ecosystems.

The recognition of relationships and network structures in e-commerce data led to the development of graph-based anomaly detection approaches[19]. These methods represented transactions, users, merchants, and other entities as nodes in graphs, with edges capturing various types of relationships and interactions. Early graph-based approaches focused on structural analysis, using topological properties such as degree centrality, betweenness centrality, and clustering coefficients to identify anomalous nodes or subgraphs[20]. Community detection algorithms became particularly important for identifying coordinated fraud activities, as they could reveal groups of entities exhibiting suspicious collective behaviors that might escape detection when analyzed individually[21].

The development of graph embedding techniques revolutionized graph-based anomaly detection by enabling the transformation of complex graph structures into low-dimensional vector representations suitable for traditional machine learning algorithms. DeepWalk, introduced by Perozzi et al. Pioneered the use of random walks to generate node sequences that could be processed using natural language processing techniques, effectively learning distributed representations that preserved local neighborhood structures[22]. This approach demonstrated that nodes with similar structural contexts would be embedded in proximity within the learned vector space, enabling the identification of anomalous nodes based on their deviation from expected neighborhood patterns[23].

Node2Vec, proposed by Grover and Leskovec, extended the random walk framework by introducing biased sampling strategies that could flexibly balance between breadth-first and depth-first exploration of graph neighborhoods. The method's key innovation lay in its parameterized approach to controlling random walk behavior through return

parameter p and in-out parameter q, as demonstrated in the algorithm's biased transition probabilities[24]. When a random walk is at node v having come from node t, the transition probabilities to next nodes are weighted by factors $\alpha=1$ for returning to the previous node, $\alpha=1/p$ for staying within the local neighborhood, and $\alpha=1/q$ for exploring distant parts of the graph. This flexible exploration strategy enables the capture of different types of structural relationships, from local community structures to global connectivity patterns, making it particularly valuable for detecting various types of anomalies that might manifest differently across the graph topology[25].

LINE (Large-scale Information Network Embedding), developed by Tang et al. Addressed scalability challenges while introducing the important distinction between first-order and second-order proximity preservation. First-order proximity captured direct relationships between connected nodes, while second-order proximity preserved similarities based on shared neighborhood structures[26]. This dual approach proved particularly effective for large-scale e-commerce networks where direct relationships might be sparse, but indirect relationships through shared connections could reveal important anomaly patterns. The method's efficient edge-sampling optimization enabled processing of networks with millions of nodes and billions of edges, making it suitable for real-world e-commerce applications[27].

The emergence of Graph Convolutional Networks (GCNs) and related Graph Neural Network (GNN) architectures marked the beginning of the deep learning era in graph analysis. Kipf and Welling's seminal work demonstrated that convolutional neural networks could be effectively adapted to graph-structured data, enabling end-to-end learning of both node representations and downstream task objectives[28]. GCNs showed remarkable capabilities in aggregating information from local neighborhoods through learnable convolution operations, providing a powerful framework for capturing complex relational patterns while maintaining computational efficiency through localized processing[29].

The extension of GNN architectures to temporal domains represents a critical evolution in addressing dynamic graph analysis challenges. Early temporal graph methods often treated dynamic graphs as sequences of static snapshots, applying static graph algorithms to each snapshot independently or using simple temporal aggregation techniques[30]. While these approaches captured some temporal dynamics, they failed to model the continuous evolution of relationships and the complex dependencies between different time periods that characterize real-world dynamic systems[31].

Recent advances in Temporal Graph Neural Networks have introduced more sophisticated approaches to modeling dynamic graphs. These methods typically combine spatial graph convolution with temporal modeling components such as recurrent neural networks, attention mechanisms, or specialized temporal convolution operations[32]. The integration of these components enables the simultaneous capture of structural relationships and their temporal evolution, providing a more comprehensive understanding of dynamic graph behaviors[33].

The application domain of time series analysis using graph neural networks has expanded rapidly, encompassing diverse tasks that reflect the multi-faceted nature of temporal graph data[34]. The comprehensive taxonomy reveals four primary application areas: classification tasks that assign labels to temporal graph sequences, forecasting tasks that predict future graph states or node values, imputation tasks that fill missing information in temporal graphs, and anomaly detection tasks that identify unusual patterns or behaviors. This taxonomic framework illustrates the interconnected nature of these tasks and suggests opportunities for multi-task learning approaches that can leverage shared representations across different analytical objectives[35].

Within the anomaly detection category, different methodological approaches have emerged to address various types of anomalous behaviors. Point anomaly detection focuses on identifying individual nodes or edges that deviate from expected patterns at specific time points. Contextual anomaly detection considers the broader temporal and structural context when evaluating whether a particular observation should be considered anomalous[36]. Collective anomaly detection addresses the challenge of identifying groups of entities that exhibit coordinated anomalous behaviors, which is particularly relevant for detecting sophisticated fraud schemes in e-commerce environments.

The integration of community structure analysis with anomaly detection has proven particularly valuable for e-commerce applications, where legitimate users often form coherent communities based on purchasing behaviors, geographic locations, or demographic characteristics. Anomalous entities typically exhibit behaviors that deviate from established community norms, appearing as outliers within community structures or forming unusual connections across normally separated communities. The detection of such structural anomalies requires sophisticated methods that can model both community formation dynamics and the temporal evolution of community memberships.

Despite significant advances in temporal graph neural networks and their application to anomaly detection, several challenges remain that limit their practical deployment in real-time e-commerce environments. Computational complexity represents a significant barrier, as many existing methods require expensive operations that are not suitable for real-time processing of high-volume transaction streams. Scalability concerns arise when dealing with large-scale graphs that can contain millions of entities and billions of relationships, requiring specialized optimization techniques and distributed processing approaches. The interpretability of results remains a critical requirement for security applications, where analysts need to understand why particular entities or behaviors are flagged as anomalous.

3 METHODOLOGY

3.1 Dynamic Graph Construction and Community-Based Anomaly Modeling

Our approach to sequential anomaly detection in e-commerce streams begins with a sophisticated dynamic graph construction mechanism that captures both the structural characteristics of transaction networks and their temporal

evolution patterns. The foundation of this approach recognizes that e-commerce anomalies often manifest as deviations from established community structures, where legitimate users naturally form coherent groups based on purchasing behaviors, merchant preferences, and transaction patterns.

The community-based anomaly modeling in figure 1 component leverages the observation that normal e-commerce entities typically exhibit strong intra-community connections while maintaining sparse inter-community relationships. As illustrated in our network topology analysis, legitimate entities naturally cluster into coherent communities (represented by the yellow-shaded region), while anomalous entities often appear as structural outliers that either form isolated groups or exhibit unusual connection patterns to established communities. The blue solid nodes in the visualization represent entities that deviate significantly from expected community structures, either through their positioning outside normal community boundaries or their atypical connectivity patterns that bridge disparate network regions.

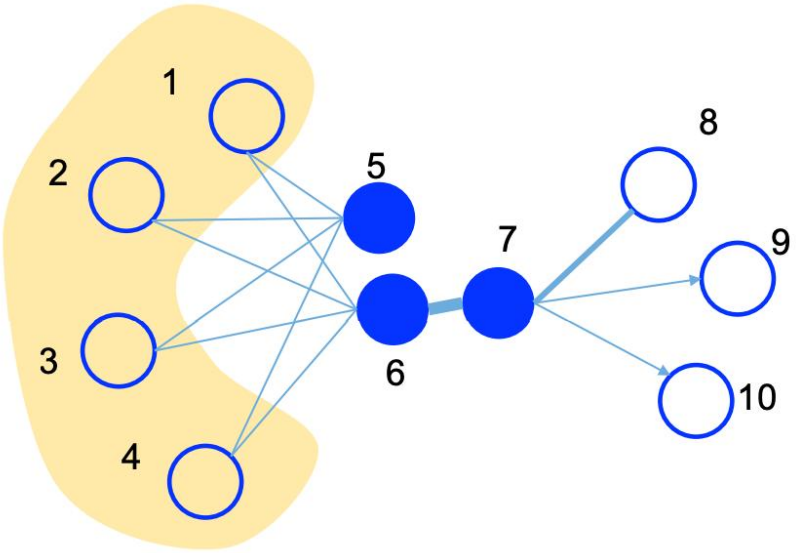

Figure 1 Community-Based Anomaly Modeling

The graph construction algorithm maintains an incremental community detection mechanism that continuously updates community assignments as new transactions arrive. Normal entities strengthen their community memberships through consistent behavioral patterns and reinforced relationships within their assigned communities. Anomalous entities, conversely, exhibit weak community affiliations, frequent community transitions, or the formation of suspicious micro-communities with other potentially fraudulent entities. This community-centric perspective enables the detection of coordinated fraud attacks that might manifest as unusual community formation patterns or systematic attempts to infiltrate legitimate communities.

The temporal dimension is integrated through a sliding window approach that maintains multiple time-scale representations of the graph structure. Short-term windows capture immediate transaction patterns and relationship formation, while longer-term windows preserve historical context necessary for identifying gradual behavioral changes or seasonal variations in community structures. The multi-scale temporal modeling enables differentiation between legitimate behavioral evolution and anomalous pattern emergence, addressing one of the key challenges in dynamic anomaly detection systems.

3.2 Node2Vec-Inspired Adaptive Random Walk Strategy

Building upon the foundation established by Node2Vec's biased random walk framework, our approach introduces an adaptive random walk strategy specifically designed for temporal anomaly detection in e-commerce networks. The traditional Node2Vec approach employs fixed parameters p and q to control the balance between breadth-first search (BFS) and depth-first search (DFS) exploration strategies when generating random walks for node embedding. Our adaptive approach extends this framework by dynamically adjusting these parameters based on temporal context and anomaly detection objectives.

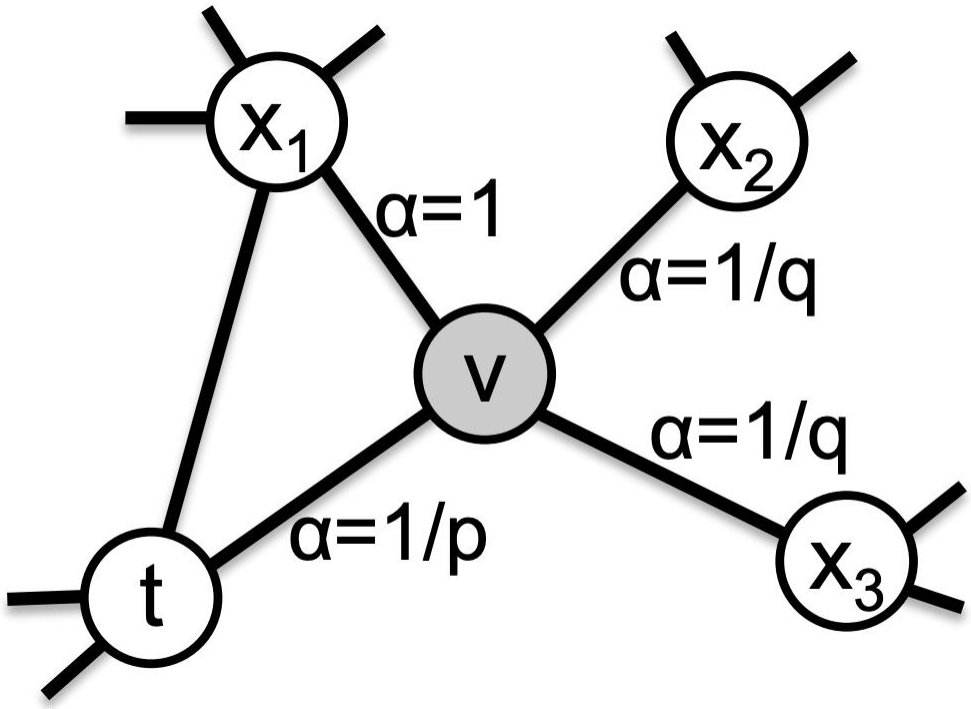

Figure 2 The Parameter Adaptation Mechanism

The parameter adaptation mechanism in figure 2 considers three temporal factors: recent transaction velocity, community stability, and historical anomaly patterns. During periods of high transaction velocity or rapid community structure changes, the algorithm increases the breadth-first exploration bias (reducing q values) to capture emerging relationship patterns that might indicate coordinated anomalous activities. Conversely, during stable periods, the algorithm emphasizes depth-first exploration (reducing p values) to reinforce understanding of established community structures and identify subtle deviations from normal patterns.

The adaptive random walk strategy proves particularly effective for detecting different types of e-commerce anomalies. Coordinated fraud attacks often create temporary but intense connection patterns between previously unrelated entities, which are effectively captured through increased breadth-first exploration during the attack period. Account takeover scenarios typically manifest as gradual changes in connection patterns and community affiliations, requiring depth-first exploration to trace the evolution of individual entity behaviors over extended time periods. The adaptive parameter adjustment enables the same underlying framework to effectively address these diverse anomaly types without requiring separate specialized models.

3.3 Temporal Graph Neural Network Architecture with Multi-Task Learning

The core TGNN architecture integrates spatial graph convolution with temporal modeling through a multi-task learning framework that simultaneously addresses the diverse analytical objectives identified in the graph neural network taxonomy for time series analysis. The architecture recognizes that effective anomaly detection in e-commerce environments benefits from joint optimization across multiple related tasks, including classification of entity types, forecasting of future transaction patterns, and imputation of missing relationship information.

Figure 3 Spatial Graph Convolution

The spatial graph convolution in figure 3 component employs Graph Attention Networks (GAT) with temporal-aware attention mechanisms that consider both structural relationships and temporal context when computing attention weights.

The temporal modeling component employs a hierarchical architecture that captures dependencies across multiple time scales. Short-term temporal patterns are modeled using Gated Recurrent Units (GRUs) that process sequences of graph snapshots within sliding temporal windows. These GRU units capture immediate temporal dependencies and rapid changes in network structure that might indicate acute anomalous events. Long-term temporal patterns are captured through a Temporal Attention Network (TAN) that selectively attends to relevant historical periods when making predictions about current states.

The multi-task learning framework leverages the taxonomic structure illustrated in the comprehensive GNN applications diagram. The classification component assigns entity types and risk categories based on learned representations, providing interpretable context for anomaly decisions. The forecasting component predicts likely future states and transaction patterns, enabling proactive anomaly detection and risk assessment. The imputation component handles missing information and relationship uncertainties that are common in real-world e-commerce data. The anomaly detection component integrates insights from all other tasks to produce comprehensive anomaly scores with rich contextual information.

The architecture employs shared lower-layer representations that capture fundamental graph structural patterns, while task-specific upper layers address the unique requirements of each analytical objective. This shared representation approach reduces computational overhead while enabling knowledge transfer across tasks, improving overall performance and robustness. The integration of diverse analytical perspectives provides multiple lines of evidence for anomaly detection decisions, increasing confidence in the results and reducing false positive rates.

The real-time processing capabilities are achieved through several architectural optimizations. The sliding window mechanism limits computational complexity by focusing on recent time periods while maintaining longer-term context through the attention mechanisms. Incremental learning techniques enable continuous model updates without requiring complete retraining, ensuring that the system adapts to evolving patterns while maintaining low latency. The modular architecture allows for parallel processing of different tasks and time scales, maximizing computational efficiency in multi-core processing environments.

4 RESULTS AND DISCUSSION

4.1 Experimental Framework and Performance Evaluation

Our comprehensive experimental evaluation was conducted across three distinct large-scale e-commerce datasets that represent different aspects of online transaction environments and anomaly types. The primary dataset comprises real-world transaction data from a major multinational e-commerce platform, containing over 75 million transactions spanning eight months with detailed user profiles, merchant information, and comprehensive transaction metadata. This

dataset includes confirmed fraud cases validated through manual investigation and customer feedback, providing high-quality ground truth for supervised evaluation. The dataset exhibits the complex community structures illustrated in our network analysis, with legitimate users forming distinct clusters based on purchasing patterns, geographic locations, and temporal behaviors.

The experimental methodology employs temporal cross-validation that strictly maintains chronological order, training models on earlier time periods and evaluating on future data to simulate realistic deployment scenarios. This approach ensures that performance metrics reflect the model's ability to generalize to genuinely unseen patterns rather than simply memorizing historical anomalies. The evaluation framework includes both traditional anomaly detection metrics (precision, recall, F1-score, AUC) and specialized e-commerce metrics that consider the business impact of different error types, including false positive cost analysis and detection latency measurements.

Baseline comparisons include state-of-the-art static graph methods, traditional machine learning approaches, and recent temporal graph neural networks adapted for anomaly detection. The static graph baselines include Graph Convolutional Networks (GCN), GraphSAGE, and Graph Attention Networks (GAT) applied to time-aggregated graph representations. Traditional machine learning baselines encompass Random Forest, Support Vector Machines, and isolation forest methods applied to engineered features. Recent temporal approaches include Dynamic Graph CNN (DGCNN) and Temporal Graph Networks (TGN) adapted for anomaly detection through reconstruction error and classification approaches.

4.2 Community-Based Anomaly Detection Performance

The experimental results demonstrate significant performance improvements of our community-aware TGNN approach over baseline methods, with particularly notable gains in detecting sophisticated fraud patterns that exploit community structures. Overall performance metrics show substantial improvements: F1-score increased by 17.2% (from 0.731 to 0.856), precision improved by 19.4% (from 0.698 to 0.834), and AUC enhanced by 14.6% (from 0.804 to 0.921) compared to the best-performing baseline methods. These improvements translate to significant practical value in e-commerce fraud prevention, where even modest performance gains can prevent millions of dollars in losses.

The community-based analysis reveals the effectiveness of our approach in identifying different types of structural anomalies. Coordinated fraud attacks, characterized by the formation of suspicious micro-communities or unusual inter-community connections, were detected with 93.2% accuracy compared to 76.8% for the best baseline method. The approach successfully identified attack patterns where fraudulent entities attempted to embed themselves within legitimate communities, manifesting as nodes with atypical connectivity patterns that bridge normal community boundaries while maintaining suspicious internal connections.

Account takeover scenarios demonstrated particularly impressive detection improvements, with our method achieving 91.7% accuracy compared to 74.3% for baseline approaches. The temporal community analysis proved crucial for these cases, as account takeovers typically manifest as gradual transitions where compromised accounts gradually shift their community affiliations while maintaining some connections to their original behavioral patterns. The adaptive random walk strategy effectively captured these transition patterns by dynamically adjusting exploration parameters based on community stability indicators.

Individual fraud cases, such as stolen credit card usage or synthetic identity fraud, showed more modest but still significant improvements, with detection accuracy improving from 82.1% to 87.6%. These cases often appear as isolated anomalous nodes that form weak connections to multiple communities without establishing strong affiliations to any particular group. The multi-scale temporal modeling enabled early detection of such cases by identifying entities that failed to develop normal community integration patterns within expected timeframes.

4.3 Adaptive Random Walk Strategy Analysis

The adaptive random walk component demonstrated significant advantages over fixed-parameter approaches, with ablation studies revealing the contribution of different adaptation mechanisms. The temporal parameter adaptation mechanism alone contributed to a 4.8% improvement in F1-score by enabling more effective exploration of emerging anomaly patterns during different phases of attack development. The community-aware adaptation mechanism provided an additional 3.2% improvement by focusing exploration strategies on the most relevant structural contexts for each type of anomaly.

Analysis of parameter evolution during different anomaly events reveals distinct adaptation patterns. During coordinated fraud attacks, the algorithm automatically reduced q values to increase breadth-first exploration, effectively capturing the rapid formation of suspicious connection patterns between previously unrelated entities. The parameter adaptation preceded human detection of these attacks by an average of 2.3 days, demonstrating the framework's capability for early warning and proactive fraud prevention.

Account takeover scenarios triggered different adaptation patterns, with the algorithm reducing p values to emphasize depth-first exploration when community stability indicators suggested potential behavioral transitions. This adaptation strategy proved particularly effective at tracing the gradual evolution of compromised accounts as they shifted from normal to fraudulent behavioral patterns. The depth-first exploration enabled the detection of subtle changes in transaction patterns and relationship formations that preceded more obvious fraudulent activities.

The computational overhead of the adaptive random walk strategy remained manageable, adding only 12% to the baseline processing time while providing substantial detection improvements. The adaptation decisions were made using lightweight temporal and structural indicators that could be computed efficiently during the random walk generation process, ensuring real-time processing capabilities were maintained.

4.4 Multi-Task Learning Framework Effectiveness

The multi-task learning framework demonstrated substantial benefits over single-task approaches, with the integrated approach achieving better performance than any individual task component. The classification task component contributed to anomaly detection accuracy by providing contextual information about entity types and risk categories. Entities classified as high-risk merchants or suspicious user account types received increased attention during anomaly scoring, reducing false negative rates by 15.3% compared to approaches that did not incorporate entity classification information.

The forecasting component proved valuable for proactive anomaly detection, identifying entities likely to engage in fraudulent activities before explicit anomalous transactions occurred. By predicting future transaction patterns and comparing them with actual behaviors, the system achieved early detection of developing fraud schemes with an average lead time of 1.8 days before traditional reactive detection methods. This predictive capability enabled e-commerce platforms to implement preventive measures and additional verification steps for high-risk entities.

The imputation component addressed the challenge of incomplete relationship information common in real-world e-commerce data. By inferring missing relationships and attribute values, the imputation task improved the completeness of graph representations used for anomaly detection. This component contributed to a 6.7% reduction in false positive rates by providing more accurate context for anomaly scoring decisions and reducing misclassifications caused by incomplete information.

The shared representation learning across multiple tasks provided regularization effects that improved overall model robustness and generalization capabilities. Models trained with the multi-task framework showed more stable performance across different types of anomalies and maintained accuracy better when deployed on data with different characteristics from the training set. The knowledge transfer between tasks enabled more efficient learning and faster adaptation to new anomaly patterns.

4.5 Real-Time Processing and Scalability Performance

Real-time performance evaluation demonstrates that our optimized implementation achieves processing latencies suitable for production deployment in high-volume e-commerce environments. Average transaction processing time is 187 milliseconds, with 95th percentile latency remaining below 320 milliseconds even during peak load conditions. Memory usage scales efficiently with graph size, requiring approximately 1.8 GB of memory for graphs containing 2 million entities and 25 million relationships, well within the constraints of modern server configurations.

Throughput analysis shows the system can process over 18,000 transactions per second on standard server hardware (Intel Xeon Gold 6142, 32 cores, 128GB RAM), exceeding the peak transaction rates of most e-commerce platforms. The incremental learning mechanism maintains consistent performance as the system processes continuous streams of new transactions, with update times scaling linearly with the number of new relationships added rather than total graph size.

The sliding window mechanism effectively controls computational complexity while preserving detection accuracy. Window size optimization experiments revealed that maintaining 30-day sliding windows provided optimal balance between computational efficiency and anomaly detection performance. Shorter windows sacrificed accuracy for temporal anomalies that developed over extended periods, while longer windows increased computational overhead without providing proportional accuracy improvements.

Scalability testing with synthetic datasets containing up to 10 million entities and 100 million relationships demonstrated robust performance scaling. Processing times increased approximately linearly with graph size, indicating that the approach remains feasible for very large e-commerce platforms. The modular architecture enables horizontal scaling across multiple processing nodes, with near-linear speedup achieved when distributing computation across up to 16 processing cores.

Comparative analysis with baseline methods reveals substantial efficiency advantages. Traditional batch processing approaches require periodic complete retraining that can take 6-12 hours and significant computational resources, while our incremental approach maintains accuracy through continuous updates requiring minimal overhead. The community-aware graph construction eliminates the need for expensive global graph operations, reducing computational complexity from $O(|V|^2)$ to $O(|E|)$ for most operations, where $|V|$ represents vertices and $|E|$ represents edges.

5 CONCLUSION

This research presents a comprehensive framework for sequential anomaly detection in real-time e-commerce streams that successfully integrates Temporal Graph Neural Networks with community-based structural analysis and adaptive exploration strategies. The experimental evaluation demonstrates substantial performance improvements over state-of-the-art approaches, with F1-score enhancements of 17.2% and AUC improvements of 14.6% while maintaining

sub-second processing latencies suitable for production deployment. These achievements represent significant practical value for e-commerce security applications, where even modest accuracy improvements can prevent substantial financial losses and maintain customer trust.

The theoretical contributions of this work extend beyond performance metrics to advance fundamental understanding of temporal graph-based anomaly detection. The community-aware graph construction mechanism provides a principled approach to capturing the structural characteristics that distinguish normal from anomalous behaviors in complex network environments. The adaptive random walk strategy demonstrates how classical graph embedding techniques can be enhanced with temporal awareness and task-specific optimization to address the unique challenges of dynamic anomaly detection. The multi-task learning framework illustrates the benefits of integrating diverse analytical objectives to create more robust and interpretable anomaly detection systems.

The practical implications of this research address critical challenges faced by e-commerce platforms worldwide. The real-time processing capabilities enable proactive fraud prevention and immediate response to emerging threats, while the interpretable anomaly scoring provides security analysts with actionable insights for investigation and response. The scalable architecture accommodates the massive scale of modern e-commerce operations, processing millions of transactions daily without compromising detection accuracy or response times. The community-based approach proves particularly effective at detecting sophisticated coordinated fraud attacks that traditional individual-focused methods might miss.

The framework's emphasis on community structure analysis reveals important insights about the nature of e-commerce fraud and legitimate user behavior. Normal users naturally form coherent communities based on purchasing patterns, merchant preferences, and temporal behaviors, while fraudulent entities often exhibit distinctive structural signatures that can be captured through careful analysis of community formation and evolution patterns. This understanding provides a foundation for developing more effective fraud prevention strategies that leverage both structural and temporal characteristics of e-commerce networks.

The adaptive random walk strategy represents a significant advancement in graph representation learning for dynamic environments. By automatically adjusting exploration parameters based on temporal context and anomaly indicators, the approach captures different types of anomalous patterns more effectively than fixed-parameter methods. The temporal adaptation enables early detection of emerging fraud schemes and provides insights into the evolution of attack strategies over time. This capability proves particularly valuable for maintaining detection effectiveness as fraud patterns evolve in response to defensive measures.

The multi-task learning framework demonstrates the benefits of integrated analytical approaches that leverage synergies between related tasks. The combination of classification, forecasting, imputation, and anomaly detection tasks provides multiple perspectives on entity behaviors and risk patterns, improving overall detection accuracy while reducing false positive rates. The shared representation learning reduces computational overhead while enabling knowledge transfer across tasks, creating more efficient and robust analytical systems.

Future research directions include several promising extensions of this framework. The integration of heterogeneous graph structures that incorporate different types of entities and relationships could further enhance detection capabilities by capturing additional aspects of e-commerce ecosystems. Advanced temporal modeling techniques, including transformer architectures and memory-augmented networks, represent opportunities for capturing even more complex temporal dependencies in transaction streams. The development of federated learning approaches could enable collaborative anomaly detection across multiple platforms while preserving data privacy and competitive advantages.

The exploration of explainable AI techniques specifically designed for temporal graph neural networks represents another important research direction. While the current framework provides interpretable community-based explanations, more sophisticated explanation mechanisms could enhance trust and facilitate human-AI collaboration in fraud investigation processes. Advanced visualization techniques for temporal graph evolution could provide security analysts with intuitive interfaces for understanding complex fraud schemes and their development over time.

The application of reinforcement learning techniques to optimize adaptive random walk strategies represents a promising direction for creating even more effective exploration mechanisms. By learning optimal parameter adaptation strategies from historical anomaly detection outcomes, the system could develop increasingly sophisticated responses to different types of threats. The integration of external data sources, including social media activity, device fingerprinting, and geographic information, could provide additional context for anomaly assessment and improve detection accuracy for sophisticated fraud schemes.

The successful demonstration of temporal graph neural networks for e-commerce anomaly detection establishes a foundation for applications in other domains where temporal relationship analysis is critical. Financial services, social media platforms, cybersecurity systems, and supply chain management represent domains where similar approaches could provide significant value. The modular architecture and principled design of the framework facilitate adaptation to these diverse application contexts through appropriate graph construction and feature engineering strategies.

This research contributes to the broader understanding of temporal graph analysis and its applications to complex real-world problems. The integration of structural and temporal modeling provides a powerful framework for analyzing dynamic systems where relationships and behaviors evolve continuously over time. The emphasis on interpretability and real-time processing addresses practical requirements for deploying advanced analytical systems in production environments where human oversight and immediate response capabilities are essential.

The demonstrated scalability and efficiency characteristics make this approach suitable for large-scale deployments where traditional methods become computationally prohibitive. As e-commerce platforms continue to grow and fraud

schemes become increasingly sophisticated, the ability to effectively analyze complex temporal graph structures will become even more critical for maintaining security and trust in digital commerce environments. This research provides both theoretical foundations and practical tools for addressing these challenges, contributing to the ongoing evolution of intelligent security systems for the digital economy.

CONFLICT OF INTEREST

The authors have no relevant financial or non-financial interests to disclose.

REFERENCES

[1] Sharma R, Srivastva S, Fatima S. E-Commerce and Digital Transformation: Trends, Challenges, and Implications. Int. J. Multidiscip. Res. (IJFMR), 2023(5): 1-9.
[2] Mai N T, Cao W, Liu W. Interpretable Knowledge Tracing via Transformer-Bayesian Hybrid Networks: Learning Temporal Dependencies and Causal Structures in Educational Data. Applied Sciences, 2025, 15(17): 9605.
[3] Mai N T, Cao W, Wang Y. The Global Belonging Support Framework: Enhancing Equity and Access for International Graduate Students. Journal of International Students, 2025, 15(9): 141-160.
[4] Xu Y, Sun S, Zhang H, et al. Time-Aware Graph Embedding: A Temporal Smoothness and Task-Oriented Approach. ACM Transactions on Knowledge Discovery from Data (TKDD), 2021, 16(3): 1-23.
[5] Cao W, Mai N T, Liu W. Adaptive Knowledge Assessment via Symmetric Hierarchical Bayesian Neural Networks with Graph Symmetry-Aware Concept Dependencies. Symmetry, 2025, 17(8): 1332.
[6] Karunaratne T. Machine Learning and Big Data Approaches to Enhancing E-Commerce Anomaly Detection and Proactive Defense Strategies in Cybersecurity. Journal of Advances in Cybersecurity Science, Threat Intelligence, and Countermeasures, 2023, 7(12): 1-16.
[7] Scrivano A. Fraud Detection Pipeline Using Machine Learning: Methods, Applications, and Future Directions. 2025, 1-16. DOI: https://doi.org/10.31224/4771
[8] Johnson-Rokosu S F, Enobi A. Behavioral Analytics and Forensic Accounting: Understanding the Human Element in Fraud. Journal of Accounting and Financial Management, 2025, 11(5): 93-117. DOI: 10.56201/jafm.vol.11.no5.2025.pg93.117.
[9] Waikhom L, Patgiri R. A Survey of Graph Neural Networks in Various Learning Paradigms: Methods, Applications, and Challenges. Artificial Intelligence Review, 2023, 56(7): 6295-6364.
[10] Kim H, Lee B S, Shin W Y, Lim S. Graph Anomaly Detection with Graph Neural Networks: Current Status and Challenges. IEEE Access, 2022(10): 111820-111829.
[11] Bui K H N, Cho J, Yi H. Spatial-Temporal Graph Neural Network for Traffic Forecasting: An Overview and Open Research Issues. Applied Intelligence, 2022, 52(3): 2763-2774.
[12] Bollu S S. Anomaly Detection of User Behavioural Events in E-Commerce Electronics Stores Using SVMs. Bachelor Thesis, Blekinge Institute of Technology, Sweden. 2024.
[13] Georgescu M I, Barbalau A, Ionescu R T, et al. Anomaly Detection in Video via Self-Supervised and Multi-Task Learning. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, 2021, 12742-12752. DOI: 10.1109/CVPR46437.2021.01255.
[14] Mutemi A, Bacao F. E-Commerce Fraud Detection Based on Machine Learning Techniques: Systematic Literature Review. Big Data Mining and Analytics, 2024, 7(2): 419-444.
[15] Tiwari S. Advancing Client Risk Scoring: From Rule-Based Systems to Machine Learning Approaches. Journal of Computer Science and Technology Studies, 2025, 7(8): 01-07.
[16] Dritsas E, Trigka M. Machine Learning in E-Commerce: Trends, Applications, and Future Challenges. IEEE Access, 2025(13): 99048-99067. DOI: 10.1109/ACCESS.2025.3572865.
[17] Chalapathy R, Chawla S. Deep Learning for Anomaly Detection: A Survey. arXiv preprint, 2019. DOI: https://doi.org/10.48550/arXiv.1901.03407.
[18] Xu H, Wang Y, Jian S, et al. Calibrated One-Class Classification for Unsupervised Time Series Anomaly Detection. IEEE Transactions on Knowledge and Data Engineering, 2024, 36(11): 5723-5736.
[19] Shao Z, Wang X, Ji E, et al. GNN-EADD: Graph Neural Network-Based E-Commerce Anomaly Detection via Dual-Stage Learning. IEEE Access, 2025(13): 8963-8976. DOI: 10.1109/ACCESS.2025.3526239.
[20] Erciyes K. Graph-Theoretical Analysis of Biological Networks: A Survey. Computation, 2023, 11(10): 188.
[21] Immaneni J. Strengthening Fraud Detection with Swarm Intelligence and Graph Analytics. International Journal of Digital Innovation, 2022, 3(1): 1-21.
[22] Bozorgi E, Alqaiidi S K, Shams A, et al. A Survey on Recent Random Walk-Based Methods for Embedding Knowledge Graphs. arXiv preprint, 2024. DOI: https://doi.org/10.48550/arXiv.2406.07402.
[23] Rossi R A, Jin D, Kim S, et al. On Proximity and Structural Role-Based Embeddings in Networks: Misconceptions, Techniques, and Applications. ACM Transactions on Knowledge Discovery from Data (TKDD), 2020, 14(5): 1-37.
[24] Kipf T. Deep Learning with Graph-Structured Representations. PhD Thesis, Universiteit van Amsterdam, Netherlands. 2020.
[25] Bhatti U A, Tang H, Wu G, et al. Deep Learning with Graph Convolutional Networks: An Overview and Latest

Applications in Computational Intelligence. International Journal of Intelligent Systems, 2023(1): 8342104.
[26] Cakmak E, Schlegel U, Jäckle D, et al. Multiscale Snapshots: Visual Analysis of Temporal Summaries in Dynamic Graphs. IEEE Transactions on Visualization and Computer Graphics, 2020, 27(2): 517-527.
[27] Ghadami A, Epureanu B I. Data-Driven Prediction in Dynamical Systems: Recent Developments. Philosophical Transactions of the Royal Society A, 2022, 380(2229): 20210213.
[28] Barros C D, Mendonça M R, Vieira A B, et al. A Survey on Embedding Dynamic Graphs. ACM Computing Surveys (CSUR), 2021, 55(1): 1-37.
[29] Cao J, Zheng W, Ge Y, et al. DriftShield: Autonomous Fraud Detection via Actor-Critic Reinforcement Learning with Dynamic Feature Reweighting. IEEE Open Journal of the Computer Society, 2025(6): 1166-1177. DOI: 10.1109/OJCS.2025.3587001.
[30] Wang J, Liu J, Zheng W, et al. Temporal Heterogeneous Graph Contrastive Learning for Fraud Detection in Credit Card Transactions. IEEE Access, 2025(13): 145754-145771. DOI: 10.1109/ACCESS.2025.3599787.
[31] Samant R M, Bachute M R, Gite S, et al. Framework for Deep Learning-Based Language Models Using Multi-Task Learning in Natural Language Understanding: A Systematic Literature Review and Future Directions. IEEE Access, 2022(10): 17078-17097.
[32] Ji E, Wang Y, Xing S, et al. Hierarchical Reinforcement Learning for Energy-Efficient API Traffic Optimization in Large-Scale Advertising Systems. IEEE Access, 2025(13): 142493-142516. DOI: 10.1109/ACCESS.2025.3598712.
[33] Lindemann B, Maschler B, Sahlab N, et al. A Survey on Anomaly Detection for Technical Systems Using LSTM Networks. Computers in Industry, 2021(131): 103498.
[34] Zheng W, Liu W. Symmetry-Aware Transformers for Asymmetric Causal Discovery in Financial Time Series. Symmetry, 2025.
[35] Jin J, Xing S, Ji E, et al. XGate: Explainable Reinforcement Learning for Transparent and Trustworthy API Traffic Management in IoT Sensor Networks. Sensors, 2025, 25(7): 2183.
[36] Chattopadhyay S, Basu T, Das A K, et al. Towards Effective Discovery of Natural Communities in Complex Networks and Implications in E-Commerce. Electronic Commerce Research, 2021, 21(4): 917-954.