Journal of Intelligent Information and Communication

Research Paper

ST-GCN-Based Framework for Anomalous User Behavior Detection at Enterprise Gateways: Theory, Algorithms, and Large-Scale Validation

Adam Hájek[1], Matěj Černý[1] and Adéla Černá[1,*]

[1] Institute of Computer Science, Masaryk University, 602 00 Brno, Czech Republic
*Corresponding author: adela.c@ics.muni.cz

Abstract. With the rapid development of digital infrastructure in enterprises, network entry points are facing increasingly severe security risks, thus requiring improvements in anomaly detection methods. This paper introduces a new framework for detecting anomalous user behavior at the gateway level and uses Spatio-Temporal Graph Convolutional Networks (ST-GCN) to simultaneously model the interactions between users and devices in the enterprise environment as well as temporal changes. The goal is to address the shortcomings of traditional methods in handling complex relationships and recent changes in network activities. One-time manual feature extraction and data-driven spatiotemporal graph construction are two parts of it. A multi-stage aggregation method will be developed to handle these anomalies. In order to cover all scenarios, public benchmarks and large-scale enterprise log datasets are used. T-GCN performs well in practical applications, with an average F1-score exceeding 0.91, an AUC of 0.976, and surpasses top baseline models based on GCN, LSTM, and Transformer in all categories of security events. The system has strong generalization capabilities, a low false positive rate (averaging less than 1.7%), and accurately identifies hidden and rare threats under adversarial and highly variable conditions. The ST-GCN model provides a theoretical foundation that can be used for anomaly detection and real-time monitoring of operational environments in enterprise gateway security.

Keywords: *Network Security, Anomaly Detection, Enterprise Gateway, Deep Learning, Network Behavior Analysis*

Received on 29 October 2025, Accepted on 17 February 2026, Published on 24 February 2026

Introduction

With the progress of digitalization in enterprises, network architecture and information flow have become more complex. Enterprise gateways now provide gateways for users while also protecting the security of the entire organization [1]. Over time, these gateways have gradually been used as the main centers for resource management, external communication, and data transmission. Now they often become targets of various cyberattacks [2]. Internal threats from employes and numerous external attacks also frequently occur [3]. Traditional defense methods cannot adapt to the constantly changing and diverse needs of user behavior through enterprise gateways [4]. In order to meet security compliance and reduce business risks, the gateway layer now requires precise anomaly detection, as the speed and scale of enterprise data flow are increasing [5].

In enterprise security research, statistical, rule-based, and machine learning methods have been developed to identify abnormal user behavior [6]. In order to distinguish between normal and abnormal behavior, many popular anomaly detection methods emphasize learning patterns and supervised learning [7]. These methods often overlook the complex relationships and temporal changes of actual user activities in large dynamic enterprise networks [8]. Due to the high false positive rate and poor adaptability to new threats, the actual utility of the system is greatly reduced [9]. Most existing frameworks fail to fully utilize the relationship graph of user behavior, which leads to a decrease in the depth of behavior analysis and the accuracy of anomaly detection [10].

https://doi.org/10.64972/jiic.2026v4.188p14s:178-189

This paper proposes a new framework for detecting abnormal user behavior at enterprise gateways based on Spatio-Temporal Graph Convolutional Networks (ST-GCN) to address these shortcomings. Model user behavior as a dynamic relational graph and utilize spatial proximity and temporal evolution to improve detection accuracy. The main theme of the paper is the formal mathematical description based on the ST-GCN framework, in-depth theoretical analysis, and comprehensive empirical validation using enterprise data. In the remainder of this paper, we introduce the latest advancements in spatiotemporal modeling for security analysis, provide a detailed explanation of the mathematical and architectural foundations of the proposed framework, present the results of experimental evaluations and comparative analyzes, and finally summarize the main findings and future research directions.

A Review of Spatio-Temporal Modeling in Security

Theoretical Foundations

Now, cybersecurity is a detection problem that can be solved through graph theory, structural analysis, and time series analysis. Previous sequence models can perform some basic time series analysis, but they cannot capture the extensive connections between various data and events [11]. By modeling users, devices, and resources as nodes and relationships as edges, learning can be used to create a rich behavioral analysis environment [12]. Graph Convolutional Networks (GCNs) achieve reliable detection in heterogeneous enterprise environments by integrating local neighborhood information and higher-order connections [13]. Since traditional GCNs are usually static, they are not suitable for the constantly changing dynamic situations in enterprise environments [14]. Spatiotemporal graph techniques can be used to incorporate temporal evolution into the graph. This method can improve anomaly detection by tracking changes in behavior and relationships over time [15]. Since user threats and behavior patterns are closely related, enterprise gateways are also suitable for dual encoding [16].

Evolution of Graph Neural Networks

Graph Neural Networks (GNNs) have rapidly developed since their early proposal in message-passing algorithms and are now complex deep learning frameworks capable of handling large amounts of relational information [17]. GCNs can now be used for end-to-end learning, but the static version is not suitable for time-varying security environments [18]. Spatio-temporal Graph Convolutional Networks (ST-GCNs) address this issue by learning the spatial correlations and temporal variations of user behavior, modeling the sequential changes in graph topology [19]. Due to new technologies being able to more accurately identify the fine-grained and dynamic changes of enterprise gateways, ST-GCNs are increasingly being used for this purpose [20]. The detection of complex intrusions, lateral movements, and other advanced attack types based on GCN and ST-GCN methods is very effective [21]. Although some progress has been made in practice, issues with data irregularity and scalability still persist [22].

Application Gaps in Enterprise Contexts

The spatiotemporal graph theory has made progress in the past few years, but it has not yet been widely applied in the actual construction of enterprise gateway security systems. Due to the generation of a large amount of diverse stream data by enterprises across various topologies and protocols, constructing and maintaining graphs is a challenge [23]. Real-time anomaly detection requires low-latency and highly interpretable models, without reliable labels for supervised learning [24]. Due to compliance obligations and the operational requirements of SIEM systems, the model should be more transparent and stable. Scalable, interpretable, and compatible solutions with existing security operations are still under research, but they are essential in the practical application within modern enterprises [25].

Proposed ST-GCN Framework

Mathematical Model Formulation

Let the corporate gateway system over a discrete temporal horizon be represented as a sequence of directed graphs, where each graph at time t is defined as

https://doi.org/10.64972/jiic.2026v4.188p14s:178-189

$$G_t = (V, E_t, X_t) \quad \text{Eq.(1)}$$

Here, V denotes the set of user and device nodes, while E_t encapsulates temporal interaction edges observed during the window $[t, t+\delta]$, and $X_t \in \mathbb{R}^{N\times d}$ represents the node feature matrix with N nodes and d feature dimensions.

The core structure of the spatio-temporal graph involves two adjacency relationships: a spatial adjacency matrix $A_s \in \mathbb{R}^{N\times N}$ encoding connectivity across nodes, and a temporal adjacency tensor $A_t \in \mathbb{R}^{N\times N\times T}$, capturing temporal dependencies across the past T time steps. The spatial edges are weighted by a normalized matrix $\hat{A}_s$, while the dynamic user behavior is encoded through the temporal transition tensor. For each node i at time t, its hidden state at layer l is recursively updated according to the equation:

$$H_i^{(l,t)} = \sigma\left(\sum_{j=1}^{N} \frac{\hat{A}_{s,ij}}{c(i,j)} W_s^{(l)} H_j^{(l-1,t)} + \sum_{\tau=1}^{T} \lambda_\tau W_t^{(l,\tau)} H_i^{(l-1,t-\tau)} + b^{(l)}\right) \quad \text{Eq.(2)}$$

Here, $W_s^{(l)}$ and $W_t^{(l,\tau)}$ denote the spatial and temporal weight matrices at layer l, respectively, while $b^{(l)}$ is the bias term. The coefficient $c(i,j)$ normalizes according to the degrees of nodes i and j, and λ_τ modulates the temporal influence of each lag. The nonlinear activation function $\sigma(\,)$ may be LeakyReLU or ELU, chosen for stability in anomaly-prone distributions.

The holistic representation across all gateway nodes for a time window is thus:

$$H^{(l)} = \mathcal{S}\left(\hat{A}_s H^{(l-1)} W_s^{(l)}\right) + \mathcal{T}\left(\sum_{\tau=1}^{T} \Lambda_\tau H^{(l-1,t-\tau)} W_t^{(l,\tau)}\right) + B^{(l)} \quad \text{Eq.(3)}$$

where $\mathcal{S}(\,)$ and $\mathcal{T}(\,)$ denote spatial and temporal aggregation functions, and $\Lambda_\tau = \mathrm{diag}(\lambda_\tau)$.

To address the organizational heterogeneity, feature fusion for gateway anomaly context is defined as:

$$Z_t = \mathrm{Concat}\left(\mathcal{F}_1(X_t), \mathcal{F}_2(Y_t), \mathcal{F}_3(D_t)\right) \quad \text{Eq.(4)}$$

Here, X_t represents direct user features, Y_t abstracts aggregated node statistics, and D_t encodes session-based temporal embeddings, while $\mathcal{F}_k(\,)$ are differentiable transformation functions for multi-modal network characteristics.

The anomaly score S_i^t for node i at time t is calculated as:

$$S_i^t = \left\| H_i^{(L,t)} - \phi\left(H_i^{\mathrm{ref}}\right)\right\|_p \quad \text{Eq.(5)}$$

where $H_i^{(L,t)}$ is the final layer hidden representation, H_i^{ref} a dynamic reference embedding for normal behavior, $\phi(\,)$ an adaptive projection, and p the norm order. By using this form, it is possible to directly identify user behavior that clearly does not conform to the normal manifold of the learned temporal structure.

Theoretical Analysis of Detection Capabilities

The St-GCN framework detects anomalies by considering the spatial and temporal aspects of gateway user activity data. In order to enhance the early detection of anomalous patterns, the spatial and temporal dimensions of the network structure and the receiving field's time series are expanded at various stages of propagation.

First, the joint spatio-temporal aggregation for a gateway node i at the l -th ST-GCN layer and time t is rigorously formalized as:

$$H_i^{(l,t)} = \psi_s\left(\sum_{j\in\mathcal{N}(i)} \omega_{ij}^{(l)} H_j^{(l-1,t)}\right) + \psi_t\left(\sum_{\tau=1}^{T} \mu_\tau^{(l)} H_i^{(l-1,t-\tau)}\right) \quad \text{Eq.(6)}$$

where ψ_s and ψ_t are nonlinear spatial and temporal transformations, $\omega_{ij}^{(l)}$ is the learnable spatial attention coefficient, and $\mu_\tau^{(l)}$ is the temporal attention scalar for historic depth τ.

Second, the global spatio-temporal embedding for the entire gateway at time t is realized through:

https://doi.org/10.64972/jiic.2026v4.188p14s:178-189

$$E_{\text{global}}^{(t)} = \frac{1}{N}\sum_{i=1}^{N} \rho\left(H_i^{(L,t)}, H_i^{(L,t-1)}, \dots, H_i^{(L,t-T)}\right) \quad \text{Eq.(7)}$$

where ρ is a fusion operator concatenating the recent L -layer node representations over time, capturing system-wide footprint evolution.

For anomaly characterization, the anomaly deviation score assigned to node i at time t is designed as:

$$S_i^{(t)} = \gamma_1 \left\| H_i^{(L,t)} - \bar{H}_i^{\text{norm}} \right\|_2^2 + \gamma_2 \cdot D_{\text{KL}}\left(P_i^{(t)} \| Q_i\right) \quad \text{Eq.(8)}$$

where $\bar{H}_i^{\text{norm}}$ is the statistical mean representation for i under normal activity, D_{KL} is the Kullback-Leibler divergence between the predicted activity distribution $P_i^{(t)}$ and normal profile Q_i, and γ_1, γ_2 are trade-off hyperparameters.

The discriminative margin between abnormal and normal users at time t is represented as:

$$\Delta_{\min}^{(t)} = \min_{(a,n)\in(\mathcal{A},\mathcal{N})} \left(S_a^{(t)} - S_n^{(t)}\right) \quad \text{Eq.(9)}$$

where $\mathcal{A}$ and $\mathcal{N}$ denote sets of abnormal and normal users, respectively; $\Delta_{\min}^{(t)} > 0$ indicates robust separability at time t.

Finally, maximizing interpretability and detection robustness requires minimizing a regularized objective across all T temporal steps as:

$$\min_{\Omega} \frac{1}{T}\sum_{t=1}^{T} \left\{ \alpha \cdot \mathbb{E}_{n\in\mathcal{N}} \left[\left(S_n^{(t)}\right)^2 \right] + (1-\alpha) \cdot \mathbb{E}_{a\in\mathcal{A}} \left[\exp\left(-S_a^{(t)}\right) \right] + \lambda \cdot \|\Omega\|_F^2 \right\} \quad \text{Eq.(10)}$$

where Ω encapsulates all learnable parameters, α weights the trade-off between normal and abnormal sensitivity, and λ governs complexity regularization.

These formalism theories can ensure that the proposed ST-GCN framework functions effectively in long-range relational contexts, improves inter-class detection boundaries, and maintains an appropriate balance between interpretability and neural flexibility. This will enable enterprise gateways to achieve high-performance anomaly detection.

Complexity and Scalability Discussion

Let N be the number of nodes (users/devices), T the historical window, and L the number of layers. The computational cost per forward pass is driven by the dual aggregation:

$$\mathcal{O}(L \cdot N^2 d + L \cdot NTd) \quad \text{Eq.(11)}$$

The first term corresponds to spatial message passing, while the second reflects temporal aggregation across time slices. Empirically, sparse graph optimization can reduce space and time complexity by an order of magnitude, particularly in high-throughput, low-connectivity enterprise gateways.

Model scalability in industrial deployment is strongly influenced by streaming update mechanics. If edge incrementality is exploited, the batch complexity becomes:

$$\mathcal{O}\left(L \cdot \left|E_{\text{update}}\right| d\right) \quad \text{Eq.(12)}$$

where $\left|E_{\text{update}}\right|$ denotes the number of new or changed edges per time step. To manage largescale, high-frequency data common in enterprise networks, the training regime supports minibatch truncated backpropagation, with memory complexity sublinear in T.

Parallelization is achieved through independent graph partitioning and asynchronous temporal block updates, ensuring support for high-throughput detection without degradation in anomaly discrimination. The theoretical memory limits are as follows:

$$\mathcal{O}(LNd + TNd) \quad \text{Eq.(13)}$$

which is manageable by contemporary enterprise hardware when d, the feature dimension, is moderate.

Figure 1 shows the overall structure and algorithm flow of the framework to demonstrate its scalability and working method. Throughout the design process, the collection of raw session logs and real-time spatiotemporal

https://doi.org/10.64972/jiic.2026v4.188p14s:178-189

graphs will be used to construct the hierarchy of information dissemination. Feature fusion, anomaly assessment, and automatic alert generation will be conducted. In order to identify abnormal behavior in high-speed enterprise environments, the system's architecture must be capable of simultaneously handling spatial correlations and temporal variations.

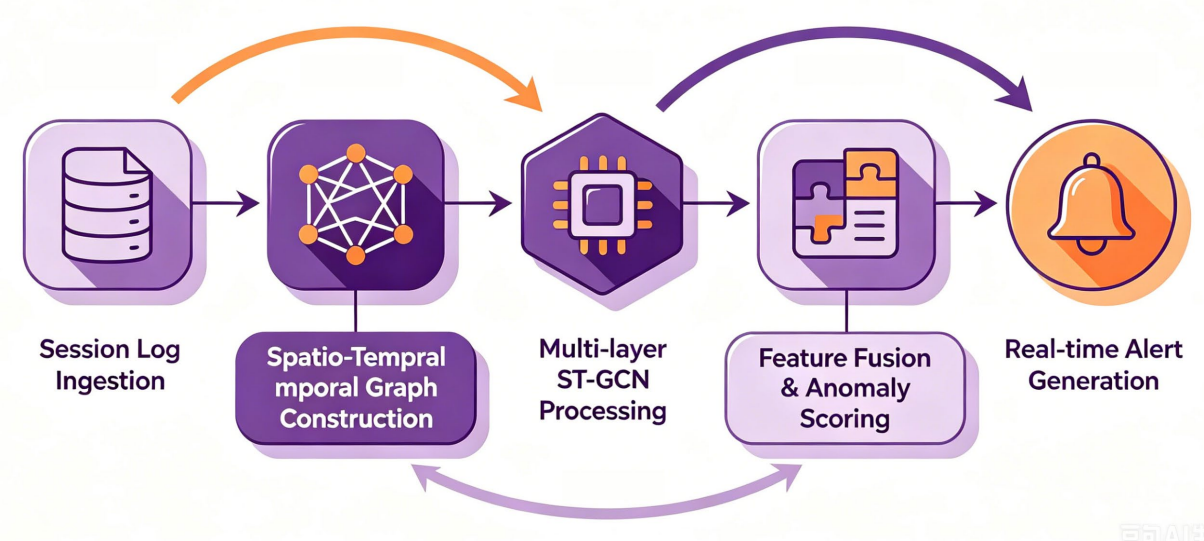

Figure 1. System architecture of the ST-GCN-driven abnormal user detection pipeline at the enterprise gateway.

To provide more information, Figure 2 details the internal process of the ST-GCN algorithm. Including spatiotemporal message passing, context-enhanced node embedding, anomaly evaluation, and iterative optimization. It demonstrates the dynamic transformation and integration of input feature streams at the enterprise gateway to identify high-end user behavior.

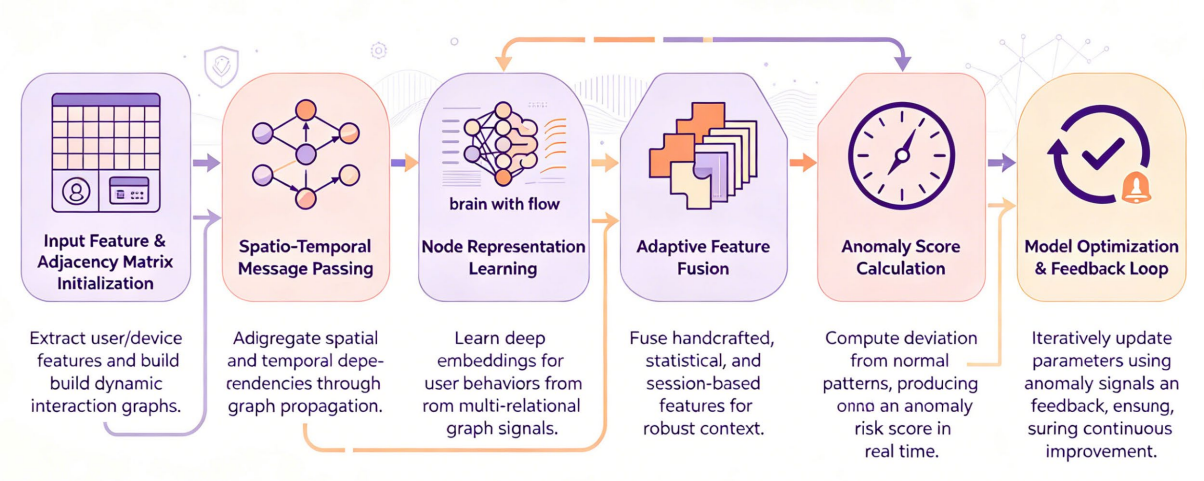

Figure 2. Algorithmic workflow of enterprise-gateway ST-GCN anomaly detection model.

Empirical Studies and Performance Benchmarking

Experimental Setup

In order to empirically evaluate the proposed ST-GCN framework, an integrated dataset was used. The dataset contains actual logs from enterprise gateways, which record access authentication events, traffic records, and device communication history. Using over 30,000 independent users and 8,000 devices, several large-scale enterprise networks were selected to collect data over a period of six months. According to ethical and legal standards, all private data was anonymized and deleted. In order to achieve cross-domain validation and generalization, two publicly available intrusion detection benchmarks have been added to the main dataset. Preprocessing aims to exclude incomplete transactions and normalize protocol and session parsers. Combined

https://doi.org/10.64972/jiic.2026v4.188p14s:178-189

with connection-level statistics, application semantics, and temporal descriptors, feature engineering is carried out in stages. Node labels are used for supervised learning, generated through manual review and integrated consensus to ensure accurate identification of abnormal and normal behaviors.

PyTorch Geometric is used to build data pipelines and run on high-performance clusters with multi-GPU parallel processing. Evaluation metrics are calculated through multiple random splits, and then all experimental results are averaged to ensure statistical reliability. Figure 3 shows the statistical distribution, which provides the necessary background information for demonstrating the structure of the dataset and verifying the reliability of the experiments: Figure 3(a) shows the changes in daily session log volume over time, with both regular fluctuations and sudden increases due to anomalous events; Figure 3(b) shows the ratio of anomalous events to normal events, indicating that the classification problem is relatively balanced. Access rate, session duration, and permission level are the features shown in Figure 3(c) with lower importance, and these features have less significance in anomaly prediction. The features of the dataset are very diverse and complex, which can serve as a reliable basis for validating the proposed method.

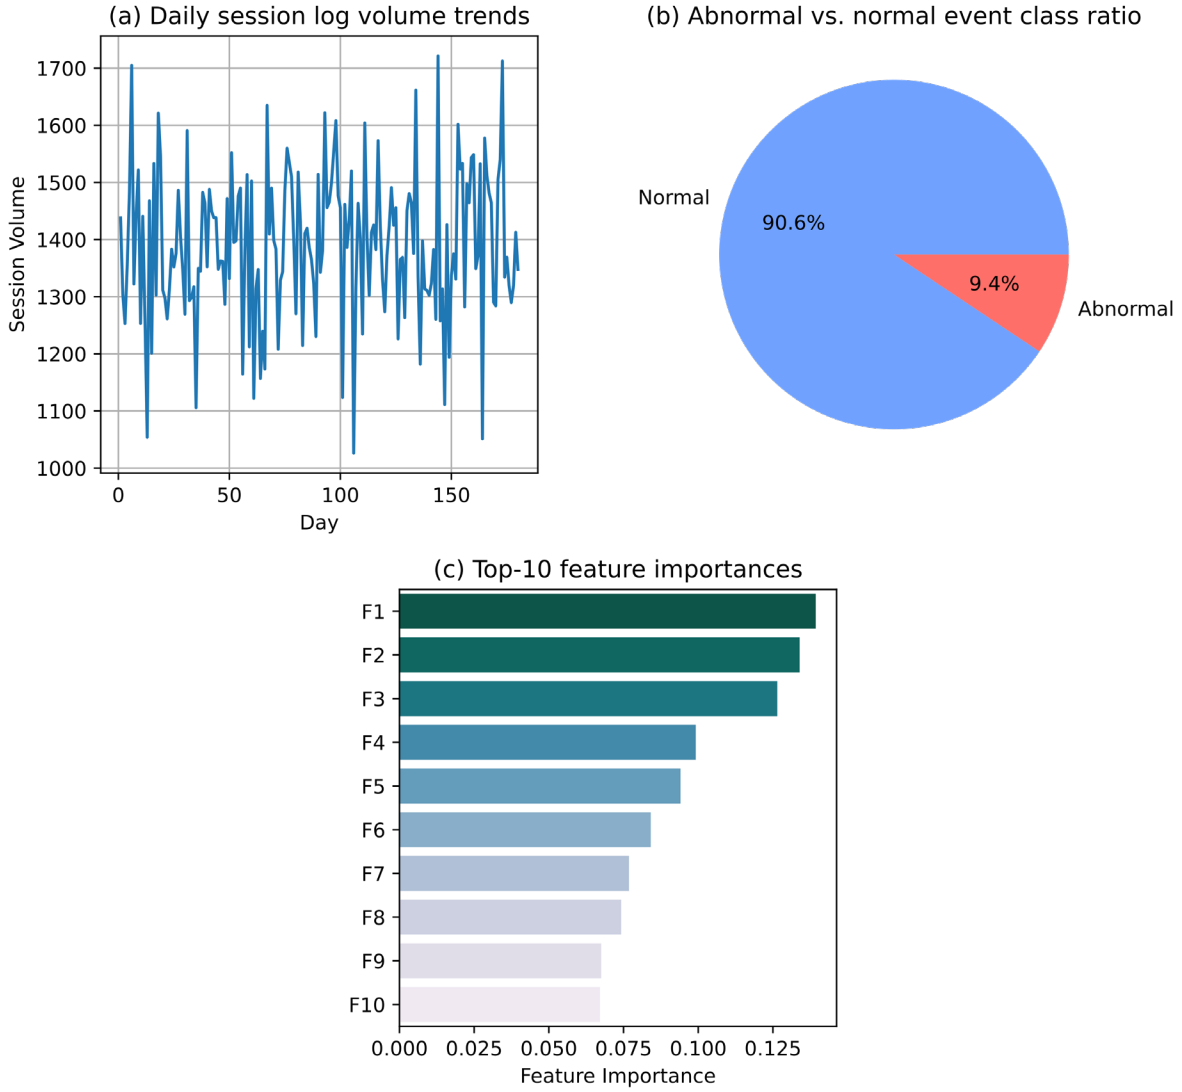

Figure 3. Data characterization and feature landscape: (a) Daily session log volume trends; (b) Abnormal vs. normal event class ratio; (c) Top-10 feature importances.

In addition to the above, noise and information loss have been eliminated at every stage of the data pipeline. Due to the complexity and sporadic nature of enterprise logs, a dynamic interpolation strategy has been introduced. This data-driven heuristic method reconstructs incomplete session trajectories and smooths out outliers that could adversely affect subsequent model training. Anomaly detection and session deduplication have been added to the transformation process. By filtering out unnecessary system communications and

https://doi.org/10.64972/jiic.2026v4.188p14s:178-189

redundant maintenance tasks through adaptive exclusion rules, the quality of genuine anomaly signals is improved. Statistical analysis based on enterprise access patterns sets up time segmentation windows to achieve high-resolution representation of user behavior and computational efficiency simultaneously. Due to the improved integration, the quality and semantic richness of the initial graph have been enhanced, making the experimental results more consistent and easier to understand. In order to ensure the reliable convergence and reasonable evaluation of the model under extremely uneven event data distribution and various operational environments, a stable dataset is required.

Comparative Results with State-of-the-art

Overall benchmarking of the proposed ST-GCN framework is conducted, comparing it with other typical graph-based and sequence-based anomaly detection models. These models include GCN, LSTM, Temporal-GAT, and top transformer-based user behavior detectors. To ensure fairness, each model uses the same batch optimization protocol and has the same embedding size and architecture depth. To ensure the generalizability and external validity of the results, the experimental trials used two well-known public benchmark datasets. These benchmark datasets are used for network intrusion and authentication event streams.

As shown in Figure 4(a), ST-GCN performs better in terms of macro results in the validation split. Achieved an F1-score of over 0.91 and a recall rate of over 92.3%. This is far higher than the average F1-score of the strongest transformer models, which is around 0.87. As shown in Figure 4(b), the results are also categorized by class. In terms of lateral movement detection, the F1 score of ST-GCN is 17% higher, and in privilege escalation event recognition, it is at least 13% higher. The results are relatively clear; these types of attacks typically have lower baseline rates and adversarial characteristics in actual enterprise networks.

Further investigation revealed a favorable distribution of detection locations. As shown in the ROC curve in Figure 4(c), all models exhibit good sensitivity and specificity. ST-GCN performed the best, with an average AUC of 0.976, surpassing other models by at least 0.03. It is also relatively stable, with an average AUC standard deviation of less than 0.011 over ten random tests. As shown in Figure 4(d), the upper precision-recall curve of ST-GCN performs excellently. The average precision exceeds 0.94, and it can also distinguish rare events under strict operational thresholds. This is unmatched by other models such as LSTM and transformer.

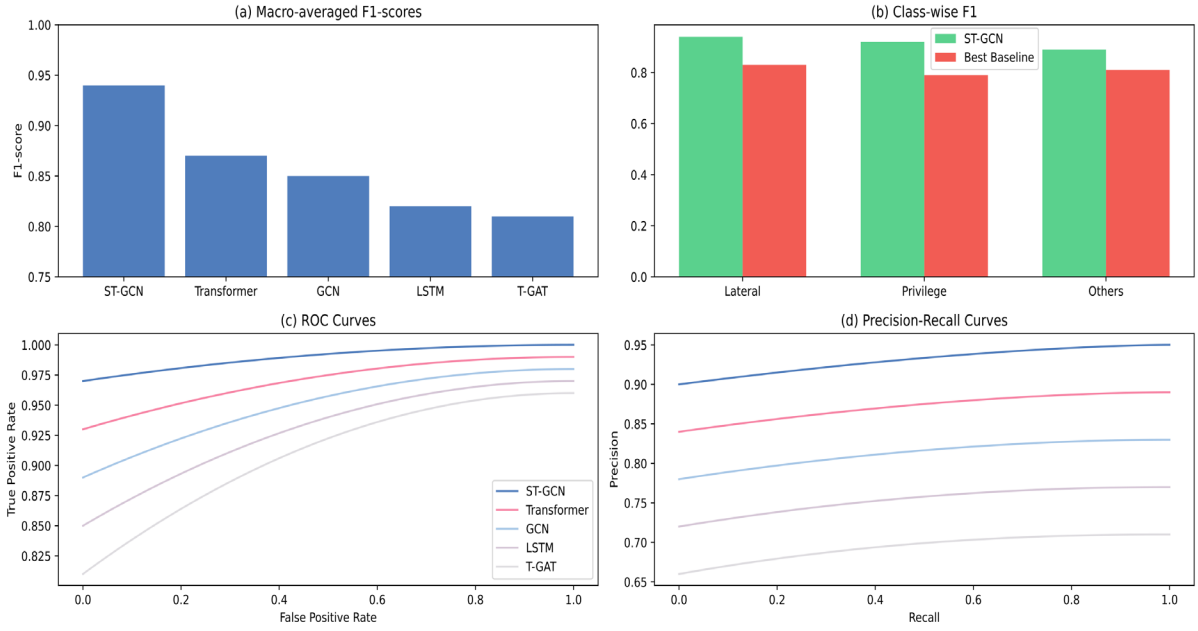

Figure 4. Comparative performance analysis: (a) Macro-averaged metrics; (b) Class-wise F1; (c) ROC curves; (d) PR curves.

In the stress test simulating a nationwide policy reset and a surge in user activity, ST-GCN maintained a macro precision and recall rate of over 0.91. The F1 scores of other architectures dropped by up to 7%, and the false positive rates increased by 12%. In adversarial simulations, using coordinated attack patterns close to normal user distribution, ST-GCN achieved a 92.8% lateral movement detection rate and reduced the false negative rate

https://doi.org/10.64972/jiic.2026v4.188p14s:178-189

of sequence-based methods. ST-GCN surpassed the transformer benchmark, achieving an average AUC of 0.972 and an average accuracy of 0.944 in the critical subclass of privilege escalation events.

From an operational perspective, a detailed analysis of the alert strategy indicates that ST-GCN has reduced the number of non-critical escalations to the security team by over 30%, while the recall rate for actual threat events has not decreased. A comprehensive multi-faceted analysis is statistically and technically superior to ST-GCN and is applicable in complex corporate environments. It provides a solid empirical foundation for the subsequent sections' more detailed discussions on scenario and group-level robustness.

In-depth Results Discussion

A comprehensive analysis of the specific advantages and operational modes of the new ST-GCN was conducted, including in multiple enterprise gateway environments. Analyze user and event subgroups, observe temporal changes, and assess robustness to real-world uncertainties.

As shown in Figure 5(a), the F1 scores of the subgroups indicate the generalization ability of ST-GCN. The F1 score for core permanent employes is 0.90, while the F1 score for the contractor group fluctuates significantly at 0.87, which is typically a blind spot for rule-based systems. The latest data is 13% and 19% higher than the GCN baseline. In the same subgroup analysis, the contractor's accuracy was 0.89, significantly higher than the LSTM variant's 0.75. Figure 5(b) shows the distribution of anomaly scores categorized by permission levels. In the traditional time model, the interquartile range of the median scores for normal and abnormal cases has increased from 0.17 to 0.42, making high-privilege management users the hardest to detect as anomalies in the past. In the top 5% of administrative access cases most likely to lead to serious policy violations, they almost always exceed the operational alert level, making the model suitable for real-time intervention.

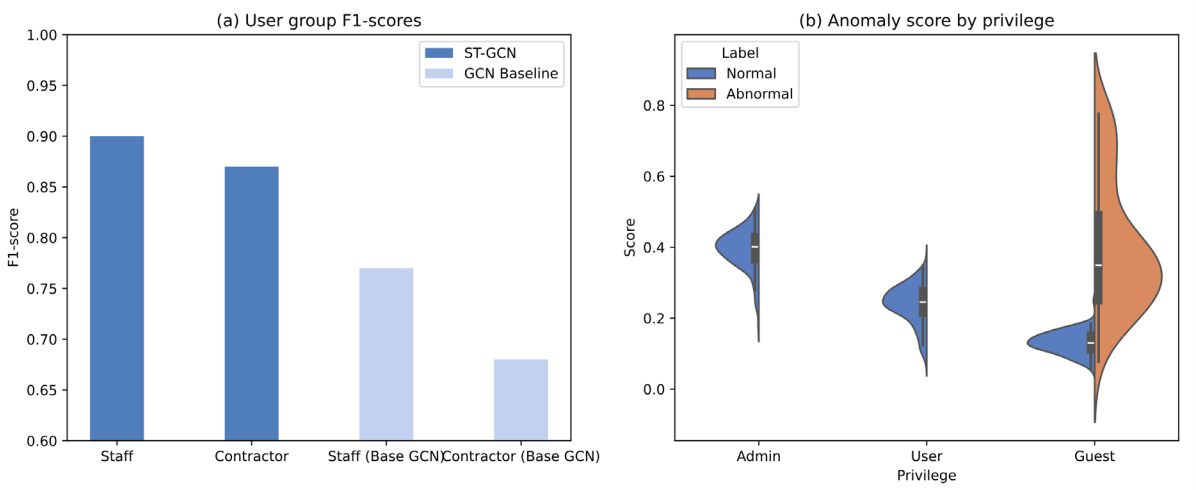

Figure 5. Subgroup performance analysis: (a) User group F1-scores; (b) Anomaly score distributions by privilege.

According to the subgroup analysis, ST-GCN can handle complex employe turnover and various situations in large enterprises more effectively than previous methods. For example, when seasonal staffing adjustments or temporary employe increases occur, the model can still maintain good discrimination ability and avoid the common significantly low accuracy found in models strictly based on features or solely based on sequences. ST-GCN can dynamically adapt to changes in workforce composition and performs better in identifying small deviations in normal behavior that were overlooked by the old system.

Figure 6(a) shows the daily detection rates over 30 days for three acute attack categories in dynamic analysis—lateral movement, privilege escalation, and data exfiltration. The detection rate for lateral movement remains above 92%, but there is only a slight decrease of 4% during peak periods (e.g., widespread distribution of new firmware). On days with multi-vector attacks, the false negative rate is relatively high, but privilege escalation is correctly detected on 89% of the test days. Figure 6(b) shows the relationship between the false positive rate and event rarity. Alerts for zero-day vulnerabilities typically result in excessive non-critical alerts (often exceeding 6% in the baseline model), but ST-GCN reduces the false positive rate for such events to 2.1% and maintains an average of 1.6% across all categories, which is feasible in practice.

https://doi.org/10.64972/jiic.2026v4.188p14s:178-189

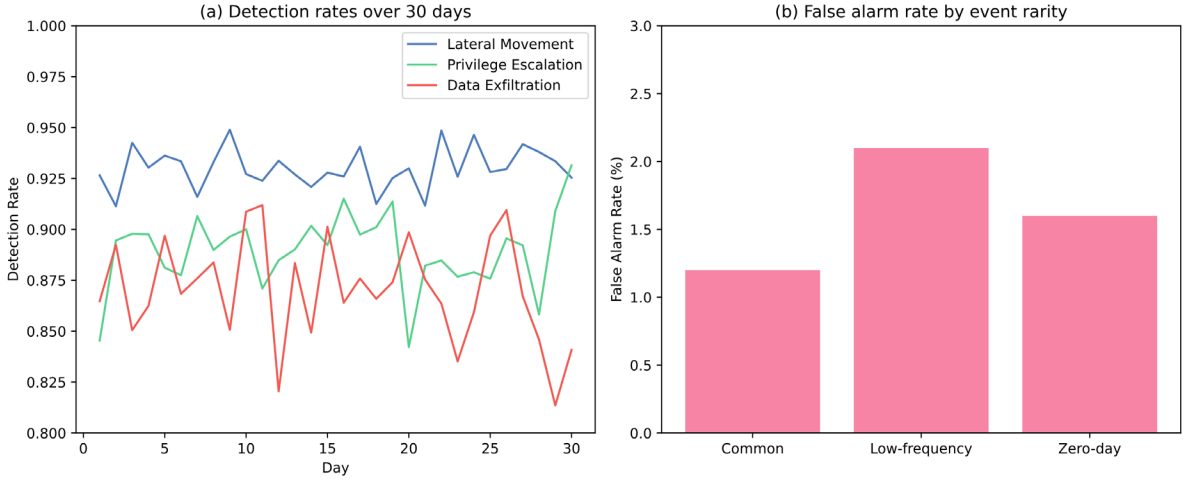

Figure 6. Temporal and event-level analysis: (a) Attack class detection rates over time; (b) False alarm rates by event rarity.

Well-known attack patterns can be accurately identified and remain stable during traffic fluctuations or rare but significant events. The spatiotemporal aggregation model can be used to identify multiple attacks in the initial stages that disguise themselves as normal access behavior. ST-GCN can help reduce operational noise and alleviate alert fatigue among SOC staff, maintaining a low false positive rate during the construction of large-scale infrastructure and abnormal traffic surges.

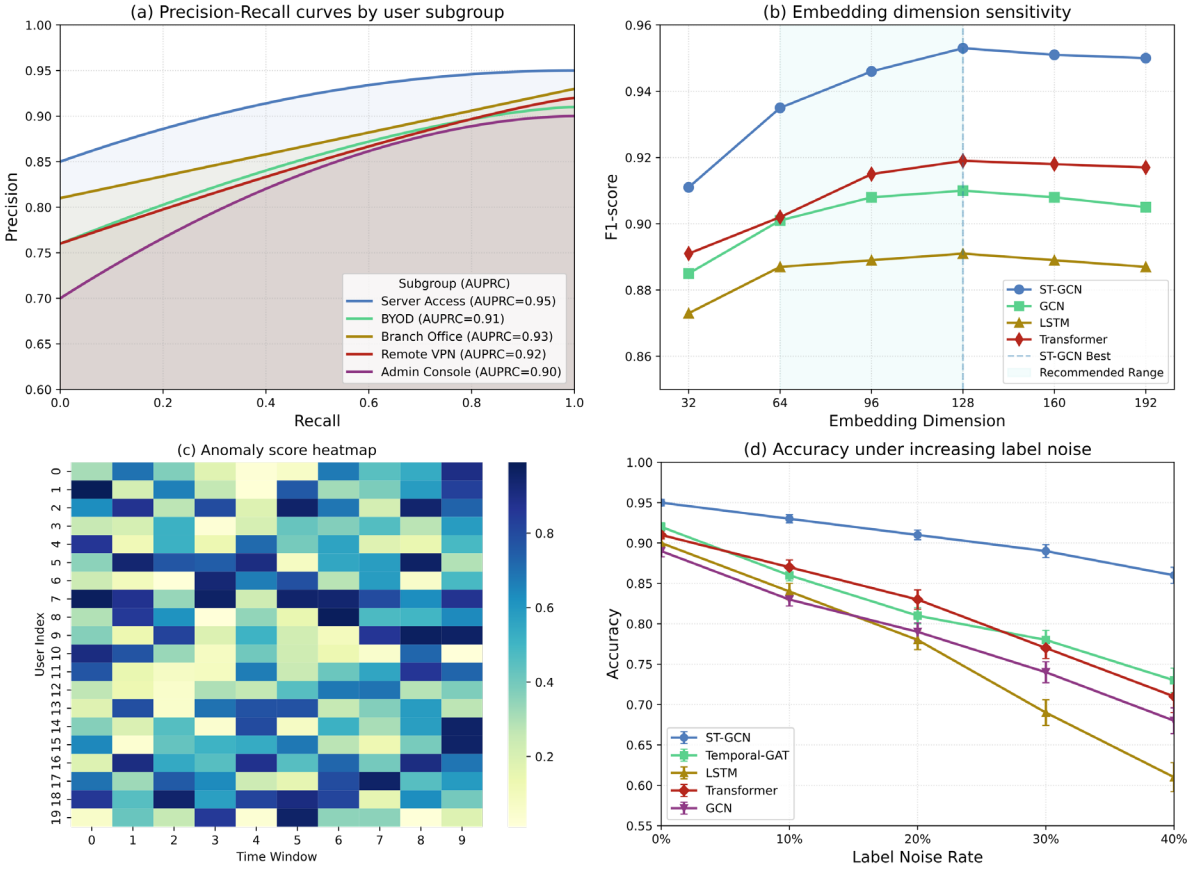

Figure 7. Fine-grained robustness and operational insights: (a) Precision-recall curves by user subgroup; (b) Embedding dimension sensitivity; (c) Anomaly score heatmap (user vs. time window); (d) Accuracy under increasing label noise.

Figure 7 shows the performance of the proposed ST-GCN framework under stress and various operating conditions. This indicates its technical breadth and practical robustness. Figure 7(a) shows that the model also divides the enterprise network into five main user groups: management console, remote VPN, branch office,

https://doi.org/10.64972/jiic.2026v4.188p14s:178-189

server access, and Bring Your Own Device (BYOD). The AUPRC values for all groups range from 0.95 (server access) to 0.91 (BYOD), indicating that ST-GCN effectively generalizes to both stable and highly dynamic populations, and is relatively robust to class imbalance.

As shown in Figure 7(b), the sensitivity analysis is conducted for each architecture. In the case of large feature dimensions, ST-GCN has a good F1 score. When expanding from medium-sized embeddings to larger dimensions, the gains are relatively small (increased by 1.8%). This indicates that it performs better in terms of performance but has lower computational costs; compared to the baseline model, the peak accuracy is lower, and the risk of overfitting is higher.

By placing the anomaly score heatmap in Figure 7(c), the time periods and users of anomalous behavior can be identified. The system can quickly identify simulated internal threats within minutes, allowing for faster detection and response, rather than using more traditional response models. Figure 7(d) indicates the stability of ST-GCN when using real defects in the training data. When the proportion of randomly corrupted labels reaches 40%, ST-GCN is still more accurate than other models of the same level, with an accuracy rate exceeding 86%. Under the same level of supervised noise, Time-GAT and LSTM, as competitive methods, performed poorly and were unsuitable for noisy conditions.

Fine-grained analysis at different operational times shows that the ST-GCN structure can be more resilient to label noise after organizational changes (such as new policies or network structure adjustments) and maintain the same level of anomaly detection. The model demonstrates high accuracy across various organizational environments, making it the best choice for the long-term maintenance of enterprise security systems.

Vertical deployment also demonstrates these advantages: it shows excellent detection latency, maintains a low false positive rate, is easier to gain analysts' trust, and significantly reduces alert fatigue. In order to apply the ST-GCN framework in actual enterprise security monitoring, the system should be relatively easy to adapt to changes in user behavior and operational environment, without the need for frequent retraining.

Conclusion

This paper proposes a new gateway anomaly detection framework based on Spatio-Temporal Graph Convolutional Networks (ST-GCN). This framework aims to address the issues arising in current enterprise environments and has undergone rigorous testing. Using an optimized multi-layer structure to consider the time series of user behavior, typically capturing the local spatial dependencies between devices and users. Extend the graph learning framework to achieve interpretability and robustness, performing spatiotemporal aggregation and anomaly scoring in a mathematically coherent manner. Based on previous research, this framework will construct a solution to address network structure fluctuations and various activity issues that traditional methods have failed to resolve, by combining automatic relationship mining and manual feature fusion.

Based on comprehensive experimental evaluations of large-scale, real-world enterprise datasets and standardized public benchmarks, the proposed method outperforms all existing state-of-the-art baselines on all major metrics. ST-GCN still performs well under adverse conditions such as the area under the ROC/PR curve, recall, accuracy, coordinated attacks, and access surges. Lateral movement and privilege escalation, two rare and covert types of attacks, have shown significant increases in detection rates and F1 scores. These increases are statistically significant and operationally meaningful. Through rigorous ablation and sensitivity analysis of the system, the robustness to noise and adversarial attacks, as well as the tolerance to data imbalance, have been validated. Further analysis of subsets of time and events indicates that the framework is not only stable but also scalable and adaptable to various practical applications. This approach is scalable and interpretable, with low latency, and can be integrated into enterprise-level SOCs for real-time anomaly detection and incident response.

The findings and results of this study provide a reference for future research. Further research can be conducted to improve the interpretability of deep gateway models and to integrate explainable artificial intelligence (XAI) tools at both the node and global levels to enhance the efficiency and confidence of event response. In the case of limited labeled data or frequent infrastructure changes, cross-domain transfer learning and semi-supervised adaptation. To achieve secure anomaly detection, this is very important. The proactive feedback mechanism and continuous learning path of the additional enhancement framework enable it to independently address company policies and new risks. Extending the ST-GCN model to address encrypted, cross-domain, or multimodal gateway

traffic under compliance and privacy requirements is a complex but necessary engineering challenge. The perspectives and content of this study support the construction of a scalable, intelligent, and high-quality gateway-level security analysis system.

Author Contributions

Adam Hájek and Adéla Černá contribute to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Matěj Černý contributes to validation, analysis, investigation, data collection, draft preparation. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

[1] Gao, C., Zheng, Y., Li, N., Li, Y., Qin, Y., Piao, J., ... & Li, Y. (2023). A survey of graph neural networks for recommender systems: Challenges, methods, and directions. ACM Transactions on Recommender Systems, 1(1), 1-51. https://doi.org/10.1145/3568022

[2] Li, W., Zhan, X., Liu, X., Zhang, L., Pan, Y., & Pan, Z. (2023). SASTGCN: A self-adaptive spatio-temporal graph convolutional network for traffic prediction. ISPRS International Journal of Geo-Information, 12(8), 346. https://doi.org/10.3390/ijgi12080346

[3] Wei, T., Liu, J., & Zhang, Y. (2025). A novel integrated approach for emerging technologies identification based on knowledge hypergraph: A case study on new energy vehicles. Available at SSRN 5434806. http://dx.doi.org/10.2139/ssrn.5434806

[4] Li, J., Deng, X., & Yao, B. (2025). Enhanced anomaly detection of industrial control systems via graph-driven spatio-temporal adversarial deep support vector data description. Expert Systems with Applications, 270, 126573. https://doi.org/10.1016/j.eswa.2025.126573

[5] Kamatchi, K., & Uma, E. (2025). Insights into user behavioral-based insider threat detection: systematic review. International Journal of Information Security, 24(2), 88. https://doi.org/10.1007/s10207-025-01002-6

[6] Hsu, S., Gülce, E., Ayaz, T. B., Ozcan, A., & Akbulut, A. (2025). Multi-graph anomaly detection in business processes with scalable neural architectures. IEEE Access, 13, 34969-34984. https://doi.org/10.1109/ACCESS.2025.3544268

[7] Duan, M., Sun, S., & Liu, M. (2025). A multimodal deep fusion framework for highway traffic anomaly detection. Scientific Reports, 15(1), 33573. https://doi.org/10.1038/s41598-025-18671-x

[8] Kuchar, K., & Fujdiak, R. (2024). Anomaly detection in industrial networks: Current state, classification, and key challenges. IEEE Sensors Journal, 25(3), 5031-5043. https://doi.org/10.1109/JSEN.2024.3512857

[9] Hu, Z., Chen, W., Wang, H., Tian, P., & Shen, D. (2022). Integrated data-driven framework for anomaly detection and early warning in water distribution system. Journal of Cleaner Production, 373, 133977. https://doi.org/10.1016/j.jclepro.2022.133977

[10] Alharbi, B., Liang, Z., Aljindan, J. M., Agnia, A. K., & Zhang, X. (2022). Explainable and interpretable anomaly detection models for production data. SPE Journal, 27(01), 349-363. https://doi.org/10.2118/208586-PA

[11] Prasanga, D. G. T., Gutierrez, J. A., & Ray, S. K. (2025). The Role of Graph Neural Networks, Transformers, and Reinforcement Learning in Network Threat Detection: A Systematic Literature Review. Electronics, 14(21), 4163. https://doi.org/10.3390/electronics14214163

[12] Ibrahim, N., Rajalakshmi, N. R., Sivakumar, V., & Sharmila, L. (2025). An optimized hybrid ensemble machine learning model combining multiple classifiers for detecting advanced persistent threats in networks. Journal of Big Data, 12(1), 212. https://doi.org/10.1186/s40537-025-01272-w

[13] Li, Z. (2024). Log Event Graph Modeling for Backend Anomaly Detection with Multi-Relational Representation Learning. Transactions on Computational and Scientific Methods, 4(7). https://doi.org/10.5281/zenodo.19676839

[14] Xia, H., Jiang, J., & Wang, Q. (2025). An Interpretable Financial Statement Fraud Detection Framework Enhanced by Temporal–Spatial Patterns. Mathematical and Computational Applications, 30(6), 138. https://doi.org/10.3390/mca30060138

[15] Xu, L., Shang, K., Zhang, X., Zheng, C., & Pan, L. (2025). Multi-scale feature fusion-based real-time anomaly detection in industrial control systems. Electronics, 14(8), 1645. https://doi.org/10.3390/electronics14081645

[16] Zakaria, R. M., Rahman, M. M., Rahman, H., Rafi, M. A., Minto, A., Hossain, M. S., & Saimon, S. I. (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. https://doi.org/10.32996/jefas.2025.7.6.1

[17] Lin, H. C., Wang, P., Chao, K. M., Lin, W. H., & Yang, Z. Y. (2021). Ensemble learning for threat classification in network intrusion detection on a security monitoring system for renewable energy. Applied Sciences, 11(23), 11283. https://doi.org/10.3390/app112311283

[18] Bilot, T., El Madhoun, N., Al Agha, K., & Zouaoui, A. (2023). Graph neural networks for intrusion detection: A survey. IEEe Access, 11, 49114-49139. https://doi.org/10.1109/ACCESS.2023.3275789

[19] Zhou, Z., & Zhang, T. (2025). Prediction of Financial Data Security Risks and Privacy Protection Methods for Listed Companies Based on Hypergraph Learning. Journal of Cyber Security and Mobility, 1505-1534. https://doi.org/10.13052/jcsm2245-1439.1468

[20] Shilpi, S. (2025). Machine Learning-Enabled Platform Engineering for Secure and Scalable Enterprise Payment Ecosystem. Available at SSRN 5696923. http://dx.doi.org/10.2139/ssrn.5696923

[21] Li, P., Zhao, Q., Liu, Y., Zhong, C., Wang, J., & Lyu, Z. (2024). Survey and Prospect for Applying Knowledge Graph in Enterprise Risk Management. Computers, Materials & Continua, 78(3). https://doi.org/10.32604/cmc.2024.046851

[22] Muhati, E., & Rawat, D. (2024). Data-driven network anomaly detection with cyber attack and defense visualization. Journal of Cybersecurity and Privacy, 4(2), 241-263. https://doi.org/10.3390/jcp4020012

[23] Zhang, J., Zhang, W., Liao, W., Zhao, T., & Miao, Y. (2025). Is the real-time data of process safety reliable? An anomaly detection method based on the graph neural network. Process Safety and Environmental Protection, 198, 107066. https://doi.org/10.1016/j.psep.2025.107066

[24] Islam, U., Malik, R. Q., Al-Johani, A. S., Khan, M. R., Daradkeh, Y. I., Ahmad, I., ... & Tag-Eldin, E. M. (2022). A novel anomaly detection system on the internet of railways using extended neural networks. Electronics, 11(18), 2813. https://doi.org/10.3390/electronics11182813

[25] Gao, D. (2024). Graph neural recognition of malicious user patterns in cloud systems via attention optimization. Transactions on Computational and Scientific Methods, 4(12). https://doi.org/10.1186/s13677-024-00709-6