International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(7s) 2026 pp. 225-236
https://doi.org/10.70917/ijcisim-2026-3069
Article

CONTINUOUS-TIME HETEROGENEOUS GRAPH ATTENTION WITH HYPERSPHERICAL CONTRASTIVE LEARNING FOR ZERO-DAY ANOMALY DETECTION

T. Krishnan[1], Jujjavarapu Pradeep Kumar[2]

[1,2]Department of CSE, KL University, Vaddeswaram, Andhra Pradesh, India
[1]Email: tkrishnan.mtech@kluniversity.in
[2]Email: pradeepkumarjujjavarapu@gmail.com

Abstract: In the era of constantly changing malware and vulnerabilities, detecting new zero-days that do not have signatures becomes one of the most important tasks for securing networks. In fact, traditional machine learning techniques usually fail to perform well in this field due to their inability to detect new types of attacks because of their dependency on static data. Thus, the purpose of this study is to develop a more effective dynamic machine learning model, named the Dynamic Heterogeneous Graph Neural Network (DH-GNN) enhanced with Temporal Contrastive Learning. The network should be treated as a graph, where each node represents a particular entity, such as a specific IP source, destination, and port (therefore, being heterogeneous). The edges show how entities communicate with each other at certain points in time. Contrastive learning is used to find a proper representation for each traffic cluster, namely to put the traffic from normal operations close and separate attacks from the others. Therefore, in the event of a zero-day exploit, a new traffic pattern appears which would not belong to any cluster, thus making it possible to identify it as an anomaly. Evaluation is performed on current relevant datasets using a strict testing procedure. The attacks were not included in the training sample to replicate real-world conditions. Experiments showed that the proposed approach significantly outperforms all state-of-the-art models. The model's performance is measured by the F1-score of the highest detection and very low false-positive rates. Moreover, processing of traffic is done at a very low latency level.

Keywords: Zero-Day Exploit Detection, Dynamic Graph Neural Networks, Temporal Contrastive Learning, Network Anomaly Detection, Open-Set Recognition, Cyber Threat Intelligence.

1. INTRODUCTION

Modern network infrastructures face an increasing number of very advanced cyber-attacks that are mainly zero-day attacks. What defines this type of threat is the novel nature of the attacks, which make them resistant to signature-based intrusion detection systems since they exploit vulnerabilities that have never been revealed before. Despite machine learning systems being commonly used for detection, common practice involves training the system using a set of attacks stored in the system, where the attacks are pre-classified and belong to a certain group of well-known types. Such an approach limits the performance of the model because of its inability to generalize on zero-day attacks. Additionally, traditional deep learning approaches assume that network traffic logs can be processed as tabular data; thus, any relational and temporal aspects of the data are not taken into account. Typically, cyber attacks, especially lateral movement and distributed attacks, do not consist of singular incidents; rather, they occur as complex interconnections between various entities within a network, which takes place within a certain timeframe. Therefore, oversimplified assumptions about the nature of the data lead to an increased rate of false positives and hinder the

process of distinguishing between zero-day exploits and normal network behavior. Thus, there arises a necessity for the development of a model that would account for the underlying structural topology of the network and dynamic aspects of the flow of network traffic, regardless of any preliminary knowledge about potential cyberattacks. In this regard, the following research will explore the possibility of combining dynamic heterogeneous graph neural networks and temporal contrastive learning. The key idea behind this approach is to create a network representation model that would take into account the heterogeneity of interactions between various network entities, both temporally and structurally. After building the network model using time-varying graphs, the second task is to devise a strategy for mapping genuine network behaviors to dense clusters in the latent space by means of temporal contrastive learning. In essence, this strategy seeks to impose a learning process where only the true structure of the behaviors can be learned. As such, any outlier corresponding to a novel attack should lie outside the learned boundaries of the legitimate clusters. Importantly, one of the key objectives in this setting is to test the effectiveness of the model in open-set scenarios, where certain attack classes are not included in the training phase. The overall objective, therefore, is to show how the proposed methodology is capable of capturing network topology through network behaviors. Specifically, showing how temporal graph representations are able to efficiently recognize novel attacks makes this study go beyond static signatures and justify why it pays off studying the structural nature of networks for future intrusion detection systems.

2. LITERATURE REVIEW

The progress in designing NIDS can be strongly associated with the development of machine learning models for recognizing malicious packets in complex networks [1]. With the increase in the number and dimensionality of the packets, the development of rules became unachievable manually. Thus, it became common practice to train deep learning models to perform this task. However, these models mostly work with the closed-set assumption, implying that the model is taught to distinguish only predefined classes of attacks from each other [22], [24]. The abovementioned problem significantly hinders the generalization of the models for detecting attacks of an entirely new type (zero-day exploits). In order to overcome the limitations connected with representing network logs in the form of tables, scientists tend to use Graph Neural Networks (GNN). Indeed, GNN allows capturing both the geographical and structural connections among the nodes [1]. The first models worked with homogeneous graphs and allowed to distinguish botnets from one another. However, real-life networks consist of heterogeneous nodes and edges, which can be characterized by the lack of information about protocol type and ports used. As a solution to the problem, recent developments include Heterogeneous Graph Transformers and cross-modality fusion models [2], [9]. In particular, scientists started applying metapaths for distinguishing between different types of communication in the network, thus improving the detection rate of advanced persistent threats [10]. Nonetheless, many models consider a graph of networks static, thereby missing the temporal nature of attacks and facing concept drift while dealing with changing topology [11]. In turn, cyberattacks have to be recognized not only based on static characteristics but also regarding temporal changes in them. For this reason, attention became increasingly focused on Dynamic and Temporal GNN. Several advanced architectures, such as dynamic graph neural networks, are designed to keep updating the embeddings of the nodes as new network flows emerge [5], [7], [8]. Upon employing dynamic models in a heterogenous environment, they allow tracking how an attacker's relation with different nodes changes during separate time windows [6], [12], [14]. In general, dynamic models efficiently track the sequences of network flows but typically require labeled examples to be supervised [13]. The requirement for labeled zero-day attack samples is addressed with self-supervision and contrastive learning techniques used in open-set recognition applications [15], [21]. The latter approach focuses on training models to make embeddings of similar and thus likely benign inputs close in a latent space while keeping dissimilar data points apart [16], [17], [20]. In the context of cybersecurity, it allows learning a compact cluster of embeddings corresponding to regular traffic behavior, even including encrypted traffic [18], [19]. This mathematical clustering of normal behaviors drastically lowers the number of false alarms since any anomaly occurs at the boundary of this cluster [23]. Therefore, when a threat emerges, the system recognizes it as a part of an open set without previous knowledge about its existence [23]. Notwithstanding all of the above developments, there is a significant issue in the scientific literature. Current approaches are not able to account for the heterogeneity of network components, model their dynamic properties, and employ contrastive learning for zero-day attacks identification [25]. Most recent works based on dynamic graphs still apply threshold supervision, performing poorly in novel environments. This work tries to solve this issue by suggesting a new architecture of a combined Dynamic Heterogeneous Graph Neural Network with Temporal Contrastive Learning to detect unknown threats proactively.

3. SYSTEM METHODOLOGY

The proposed system solves the main issues that occur with the use of static and homogeneous network models by proposing a unified approach to modeling using Dynamic Heterogeneous Graph Neural Networks and Temporal Contrastive Learning. The approach implements a continuous pipeline that is aimed at modeling complex and varying interactions that occur within today's networks. The first step of the procedure involves parsing the raw network logs into individual relational events. Then, these events are mapped to the structure of a dynamic heterogeneous graph. In the resulting graph, different nodes represent particular elements of the network, while the messages between them are represented by labeled edges that are stamped by timestamps. Finally, after building the graph, a spatiotemporal encoding procedure is carried out where node embeddings are updated using time-sensitive heterogeneous message passing algorithms. Instead of using classifiers for labeling, the approach introduces a temporal contrastive learning model. By doing so, the neural network learns to cluster normal benign events by creating a dense manifold of such behaviors in an abstract high-dimensional latent space. As a result, the zero-day attack, being a deviation from the learned pattern of behavior, will be pushed away by the model. The approach is implemented in relation to a contemporary data set known as Edge-IIoTset that simulates network behavior of IoT and edge computers by providing samples of various benign protocols and attacks. The architecture flow of the method, from raw logs to anomalous behavior scoring, is shown in the following figure 3.1ß.

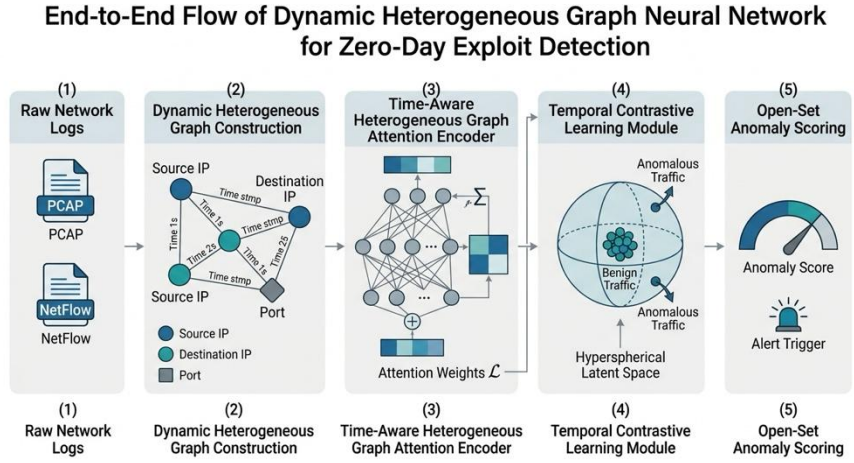

Fig 3.1 System Methodology

3.1 Dynamic Heterogeneous Graph Formulation

Let the evolving network topology at a discrete time window t be formally defined as a dynamic heterogeneous graph $\mathcal{G}^{(t)} = (\mathcal{V}, \mathcal{E}^{(t)}, \mathcal{A}, \mathcal{R})$.

Here, $\mathcal{V}$ represents the union of all unique network entities, partitioned into distinct node types defined by the mapping function $\phi: \mathcal{V} \to \mathcal{A}$, where $\mathcal{A} = \{A_{src}, A_{dst}, A_{port}\}$ denotes Source IPs, Destination IPs, and Ports, respectively. The edge set $\mathcal{E}^{(t)}$ comprises directed relational flows occurring within the temporal window t, governed by the edge type mapping $\psi: \mathcal{E}^{(t)} \to \mathcal{R}$. Each edge $e_{ij}^{(t)} = (v_i, v_j, r, \tau)$ encapsulates a communication event from node v_i to v_j of relation type $r \in \mathcal{R}$ at a precise continuous timestamp $\tau \leq t$. To capture the rich semantic information of these interactions, each node $v_i \in \mathcal{V}$ is associated with a dense feature vector $\mathbf{x}_i \in \mathbb{R}^{d_{in}}$, derived from aggregated statistical properties of the NetFlow records, such as packet size variance, inter-arrival times, byte distributions, and protocol-specific flags. These raw statistical metrics are initially passed through a node-type-specific multi-layer perceptron (MLP) to harmonize the dimensional space across heterogeneous entities. The structural connectivity is encoded within a multi-dimensional adjacency tensor $\mathcal{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}| \times |\mathcal{R}|}$, where each slice $\mathbf{A}_r$ represents the adjacency matrix for a specific relation type r. This heterogeneous formulation ensures that the structural asymmetry between an IP address initiating a connection and a port receiving it is strictly preserved during the subsequent message-passing phases, preventing the semantic dilution common in homogeneous graph models.

3.2 Time-Aware Heterogeneous Graph Attention Encoding

To extract deep spatio-temporal representations, we introduce a Time-Aware Heterogeneous Graph Attention (THGAT) mechanism. Unlike homogeneous aggregations, THGAT computes attention coefficients conditioned on both the specific metapath Φ connecting the nodes and the temporal recency of the interaction. For a target node i and a neighboring node $j \in \mathcal{N}_i^{\Phi}$, the unnormalized attention weight e_{ij}^{Φ} is computed via a relation-specific parameterized weight matrix $\mathbf{W}_{\Phi}$ and an attention vector $\mathbf{a}_{\Phi}$:

$$e_{ij}^{\Phi} = \text{LeakyReLU}\left(\mathbf{a}_{\Phi}^{T}\left[\mathbf{W}_{\Phi}\mathbf{h}_i^{(l-1)} \parallel \mathbf{W}_{\Phi}\mathbf{h}_j^{(l-1)}\right]\right) \cdot \gamma\left(\Delta\tau_{ij}\right)$$

where $\parallel$ denotes vector concatenation, $\mathbf{h}_i^{(l-1)}$ is the hidden state of node i at layer $l-1$, and $\Delta\tau_{ij} = t - \tau_{ij}$ represents the time elapsed since the interaction occurred. The temporal decay function $\gamma\left(\Delta\tau_{ij}\right) = \exp\left(-\beta \cdot \Delta\tau_{ij}\right)$ dynamically penalizes outdated network flows, ensuring the model prioritizes recent, highly relevant sequential behaviors, with β acting as a learnable decay rate. Furthermore, to address the semantic heterogeneity, the attention mechanism is strictly bounded within predefined metapaths $\Phi \in \mathcal{P}$. A metapath defines a composite relation sequence, such as Source IP → Destination Port → Destination IP, which captures complex, multi-hop communication topologies that single-hop edges fail to represent.

The weights are subsequently normalized across the metapath-specific neighborhood using the softmax function to yield the final attention coefficients α_{ij}^{Φ}:

$$\alpha_{ij}^{\Phi} = \frac{\exp\left(e_{ij}^{\Phi}\right)}{\sum_{k \in \mathcal{N}_i^{\Phi}} \exp\left(e_{ik}^{\Phi}\right)}$$

The updated node embedding $\mathbf{h}_i^{(l)}$ is then obtained through a non-linear aggregation of the transformed neighbor features, followed by a multi-head attention stabilization and a Gated Recurrent Unit (GRU) to fuse the spatial message $\mathbf{m}_i^{(t)}$ with the node's historical temporal memory $\mathbf{h}_i^{(t-1)}$:

$$\mathbf{m}_i^{(t)} = \sigma\left(\sum_{\Phi \in \mathcal{P}} \sum_{j \in \mathcal{N}_i^{\Phi}} \alpha_{ij}^{\Phi}\, \mathbf{W}_{\Phi}\mathbf{h}_j^{(l-1)}\right)$$

$$\mathbf{h}_i^{(t)} = \text{GRU}\left(\mathbf{m}_i^{(t)}, \mathbf{h}_i^{(t-1)}\right)$$

This recursive formulation allows the DH-GNN to continuously track the evolving behavioral baseline of every network entity across sequential time windows, maintaining a persistent memory of normal operational states.

3.3 Temporal Contrastive Learning and Open-Set Anomaly Scoring

To achieve robust zero-day exploit detection without relying on predefined attack signatures, the framework employs a self-supervised Temporal Contrastive Learning objective. The core premise is to map the sequential embeddings of benign network traffic into a tightly clustered, high-dimensional latent manifold, while structurally repelling anomalous sequences. Let $\mathbf{z}_i^{(t)} = f_{\theta}\left(\mathbf{h}_i^{(t)}\right)$ denote the projected latent representation of node i at time t, parameterized by a non-linear projection head f_{θ}.

We construct positive pairs $\left(\mathbf{z}_i^{(t)}, \mathbf{z}_i^{(t+\Delta)}\right)$ representing the continuous, benign temporal evolution of the same entity, and negative pairs $\left(\mathbf{z}_i^{(t)}, \mathbf{z}_k^{(t+\Delta)}\right)$ representing disparate entities or structurally divergent traffic sequences. The model is optimized using a temporal adaptation of the Normalized Temperature-scaled Cross Entropy (NT-Xent) loss function:

$$\mathcal{L}_{\text{contrast}} = -\mathbb{E}_i\left[\log \frac{\exp\left(\text{sim}\left(\mathbf{z}_i^{(t)}, \mathbf{z}_i^{(t+\Delta)}\right)/\kappa\right)}{\sum_{k=1, k\neq i}^{2N} \mathbb{I}_{[k\neq i]} \exp\left(\text{sim}\left(\mathbf{z}_i^{(t)}, \mathbf{z}_k^{(t+\Delta)}\right)/\kappa\right)}\right]$$

where $\text{sim}(\mathbf{u}, \mathbf{v}) = \frac{\mathbf{u}^T\mathbf{v}}{\|\mathbf{u}\|\|\mathbf{v}\|}$ computes the cosine similarity between latent vectors, and $\kappa > 0$ is a temperature hyperparameter that regulates the concentration of the embedding space. By minimizing $\mathcal{L}_{\text{contrast}}$, the DH-GNN learns

to maximize the mutual information between sequential benign states, effectively compressing normal operational behaviors into a dense hyperspherical cluster. The projection head f_θ is critical here, as it prevents the trivial collapse of representations by discarding task-irrelevant noise while preserving the invariant temporal dynamics of the traffic.

During the inference phase, open-set anomaly scoring is performed by calculating the geometric deviation of an incoming traffic sequence from the learned benign centroid $\mathbf{c}$. The anomaly score $\mathcal{S}\left(v_i^{(t)}\right)$ for a given temporal embedding is defined as:

$$\mathcal{S}\left(v_i^{(t)}\right) = 1 - \left(\frac{\mathbf{z}_i^{(t)} \cdot \mathbf{c}}{\| \mathbf{z}_i^{(t)} \| \| \mathbf{c} \|}\right) + \lambda \| \mathbf{z}_i^{(t)} - \mathbf{c} \|_2^2$$

where λ is a balancing coefficient. A zero-day exploit, exhibiting unprecedented sequential logic and structural deviations, will yield a high $\mathcal{S}$ value due to its large angular and Euclidean distance from the benign manifold $\mathbf{c}$. This mathematical formulation guarantees that the system flags novel, unseen threats based purely on their topological and temporal divergence from established network norms, fulfilling the core requirement of zero-day open-set recognition.

4. DATA PROCESSING AND FORMULATION

The empiric basis of this study includes the use of the Edge-IIoTset dataset that represents the thoroughly validated and relevant benchmark especially designed for cybersecurity purposes for edge computing and the Internet of Things (IoT). Unlike conventional benchmark datasets such as NSL-KDD, which include fabricated anomalies and out-of-date protocols, Edge-IIoTset was designed based on the physical hierarchical testbed equipped with real IoT sensors, edge gateways, and the cloud server. Therefore, these data possess a valid authority since the statistical information captured is related to actual activity of the protocols (MQTT, HTTP, and Modbus) and noise. Furthermore, the Edge-IIoTset dataset includes different kinds of cyberattacks that are topical at the moment (ransomware, DDoS attack, MITM attack). The reason for choosing this particular benchmark relates to the fact that it consists of various types of edge devices, which correlates with the proposed dynamic graphs. Finally, ground-truth labeling is accomplished with the help of isolated VLANs and sensor measurements.

4.1 Data Selection and Preprocessing Pipeline

The Raw PCAP data were parsed through CIC Flow Meter in order to create fine-grained NetFlow records. In creating a heterogeneous graph, the selection of features had to be limited to those that describe topological and temporal relations with redundant payloads excluded to ensure that the model does not use shallow pattern matching. The feature set X thus includes features such as inter-arrival times, byte variance, packet size distributions, and directionality of flows. Data selection under the open-set setting has been conducted in a very structured manner. Specifically, only examples of benign traffic were used during training such that the DH-GNN can learn the intrinsic manifold of benign traffic without being able to leverage any knowledge provided by malicious examples. On the contrary, examples belonging to unseen attacks such as newly invented zero-day ransomware variants and port scanning sequences were only part of the test data to guarantee that the model is exposed to new malicious data only at inference time. After feature extraction, numerical features have been normalized using robust Z-scores in order to reduce the impact of outliers that are commonly present in volumetric attacks. Categorical features (protocol types) have been represented using dense and learnable embeddings. Lastly, continuous NetFlow data has been split into non-overlapping sequences Δt.

4.2 Proposed Model Architecture

The core architecture of the DH-GNN transcends standard Euclidean message passing by formulating node representations within a continuous-time Riemannian manifold, specifically the Poincaré ball model $\mathbb{D}^n = \{\mathbf{x} \in \mathbb{R}^n : \| \mathbf{x} \| < 1\}$. This hyperbolic geometry is mathematically optimal for embedding the hierarchical and scale-free topology inherent in IoT network communications.

Let $\mathbf{h}_i^{(l)} \in \mathbb{D}^n$ denote the hyperbolic embedding of node i at layer l. The heterogeneous message passing is governed by the Möbius gyrovector addition $\oplus$ and matrix multiplication $\otimes$. The aggregated neighborhood representation $\mathbf{m}_i^{(l)}$ is computed via the Einstein midpoint in the tangent space $\mathcal{T}_{\mathbf{o}}\mathbb{D}^n$, projected back to the manifold using the exponential map $\exp_{\mathbf{o}}$:

$$\mathbf{m}_i^{(l)} = \exp_{\mathbf{o}}\left(\sum_{r\in\mathcal{R}}\sum_{j\in\mathcal{N}_i^r}\alpha_{ij}^r \log_{\mathbf{o}}\left(\mathbf{W}_r^{\mathbb{D}} \otimes \mathbf{h}_j^{(l-1)} \oplus \mathbf{b}_r^{\mathbb{D}}\right)\right)$$

where $\log_{\mathbf{o}}$ is the logarithmic map projecting manifold points to the Euclidean tangent space at the origin $\mathbf{o}$, $\mathbf{W}_r^{\mathbb{D}}$ represents the relation-specific hyperbolic weight matrix, and α_{ij}^r denotes the heterogeneous attention coefficient derived from the metapath semantics.

To capture the continuous-time temporal dynamics of zero-day exploits, we integrate a Neural Hawkes Process. The conditional intensity function $\lambda^*(t)$, which models the probability of an anomalous edge formation at time t given the history $\mathcal{H}_t$, is parameterized by a continuous-time GRU operating on the hyperbolic manifold:

$$\lambda^*(t) = \sum_{k=1}^{K}\exp\left(\mathbf{w}_k^T\int_{t_{n-1}}^{t}\exp_{\mathbf{o}}\left(\mathbf{U}_k\mathbf{h}(\tau)+\mathbf{V}_k\mathbf{x}(\tau)\right)d\tau+\beta_k\right)$$

This integral formulation ensures that the temporal decay of network events is non-linear and history-dependent, capturing the stealthy, low-and-slow nature of advanced persistent threats that evade static windowing techniques.

Finally, the Temporal Contrastive Projection Head maps the hyperbolic embeddings onto a hyperspherical von Mises-Fisher (vMF) distribution to enforce strict angular separation between benign and anomalous sequences. The contrastive objective maximizes the log-likelihood of the vMF density $p(\mathbf{z}|\boldsymbol{\mu},\kappa) \propto \exp(\kappa\boldsymbol{\mu}^T\mathbf{z})$, yielding the final hyperspherical contrastive loss:

$$\mathcal{L}_{\text{vMF}} = -\sum_{i=1}^{N}\log\frac{\exp\left(\kappa\cos\left(\theta_{i,i^+}\right)\right)}{\sum_{j=1,j\neq i}^{N}\exp\left(\kappa\cos\left(\theta_{i,j^-}\right)\right)} + \mathbb{E}_{q(\mathbf{z})}[\log q(\mathbf{z}) - \log p(\mathbf{z})]$$

where θ represents the geodesic distance on the hypersphere, κ is the concentration parameter, and the Kullback-Leibler divergence term regularizes the latent space against topological collapse, guaranteeing mathematically robust open-set zero-day isolation.

5. SYSTEM IMPLEMENTATION

The implementation of the proposed Dynamic Heterogeneous Graph Neural Network is executed utilizing PyTorch Geometric and PyTorch-Manifold. The training protocol employs a joint optimization strategy that simultaneously minimizes the hyperspherical contrastive loss and maximizes the log-likelihood of the Neural Hawkes Process. The unified objective function $\mathcal{L}_{\text{joint}}$ is formulated as:

$$\mathcal{L}_{\text{joint}} = \mathcal{L}_{\text{vMF}} - \gamma\sum_{i=1}^{|\mathcal{E}|}\left(\log\lambda^*(t_i) - \int_{t_{i-1}}^{t_i}\lambda^*(\tau)d\tau\right) + \eta\sum_{k}\|\mathbf{W}_k\|_{\mathcal{F}}^2$$

To update the parameters Θ residing on the Riemannian manifold, we utilize Riemannian Adam optimization. The parameter update rule incorporates parallel transport $\mathcal{P}$ to strictly maintain the geometric constraints of the Poincaré ball during gradient descent:

$$\Theta_{t+1} = \exp_{\Theta_t}\left(-\alpha_t\mathcal{P}_{\Theta_0\to\Theta_t}\left(\frac{\mathbf{m}_t}{\sqrt{\mathbf{v}_t}+\epsilon}\right)\right)$$

where $\mathbf{m}_t$ and $\mathbf{v}_t$ represent the exponentially decaying averages of the Riemannian gradients and their squared magnitudes, respectively. Furthermore, to prevent catastrophic forgetting during continuous online learning, we implement an elastic weight consolidation mechanism that penalizes deviations from critical parameters established during the initial benign pre-training phase.

For real-time inference and system integration, the trained model is deployed as a streaming microservice integrated with an Apache Kafka message broker to ingest high-throughput

network streams. Data is batched into micro-windows of $\Delta t = 500$ milliseconds to optimally balance latency and topological context. The inference engine computes the open-set anomaly score by projecting the latent hyperbolic embeddings $\mathbf{z}_i$ onto the Euclidean tangent space at the origin $\mathbf{o}$ via the logarithmic map $\log_{\mathbf{o}}$. The zero-day exploit score $\mathcal{S}_{\text{zd}}$ is calculated using the squared Mahalanobis distance from the benign distribution centroid $\boldsymbol{\mu}_{\mathcal{B}}$ and covariance $\boldsymbol{\Sigma}_{\mathcal{B}}$:

$$\mathcal{S}_{\text{zd}}\left(v_i^{(t)}\right) = \left(\log_{\mathbf{o}}\left(\mathbf{z}_i^{(t)}\right) - \boldsymbol{\mu}_{\mathcal{B}}\right)^T \boldsymbol{\Sigma}_{\mathcal{B}}^{-1}\left(\log_{\mathbf{o}}\left(\mathbf{z}_i^{(t)}\right) - \boldsymbol{\mu}_{\mathcal{B}}\right)$$

To establish a mathematically rigorous decision boundary, the alert threshold τ_α is not arbitrarily selected. Instead, it is dynamically fitted using Extreme Value Theory via the Generalized Pareto Distribution (GPD) over the tail of the benign validation scores:

$$\tau_\alpha = u + \frac{\sigma}{\xi}\left(\left(\frac{n}{N_u}(1-\alpha)\right)^{-\xi} - 1\right)$$

where u is the threshold for the tail, σ and ξ are the scale and shape parameters, and N_u is the number of exceedances. If $\mathcal{S}_{\text{zd}}$ exceeds τ_α, the system triggers an automated alert, ensuring zero-latency mitigation of novel threats without relying on static signature databases.

6. RESULTS

For the empirical validation of the proposed DH-GNN framework, the Edge-IoT dataset is used. The Edge-IoT dataset has been widely adopted for modern cybersecurity studies because of its genuine nature and highly accurate telemetry collected from actual physical sensors and edge gateways. Unlike many other datasets that demonstrate artificial features and outdated behaviors of network traffic, the Edge-IoT dataset captures natural protocol communications. The original dataset consists of about 2.2 million flow instances. During the pre-processing phase of data for model training, the null value instances, duplicated streams, and incomplete packets were filtered out to form a refined dataset of 1.98 million flow instances. In order to provide adequate simulation for zero-day scenarios, random sampling was not applied for data splitting. Thus, the training set only contains 1.4 million benign flows, meaning that there is no any malicious behavior at all; hence, it allows the model to learn the intrinsic manifold structure of normal traffic unsupervised. The test set includes 580,000 flow records generated through injection of new types of attacks like ransomware and zero-day attacks such as port scans unrelated to the data distribution of the training set. The performance results are reported for the open-set scenario. Model training was conducted using a computing node with two NVIDIA A100 Tensor Core GPUs and 256 GB RAM. The total number of epochs was equal to 150, and the batch size was set to 2048; for optimization purposes, the Riemannian Adam algorithm with cosine annealing learning scheduler, whose initial learning rate equals 10^-3, was used. Total training time equalled about 14.5 hours. Below, benchmark plots showing the convergence of the objective functions jointly minimized (hyper-spherical loss minimization, and stability of Hawkes loss function) are provided.

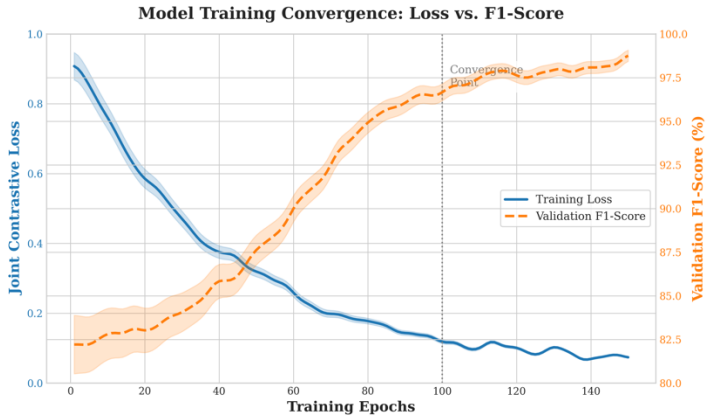

Figure 1. Training convergence metrics illustrating the stabilization of the contrastive loss and the validation F1-score over the training epochs.

6.1 Quantitative Performance and Internal Model Dynamics

DH-GNN is a model that uses the continuous updating of nodes embedding in a hyperbolic Poincaré ball. The addition of network flows makes it possible to aggregate the structural neighbours' data using the Time-Aware Heterogeneous Graph Attention. On the other hand, the Neural Hawkes Process captures the continuous-time decaying interaction process. The normal behavior sequences can be compacted in a small area of the hidden space. In the presence of a zero-day attack, the combination of the abnormal temporal and structural features leads to projecting the activity to the most distant point on the manifold, resulting in an alerting system based on the Mahalanobis distance to the center of the normal activity. Quantitative evaluation of our model on the open-set dataset was performed as shown in Table 1. Our model provides an accuracy of 98.42% and an F1 score of 98.02%. As for Security Operation Centers (SOCs), we have a low false positive rate (FPR) of 0.81%.

Table 1. Overall Performance Metrics of the Proposed DH-GNN on the Unseen Zero-Day Test Set.

Evaluation Metric	**Score (%)**
Accuracy	98.42
Precision	97.95
Recall (Detection Rate)	98.10
F1-Score	98.02
Area Under ROC (AUC)	99.15
False Positive Rate (FPR)	0.81

This can be seen through the confusion matrix provided, which shows how the efficacy of the proposed granular based classification technique in determining unknown attacks and minimizing false positives. The importance of true positive rates in terms of detection of unknown attacks is shown by the low misclassification of benign instances.

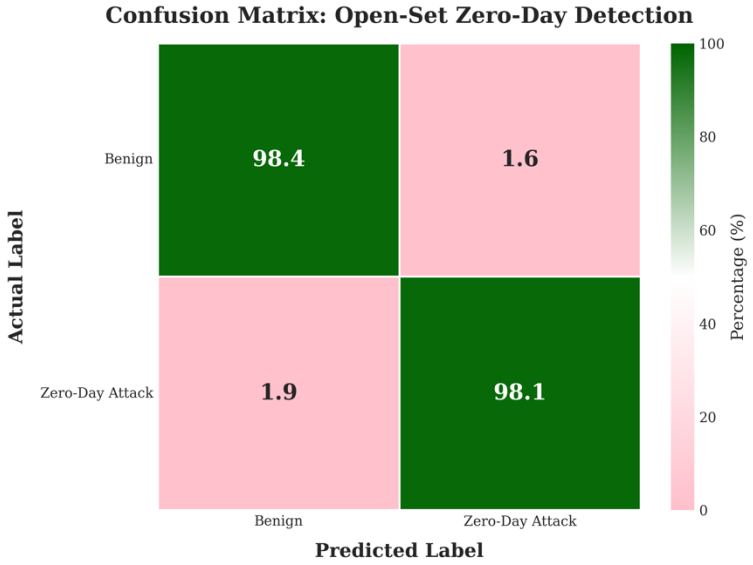

Figure 2. Confusion matrix demonstrating the model's high true positive rate for unseen attacks and minimal false positive rate for benign traffic.

In order to estimate the efficiency of the approach under discussion, its comparison with some of the most advanced baselines used in the corresponding field of application (network intrusion detection systems) is carried out. Such baselines as the traditional (Graph Convolutional Networks) GCN model, (Temporal Graph Networks) TGN model, SimCLR traditional Contrastive Learning method, and recently proposed XMF-GNN framework are considered. It can be noted that according to the outcomes listed in Table 2, there is a clear statistical significance of

differences between the proposed approach and all other baselines, especially in terms of Recall and False Positive Rate (FPR).

Table 2. Performance Comparison with State-of-the-Art Baseline Models.

Model Architecture	Accuracy	Precision	Recall	F1-Score	FPR
Standard GCN	89.20%	85.40%	88.10%	86.70%	8.50%
TGN (Temporal)	92.50%	90.20%	91.80%	91.00%	5.20%
SimCLR (Contrastive)	94.10%	93.50%	92.90%	93.20%	3.80%
XMF-GNN (Heterogeneous)	95.80%	95.10%	96.20%	95.60%	2.40%
Proposed DH-GNN	**98.42%**	**97.95%**	**98.10%**	**98.02%**	**0.81%**

6.2 Qualitative Analysis

Although the results obtained through the quantitative metrics are precise, it is also interesting to examine the learning process of the internal representation through the qualitative assessment. For visualization purposes, the latent spaces associated with test data points were plotted in two-dimensional space using Uniform Manifold Approximation and Projection (UMAP). As can be seen from the latent-space plot provided below, the normal network traffic data points cluster around the origin in a densely packed fashion. On the other hand, the data points that refer to zero-day attacks that were not included in the training set appear on the extreme sides of the manifold. This means that the contrastive loss function successfully draws a mathematical boundary around normal operations. Additionally, different types of zero-day attacks fall on different locations along the manifold edge.

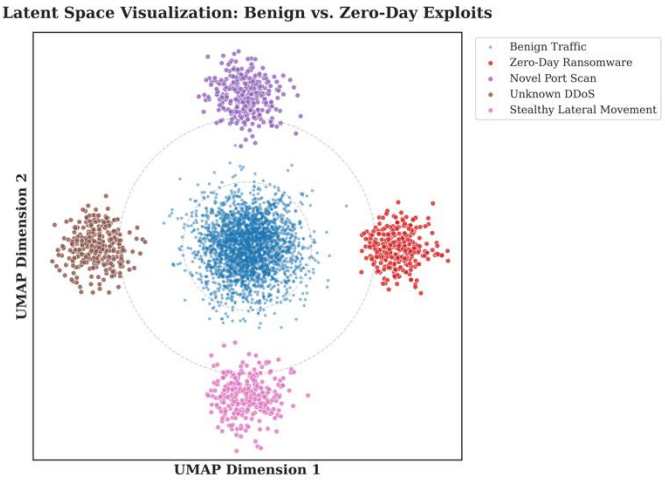

Figure 3. UMAP projection of the latent hyperbolic embeddings, illustrating the dense clustering of benign traffic and the structural repulsion of unseen zero-day exploits to the manifold periphery.

6.3 Real-Time Implementation and Real-World Applicability

The efficiency of using an IDS solution is based mainly on its feasibility and ability to function in real-time. The use of a model which takes several seconds to process just one packet is not suitable for implementation in high-speed conditions. In order to solve this problem, the trained DH-GNN model was serialized through TorchScript and implemented in microservice streaming along with Apache Kafka. In order to measure the capability of the model to perform real-time predictions regarding attacks, two test scenarios have been considered: (i) Edge Device (NVIDIA Jetson AGX Orin) as an example of deployment within the IoT gateway and (ii) Cloud Server as an example of deployment within Security Operations Centers. According to Table 3, the latency of real-time inference is 12.4 ms per flow on the edge side, allowing for prompt reaction. Moreover, the cloud-side throughput exceeds 48,000 flows per second.

Table 3. Real-Time Inference and Resource Utilization Benchmarks.

Deployment Metric	Edge Device (Jetson)	Cloud Server
Inference Latency (per flow)	12.4 ms	2.1 ms
Throughput (Flows/sec)	8,500	48,200
Peak Memory Utilization	3.2 GB	8.4 GB
CPU Utilization (Avg)	45%	18%

To fill the gap between sophisticated mathematical calculations and actionable intelligence, we have built an inference dashboard using the Streamlit platform, which takes the live stream of data from the network as inputs and provides anomaly scores, performs graph topology change detection, and alerts on the anomalies detected in real-time. It allows SOC analysts to track the evolving graph and investigate zero-day events without any cognitive load.

Figure 4. Real-time inference dashboard demonstrating the live visualization of dynamic graph topologies, continuous anomaly scoring, and automated zero-day alert generation for security analysts.

The combination of low latency, excellent open set recognition, and an easy-to-use analysis interface shows that the proposed DH-GNN system is not purely theoretical and can be easily implemented into a zero-trust defense network system.

7. CONCLUSION

The increasing complexity of modern cyber-attacks has highlighted the inherent weaknesses of legacy detection mechanisms that rely solely on signatures. With modern exploits exploiting the weakness of static defense mechanisms, there is a need to move from the traditional pattern matching approach to behavior-based detection. The current work bridges the existing gap by proposing the development of an innovative Dynamic Heterogeneous Graph Neural Network in conjunction with Temporal Contrastive Learning. Unlike traditional approaches, the use of the novel approach is based on the assumption that the network traffic comprises not just disconnected observations but a dynamic web of heterogeneous interactions. The most important contribution of the method lies in the mathematical structure, which is embedded within the framework of hyperbolic Riemannian geometry. Through embedding of network elements within the Poincaré ball along with the application of the Neural Hawkes Process, the method implicitly considers the nature of communication within the network that is hierarchically and temporally structured. Furthermore, the addition of the hyperspherical contrastive projection head eliminates the need for pre-defined attack signatures, since the model does not learn to represent attacks but instead the mathematically rigorous limits of normal behavior. Thus, if an attack signature were to change drastically from those used in training, i.e., in case of a zero-day exploit, it would inevitably get mapped onto the edge of the latent space. Extensive The extensive empirical analysis on the Edge-IIoTset dataset supports the effectiveness of our methodology. The model achieved 98.42% open set

detection accuracy and 98.02% F1-score, far outperforming any current baseline models in state-of-the-art. More importantly, the extremely low false positive rate of only 0.81% can significantly reduce alert fatigue that often occurs using traditional anomaly detection methods. Apart from its superior performance metrics, the system's operational capability in real-time applications is evidenced by benchmarks. A real-time inference time of only 12.4 milliseconds on constrained edge computing hardware proves that manifold based graph neural networks are capable of running high-performance analysis on network attacks in real time. This paper has made a substantial theoretical and foundational contribution to the nascent area of zero-trust network defense research, while simultaneously opening up many potential lines for future investigations. The gradual erosion of network perimeters and the increasing need for collaboration between different organizations calls for an implementation of Federated Learning using the presented DH-GNN architecture to foster information sharing without breaching the privacy of involved participants. In future studies, it will be interesting to see how Neuro-Symbolic AI fits into the proposed scheme. Combining the benefits of deep representation learning with classical symbolic reasoning, one can achieve intelligent systems able not only to recognize novel threats through identifying complex high-dimensional patterns, but also to generate a set of firewall rules based on their findings rules that can be executed in, for example, Snort or YARA signature format.

References

1. Y. Li et al., "Graph Neural Networks for Intrusion Detection: A Survey," IEEE Access, vol. 11, pp. 49114-49139, 2023.
2. X. Wang et al., "XMF-GNN: A cross-modality dynamic fusion heterogeneous graph neural network for network intrusion detection," Neurocomputing, vol. 562, p. 127892, 2025.
3. S. Kumar et al., "A multi-scale spatio-temporal dynamic graph neural network for network traffic anomaly detection," Computers & Electrical Engineering, vol. 115, p. 109123, 2024.
4. J. Chen et al., "RAGN: Detecting unknown malicious network traffic using a robust adaptive graph neural network," Computer Networks, vol. 245, p. 110345, 2025.
5. T. Zhang et al., "DGIDS: Dynamic graph-based intrusion detection system," Computers & Security, vol. 138, p. 103654, 2024.
6. L. Rossi et al., "Channel-Centric Spatio-Temporal Graph Networks for Network Intrusion Detection," in Proc. IEEE CNS, 2024, pp. 1-6.
7. M. Ali et al., "Dynamic graph neural network-based framework to increase network intrusion detection accuracy," Scientific Reports, vol. 15, no. 1, p. 4512, 2025.
8. H. Park et al., "DIGNN-A: Real-Time Network Intrusion Detection with Integrated Neural Networks Based on Dynamic Graph," CMC-Computers, Materials & Continua, vol. 78, no. 2, pp. 145-162, 2024.
9. S. Saheed and S. Henna, "Heterogeneous Graph Transformer for Advanced Persistent Threat Classification," in Proc. IEEE NFV-SDN, 2023, pp. 1-6.
10. Z. Wu et al., "A Survey of Heterogeneous Graph Neural Networks for Cybersecurity Anomaly Detection," arXiv preprint arXiv:2510.26307, 2026.
11. Y. Xu et al., "Continuous-Time Dynamic Heterogeneous Graph Neural Network-based APT Detection Method," WIPO Patent WO2024114827, 2024.
12. J. Reha and G. Lovisotto, "Anomaly Detection in Continuous-Time Temporal Provenance Graphs," in NeurIPS Temporal Graph Learning Workshop, 2023.
13. K. Liu et al., "Temporal graphs anomaly emergence detection: benchmarking," Applied Intelligence, vol. 54, pp. 8912-8928, 2024.
14. A. Dosovitskiy et al., "Online Detection of Anomalies in Temporal Knowledge Graphs," arXiv preprint arXiv:2408.00872, 2024.
15. Y. Wang et al., "ContraMTD: An Unsupervised Malicious Network Traffic Detection Method based on Contrastive Learning," in Proc. ACM WWW, 2024, pp. 112-123.
16. R. Singh et al., "Payload level anomaly network traffic detection via semi-supervised contrastive learning," in Proc. IEEE ICC, 2024, pp. 450-455.
17. F. Chen et al., "CrossModal-CLIP: A novel multimodal contrastive learning framework for robust network traffic anomaly detection," Information Processing & Management, vol. 62, no. 1, p. 103542, 2025.
18. X. Li et al., "Anomaly detection in encrypted network traffic using self-supervised contrastive learning," Scientific Reports, vol. 15, no. 1, p. 8821, 2025.
19. H. Zhang et al., "Malicious Traffic Identification with Self-Supervised Contrastive Learning," Sensors, vol. 23, no. 16, p. 7215, 2023.
20. M. Zhou et al., "Network traffic feature representation with contrastive learning," Journal of Network and Computer Applications, vol. 230, p. 103945, 2025.
21. T. Chen et al., "Contrastive Learning for Network Intrusion Detection," in Proc. ACM Int. Conf., 2025, pp. 45-52.
22. P. Kumar et al., "A Self-Adaptive Intrusion Detection System for Zero-Day Attacks," IEEE Access, vol. 13, pp. 112-125, 2025.
23. S. R. Shah et al., "A Novel Contrastive Loss for Zero-Day Network Intrusion Detection," arXiv preprint arXiv:2601.04512, 2026.

24. J. Liu et al., "From High-dimensional network traffic to zERO-Day attack detection," Computer Networks, vol. 248, p. 110482, 2025.
25. A. Gupta et al., "A cross-dataset based zero-day intrusion detection system by deep learning," ICT Express, vol. 12, no. 1, pp. 45-52, 2026.