

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

---

10-2023

### HRGCN: Heterogeneous graph-level anomaly detection with hierarchical relation-augmented graph neural networks

Jiaxi LI

Guansong PANG

Singapore Management University, [gspang@smu.edu.sg](mailto:gspang@smu.edu.sg)

Ling CHEN

Mohammad-Reza NAMAZI-RAD

Follow this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [Databases and Information Systems Commons](#), and the [Numerical Analysis and Scientific Computing Commons](#)

---

#### Citation

LI, Jiaxi; PANG, Guansong; CHEN, Ling; and NAMAZI-RAD, Mohammad-Reza. HRGCN: Heterogeneous graph-level anomaly detection with hierarchical relation-augmented graph neural networks. (2023). *2023 IEEE 10th International Conference on Data Science and Advanced Analytics (DSAA): Thessaloniki, Greece, October 9-13: Proceedings*. 1-10.

Available at: [https://ink.library.smu.edu.sg/sis\\_research/8412](https://ink.library.smu.edu.sg/sis_research/8412)

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylds@smu.edu.sg](mailto:cherylds@smu.edu.sg).

# HRGCN: Heterogeneous Graph-level Anomaly Detection with Hierarchical Relation-augmented Graph Neural Networks

Jiayi Li  
*University of Technology Sydney*  
 Sydney, Australia  
 jiayi.li-1@student.uts.edu.au

Guansong Pang  
*Singapore Management University*  
 Singapore 178902, Singapore  
 gspang@smu.edu.sg

Ling Chen  
*University of Technology Sydney*  
 Sydney, Australia  
 ling.chen@uts.edu.au

Mohammad-Reza Namazi-Rad  
*University of Wollongong*  
 Wollongong, Australia  
 mrad@uow.edu.au

**Abstract**—This work considers the problem of heterogeneous graph-level anomaly detection. Heterogeneous graphs are commonly used to represent behaviours between different types of entities in complex industrial systems for capturing as much information about the system operations as possible. Detecting anomalous heterogeneous graphs from a large set of system behaviour graphs is crucial for many real-world applications like online web/mobile service and cloud access control. To address the problem, we propose HRGCN, an unsupervised deep heterogeneous graph neural network, to model complex heterogeneous relations between different entities in the system for effectively identifying these anomalous behaviour graphs. HRGCN trains a hierarchical relation-augmented Heterogeneous Graph Neural Network (HetGNN), which learns better graph representations by modelling the interactions among all the system entities and considering both source-to-destination entity (node) types and their relation (edge) types. Extensive evaluation on two real-world application datasets shows that HRGCN outperforms state-of-the-art competing anomaly detection approaches. We further present a real-world industrial case study to justify the effectiveness of HRGCN in detecting anomalous (e.g., congested) network devices in a mobile communication service. HRGCN is available at <https://github.com/jiayililearn/HRGCN>.

**Index Terms**—GNN, Heterogeneous Graph, Anomaly Detection, Industrial Application

## I. INTRODUCTION

Anomaly detection is an important area of data analytics and machine learning which involves the development of algorithms and techniques to identify deviations from normal behaviours. By distinguishing between normal and abnormal patterns, anomaly detection empowers organisations to proactively address emerging issues, prevent system failures, and enhance reliability [1]. In the realm of large and complex systems, such as cloud platforms and mobile/IoT networks, anomaly detection has garnered increasing attention in recent years [2]. These systems are characterised by their extensive data volumes, diverse services and their relations, presenting formidable challenges in detecting anomalous patterns,

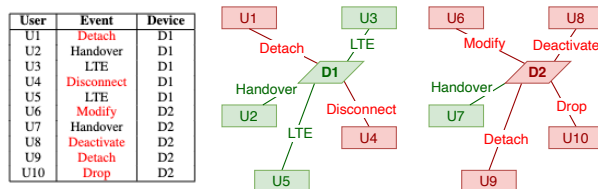


Fig. 1. Example mobile network events with heterogeneous entities and relations. Regular (green) vs. irregular (red) events and devices.

such as system performance degradation, unauthorised access, or malicious activities. Addressing these challenges requires comprehensive evaluations of behaviours involving diverse entities in the system. In this context, behaviour data is typically represented by graphs to capture the entity interactions. Figure 1 illustrates an example of normal and abnormal devices based on heterogeneous mobile networks/graphs involving two types of entities (users and devices) and multiple types of interactions among the entities. Although the users may occasionally encounter intermittent “bad” events, it does not necessarily imply that the entire graph exhibits abnormal behaviours. Therefore, a nuanced understanding of the complex inter-dependencies and interactions among the heterogeneous entities within the system is crucial for accurately detecting the anomalies.

There are three main streams of related anomaly detection approaches on graphs, including homogeneous-graph-based approaches [3]–[5], heterogeneous-graph-based approaches [6]–[8], and hybrid approaches considering both graph and non-graph data (e.g., system logs) [9]. Their common objective is to identify graph-level anomalies, i.e., exceptional graphs that deviate significantly from the other graphs. Compared with homogeneous-graph-based approaches that focus on dealing with graphs consisting of a single type of nodes

and edges, heterogeneous-graph-based approaches generally perform better on complex-structured datasets. This is mainly because heterogeneous graphs, which can be used to represent behaviours between different types of entities in complex industrial systems, excel at capturing a variety of important information about the system operations for anomaly detection. On the other hand, hybrid approaches like DeepTraLog [9] can learn normal patterns from more sources of data, such as heterogeneous graphs based on the system trace data and log embeddings based on the log texts. The hybrid approaches are generally more effective, but they rely on the availability of extra data sources, in addition to the heterogeneous graphs. Since these additional data sources are often unavailable or difficult to obtain, this work focuses on exploring the heterogeneous-graph-based approaches.

Most current approaches [8], [10]–[14] based on heterogeneous graphs primarily focus on learning representations in a supervised or semi-supervised manner. However, in practical anomaly detection scenarios, acquiring high-quality labeled anomaly data is prohibitively costly. There are two quick solutions to this problem. One is to apply an unsupervised anomaly detection optimisation objective, such as SVDD, to replace the loss layer of the traditional HetGNNs [8] so as to avoid the requirement for labelled anomaly data. Another solution is to adapt an unsupervised homogeneous-graph-based anomaly detection model, such as the recently proposed GLocalKD [3], by replacing the used homogeneous GNNs with HetGNNs layers. Nonetheless, both of the approaches could not produce competitive results in the context of anomaly detection within complex systems (see Sec. IV-C).

To this end, we propose Hierarchical Relation-augmented Graph Convolution Neural Network (HRGCN), which is an unsupervised approach that extends heterogeneous HetGNNs by exploiting more complex heterogeneous relation hierarchies for identifying abnormal graphs that are difficult to spot by considering only shallow heterogeneous relations. Figure 2 provides a simple example to illustrate how anomalous and normal behaviours can be distinguished with fine aggregation based on node types. Particularly, we first introduce a hierarchical relation module that considers deeper interactions within the source-destination neighbourhood types and edge types rather than considering all neighbour types and edge types simultaneously. Then, we introduce a self-supervised HetGNN heterogeneous relation prediction module that does not require human-annotated anomaly information. Furthermore, to support the self-supervised task, we introduce a Heterogeneous Graph Data Augmentation (HetGDA) method to enhance the data coverage during the training process and produce a more generalised model. Finally, the model is trained by using a joint objective of one-class classification and self-supervised heterogeneous relation prediction.

In summary, this work makes the following contributions:

- We propose a novel hierarchical heterogeneous graph convolutional network model, i.e., HRGCN, for improved representation learning on heterogeneous graph-level anomaly detection. HRGCN can capture the general

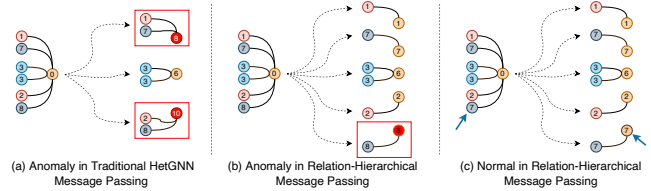


Fig. 2. Simple Example Illustrates the Architectural Difference in Message Passing (MP). Each node is labelled with node attribute value which is passed onto its neighbour during MP. Traditional HetGNN (a) considers all the neighbour nodes with mixed node types at once without further specification. In contrast, the proposed model (b, c) considers each source-destination pair individually during MP. Assume simple criteria where a value greater than 7 is anomalous. (a) and (b) identify the graph as an anomaly since anomalous behaviours are detected. (c) is a normal graph introduced with a different node in the neighbourhood, the traditional HetGNN would still identify it as an anomaly, while the proposed architecture can effectively avoid such errors.

structural and attributed heterogeneity and is extended to the closer relations within source-to-destination neighbourhood types and edge types, which makes it effective in learning complex system-flow-like graphs for anomaly detection.

- We introduce a self-supervised prediction module and heterogeneous graph data augmentation method for training HRGCN more effectively in unsupervised settings.
- We conduct extensive experiments on two system-focused graph datasets with distinct structural properties and attributes. The result demonstrates significant performance improvement in HRGCN compared to other state-of-the-art baseline models.
- We conduct extensive case studies over real-world-like mobile network event data. HRGCN demonstrates a superior performance running in a real-world industrial setting compared to other state-of-art baseline models.

## II. RELATED WORK

**Heterogeneous Graph Neural Networks.** Heterogeneous Graphs (HGs) have been widely studied in representation learning [15]. Dong et al. introduced a heterogeneous graph model, metapath2vec, to produce scalable network representations [16]. Shi et al. created the HERec model specifically for recommendation systems with heterogeneous graph data and introduced an extended matrix factorisation (FM) module to optimise the rating prediction tasks [17]. Hu et al. introduced a three-way characterisation to combine object embeddings with meta-path-based context embedding for the recommendation tasks [18], and further introduced an adversarial learning model, HeGAN, on heterogeneous graphs in 2019 [19] which is inspired by the generative adversarial networks (GANs).

Graph Neural Networks (GNNs) [20] are neural network models that can be used directly to analyse graph structure data and focuses on revealing complicated relationships between objects in a network. In the past few years, research has been conducted on improving GNNs and applying them to different areas. A standard learning methodology from GNNs generates meaningful representations describing the input graphs. To this extent, graph-structured data demonstrates the capability

to easily model more complicated systems, such as social networks [21], biomedical networks [22], the Internet [23], and other complex information networks with attributes [24]. Therefore, producing helpful network representations becomes a continuous research topic. Li et al. proposed the Gated Graph Neural Network (GGNN), which uses a gated recurrent unit (GRU) to remember or forget messages during each message passing in GNNs [25]. Zhang et al. proposed the Heterogeneous Graph Neural Network (HetGNN), which incorporates heterogeneous structural information, attributes, and contents [8]. The HetGNN model is aimed to produce helpful node embeddings within the networks by introducing multiple aggregation modules based on various heterogeneous information. Wang et al. proposed the Heterogeneous Graph Attention Network (HAN) and introduced the hierarchical attention mechanism into the GNNs [10]. Specifically, it includes both node-level and semantic-level attention, which considers the importance between the nodes and meta-paths and the importance within different meta-paths. Fu et al. proposed a meta-path aggregated Heterogeneous GNN model (MAGNN) to overcome the limitations of traditional meta-path-based GNNs [11]. It introduces the additional node content transformation, intra-meta-path aggregation, and inter-meta-path aggregation to accommodate input node attributes and the message combination between meta-paths. Hu et al. proposed the Heterogeneous Graph Transformer to model heterogeneity over each edge and feed the sampled graph features into a Transformer before the message passing [12]. Zhao et al. proposed a Heterogeneous Network Embedding based on network schema (NSHE) [13]. Instead of using the meta-paths, utilising the network schema directly could preserve rich semantics from the high-order structure when learning the network representations. Similarly, Sun et al. introduced a graph classification method using hypergraph embeddings, which is a network-based learning framework for heterogeneous hypergraphs [14].

**Graph-level Anomaly Detection.** There are limited research on anomaly detection via Heterogeneous Graph representation learning, particularly detecting anomalous behaviours from running systems and applications. Manzoor et al. proposed a streaming low-latency model to detect anomalous actions from the heterogeneous graphs constructed from system APIs [7]. It introduces a light-weighted hashing mechanism in real-time to update and compute the similarities of the hashed graphs. MultimodalTrace [26] is a graph-based anomaly detection model and leverages a multi-modal LSTM module to learn the graph representations from only the normal graphs. TraceAnomaly [27] is an unsupervised model proposed to detect anomaly traces from microservice applications. It considers only service-level graphs and uses the variational auto-encoders (VAE) to detect anomaly patterns. Zhang et al. [9] proposed a DeepTraLog model to combine both log data and graph-structured data to detect anomalies from microservice applications more confidently. Another closely related line of research is homogeneous graph anomaly detection [3]–[5]. They also perform graph-level anomaly detection, but

TABLE I  
NOTION SUMMARY

Notions	Definitions
$V, E, T_v, T_e$	Collection of node, edge, node type, edge type
$v_i, e_{j,i}$	Node $i$ , edge between node $j$ and $i$
$T_{v_j}$	Node type of node $j$
$t_{src}, t_{dst}$	Source node type, destination node type
$t_{e_{j,i}}; t \in T_e$	Edge type of the edge between node $j$ and $i$
$\mathbf{c}$	SVDD center
$Batch(i)$	A collection of graphs in the $i$ -th batch
$n_B$	The batch size
$\mathcal{N}(i)$	Neighbours of node $i$
$\mathcal{G}_i$	Embedding of graph $i$
$s(\mathcal{G}_i)$	Anomaly score of graph $i$
$\tilde{T}_v$	A collection of augmented node types of graph
$\tilde{T}_e$	A collection of augmented edge types of graph
$\tilde{E}$	A collection of augmented edges

they are focused on homogeneous graphs. Thus, they cannot work effectively when applied to heterogeneous graphs. We also demonstrate that simple adaptation of these homogeneous graph approaches cannot work well in our experiments.

### III. METHODOLOGY

We first introduce a baseline model that adapts a popular deep one-class classifier, Deep SVDD [28], for heterogeneous graph-level anomaly detection. We then introduce the proposed framework and methodologies which are related to the proposed HRGCN model in details. A list of notations used is summarised in Table I.

#### A. Baseline Model: One-Class Heterogeneous Graph Convolution Network

The original HetGNN model [8] was built to learn on a large graph to generate each node’s representations. It was implemented based on the neighbour sampling mechanisms and LSTM layers for encoding node/neighbourhood contents. Furthermore, the loss function is computed by calculating the cross entropy loss between the triplet outputs, including positive and negative representations. To comply with our settings, we have extended the HetGNN model with Message Passing (MP) [29] and One-Class SVDD loss [28].

##### 1) Heterogeneous Convolution Neural Network (GCN):

An attributed heterogeneous graph can be defined as  $G = (V, E, T_V, T_E)$ , where  $v_i \in V$  denotes the nodes,  $e_{j,i} \in E$  denotes the edges, and  $T_V$  and  $T_E$  represent the sets of node types and edge types respectively. For each node  $v_i$ , it associates with a set of node attributes  $\mathbf{x}_i \in \mathbb{R}$  for node  $i$ , and its neighbourhood can be defined as  $\mathcal{N}(i)$ . On the other hand, message passing can also be expressed as neighbourhood aggregation. Under this implementation, message passing at  $k$ th message propagation in graph neural network can be described as Eq. 1:

$$\mathbf{x}_i^{(k)} = \gamma^{(k)} \left( \mathbf{x}_i^{(k-1)}, \sum_{j \in \mathcal{N}(i)} \phi^{(k)} \left( \mathbf{x}_i^{(k-1)}, \mathbf{x}_j^{(k-1)}, \mathbf{e}_{j,i} \right) \right). \quad (1)$$

where  $\phi$  and  $\gamma$  are differentiable functions for constructing messages and updating node embeddings. Furthermore, the

implemented Graph Convolution Network (GCN) layer can be mathematically described as below

$$\mathbf{x}_i^{(k)} = \sum_{j \in \mathcal{N}(i) \cup \{i\}} \frac{1}{\sqrt{\deg(i)} \cdot \sqrt{\deg(j)}} \cdot (\mathbf{w}^T \cdot \mathbf{x}_j^{(k-1)}) + \mathbf{b}. \quad (2)$$

That is, the node features of all the neighbour nodes are first transformed by a weight matrix  $\mathbf{w}$ , then normalised by the degrees of the node pairs, and the final node feature is acquired by summing up the transformed values of all the neighbours at the end.

When considering the heterogeneous node types, each node type  $t$  from the neighbourhood is transformed using a different neighbourhood encoder  $\phi_t$  to transform the contents from all the nodes in the same neighbourhood type. Eq. 3 and Eq. 4 below describe the masked message passing processes with different node types:

$$\mathcal{F}(x_i, x_j, t) = \begin{cases} 0, & \text{if } t_{v_j} \neq t \\ \frac{1}{\sqrt{\deg(i)} \cdot \sqrt{\deg(j)}} \cdot (\mathbf{w}^T \cdot \mathbf{x}_j^{(k-1)}) + \mathbf{b}, & \text{else} \end{cases} \quad (3)$$

$$\mathbf{x}_i^{(k)} = \text{Concat}_{t \in T_{v_j}} \left( \sum_{j \in \mathcal{N}(i) \cup \{i\}} \phi_t(\mathcal{F}(x_i, x_j, t)) \right). \quad (4)$$

The neighbour nodes of the different types are fed into different differentiable functions. Then concatenation is performed on all the transformed neighbourhood types to get resulting node representations.

In the end, to produce the embedding of the graph  $\mathcal{G}$ , max pooling is performed on all the node representations from the last message passing iteration  $K$  to produce a single vector for the final representation:  $\mathcal{G} = \max_{i \in V} (x_i^{(K)})$ .

2) *One-Class Deep SVDD*: Deep Support Vector Data Description (Deep SVDD) is an unsupervised loss function used in the area of Anomaly Detection. SVDD can be described as a quadratic loss between the network representation  $\phi(\mathcal{G}_i; \mathcal{W})$  and a hyper-sphere center  $\mathbf{c}$ . The specific objective is to minimise the distance of every network representation to the hyper-sphere center:

$$\min_{\mathcal{W}} \frac{1}{n} \sum_{i=1}^n \|\phi(\mathcal{G}_i; \mathcal{W}) - \mathbf{c}\|^2 + \frac{\lambda}{2} \sum_{\ell=1}^L \|\mathbf{W}^\ell\|_F^2. \quad (5)$$

where  $\lambda$  is a regularisation parameter and  $\mathbf{W}^l$  denotes the weight of the layer  $l$  in the HetGNN model. After training the one-class HetGNN, given a test graph  $\mathbf{g} \in \mathcal{G}$ , the anomaly score  $s$  is then defined as the distance to the fixed SVDD center  $\mathbf{c}$ , i.e.,  $s(\mathcal{G}_i) = \|\phi(\mathcal{G}_i; \mathcal{W}) - \mathbf{c}\|^2$ . Note that, the center  $\mathbf{c}$  is computed as an averaged vector of the graph embeddings within the first data batch:

$$\mathbf{c} = \frac{\sum_{i=1}^{n_B} (\phi(\mathcal{G}_i))}{n_B}, \{\mathcal{G}_i \in \text{Batch}(1)\}. \quad (6)$$

## B. Modelling Hierarchical Heterogeneous Relations in Graphs of Normal Class

The above-implemented HetGCN model aggregates all the neighbourhood types together regardless of the source node type. In scenarios where there are numerous source node types, the exchange of information between specific source-destination node types may diminish in subsequent transformations and aggregations within the network architecture. To address this issue and preserve more information during the learning process, we propose the HRGCN, which introduces a segregation of the architecture based on individual source nodes and edge types. This approach allows for the incorporation of deeper hierarchical relations within the heterogeneous graph. Figure 3 illustrates the overall framework of HRGCN in a detailed example, including the relation-hierarchy feature convolution module and self-supervised module.

1) *Source-Destination Node Type and Edge Type Relations*: To establish a more distinctive hierarchy during the message passing process from various source nodes, we introduce the inclusion of source node type  $t_{src}$  and destination node type  $t_{dst}$ , where  $t_{src} \in T_v$  and  $t_{dst} \in T_v$ , when iterating within the network.

The source-destination hierarchical masking function  $\mathcal{F}^1$  can then be extended as below

$$\mathcal{F}^1(i, j, T_v) = \sum_{t_{src} \in T_v, t_{dst} \in T_v} \left( \begin{cases} 0, & \text{if } (t_{v_i}, t_{v_j}) \neq (t_{src}, t_{dst}) \\ \text{Norm}(x_i, x_j), & \text{else} \end{cases} \right), \quad (7)$$

where  $\text{Norm}(x_i, x_j)$  is a normalisation step of the node pairs for each source-destination type which is defined as follows

$$\text{Norm}(x_i, x_j) = \frac{1}{\sqrt{\deg(i)} \cdot \sqrt{\deg(j)}} \cdot (\mathbf{w}^T \cdot \mathbf{x}_j^{(k-1)}) + \mathbf{b}. \quad (8)$$

With an additional hierarchy introduced for source node types, the differentiable functions are expanded to accommodate the resulting changes. The message passing graph can then be defined as

$$\mathbf{x}_i^{(k)} = \text{Concat}_{t \in T_v} \left( \sum_{j \in \mathcal{N}(i) \cup \{i\}} \phi_{t_{i,j}}(\mathcal{F}^1(i, j, T_v)) \right). \quad (9)$$

where  $\phi_{t_{i,j}}$  denotes as individual differentiable function of each source-destination node type pair.

A similar hierarchical enhancement can also be applied to each edge type relation, further segregating the hierarchy by combining edge types and source-destination node type pairs via Eq. 11.

$$\mathbf{x}_i^{(k)} = \text{Concat}_{t \in T_v, T_e} \left( \sum_{j \in \mathcal{N}(i) \cup \{i\}} \phi_{t_{i,j}}(\mathcal{F}^2(i, j, T_e)) \right), \quad (10)$$

$$\mathcal{F}^2(i, j, T_e) = \sum_{t \in T_e} \left( \begin{cases} \phi_{t_e}(\mathcal{F}^1(i, j, T_v)), & \text{if } t_{e_{i,j}} = t \\ 0 & \text{else} \end{cases} \right). \quad (11)$$

where  $\phi_{t_e}$  is the differentiable function for each edge type in the graph. Furthermore, the final node representations

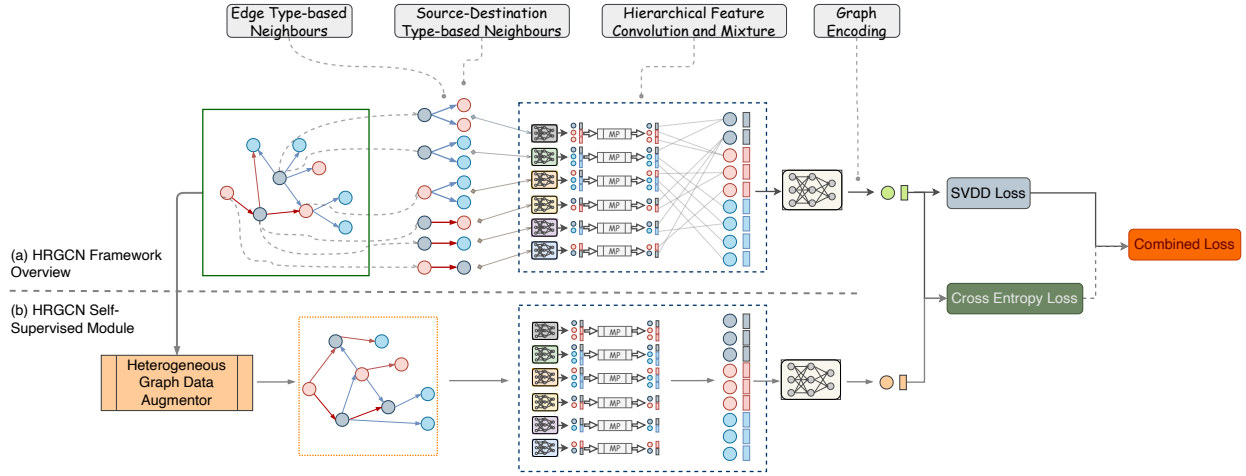


Fig. 3. (a) Overall Architecture of HRGCN. It first identifies all the edge types and source-to-destination node types in the input graph and encodes the segregated node features via hidden NN layers. Then the encoded node features in each edge and source-to-destination type are fed to the message-passing network to aggregate the neighbour information. Next, the aggregated neighbour features are then mixed and aggregated for each node in the graph. The final node embeddings are aggregated to create graph representation via the final NN layers and the SVDD loss. (b) The Self-Supervised module framework. It augments the input HGs and produces the graph representation with the HRGCN architecture, and then both the input graph and augmented graph are combined via Cross-Entropy (CE) Loss. In the end, the final loss of the input graph data is calculated based on the SVDD error and weighted CE errors.

are updated as Eq. 10 which aggregates on all the source-destination node type pairs  $T_{v_i, j}$  and edge types  $T_e$ . This indicates that for edges in each edge type, we will perform the message passing for all the source-destination node type pairs.

### C. Heterogeneous Graph Data Augmentation for Self-Supervised Normality Learning

Graph Data Augmentation (GraphDA) methods are employed to augment graph datasets and enhance the representation learning process. Existing graph augmentation approaches [30] primarily focus on homogeneous graphs and simple heterogeneous graphs. However, to cater to attributed heterogeneous graphs, we introduce four heterogeneous graph augmentation methods specifically designed to complement the proposed HRGCN self-supervised training process.

1) *Heterogeneous Edge Perturbation*: Edge Perturbation is widely used in GraphDA methods in different Deep Graph Learning tasks [31]–[33]. It is the technique of randomly adding or removing edges within the graph. Given an adjacency matrix, edge perturbation preserves the original node order and overwrites a portion of the edges. For example, Eq. 12 describes the process of perturbation of edges. By obtaining a randomly sampled corruption matrix  $C$ , then XOR operation  $\oplus$  is carried out between the adjacency matrix  $A$  and the corrupted matrix  $C$  to create a new adjacency matrix  $\tilde{A}$ :

$$\tilde{A} = A \oplus C. \quad (12)$$

To also preserve the heterogeneity from the original graphs and generate better quality heterogeneous graphs, the matrix  $C$  is generated based on the prior information of the source-destination node type pair and edge type distributions:

$$C(T_{v_i}, T_{v_j}, T_e) \sim \{P(T_{v_i}, T_{v_j}), P(T_e)\}. \quad (13)$$

2) *Heterogeneous Edge Replacement*: Edge Replacement is an alternative edge augmentation method involves addition and removal of edges from the original graph. In contrast to the randomness of Edge Perturbation, which applies to all possible node pairs, edge replacement maintains the total number of edges in the specific neighbourhood while introducing new edges to the graph. Moreover, the selection of edges to add follows a criterion based on their rarity within the graphs. Edges with infrequent occurrences in the neighborhood are more likely to be generated in the augmented graphs. The replacement can be described as follows:

$$\mathbf{R} \sim P(e_{i,j}) \propto \frac{1}{P(T_{v_i}, T_{v_j}, T_e)}. \quad (14)$$

In addition, some existing edges in the same neighbourhood are removed to form similar-sized resulting graphs, as described in Eq. 15. The new set of edges  $\tilde{\mathbf{E}}$  is acquired by replacing a subset of edges  $\mathbf{E}'$  in the neighbourhood with the newly sampled edges  $\mathbf{R}$ .

$$\tilde{\mathbf{E}} = \mathbf{E} \cup \mathbf{R} \setminus \{\mathbf{E}' \subseteq \mathbf{E}(T_{v_i}, T_{v_j}, T_e)\}. \quad (15)$$

3) *Heterogeneous Node and Edge Type Swapping*: Another approach to enhance the training process of HRGCN is by creating augmented graphs through node and edge type swapping. In the node type swap augmentation step, a random node type pair  $(T_{v_a}, T_{v_b})$  is selected for the swap operation, but only if  $T_{v_a} \neq T_{v_b}$ . Consequently, a new set of node types, denoted as  $\tilde{T}_v$ , can be defined after the swapping process.

$$\tilde{T}_v \sim \{\mathbf{V}_a' \subseteq \mathbf{V}(T_{v_a}), \mathbf{V}_b' \subseteq \mathbf{V}(T_{v_b})\}. \quad (16)$$

On the other hand, edge type swap can be performed using a similar method as shown in Eq. 17 below. This operation results in a new edge type list denoted as  $\tilde{T}_e$ , which is generated based on the edge type pair  $(T_{e_a}, T_{e_b})$ .

$$\tilde{\mathbf{T}}_e \sim \{\mathbf{E}'_a \subseteq \mathbf{E}(\mathbf{T}_{e_a}), \mathbf{E}'_b \subseteq \mathbf{E}(\mathbf{T}_{e_b})\}. \quad (17)$$

4) *Self-Supervised Loss with HetGDA*: The augmented graph data is subsequently employed in the self-supervised module as depicted in Figure 3 (b). Within the self-supervised context, the augmented graph data is treated as abnormal graphs when computing the supervised loss and passed through the same HRGCN architecture during the training process.

#### D. Training and Inference

By incorporating both SVDD and self-supervised modules into the model, the overall objective is defined as minimizing the following function:

$$\left( \frac{1}{n} \sum_{i=1}^n \|\phi(\mathcal{G}) - \mathbf{c}\|^2 + \frac{\lambda}{2} \sum_{\ell=1}^L \|\mathbf{W}^\ell\|_F^2 + \alpha \sum_{i=1}^n \mathcal{S}(\mathcal{G}, \tilde{\mathcal{G}}) \right). \quad (18)$$

where  $\mathcal{S}(\mathcal{G}, \tilde{\mathcal{G}})$  represents the self-supervised objective using the augmented graphs and labels, and  $\alpha$  represents the weight assigned to the self-supervised loss in the module. I.e.,  $\mathcal{S}(\mathcal{G}, \tilde{\mathcal{G}}) =$

$$-\frac{1}{N} \sum_{g \in \{\mathcal{G} \cup \tilde{\mathcal{G}}\}} (y \log(\phi(g)) + (1 - y) \log(1 - \phi(g))). \quad (19)$$

The training iteration comprises 3 steps, 1) compute SVDD loss from the graph representation and centre  $\mathbf{c}$  as described in Section III-A2 ; 2) compute self-supervised loss between input graph  $\mathcal{G}$  and augmented graph  $\tilde{\mathcal{G}}$  as Eq. 19; 3) sum SVDD and self-supervised losses adjusted by a weight parameter  $\alpha$  as equation 18. In the end, the anomaly score is obtained as below where output graph representation is denoted as  $\phi_1$  while the self-supervised prediction score is represented by  $\phi_2$ .

$$s(\mathcal{G}_i) = \|\phi_1(\mathcal{G}_i) - \mathbf{c}\|^2 * \phi_2(\mathcal{G}_i) \quad (20)$$

## IV. EXPERIMENTS

### A. Datasets

Despite its broad application domains, there are limited publicly accessible datasets for heterogeneous graph-level learning tasks. Our experiments use two real-world datasets with heterogeneous system and application API entities. Even though labels are provided with both datasets, an unsupervised learning setting where training data contains only normal graphs is conducted for all the experiments to mimic real-life application scenarios where abnormal data is difficult and/or costly to obtain.

**Train Ticket Graph Dataset (TraceLog)**. TraceLog is a large-scale heterogeneous graph dataset derived from a train ticket-based microservice system [34] for booking training tickets. It is widely used for evaluating anomaly detection in system trace log sequences [27], [35], [36]. TraceLog has four fault classes and 14 fault subclasses for different faulty types, such as asynchronous interaction, multi-instance,

configuration-related, and monolithic [9]. All 14 system faults are used as anomalies against normal system traces.

**System Flow Graph Dataset (FlowGraph)**. FlowGraph is a system flow graph dataset created and used in [7], where each graph consists of heterogeneous user web browsing activities, such as watching YouTube videos, downloading files and artifacts, opening *cnm.com* webpage, opening Gmail, and playing video games. The dataset contains 600 graphs, including one attack and five normal scenarios, with 100 graphs in each scenario. The attack scenario is categorised as drive-by download actions triggered by visiting unsecured URLs that exploit malicious software to gain root access.

A detailed introduction of the two datasets, including graph statistics and train/test configurations, can be found in Appendix A.

### B. Competing Methods and Evaluation Metrics

Our proposed model HRGCN is compared with four state-of-the-art graph-level anomaly detection models from diverse related research lines, including StreamSpot [7], DeepTraLog [9], GLocalKD [3], and one-class HetGCN (OCHetGCN) [8], [28].

**StreamSpot** is a seminal work on heterogeneous graph anomaly detection. It is a shallow method that uses an efficient hashing algorithm to encode each graph and defines anomaly score based on the similarity to graph clusters.

**DeepTraLog** is a recently proposed deep heterogeneous graph-level anomaly detection method, which leverages request logs and heterogeneous graph data to learn an anomaly detection model. The log data is first modelled with natural language processing (NLP) model to acquire log embeddings and then combined with attributed trace graph data. Since our problem setting does not involve additional log embeddings, the DeepTraLog model using the graph-structured data only is taken in our experiment.

**GLocalKD** is a recent graph-level anomaly detection model on homogeneous graphs. It is adapted to heterogeneous graph data by replacing its homogeneous GCN architecture with the HetGCN architecture [8], with all other components unchanged.

**OCHetGCN** is the baseline model we implement, in which the SVDD objective [28] is applied to the HetGCN model to learn a one-class HetGCN model for heterogeneous graph-level anomaly detection.

The implementation details of our method and the four competing methods are presented in Appendix B which could be used for reproducing the experiment results.

The model performance evaluation is carried out in two main categories, including the effectiveness of the model and the computational time cost. In terms of effectiveness, following the literature [7], [37], [38], two popular and complementary metrics, Area Under ROC Curve(AUC) and Average Precision (AP, also known as the area under the precision-recall curve), are used. Both metrics have a range of [0, 1]; a larger AUC (AP) indicates better performance.

TABLE II  
MODEL EVALUATION AND COMPARISON RESULTS.

Model	AUC		AP	
	TraceLog	FlowGraph	TraceLog	FlowGraph
OCHetGCN	0.628	0.952	0.518	0.954
GLocalKD	0.467	0.618	0.594	0.612
StreamSpot	0.634	<b>1.0</b>	0.517	0.992
DeepTraLog	0.657	0.928	0.498	0.821
HRGCN (Ours)	<b>0.864</b>	<b>1.0</b>	<b>0.754</b>	<b>1.0</b>

### C. Effectiveness on Real-world Datasets

The model performance comparison details are described in Table II, and the best-performed results are highlighted in bold. According to this table, most of the baseline models performed quite similarly, where StreamSpot performed well in the FlowGraph dataset and the Simplified DeepTraLog performed well in the TraceLog dataset among all the baseline models. Furthermore, our proposed model HRGCN outperforms all the baselines in both AUC and AP and on both datasets. The average relative improvements in % over the baseline models are around 20% for the TraceLog dataset and 16.5% for the FlowGraph dataset, respectively. It demonstrates that the proposed HRGCN framework is practical and produces better graph embeddings than the baselines.

### D. Computational Efficiency

Table III describes the training and predicting time cost in each model setup on both datasets. HetGCN took the least time to train a batch of instances on both TraceLog (batch of 8) and FlowGraph datasets (batch of 25). During predicting phase, StreamSpot performs the fastest in running a single instance on the TraceLog Dataset, while the Simplified DeepTraLog performs the fastest on the FlowGraph dataset. On the other hand, our proposed HRGCN model performs the slowest in both experiments. The HRGCN with Self-Supervised Module took 8.61 seconds and 54.71 seconds in training on TraceLog and FlowGraph datasets, respectively. However, during the predicting phase, HRGCN took only 1 second on TraceLog and 0.4 seconds on FlowGraph to run one instance. In addition, the training time for StreamSpot is skipped in this comparison because it uses a predefined list of bootstrap clusters and takes 0.072s and 0.006s to compute the initial cluster center for TraceLog and FlowGraph datasets. It should be noted that HRGCN demonstrates faster inference on larger graphs within the FlowGraph dataset compared to smaller graphs in the TraceLog dataset. Despite the significantly larger graph sizes in the FlowGraph dataset, it only consists of 16 source-destination node type pairs (2 source node types and 8 destination node types), whereas the TraceLog graphs contain 64 pairs (8 for both source and destination node types). This highlights the dependence of computational efficiency on the number of source-destination node type pairs. This situation could lead to a protracted training/inference phase when dealing with an extensive array of source-destination pairs within large-scaled graphs, unless effectively managed. One potential improvement involves increasing parallelism by

TABLE III  
MODEL TRAINING AND PREDICTING TIME RESULTS. TRAINING AND PREDICTING TIME ARE THE TIME SPENT RUNNING A SINGLE BATCH.

Model	Training Time		Inference Time	
	TraceLog	FlowGraph	TraceLog	FlowGraph
OCHetGCN	<b>0.78s</b>	<b>0.69s</b>	0.054s	0.188s
GLocalKD	0.182s	9.95s	0.054s	0.477s
StreamSpot	NA	NA	<b>0.007s</b>	6.212s
DeepTraLog	0.80s	0.727s	0.032s	<b>0.012s</b>
HRGCN (Ours)	8.61s	54.71s	1.057s	0.395s

processing different relation type pairs simultaneously before the final aggregation step.

### E. Ablation Study

The proposed HRGCN introduces additional graph hierarchies into the framework and improves graph representation learning. To understand the effectiveness of improving the model capability, we conduct ablation studies to evaluate the model performances of several variant settings, including a) baseline HetGCN model with introduced Heterogeneous Graph Data Augmentation (HetGDA); b) HRGCN model with only edge relation hierarchy, which removes source-destination relation hierarchy and HetGDA (HRGCN-ER); c) HRGCN model with only source-destination relation hierarchy, which removes edge relation hierarchy and HetGDA (HRGCN-SDR); d) HRGCN model with both edge and source-destination hierarchies but removes HetGDA (HRGCN-R2). Table IV describes the performance results of each variant model. The results of the plain HetGCN model and the proposed HRGCN model are also logged to the table to make it easier to compare.

Table IV demonstrates that the proposed HRGCN performs better than the models without considering the source-destination node type hierarchy. On the FlowGraph dataset, the HRGCN model with source-destination node type hierarchy (HRGCN-SDR) achieved optimal performance in both AUC and AP. On the TraceLog dataset, performance has improved significantly from the HRGCN-SDR model variant, which achieved around 25% boost in both AUC and AP (compared to the plain HetGCN model). HRGCN model with edge-type relation hierarchy (HRGCN-ER) has a limited effect over the FlowGraph dataset, which increases AUC by 3% while no improvement in AP. On the other hand, the HRGCN-ER slightly improves the performance on the TraceLog dataset by 6% and 2% in AUC and AP, respectively. Although both individual relation hierarchies made some improvements compared to the plain HetGCN, combining both relation hierarchies (HRGCN-R2) did not receive a significant performance boost. This is also reflected in the full HRGCN model performance. Nonetheless, these three model variants demonstrate performance trade-offs between AUC and AP, i.e., *HRGCN-SDR* performs the best in AUC but has lower AP; *HRGCN-R2* performs the best in AP but has lower AUC; *HRGCN(full)* shows a balanced performance of AUC and AP. In addition, applying Heterogeneous Graph Data Augmentation alone with the plain HetGCN model did not help the TraceLog data. Still,



TABLE IV  
MODEL ABLATION ANALYSIS ON MULTIPLE VARIANTS.

Model	AUC		AP	
	TraceLog	FlowGraph	TraceLog	FlowGraph
OCHetGCN (base)	0.628	0.952	0.518	0.954
OCHetGCN-HetGDA	0.573	<b>1.0</b>	0.452	<b>1.0</b>
HRGCN-ER	0.683	0.988	0.534	0.953
HRGCN-SDR	<b>0.871</b>	<b>1.0</b>	0.743	<b>1.0</b>
HRGCN-R2	0.858	<b>1.0</b>	<b>0.761</b>	<b>1.0</b>
HRGCN (full)	0.864	<b>1.0</b>	0.754	<b>1.0</b>

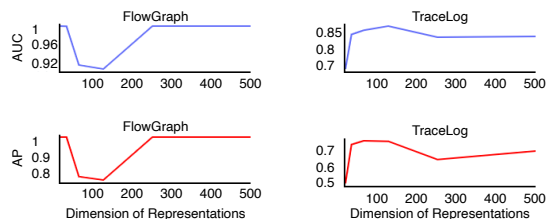


Fig. 4. Results of Varying Representation Dimensions.

it significantly improved the model performance by 5% when applied to the FlowGraph dataset.

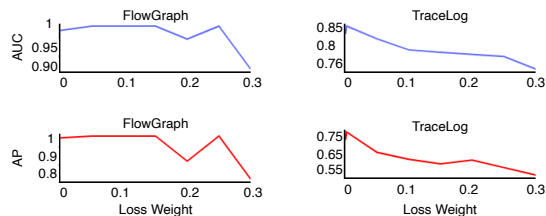


Fig. 5. Impact of Self-Supervised Loss in HRGCN Model.

## F. Hyperparameter Analysis

Hyperparameters play an essential role when training HRGCN on different datasets, such as the size of the hidden layer and the weights of self-supervised loss.

Figure 4 illustrates the impact of the representation dimension sizes. Both evaluation metrics are more consistent when the dimension size increases or is set to a relatively smaller number. From the FlowGraph dataset, the model returns a more consistent performance when the hidden layer size is larger than 256. In comparison, the model performs the best when the dimension size is set to 128 on the TraceLog dataset.

The results of varying self-supervised loss weights is shown in Figure 5. On the FlowGraph dataset, the model achieves the optimal performance when the loss weight is smaller than 0.15, while the TraceLog dataset requires a narrower range between 0 and 0.01.

## V. INDUSTRIAL CASE STUDY

### A. Problem Statement

This section will introduce the details of the problem we try to tackle in real-world industrial settings, including the problem definition, hypothesis, and the dataset used for conducting this study.

1) *Problem Definition*: A typical network behaviour can be described in combination with entities of Users, Cells, and Sectors. Users are the objects that perform or receive certain network actions that happen to themselves. Cells are the most granular network devices (in this scenario) that serve certain network actions to individual users. Sectors describe a collection of Cell devices and contain the logical grouping information among every Cell device (see Figure 10 (a) in Appendix C for a simplified example of how the network graph would look).

Serving good connections and networks to customers is always one of the goals of a telecommunication company. Therefore, in this study, we are trying to detect if any anomalous behaviours are happening within the network that could impact user’s experience. In this case, we can identify the specific network devices and times they happened and provide a means for monitoring and post-event analysis of the network behaviour.

Given a collection of the network event data  $\mathcal{D} \sim \{n_0, n_1, \dots, n_n\}$  that happened in a certain period  $T$ . We expect that there exist certain relations within the network events which could be modelled by a differentiable function  $\mathcal{S}$ , which outputs an anomaly score  $P$  for the input network events. Note that an anomaly score of 1 means highly anomalous, and 0 indicates no detected anomalous patterns:

$$P(1|d, t) = \mathcal{S}(d, t), \{d \in \mathcal{D}; t \in T\} \quad (21)$$

2) *Trace Event Dataset*: The Trace Event Dataset is populated based on relevant mobile event network data at TPG Telecom in Australia. We have simulated and resampled the new Trace Event Dataset with similar data statistics and properties so that this could reflect closely to real-world events.

To keep an appropriate data size to conduct a case study, we have sampled 20% of simulated network event data in a specific period. The sampled network data contains in total of 79,446,610 raw network events. Each network event can be categorised into one of 30 individual network event codes, and the frequency of each event that happened throughout the day could vary significantly. For example, the most frequent event could occur 49 million times, while the least frequent event may occur only two times. Moreover, some graphs may contain many vertices in a certain node and neighbourhood type (see Figure 10 (b) in Appendix C for an example of how the relation between User and Cell would look like).

After processing the raw data into graph structures, we formed the dataset containing 21,849 traces in formatted graph data around the network events that happened to users on each network device to track the Sector behaviours. There are 876 abnormal graphs in the dataset, with around 4% occurrence

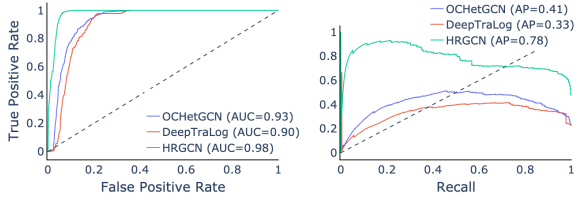


Fig. 6. Testing Results on the Industrial Trace Event Data

in this dataset. Besides, the average-sized traces consist of around 478 nodes and 477 edges individually. The dataset has three node types: User, Cell, and Sector. There are around 2.5 million user nodes, 26 thousand distinct Cell nodes, and 9 thousand distinct Sector nodes. Because the dataset is formed based on different periods, the total number of nodes is around 10 million. A detailed summary of the data statistics can be found in Table VI in Appendix C.

Furthermore, although we are conducting the unsupervised training methodology, we still need predefined labels to evaluate the model performances. In this case, we also introduce SME (Subject Matter Expert) knowledge onto the Trace Event Dataset to help with the evaluation purposes.

The following experiment divides the dataset into training, validation, and testing sets. The training and validation set contains 60% and 20% of normal traces, and the rest normal and abnormal traces are grouped into the testing set.

### B. Experiment & Results

Figure 6 illustrates the model performance comparison on the test dataset. The proposed HRGCN model outperforms the other two models by 5% in AUC and 37% in AP. We found that the model could converge at an early stage and performed stably for the rest of the training iterations.

To understand the effectiveness of the produced model, we have also applied the model to an unseen set of traces for the network event. In this section, we'll dive deep into the model outputs and the reflections from the results.

1) *Trace Representation & Anomaly Scores*: Since the model's output produces vectors in 300 dimensions, we applied additional post-processing to the model outputs to visualise the trace representation onto a 2D plane which projects each trace representation into lower dimensions. Figure 7 illustrates the scatter view of each trace in test sets, and it shows that most abnormal trace representations are closely grouped into a simple space throughout the 2D plane. . most normal traces have a score very close to 0.0025, while most abnormal traces have a score of 0.15.

2) *Network Analysis*: Figure 8 (a) and (b) illustrate a simplified network structure of how a normal/abnormal trace would look like. Although the trace graph structure may be similar, individual node features may be evolved over time and certain devices may become anomalous.

To investigate deeper onto the network analysis around the device performances and experiences, we have sampled a collection of devices and started to look closer at their

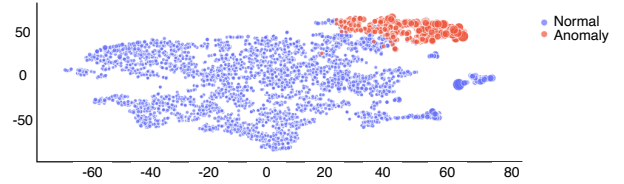


Fig. 7. t-SNE visualisation of graph representation on the unseen trace data.

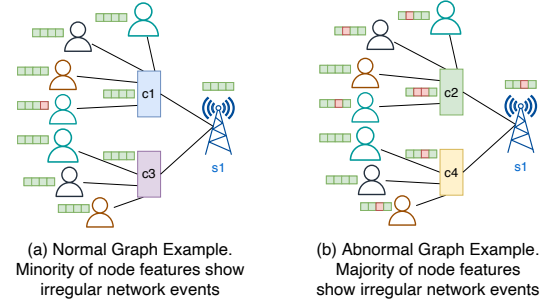


Fig. 8. Example Anomalies in Trace Event Graph. (a) and (b) demonstrate simplified examples of normal/abnormal graphs. Individual node features may reflect on the overall performance of the network device. In simple words, high volume of users received irregular network events demonstrates there is a congested network device, low volume of users received irregular events may not indicate a malfunctioning device.

results for monitoring and alerting purposes. To do so, we have rearranged the view to demonstrate the anomaly score changes within certain sequential time windows. Figure 9 demonstrates that for the selected abnormal devices that experienced congested behaviours. From the results, we can see that the model is adaptive to behaviour changes on individual devices over time and can detect and identify intermittent normal/abnormal behavioural changes throughout network devices.

## VI. CONCLUSION

In this work, we propose an unsupervised model, namely Hierarchical Relation-augmented Heterogeneous Graph Convolutional Neural Network (HRGCN) to improve the representation learning task on heterogeneous graph (HG) data. We improve the effectiveness of modelling additional relations between edges and source-destination node neighbourhoods to discover deep latent information within the complex heterogeneous graph networks without labels. Furthermore, a complementary self-supervised learning module and associated Heterogeneous Graph Data Augmentation (HetGDA) method are also introduced to enhance the model's performance during training. Extensive experiments on multiple system-focused

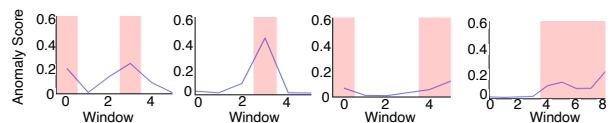


Fig. 9. Anomaly Scores Over Time on Abnormal Devices

anomaly detection datasets demonstrate that the HRGCN model can significantly outperform state-of-the-art methods. In addition, the impact of each proposed module, i.e., the relation-hierarchical module and self-supervised module, is unravelled to demonstrate the performance boosts.

On the other hand, an industrial case study is conducted over a real-world industrial setting in telecommunication. We demonstrate through the case study that our proposed HRGCN model can work effectively in real-world settings and significantly outperforms current state-of-the-art models.

#### ACKNOWLEDGMENT

This research is supported by TPG Telecom Australia. Any conclusions drawn in this study are of the authors and do not necessarily reflect the views and perspectives of TPG Telecom.

#### REFERENCES

- [1] G. Pang, C. Shen, L. Cao, and A. V. D. Hengel, "Deep learning for anomaly detection: A review," *ACM computing surveys (CSUR)*, vol. 54, no. 2, pp. 1–38, 2021.
- [2] B. Lindemann, B. Maschler, N. Sahlab, and M. Weyrich, "A survey on anomaly detection for technical systems using lstm networks," *Computers in Industry*, vol. 131, p. 103498, 2021.
- [3] R. Ma, G. Pang, L. Chen, and A. van den Hengel, "Deep graph-level anomaly detection by glocal knowledge distillation," in *The Fifteenth ACM International Conference on Web Search and Data Mining*, 2022.
- [4] G. Zhang, Z. Yang, J. Wu, J. Yang, S. Xue, H. Peng, J. Su, C. Zhou, Q. Z. Sheng, L. Akoglu *et al.*, "Dual-discriminative graph neural network for imbalanced graph-level anomaly detection," in *Advances in Neural Information Processing Systems*, 2022.
- [5] C. Niu, G. Pang, and L. Chen, "Graph-level anomaly detection via hierarchical memory networks," *arXiv preprint arXiv:2307.00755*, 2023.
- [6] S. Tang, L. Jin, and F. Cheng, "Fraud detection in online product review systems via heterogeneous graph transformer," *IEEE Access*, vol. 9, pp. 167 364–167 373, 2021.
- [7] E. Manzoor, S. M. Milajerdi, and L. Akoglu, "Fast memory-efficient anomaly detection in streaming heterogeneous graphs," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 1035–1044.
- [8] C. Zhang, D. Song, C. Huang, A. Swami, and N. V. Chawla, "Heterogeneous graph neural network," in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 793–803.
- [9] C. Zhang, X. Peng, C. Sha, K. Zhang, Z. Fu, X. Wu, Q. Lin, and D. Zhang, "Deeptralog: Trace-log combined microservice anomaly detection through graph-based deep learning," 2022.
- [10] X. Wang, H. Ji, C. Shi, B. Wang, Y. Ye, P. Cui, and P. S. Yu, "Heterogeneous graph attention network," in *The world wide web conference*, 2019, pp. 2022–2032.
- [11] X. Fu, J. Zhang, Z. Meng, and I. King, "Magnn: Metapath aggregated graph neural network for heterogeneous graph embedding," in *Proceedings of The Web Conference 2020*, 2020, pp. 2331–2341.
- [12] Z. Hu, Y. Dong, K. Wang, and Y. Sun, "Heterogeneous graph transformer," in *Proceedings of the web conference 2020*, 2020, pp. 2704–2710.
- [13] J. Zhao, X. Wang, C. Shi, Z. Liu, and Y. Ye, "Network schema preserving heterogeneous information network embedding," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- [14] X. Sun, H. Yin, B. Liu, H. Chen, J. Cao, Y. Shao, and N. Q. Viet Hung, "Heterogeneous hypergraph embedding for graph classification," in *Proceedings of the 14th ACM international conference on web search and data mining*, 2021, pp. 725–733.
- [15] X. Wang, D. Bo, C. Shi, S. Fan, Y. Ye, and S. Y. Philip, "A survey on heterogeneous graph embedding: methods, techniques, applications and sources," *IEEE Transactions on Big Data*, 2022.
- [16] Y. Dong, N. V. Chawla, and A. Swami, "metapath2vec: Scalable representation learning for heterogeneous networks," in *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining*, 2017, pp. 135–144.
- [17] C. Shi, B. Hu, W. X. Zhao, and S. Y. Philip, "Heterogeneous information network embedding for recommendation," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 2, pp. 357–370, 2018.
- [18] B. Hu, C. Shi, W. X. Zhao, and P. S. Yu, "Leveraging meta-path based context for top-n recommendation with a neural co-attention model," in *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 2018, pp. 1531–1540.
- [19] B. Hu, Y. Fang, and C. Shi, "Adversarial learning on heterogeneous information networks," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 120–129.
- [20] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE transactions on neural networks*, vol. 20, no. 1, pp. 61–80, 2008.
- [21] W. Fan, Y. Ma, Q. Li, Y. He, E. Zhao, J. Tang, and D. Yin, "Graph neural networks for social recommendation," in *The world wide web conference*, 2019, pp. 417–426.
- [22] X. Yue, Z. Wang, J. Huang, S. Parthasarathy, S. Moosavinasab, Y. Huang, S. M. Lin, W. Zhang, P. Zhang, and H. Sun, "Graph embedding on biomedical networks: methods, applications and evaluations," *Bioinformatics*, vol. 36, no. 4, pp. 1241–1251, 2020.
- [23] Q. Zhang, K. Yu, Z. Guo, S. Garg, J. J. Rodrigues, M. M. Hassan, and M. Guizani, "Graph neural network-driven traffic forecasting for the connected internet of vehicles," *IEEE Transactions on Network Science and Engineering*, vol. 9, no. 5, pp. 3015–3027, 2021.
- [24] J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun, "Graph neural networks: A review of methods and applications," *AI Open*, vol. 1, pp. 57–81, 2020.
- [25] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel, "Gated graph sequence neural networks," *arXiv preprint arXiv:1511.05493*, 2015.
- [26] S. Nedelkoski, J. Cardoso, and O. Kao, "Anomaly detection from system tracing data using multimodal deep learning," in *2019 IEEE 12th International Conference on Cloud Computing (CLOUD)*. IEEE, 2019, pp. 179–186.
- [27] P. Liu, H. Xu, Q. Ouyang, R. Jiao, Z. Chen, S. Zhang, J. Yang, L. Mo, J. Zeng, W. Xue *et al.*, "Unsupervised detection of microservice trace anomalies through service-level deep bayesian networks," in *2020 IEEE 31st International Symposium on Software Reliability Engineering (ISSRE)*. IEEE, 2020, pp. 48–58.
- [28] L. Ruff, R. Vandermeulen, N. Goernitz, L. Deecke, S. A. Siddiqui, A. Binder, E. Müller, and M. Kloft, "Deep one-class classification," in *International conference on machine learning*. PMLR, 2018, pp. 4393–4402.
- [29] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *International conference on machine learning*. PMLR, 2017, pp. 1263–1272.
- [30] K. Ding, Z. Xu, H. Tong, and H. Liu, "Data augmentation for deep graph learning: A survey," *arXiv preprint arXiv:2202.08235*, 2022.
- [31] Y. You, T. Chen, Y. Shen, and Z. Wang, "Graph contrastive learning automated," in *International Conference on Machine Learning*. PMLR, 2021, pp. 12 121–12 132.
- [32] Y. You, T. Chen, Y. Sui, T. Chen, Z. Wang, and Y. Shen, "Graph contrastive learning with augmentations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 5812–5823, 2020.
- [33] Y. Zhu, Y. Xu, F. Yu, Q. Liu, S. Wu, and L. Wang, "Graph contrastive learning with adaptive augmentation," in *Proceedings of the Web Conference 2021*, 2021, pp. 2069–2080.
- [34] X. Zhou, X. Peng, T. Xie, J. Sun, C. Ji, W. Li, and D. Ding, "Fault analysis and debugging of microservice systems: Industrial survey, benchmark system, and empirical study," *IEEE Transactions on Software Engineering*, vol. 47, no. 2, pp. 243–260, 2018.
- [35] X. Zhou, X. Peng, T. Xie, J. Sun, C. Ji, D. Liu, Q. Xiang, and C. He, "Latent error prediction and fault localization for microservice applications by learning from system trace logs," in *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2019, pp. 683–694.
- [36] N. Zhao, J. Chen, Z. Yu, H. Wang, J. Li, B. Qiu, H. Xu, W. Zhang, K. Sui, and D. Pei, "Identifying bad software changes via multimodal anomaly detection for online service systems," in *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2021, pp. 527–539.

- [37] G. Pang, A. van den Hengel, C. Shen, and L. Cao, "Toward deep supervised anomaly detection: Reinforcement learning from partially labeled anomaly data," in *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, 2021, pp. 1298–1308.
- [38] H. Qiao and G. Pang, "Truncated affinity maximization: One-class homophily modeling for graph anomaly detection," *arXiv preprint arXiv:2306.00006*, 2023.

## APPENDIX

### A. TraceLog and FlowGraph Dataset Overview

The detailed statistics of the TraceLog and FlowGraph datasets are summarised in Table V.

TABLE V  
TRACELOG AND FLOWGRAPH DATASET OVERVIEW

TraceLog	FlowGraph
#node-types: 8	#node-types: 8
#edge-types: 4	#edge-types: 26
#node-attributes: 7	#node-attributes: 26
#graphs: 132,485	#graph: 600
#anomaly-graph: 23,334	#anomaly-graph: 100
#mean-size: 205	#mean-size: 8,411
#mean-edges: 224	#mean-edges: 12,730
#mean-node-types: 8	#mean-node-types: 8
#mean-edge-types: 2.87	#mean-edge-types: 25.5

1) *Details of TraceLog Dataset:* Specifically, the dataset has 132,485 trace graphs, including 23,334 anomalous graphs and 109,151 normal graphs. 60% and 20% normal graphs are used as training and validation datasets. The remaining normal graphs are combined with abnormal graphs as the test data. The dataset contains 8 node types and 4 edge types in total. On average, each graph consists of around 205 nodes, 224 edges, 8 node types and 2.87 edge types.

2) *Details of FlowGraph Dataset:* There are 8 node types and 26 edge types in total. On average, each graph contains around 8,411 nodes, 12,730 edges, 8 node types, and 25.5 edge types. FlowGraph is also split into training, validation, and testing. Following [7], 60% and 15% normal graphs are used for training and validation, and the rest is combined for the testing set.

### B. Implementation Details

The following hyperparameters are used in the proposed HRGCN. On the TraceLog dataset, the learning rate is 0.0001, the representation dimension is 300, the number of hidden layers is 2, the self-supervised loss weight is 0.001, and the batch size is 8. In addition, the parameters used for heterogeneous graph data augmentations are: edge perturbation probability is 0.84, edge replacement percentage is 0.13, swap node type percentage is 0.1, and swap edge type percentage is 0.17. On the FlowGraph dataset, the learning rate is 0.01, the representation dimension is 32, the number of hidden layers is 2, the self-supervised loss weight is 0.21, the batch size is 25, the edge replacement percentage is 0.39, and swap node type percentage is 0.52. During the training process, the model is trained on the training set, and tuned on the validation set using the self-supervised classification accuracy and SVDD distance. A network architecture of the same depth and breadth is used in all deep competing methods. The implementation of StreamSpot is taken from its authors.

The experiment is conducted in AWS using SageMaker Training Jobs<sup>1</sup>. The instance type used for training is g4dn.2xlarge, which has 8 vCPU, 32 GiB memory, and 1 NVIDIA T4 GPU.

<sup>1</sup><https://aws.amazon.com/sagemaker/>

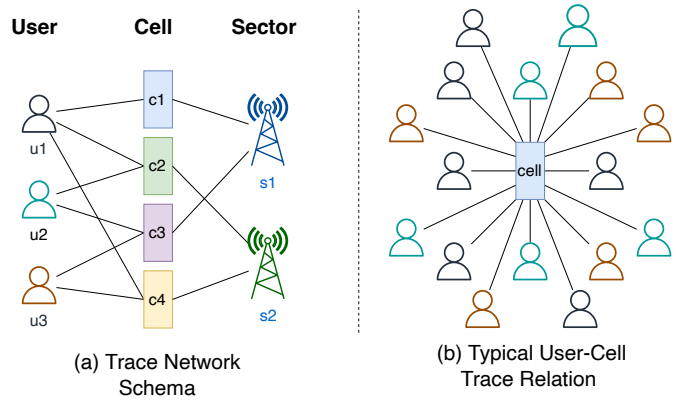


Fig. 10. Trace Event Graph Example and Schema. (a) demonstrates the basic relationships among User, Cell, and Sector. (b) illustrates how a typical Cell would look like in the abstract when there is a lot of network traffic.

### C. Case Study – Overview of the Trace Event Dataset

The properties and statistics of the case study dataset is summarised in Table VI.

An example of trace event graph example and schema is presented in Figure 10.

TABLE VI  
CASE STUDY DATASET PROPERTIES AND STATISTICS

Trace	Node	Edge
#traces: 21,849	#node-types: 3	#edge-types: 30
#normal: 20,973	#total: 10,423,149	#total: 10,423,149
#abnormal: 876	#user: 2,597,531	
mean-size: 478	#cell: 26,761	
#mean-edges: 477	#sector: 8,724	

### D. Case Study Implementation details

We have selected a set of hyperparameters for the Trace Event Datasets based on the previous experiments on the FlowGraph and TraceLog datasets. The learning rate is set to 0.01, the hidden layer size and output embedding size are 300, the number of hidden HetGCN layers is 2, and the batch size is 256. The experiments are conducted in the same environment as that on the other two datasets.