



HAL
open science

Few Edges Are Enough: Few-Shot Network Attack Detection with Graph Neural Networks

Tristan Bilot, Nour El Madhoun, Khaldoun Al Agha, Anis Zouaoui

► **To cite this version:**

Tristan Bilot, Nour El Madhoun, Khaldoun Al Agha, Anis Zouaoui. Few Edges Are Enough: Few-Shot Network Attack Detection with Graph Neural Networks. The 19th International Workshop on Security (IWSEC 2024), Sep 2024, Kyoto, Japan. pp.257–276, <10.1007/978-981-97-7737-2_15>. <hal-04701979>

HAL Id: hal-04701979

<https://hal.science/hal-04701979v1>

Submitted on 19 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Few Edges Are Enough: Few-Shot Network Attack Detection with Graph Neural Networks

Tristan Bilot^{1,2,3}[0000-0003-4153-735X], Nour El Madhoun^{3,4}[0000-0001-7742-7748], Khaldoun Al Agha¹[0000-0002-6595-5866], and Anis Zouaoui²[0009-0001-0273-4601]

¹ Université Paris-Saclay, CNRS, Laboratoire Interdisciplinaire des Sciences du Numérique, Gif-sur-Yvette, France tristan.bilot@universite-paris-saclay.fr
alagha@liscn.fr

² Iriguard, Puteaux, France anis.zouaoui@adservio.fr

³ LISITE Laboratory, ISEP (Institut Supérieur d'Electronique de Paris), Issy-les-Moulineaux, France nour.el-madhoun@isep.fr

⁴ Sorbonne Université, CNRS, LIP6, Paris, France

Abstract. Detecting cyberattacks using Graph Neural Networks (GNNs) has seen promising results recently. Most of the state-of-the-art models that leverage these techniques require labeled examples, hard to obtain in many real-world scenarios. To address this issue, unsupervised learning and Self-Supervised Learning (SSL) have emerged as interesting approaches to reduce the dependency on labeled data. Nonetheless, these methods tend to yield more anomalous detection algorithms rather than effective attack detection systems. This paper introduces *Few Edges Are Enough* (FEAE), a GNN-based architecture trained with SSL and Few-Shot Learning (FSL) to better distinguish between false positive anomalies and actual attacks. To maximize the potential of few-shot examples, our model employs a hybrid self-supervised objective that combines the advantages of contrastive-based and reconstruction-based SSL. By leveraging only a minimal number of labeled attack events, represented as attack edges, FEAE achieves competitive performance on two well-known network datasets compared to both supervised and unsupervised methods. Remarkably, our experimental results unveil that employing only 1 malicious event for each attack type in the dataset is sufficient to achieve substantial improvements. FEAE not only outperforms self-supervised GNN baselines but also surpasses some supervised approaches on one of the datasets.

Keywords: Attack Detection · Network Security · Few-shot Learning · Self-Supervised Learning · Graph Neural Networks.

1 Introduction

In our interconnected and digitalized world, cybersecurity is of utmost importance, impacting our society and technological infrastructure. As technology advances rapidly, cyber threats become more sophisticated, emphasizing the need for strong defense systems to combat these evolving challenges.

With the recent advancements in Machine Learning and Deep Learning, researchers are increasingly turning to these techniques to reinforce computer systems against complex attacks. The abundance of data in network environments makes Deep Learning an attractive approach to uncover hidden patterns within malicious communications occurring in the network. The emergence of Graph Deep Learning and Graph Neural Networks (GNNs) has generated significant interest in fields where interconnected data is prevalent. Cybersecurity researchers have notably adopted these graph-based models to detect attacks in network data [2,21,25], which inherently possess interconnected structures. Consequently, multiple variants of message-passing GNNs [11] have been successfully applied to diverse network datasets, including network flows, packets and authentication logs, to detect specific attacks such as DDoS attacks [5, 18] and Botnet attacks [21, 38], relying heavily on labeled training examples. However, the reliance on numerous labeled attacks presents significant challenges for real-world scenarios, as reliable labels are often difficult to obtain from enterprise networks. Furthermore, training a model to detect specific attacks in a supervised manner inherently limits its ability to generalize to unseen attacks [8]. Consequently, researchers are increasingly exploring unsupervised and self-supervised techniques that do not require pre-labeled attacks [2]. While these methods are effective at detecting anomalies and unusual events, their application to attack detection can result in a high number of false positives, particularly in highly imbalanced real-world data. The model might struggle to distinguish between an actual attack and an unprecedented but benign action, such as a first-time SSH connection between two hosts, in the absence of labeled attack data for training.

Therefore, we encounter a dual trade-off between the complexity of obtaining a multitude of labeled attacks and the requirement of having some labels to effectively distinguish between anomalies and attacks. To address this trade-off, we propose a viable solution using Few-Shot Learning (FSL), employing a minimal number of labeled examples in conjunction with the Self-Supervised Learning (SSL) loop. Motivated by this idea, we introduce *Few Edges Are Enough (FEAE)*, a GNN-based detection system that leverages FSL to more accurately identify attacks while relying on only a very small number of labeled examples. FEAE employs contrastive learning along with a few-shot aware reconstruction-based objective, allowing it to cluster similar unlabeled attack edges while requiring only few malicious edges.

This paper introduces the following contributions:

- We propose FEAE, which, to the best of our knowledge, is the first approach that leverages GNNs with few-shot learning and hybrid SSL for network attack detection. Other similar works typically rely on either supervised, self-supervised, or fully unsupervised methods.
- We introduce a few-shot aware reconstruction loss that successfully integrates the knowledge of few-shot malicious samples within the self-supervised training of the GNN encoder. This capability allows the model to directly cluster various attack families within the embedding space, enabling the downstream decoder to distinctly separate them from benign activities.

- We evaluate the detection capabilities of FEAE on two well-known network datasets and across a variety of attack families.

The main achievement of this research reveals that the performance of fully-supervised methods can be approached or even exceeded by using merely one few-shot malicious edge per attack family. This result is encouraging as it reduces the need for extensive labeling and reliance on noisy dataset labels.

The paper is organized as follows: Section 2.2 presents some background knowledge and related works. Section 3 introduces the architecture of FEAE. Section 4 presents the model configurations, used datasets and baselines, along with a performance analysis of FEAE. Section 5 benchmarks the performance of the model compared to the baselines and discusses future research directions, whereas Section 6 concludes this work.

2 Background

We introduce the background knowledge and essential definitions in Section 2.1, and present related works in Section 2.2.

2.1 Definitions

Network Graph. Networks can be inherently represented using graph structures. We define a network graph as a type of graph where nodes represent hosts characterized by their IP addresses, and edges represent network flows, typically associated with features obtained through monitoring tools. Mathematically, a graph is represented as $G = (V, E)$, where V denotes the set of nodes (or hosts), and E denotes the set of edges (or events) that connect pairs of nodes within the set V . Network graphs are represented in memory as adjacency lists of shape $(2, |E|)$, constructed from netflow logs by creating an edge between any pair of connected nodes and integrating optional flow features as edge attributes.

Graph Neural Networks. Graph Neural Networks (GNNs) [27, 34] are a class of neural networks designed to operate on graph-structured data. They leverage the graph structure along with integrated features to learn robust structural representations of nodes and edges. GNNs iteratively aggregate and transform node features from their neighbors, enabling tasks such as the detection of malicious nodes and edges in network graphs [2, 37], provenance graphs [8, 14] or code-based graphs [3, 4, 9].

Graph Self-Supervised Learning. Graph Self-Supervised Learning (Graph SSL) aims to learn node and edge representations using the graph structure itself as label, eliminating the need for manually labeled data [20]. A prominent approach is graph contrastive learning [1, 36], which generates multiple augmented views $\{G_i\}$ of a graph G and trains a model f to maximize the mutual information between the representations $f(G_i)$ and $f(G_j)$ of these views for positive pairs (i, j) , while minimizing it for negative pairs generated through augmentation methods.

Few-Shot Learning. Few-shot learning (FSL) aims to train models that perform well with only a limited number of labeled instances per class [29,33]. Given a support set $S = \{(x_i, y_i)\}_{i=1}^k$ with k labeled examples per class, the goal is to accurately classify instances in a query set Q . Unlike traditional supervised learning, which requires a large labeled dataset, FSL leverages the few available examples to generalize across new tasks.

2.2 Related Works

Since the emergence of Graph Neural Networks (GNNs) across various disciplines, these methods have demonstrated significant efficacy in the domain of network attack detection. Consequently, there has been a substantial increase in research efforts dedicated to this area.

E-GraphSAGE [22] is a model inspired by GraphSAGE [12], specifically designed for edge-level tasks on network graphs. The message-passing function has been adapted to incorporate edge features before updating the node embeddings, with edge embeddings obtained by concatenating connected pairs of nodes. Due to its simplicity and remarkable performance, this model has become a fundamental building block for numerous other works.

In another study [17], authors introduced an improvement to E-GraphSAGE by implementing a pre-sampling step before training. This technique reduces graph size, enhancing scalability, and has shown slight performance improvements on the UNSW-NB15 dataset [23].

Another innovative approach, E-ResGAT [7], operates on a line graph instead of an edge-attributed graph. Consequently, the detection task transforms into a node classification problem, where edges and their features are converted into nodes. The GAT model [31] serves as the GNN encoder to derive node embeddings, and residual connections are incorporated for propagating original edge features. To achieve greater scalability, the authors propose a neighbor sampling strategy, which uniformly samples neighboring nodes during aggregation.

Beyond these advancements, E-GraphSAGE has also been adapted into a self-supervised method known as Anomal-E [6]. Anomal-E utilizes the same graph structure but does not require any labels during the training of the GNN encoder. Deep Graph Infomax (DGI) [32] is employed to generate embeddings for positive and negative graph samples, maximizing local-global mutual information by comparing embeddings to a summarized version of the graph. Negative graphs are created using a corruption function, involving shuffling all edges. E-GraphSAGE functions as the encoder and is trained in a self-supervised manner to distinguish between positive and negative edges. Following training, edge embeddings are classified using an Isolation Forest (IF) [19], using all benign embeddings as the training set. We elaborate further in Section 3 on the challenges posed by the requirement of all these benign labels in the classifier and how FSL can mitigate this problem.

Authors in [30] introduce a few-shot malicious IoT traffic detection system based on a GNN. This method visualizes binary traffic flows as color images and extracts features of nodes using a GCN model [16]. By employing FSL, the

model is trained to generalize and detect new types of attacks with minimal labeled samples. The evaluation demonstrates that this approach achieves high F1 scores on a dataset containing benign traffic and six different types of attack, outperforming traditional detection methods.

3 Design of FEAE

In this section, we provide a comprehensive overview of our proposed FEAE model. Section 3.1 outlines the challenges associated with current models and introduces the innovative solutions implemented in FEAE. Section 3.2 offers a comparative analysis with other state-of-the-art approaches. Sections 3.3 to 3.6 detail the technical design and implementation of the model.

3.1 Intuition

Our proposed framework leverages a GNN-based SSL method enhanced with FSL for the detection of network attacks, such as Distributed Denial of Service (DDoS), brute force, and botnets. The primary objective of our framework is to achieve high precision in detecting these attacks using minimal labeling. To this end, the input network flow data are first converted into a large graph representation, where nodes correspond to hosts, identified through their IP addresses, and edges represent the network flows. This graph is enriched with additional flow features, such as the count of packets within a Netflow record and the mean packet size, represented as a vector of edge features.

FEAE comprises three key components illustrated in Figure 1: (i) the **GNN encoder**, (ii) the **SSL module**, and (iii) the **few-shot decoder**.

The GNN encoder, presented in Section 3.4, is designed to compute node and edge embeddings that capture the intrinsic relationships between hosts by leveraging the graph topology and the flow features within the network. It operates by training in a self-supervised manner through the SSL module, which ensures that the generated embeddings preserve the original network structure and semantics.

The SSL module, introduced in Section 3.5, employs hybrid self-supervised strategies, making use of both contrastive-based and reconstruction-based objectives. The contrastive objective is designed to differentiate between original (positive) and augmented (negative) edges, while the reconstruction objective, combined with few-shot samples, aims to produce dissimilar embeddings for benign and malicious edges identified in the few-shot context, also called malicious few-shot edges. The training of the encoder is conducted end-to-end with the SSL module, enabling the generation of embeddings that are aware of the topology and features associated to malicious activity.

Lastly, the few-shot decoder is trained separately on the learned edge embeddings. It uses a supervised learning approach focused on edge classification and takes advantage of the few-shot labels related to network attacks. The design of the decoder is explained in Section 3.6.

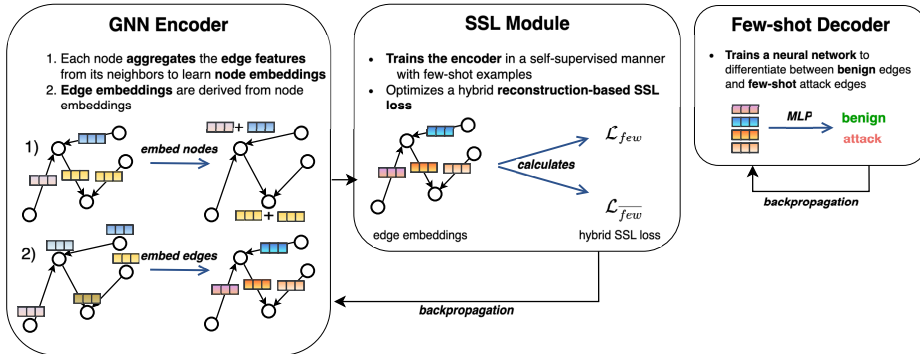


Fig. 1: Overall architecture of FEAE.

3.2 Comparison with Other Works

Other SSL techniques, such as Anomal-E [6], necessitate the inclusion of both benign and attack samples to train the GNN encoder, and exclusively benign samples to train the Isolation Forest decoder. This indirectly leads to a supervised learning paradigm as it requires prior identification of which samples are benign or malicious. In contrast, our few-shot approach requires only a minimal number of malicious samples, thereby obviating the need for identifying benign samples beforehand.

Moreover, fully self-supervised or unsupervised approaches primarily serve as anomaly detection mechanisms rather than cyberattack detection systems [2, 10, 15, 24]. The lack of training labels prevents these models from accurately differentiating between legitimate cyberattacks and false positive anomalies. Conversely, models trained solely under supervised conditions may become excessively tailored to the attack patterns observed during training, limiting their generalizability to unseen attacks. Therefore, FSL is proposed as a balanced intermediary, merging the advantages of both paradigms to enhance the precision in detecting diverse network attacks.

3.3 Notations

The subsequent sections of this paper will elaborate on each of the components within FEAE using the scientific notations summarized in Table 1.

3.4 GNN Encoder

In this paper, we propose a simple and lightweight GNN encoder to perform message-passing between nodes in the network graph. Our proposed GNN encoder first computes node embeddings by leveraging the flow numerical flow features attributed to neighboring edges. Precisely, each node aggregates its neighboring edge features in such a way that both the local neighborhood topology

Table 1: Notations

E	Training edges
X	Training edges' features
k	Number of selected few-shot edges in each attack family
m	Number of attack families
\mathcal{E}	Few-shot edges (benign and malicious)
\mathcal{E}_{mal}	Few-shot edges (only malicious)
$Y_{\mathcal{E}}$	Few shot edges' labels
\mathbf{H}	Positive edge embeddings matrix
$\tilde{\mathbf{H}}$	Negative edge embeddings matrix

and the flow features are captured. Formally, we compute the aggregation of neighboring edges such that:

$$h_{\mathcal{N}(u)} = \sum_{v \in \mathcal{N}(u)} e_{uv}, \quad u \in N, \quad (1)$$

where e_{uv} denotes the feature vector of the edge (u, v) , $\mathcal{N}(u)$ represents the neighboring nodes of node u , and $h_{\mathcal{N}(u)}$ represents the sum aggregation u 's neighboring edges. We employ here a sum aggregation as it offers interesting injective capabilities. As detailed in the original Graph Isomorphism Network (GIN) paper [35], the sum aggregation is characterized by its injectivity concerning node features, meaning that a unique combination of features will produce a distinctive sum. Through the use of a similar sum aggregation on the neighboring edge features, the model effectively preserves a greater amount of the structural features within the graph when contrasted with other aggregations like the mean or maximum.

Following the aggregation of node features, these features are then processed by a linear layer with activation function, as described by the following formula:

$$h_u = \sigma(h_{\mathcal{N}(u)} \mathbf{W}_{\text{agg}}), \quad (2)$$

where h_u represents the embedding of node u , σ is the ReLU activation function, and \mathbf{W}_{agg} is a trainable weight matrix.

Following the calculation of node embeddings, our objective is to generate specific edge embeddings from these node representations. These edge embeddings are intended for the subsequent identification of network attacks, characterized here by malicious edges. To achieve this, embeddings from each pair of connected nodes are concatenated. This concatenated output is then multiplied by a distinct trainable matrix, which aims to learn edge embeddings given the concatenated source and destination node embeddings.

$$h_{uv} = [h_u, h_v] \mathbf{W}_{\text{edge}}, \quad (3)$$

where $[,]$ designates the concatenation operation and \mathbf{W}_{edge} is the edge-level trainable matrix.

Our experiments have shown that a single-layer of this GNN encoder yields superior results in learning representations from positive and negative edges during the self-supervised phase.

3.5 SSL Module

The SSL technique employed to train the edge encoder plays a crucial role in obtaining meaningful embeddings, essential for effective classification by the few-shot decoder. To address this, we propose a hybrid SSL objective that merges contrastive- and reconstruction-based losses. By integrating both methods, the model gains the ability to distinguish between positive edges and negatively augmented edges, while maximizing the reconstruction error for the malicious few-shot edges. This hybrid approach enhances the discriminative power of the embeddings and contributes to the overall detection performance.

Contrastive-based Loss The recent success of contrastive learning in network intrusion detection [2, 6] and many other domains [20], motivated us to adopt it in the SSL training. Deep Graph Infomax (DGI) is the first attempt to apply contrastive learning to graphs, aiming to maximize the mutual information between patch and global representations of the original graph while minimizing it with negative augmentations. Initially designed for node-level tasks, DGI leverages node features when creating node embeddings. However, Anomal-E has demonstrated the successful application of DGI to edge-level tasks, utilizing edge features. In this version, an edge-level encoder such as E-GraphSAGE, computes the edge embeddings for the original graph G and its negatively augmented version \tilde{G} :

$$\mathbf{H} = \text{enc}(G), \tag{4}$$

$$\tilde{\mathbf{H}} = \text{enc}(\tilde{G}), \tag{5}$$

where \mathbf{H} and $\tilde{\mathbf{H}}$ correspond to the edge embedding matrix for the original graph and its negative augmentation, respectively. Graph \tilde{G} is defined by $\tilde{G} = \mathcal{A}(G)$ with \mathcal{A} an augmentation function that generates a modified version of G . The edge embeddings will then be compared to a compact version of the original graph to measure the similarity. This compact version, also known as global summary, is represented by a single vector \mathbf{s} that preserves global graph information, achieved through a readout function given by:

$$\mathbf{s} = \sigma(\mathcal{R}(\mathbf{H})), \tag{6}$$

where \mathcal{R} is the mean readout operation and σ is the sigmoid function. For each edge in both the original and augmented graphs, the local-global similarity is measured by calculating the dot product between the corresponding edge and the

global summary. Additionally, a weight matrix \mathbf{W} is utilized during training to either maximize or minimize these similarities. The sigmoid activation function σ is then applied to convert the similarity scores into probabilities, indicating whether the input edge is positive or negative:

$$\mathcal{D}(\mathbf{H}_{uv}, \mathbf{s}) = \sigma(\mathbf{H}_{uv} \mathbf{W} \mathbf{s}), \quad (7)$$

$$\mathcal{D}(\tilde{\mathbf{H}}_{uv}, \mathbf{s}) = \sigma(\tilde{\mathbf{H}}_{uv} \mathbf{W} \mathbf{s}), \quad (8)$$

where \mathcal{D} is called the discriminator function, which returns the probability of an edge being either positive or negative, whereas \mathbf{H}_{uv} and $\tilde{\mathbf{H}}_{uv}$ respectively represent the positive and negative embeddings for the edge uv .

Ultimately, the encoder is trained using the Binary Cross-Entropy (BCE) loss function on the positive and negative edges:

$$\mathcal{L}_{\text{DGI}} = -\frac{1}{|E| + |\tilde{E}|} \sum_{uv \in E} \mathbb{E}_G [\log \mathcal{D}(\mathbf{H}_{uv}, \mathbf{s})] + \sum_{uv \in \tilde{E}} \mathbb{E}_{\tilde{G}} [\log (1 - \mathcal{D}(\tilde{\mathbf{H}}_{uv}, \mathbf{s}))], \quad (9)$$

where $|E|$ and $|\tilde{E}|$ represent the number of edges in the positive graph and negative graph, respectively. The loss is optimized by comparing the probabilities of positive and negative edges to all-ones and all-zeros vectors. Minimizing this loss enables the model to learn to distinguish between original and fake edges, thus producing meaningful edge embeddings that preserve the original graph information. These embeddings can be used as valuable input vectors for the downstream task, which here is a few-shot classifier.

Reconstruction-based Loss While the contrastive-based loss method effectively learns edge representations through self-supervision, it overlooks the potential of leveraging the few-shot labeled samples available in the FSL context. Recognizing the value of these limited labeled edges, we introduce a novel loss function intended to capitalize on these few labeled instances during the SSL phase. Our goal is to create distinct embeddings specifically for the labeled few-shot examples. To achieve this, we propose an approach inspired by reconstruction-based SSL methods, where we aim to reconstruct edge features from edge embeddings, thereby ensuring that the embeddings of few-shot labeled edges can be easily identified in embedding space. Formally, given an edge embedding \mathbf{H}_{uv} , the reconstructed edge feature of (u, v) is defined as:

$$\hat{\mathbf{X}}_{uv} = \sigma(\mathbf{H}_{uv} \mathbf{W}_{\text{rec}}) \quad (10)$$

where σ is the sigmoid function and \mathbf{W}_{rec} is a weight matrix with same output dimension as the number of original edge features. The Mean Squared Error (MSE) loss is used to measure the reconstruction error between the reconstructed and the original edge features. We specifically divide this problem in two separate

loss functions \mathcal{L}_{few} and $\mathcal{L}_{\overline{\text{few}}}$, which correspond to the reconstruction loss of the few-shot and non-few-shot examples, respectively:

$$\mathcal{L}_{\text{few}} = \sum_{uv \in \mathcal{E}_{\text{mal}}} (\mathbf{X}_{uv} - \hat{\mathbf{X}}_{uv})^2, \quad (11)$$

$$\mathcal{L}_{\overline{\text{few}}} = \sum_{uv \in E \setminus \mathcal{E}_{\text{mal}}} (\mathbf{X}_{uv} - \hat{\mathbf{X}}_{uv})^2, \quad (12)$$

where \mathcal{E}_{mal} is a set of k malicious edges selected among each of the m attack families present in the dataset, $E \setminus \mathcal{E}_{\text{mal}}$ is the set of remaining (unlabeled) non-few-shot edges, and \mathbf{X}_{uv} represents the original features of edge (u, v) . As illustrated in Fig. 2, our aim is to maximize the reconstruction loss of the $k * m$ malicious few-shot edges while simultaneously minimizing the reconstruction loss of the other non-few-shot edges. This dynamic of maximizing loss for a few malicious instances while minimizing loss for many benign edges fosters the creation of edge embeddings that are distinctively separable in the embedding space. This separation is achievable provided the dataset exhibits significant imbalance, with the majority of edges being benign, thereby statistically ensuring that non-few-shot edges predominantly consist of benign examples.

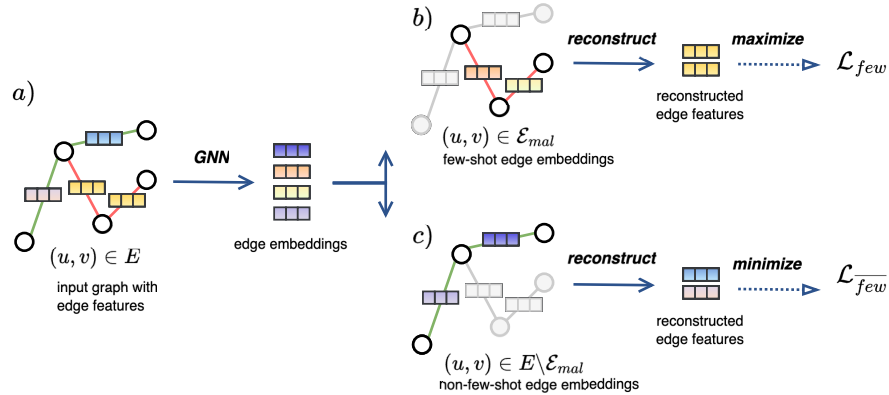


Fig. 2: Illustration of the few-shot aware reconstruction-based loss. a) The GNN encoder first compiles features from the local neighborhood edges to produce edge embeddings. Here, red edges represent malicious few-shot edges, whereas green edges symbolize non-few-shot edges, presumed to contain a high rate of benign edges. b) The SSL module leverages the few-shot edges by maximizing the loss associated with these malicious events. This action compels the encoder to create dissimilar edge embeddings for the malicious few-shot edges, ensuring they are easily distinguishable from benign edges. c) The loss function is also designed to minimize the loss for all non-malicious edges.

The resulting loss function presented in Eq. 13, integrates the contrastive-based and reconstruction-based objectives presented previously:

$$\mathcal{L}_{\text{FEAE}} = \mathcal{L}_{\text{DGI}} + \alpha \mathcal{L}_{\overline{\text{few}}} - \beta \mathcal{L}_{\text{few}} \quad (13)$$

where \mathcal{L}_{DGI} is the loss from DGI presented in Eq. 9. α and β are trade-off coefficients to balance the reconstruction error of few-shot and non-few-shot examples. We recommend to set $\alpha < \beta$, particularly when the dataset contains a significant number of malicious edges. Indeed, $\mathcal{L}_{\overline{\text{few}}}$ should be controlled by α to avoid minimizing the reconstruction error of unlabeled malicious edges present in the non-few-shot edges $E \setminus \mathcal{E}_{\text{mal}}$. Conversely, special attention should be directed towards maximizing \mathcal{L}_{few} over the malicious few-shot edges, as they have a more significant impact than all other unlabeled edges.

3.6 Few-Shot Decoder

The decoder is responsible for the classification of the edge embeddings computed by the encoder and the SSL module; therefore, its training is separate from theirs. It consists of a 2-layer Multi-Layer Perceptron (MLP) trained using BCE in a supervised manner on the few-shot edge embeddings. The decoder outputs a prediction $\hat{\mathbf{y}}$ for every edges such that:

$$\hat{\mathbf{y}} = \sigma(\text{MLP}(\mathbf{H})), \quad (14)$$

where σ is the sigmoid function. To mitigate overfitting on the malicious few-shot edges, we introduce a new set \mathcal{E} that supplements \mathcal{E}_{mal} such that $\mathcal{E}_{\text{mal}} \subset \mathcal{E}$, by randomly selecting benign edges from the dataset while attempting to maintain the same class distribution. However, in the few-shot scenario, information on labels and their distribution is considered unavailable. As a result, we only select a fixed percentage of benign edges, which varies based on the original dataset distribution. For datasets recognized to be highly imbalanced, primarily consisting of benign samples, as exemplified in this paper, we assume the selection of random edges as the benign few-shot edges for the encoder.

The decoder is trained over all edges in \mathcal{E} by computing the BCE loss between the prediction $\hat{\mathbf{y}}$ and the actual few-shot label such that:

$$\mathcal{L}_{\text{DEC}} = \frac{\sum_{uv \in \mathcal{E}} \text{BCE}(\hat{\mathbf{y}}_{uv}, \mathbf{y}_{uv})}{|\mathcal{E}|}, \quad (15)$$

where $|\mathcal{E}|$ represents the number of select few-shot edges and $\mathbf{y}_{uv} \in Y_{\mathcal{E}}$ denotes the label of an edge (u, v) , which is set to a malicious label if $(u, v) \in \mathcal{E}_{\text{mal}}$ and a benign label otherwise.

4 Experiments

In this section, we present the configuration and datasets used to evaluate FEAE, along with a comprehensive evaluation of the model using different numbers of few-shot edges and comparing it against multiple baselines. We also discuss the scalability of the proposed method.

4.1 Configurations and Datasets

The FEAE architecture consists of one encoder layer, described in Section 3), trained for a maximum of 600 epochs, with early stopping set to 150 epochs. For all experiments, we utilized the sum aggregation method, as it provided better performance than the mean aggregation, which led to smoother embeddings. Both the encoder and SSL module were implemented with a hidden size of 128 neurons. To train the encoder, we used a learning rate of 0.001 with the Adam optimizer. Additionally, we applied a weight decay of 0.0001 to the optimizer to better control large gradients generated by the sum aggregation. Regarding the SSL module, the trade-off parameters α and β were set to 0.2 and 0.8, respectively, and 5% of edges in E were randomly selected as benign samples to create \mathcal{E} , the full set of selected few-shot edges. The few-shot decoder was trained for a maximum of 4000 epochs, and we applied early stopping at 1500 epochs. We used the Adam optimizer for training, setting the learning rate to 1×10^{-3} and the weight decay to 1×10^{-5} .

Our experiments were conducted on two network datasets [26] commonly used in GNN-based detection methods [2].

NF-CSE-CIC-IDS2018-v2. This dataset is a Netflow version of the original CSE-CIC-IDS2018 dataset [28], containing approximately 18.9 million network flows. Among these flows, around 12% correspond to attack samples, which are divided into 6 attack families including BruteForce, Bot, DoS, DDoS, Infiltration, Web attacks.

NF-UNSW-NB15-v2. Also converted to Netflow format, this version of the UNSW-NB15 dataset [23] comprises 2.3 million flows, with attack samples accounting for 4% of the dataset, distributed across 9 attack families including Fuzzers, Analysis, Backdoor, DoS, Exploits, Generic, Reconnaissance, Shellcode, Worms.

Both datasets have been standardized using the Netflow format to facilitate ease of use and benchmarking across different detection methods. They offer 43 standardized network flow features [26] such as the duration of a flow or the number of packets and bytes within a flow, which we employed as edge features in our experiments. Similar to the approach used in Anomal-E, we used only 10% of each dataset for scalability concerns, with 70% used as the train set and 30% as the test set. The GNN encoder and SSL module are trained end-to-end on a graph containing both benign and malicious edges. Subsequently, the few-shot decoder is trained separately using these embeddings to classify edges.

The forward and backward steps were performed on an NVIDIA Tesla V100 GPU with 32GB of memory and an Intel Xeon Gold 6148 CPU with 20 cores and 60GB of memory. The FEAE architecture along with the different baseline models were developed in PyTorch using the DGL library.

4.2 Baselines

To evaluate the performance of FEAE, we conducted a benchmark using multiple GNN baselines that follow different learning paradigms.

Supervised Baselines We introduced four supervised baselines to compare the performance of fully-supervised and few-shot approaches. The first baseline chosen was E-GraphSAGE, as detailed in Section 2.2. This model is foundational for malicious edge classification in network graphs. The remaining three baselines—LineGAT, LineGCN, and LineSAGE—are adaptations of the GAT, GCN, and GraphSAGE models, respectively, applied to a line graph [13]. Given that these GNNs are originally designed for node features, their direct application to scenarios leveraging edge features is impractical. To accommodate this, we converted each edge into a node within a line graph, thus framing it as a node-classification problem where edge features are embodied within nodes. A comparable approach employing a line graph was also utilized in E-ResGAT. However, the complex architecture of this model precluded experiment replication due to memory exhaustion, even when employing the sampling strategy recommended in the original publication. For E-GraphSAGE, a single layer configuration with 128 hidden neurons yielded best results, whereas all line graph-based baselines were implemented with two layers, each comprising 128 hidden neurons.

Benign-Supervised Baselines We propose multiple variants of the Anomal-E model, as it is the main framework leveraging self-supervised learning to detect network attacks using flow features. Anomal-E consists of the E-GraphSAGE model as the encoder, DGI as the SSL module, and an Isolation Forest (IF) as the decoder. Anomal-E falls under the benign-supervised category as it requires knowledge of all benign edges to train the IF classifier. We evaluate two variants of this model, using different positive and negative augmentations in the SSL training. Findings in [1] reveal that employing different augmentation techniques for the positive and negative graphs can improve the prediction accuracy. DGI and Anomal-E both use the original graph as the positive graph and the original graph with randomly permuted edges as the negative graph. However, using different augmentations can improve detection performance of such self-supervised techniques depending on the underlying dataset. Consequently, we propose a variant of this baseline denoted as aug_1 that uses the same random edge permutation as the negative graph, but builds a positive graph by randomly dropping 30% of the nodes along with all connected edges. This augmentation aims to simulate a graph with fewer hosts while conserving the same topology between the remaining hosts.

Few-Shot Baselines Given the absence of established few-shot GNN baselines for attack detection, we introduce a modified version of Anomal-E. In this adaptation, the IF classifier is substituted with the few-shot classifier from FEAE. The SSL module incorporating DGI remains unaltered to enable a direct comparison with the SSL module employed in our methodology. Moreover, we assess two variations of FEAE: one adheres to the same positive and negative data augmentations as used in DGI, while the other applies the augmentation discussed earlier. For each baseline, we limit the use of $k = 1$ malicious few-shot edge per attack family, as supported by the data in Table 3. This decision to use a

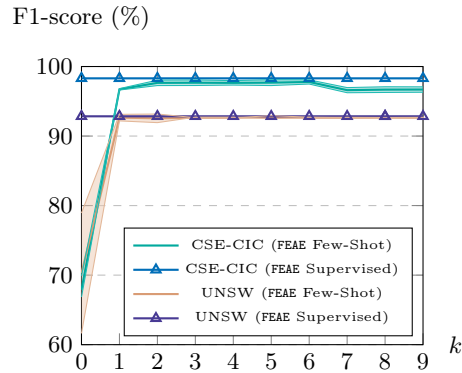


Fig. 3: FEAE performance with respect to k . Setting $k = 0$ indicates that only benign edges are used for training, without any labeled malicious edge.

single attack edge among several potential edges aims to circumvent the need for extensive labeling and analysis of attacks within the training set by analysts, who may not have prior knowledge of the specific edges involved in the attack. Consequently, this approach results in the use of 6 labeled malicious edges for the NF-CSE-CIC-IDS2018-v2 dataset and 9 labeled malicious edges for the NF-UNSW-NB15-v2 dataset.

4.3 Experimental Results

Evaluation with Respect to k The effectiveness of FEAE across the two presented datasets, with respect to varying k values, is showed in Fig. 3. Here, k denotes the count of malicious few-shot edges per each attack family used in training both the SSL module and the few-shot decoder. Additionally, we compare the performance to the supervised variants of FEAE that leverage all available supervised examples for training. The evaluation metrics chosen for this analysis are the precision and the macro F1-score, presented in Eq. 16. The F1 metric is particularly suitable as it provides a balanced measure of the precision and recall across all classes, which is crucial in imbalanced datasets where attack events are significantly less frequent than benign events.

$$F1 = 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}} \quad (16)$$

This metric takes into account both false positives and false negatives, providing a comprehensive evaluation that is particularly valuable in scenarios where class distributions are skewed.

On the NF-CSE-CIC2018-v2 dataset, the results in the figure indicate a remarkable 96.40% F1-score with only $k = 1$, and the performance further improves to 97.65% with $k = 4$. This demonstrates the ability of the few-shot

aware reconstruction loss function to learn effective representations using a minimal number of malicious labels.

For the NF-UNSW-NB15-v2 dataset, a 92.60% F1-score is reached starting from $k = 1$, and the performance remains nearly linear for $k > 1$.

In the two outlined scenarios, employing merely one or two malicious samples within this model’s framework suffices to reach performance levels comparable to those of fully supervised baselines, which leverage the entirety of available labels. To guarantee a detailed evaluation, the ensuing sections will detail a comparative analysis against various baseline models.

Evaluation Against Baselines The experimental results are summarized in Table 2, highlighting the performance of the different approaches described in Section 4.

Table 2: Performance benchmark between FEAE and other baselines.

Data	Model	NF-CSE-CIC-IDS2018-v2			NF-UNSW-NB15-v2		
		F1	Precision	Time	F1	Precision	Time
A, X, Y	E-GraphSAGE	96.02	98.82	0.31	95.35	92.49	0.32
A, X, Y	LineGAT	93.84	96.84	4.3	95.33	91.81	14.2
A, X, Y	LineGCN	89.29	95.42	0.43	95.35	91.83	0.58
A, X, Y	LineSAGE	94.94	97.10	1.00	95.90	93.11	2.08
A, X, Y_{ben}	Anomal-E (IF)	94.46	96.86	85.1	91.14	85.78	9.2
A, X, Y_{ben}	Anomal-E (IF) + aug_1	96.53	98.84	81.3	87.38	84.13	7.9
A, X, Y_{few}	Anomal-E (Few-Shot)	95.3	97.28	24.5	92.47	86.42	1.45
A, X, Y_{few}	FEAE	96.40	99.05	19.6	92.64	87.69	1.22
A, X, Y_{few}	FEAE + aug_1	97.44	98.71	18.4	92.64	87.69	1.19

Data refers to the input data required to train a specific **Model**. The first group of baselines represents supervised approaches, whereas the second group represents benign-supervised methods (with Y_{ben} corresponding to all benign labels) and the last group is dedicated to the few-shot baselines (with Y_{few} corresponding to the malicious few-shot labels). **Time** corresponds to the overall training time in minutes, to reach the score of the corresponding baseline. For both datasets, the performance is measured using the macro **F1** score and the **Precision** over 5 iterations, with the notation $\text{mean} \pm \text{standard deviation}$.

On the NF-CSE-CIC-IDS2018-v2 dataset, the variant of FEAE employing the augmentation strategy variant achieves the highest performance. This result demonstrates the effectiveness of FSL methods on this dataset, as it surpasses the results of fully supervised approaches by using merely one label per attack family. Additionally, FEAE exhibits enhanced performance compared to Anomal-E in the few-shot scenario, highlighting the capabilities of the hybrid SSL objective incorporated in FEAE. For a deeper insight into the learning mechanisms of the model, we direct attention to Fig. 4, which illustrates 5000 edge

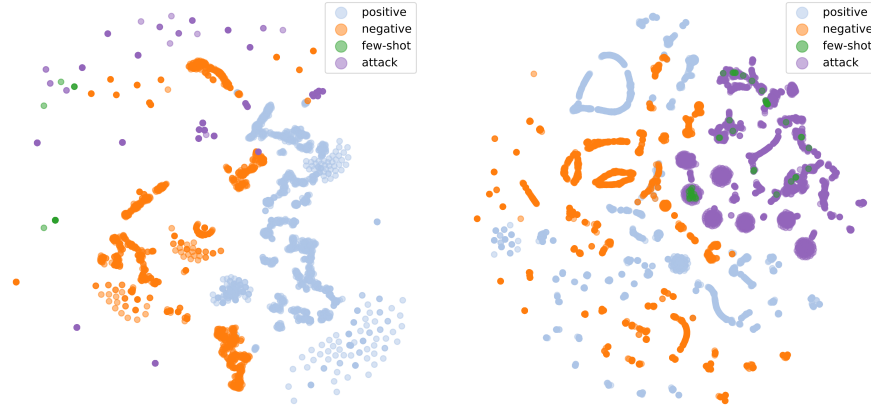


Fig. 4: **Left:** Some edge embeddings produced by Anomal-E. Note that the few-shot edges are just for comparison as they are not leveraged in the original Anomal-E. **Right:** Edge embeddings generated by FEAE.

embeddings produced by the encoder within FEAE. On the left are embeddings from E-GraphSAGE trained with DGI in Anomal-E, highlighting the model’s ability to separate positive and negative edges without specifically focusing on attack edges (purple). On the right, FEAE’s embeddings using DGI’s augmentations are displayed, offering insights into its differentiation strategy. In the case of Anomal-E, the model successfully separates positive and negative edges without giving special attention to attack edges (denoted in purple). Furthermore, attack edges do not naturally cluster within a common region in the embedding space since the loss function does not specifically enforce this clustering. However, when examining the FEAE edge embeddings, we observe a distinct cluster of attack edges, including the labeled few-shot edges used during the training. By directly clustering these attack edges in the SSL training, FEAE assists the downstream classifier in more effectively distinguishing between benign and malicious edges. Furthermore, in the embeddings from Anomal-E, attack edges do not aggregate into a unified region within the embedding space, due to the absence of a specific requirement in the loss function for such clustering. In contrary, the FEAE edge embeddings exhibit a notable clustering of attack edges. This cluster encompasses the few-shot labeled edges used during the training of the encoder. By directly clustering these attack edges in the SSL training phase, FEAE successfully clusters benign and malicious edges, which facilitates the classification by the downstream decoder.

On the NF-UNSW-NB15-v2 dataset, the observed performance surpasses that of Anomal-E but falls short of the results achieved by supervised baselines. This suggests that the model may not have been optimally trained on this dataset. It underscores the need for further experimentation and exploration of SSL methodologies specific to this dataset.

The experiments also highlight the importance of using different graph augmentations, as these techniques lead to an improvement of up to 1% in the F1-score for FEAE and up to 2% for Anomal-E. However, we also notice that one augmentation pair may not perform well on both datasets, suggesting that this strategy requires a prior understanding of the underlying data to be applied.

4.4 Scalability

The FEAE model demonstrates a significant reduction in computation time compared to the computationally intensive IF classifier used in Anomal-E baselines. Indeed, the IF classifier used in Anomal-E accounts for a significant portion of the computational time, thus serving as a primary bottleneck. Furthermore, experimental evidence suggests that the integration of few-shot examples within the learning process reduces the total number of training epochs, as an extended training on the few-shot examples may induce overfitting. This effectively reduces training time compared to fully self-supervised methods like Anomal-E.

Despite their advantages, both FEAE and Anomal-E incur higher computational costs than simpler supervised GNN models, which benefit from simpler end-to-end training. Indeed, self-supervised models based on DGI are inevitably slower than the supervised ones, which do not require any SSL module. Nonetheless, the trade-off between computational expense and the ability to efficiently handle few-shot scenarios with FEAE justifies its application, particularly in the cybersecurity domain where dataset labeling is scarce.

5 Discussion

This section discusses the advantages, limitations and research directions of FEAE and FSL methods for network attack detection.

5.1 Benefits of Few-Shot Learning

The conducted experiments have demonstrated the interesting capability of detecting attacks using only one malicious example per attack class through few-shot learning. In a real-world scenario, this labeling strategy could reduce the reliance on benign-supervised methods, which require the assumption of training the model solely using benign examples. However, it is often challenging to ensure that the labeled examples are genuinely benign, which can be a costly and time-consuming process. Few-shot learning, on the other hand, offers a promising alternative, as it requires only a small number of labeled malicious examples to be effective. Furthermore, using graphs in network attack detection offers a significant benefit: cybersecurity analysts can leverage the predictions of the FEAE model to approximate the initial structure of an attack graph. This is achieved by gathering and analyzing the edges that FEAE identifies as malicious. Through this process, analysts can gain insights into the attack’s topology and understand how different nodes (potentially representing users, servers, endpoints, or

other network entities) are interconnected in a malicious campaign. This capability not only aids in understanding the attack but also in devising targeted countermeasures by revealing the path of the attack within the network.

5.2 Limitations and Further Works

A limitation of FEAE and similar few-shot learning approaches is their dependency on requiring some malicious activity data, which may not exist in the historical data of many enterprise networks. To address this challenge, future research might focus on integrating malicious activities from synthetic sources or datasets with known malicious content into real-world datasets. Such an approach could reduce the necessity for labeling, as labels from these enriched datasets would suffice. This strategy not only promises to alleviate the labeling burden but also verifies that the performance of the model remains unaffected by new attack types. Moreover, exploring the influence of flow features on prediction accuracy and conducting studies on datasets limited to topological data, devoid of additional features, could provide deeper insights into the model’s operational dynamics in varied environments.

Moreover, substantial effort should be dedicated to enhancing the scalability of such models, especially when dealing with large graphs that may undergo structural changes over time. As networks evolve and grow, the model’s ability to adapt and handle dynamic structures efficiently becomes crucial. Research into developing scalable approaches for few-shot learning on large and dynamic graphs would greatly benefit practical deployment in real-world environments.

6 Conclusion

In this study, we undertake the first analysis of applying few-shot learning to network-based attack detection with GNNs. Through the introduction of the FEAE architecture and the conducted experiments, it has been demonstrated that merely one labeled attack example per attack family is sufficient to achieve competitive detection capabilities when compared to fully-supervised or benign-supervised approaches. These encouraging results open up exciting possibilities for reducing the reliance on fully supervised methods and addressing the challenges of limited labeled data in the field of attack detection. The integration of few-shot learning in attack detection represents a promising research direction, offering more efficient solutions that can adapt well to real-world scenarios.

References

1. Bilot, T., El Madhoun, N., Al Agha, K., Zouaoui, A.: A benchmark of graph augmentations for contrastive learning-based network attack detection with graph neural networks. In: 2023 7th Cyber Security in Networking Conference (CSNet). pp. 53–56. IEEE (2023)

2. Bilot, T., El Madhoun, N., Al Agha, K., Zouaoui, A.: Graph neural networks for intrusion detection: A survey. *IEEE Access* (2023)
3. Bilot, T., El Madhoun, N., Al Agha, K., Zouaoui, A.: A survey on malware detection with graph representation learning. *ACM Computing Surveys* (2023)
4. Bilot, T., Geis, G., Hammi, B.: Phishgnn: A phishing website detection framework using graph neural networks. In: 19th International Conference on Security and Cryptography. pp. 428–435. SCITEPRESS-Science and Technology Publications (2022)
5. Cao, Y., Jiang, H., Deng, Y., Wu, J., Zhou, P., Luo, W.: Detecting and mitigating ddos attacks in sdn using spatial-temporal graph convolutional network. *IEEE Transactions on Dependable and Secure Computing* **19**(6), 3855–3872 (2021)
6. Caville, E., Lo, W.W., Layeghy, S., Portmann, M.: Anomal-e: A self-supervised network intrusion detection system based on graph neural networks. *Knowledge-Based Systems* **258**, 110030 (2022)
7. Chang, L., Branco, P.: Graph-based solutions with residuals for intrusion detection: The modified e-graphsage and e-resgat algorithms. *arXiv preprint arXiv:2111.13597* (2021)
8. Cheng, Z., Lv, Q., Liang, J., Wang, Y., Sun, D., Pasquier, T., Han, X.: Kairos:: Practical intrusion detection and investigation using whole-system provenance. *arXiv preprint arXiv:2308.05034* (2023)
9. Fang, Y., Huang, C., Zeng, M., Zhao, Z., Huang, C.: Jstrong: Malicious javascript detection based on code semantic representation and graph neural network. *Computers & Security* **118**, 102715 (2022)
10. Fang, Y., Wang, C., Fang, Z., Huang, C.: Lmtracker: Lateral movement path detection based on heterogeneous graph embedding. *Neurocomputing* **474**, 37–47 (2022)
11. Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural message passing for quantum chemistry. In: International conference on machine learning. pp. 1263–1272. PMLR (2017)
12. Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. *Advances in neural information processing systems* **30** (2017)
13. Harary, F., Norman, R.Z.: Some properties of line digraphs. *Rendiconti del circolo matematico di palermo* **9**, 161–168 (1960)
14. Jia, Z., Xiong, Y., Nan, Y., Zhang, Y., Zhao, J., Wen, M.: Magic: Detecting advanced persistent threats via masked graph representation learning. *arXiv preprint arXiv:2310.09831* (2023)
15. King, I.J., Huang, H.H.: Euler: Detecting network lateral movement via scalable temporal link prediction. *ACM Transactions on Privacy and Security* (2023)
16. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016)
17. Lan, J., Lu, J.Z., Wan, G.G., Wang, Y.Y., Huang, C.Y., Zhang, S.B., Huang, Y.Y., Ma, J.N.: E-minbatch graphsage: An industrial internet attack detection model. *Security and Communication Networks* **2022** (2022)
18. Li, Y., Li, R., Zhou, Z., Guo, J., Yang, W., Du, M., Liu, Q.: Graphddos: Effective ddos attack detection using graph neural networks. In: 2022 IEEE 25th International Conference on Computer Supported Cooperative Work in Design (CSCWD). pp. 1275–1280. IEEE (2022)
19. Liu, F.T., Ting, K.M., Zhou, Z.H.: Isolation forest. In: 2008 eighth IEEE international conference on data mining. pp. 413–422. IEEE (2008)

20. Liu, Y., Jin, M., Pan, S., Zhou, C., Zheng, Y., Xia, F., Philip, S.Y.: Graph self-supervised learning: A survey. *IEEE Transactions on Knowledge and Data Engineering* **35**(6), 5879–5900 (2022)
21. Lo, W.W., Kulatilleke, G., Sarhan, M., Layeghy, S., Portmann, M.: Xg-bot: An explainable deep graph neural network for botnet detection and forensics. *Internet of Things* **22**, 100747 (2023)
22. Lo, W.W., Layeghy, S., Sarhan, M., Gallagher, M., Portmann, M.: E-graphsage: A graph neural network based intrusion detection system for iot. In: *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*. pp. 1–9. IEEE (2022)
23. Moustafa, N., Slay, J.: Unsw-nb15: a comprehensive data set for network intrusion detection systems (unsw-nb15 network data set). In: *2015 military communications and information systems conference (MilCIS)*. pp. 1–6. IEEE (2015)
24. Paudel, R., Huang, H.H.: Pikachu: Temporal walk based dynamic graph embedding for network anomaly detection. In: *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*. pp. 1–7. IEEE (2022)
25. Pujol-Perich, D., Suárez-Varela, J., Cabellos-Aparicio, A., Barlet-Ros, P.: Unveiling the potential of graph neural networks for robust intrusion detection. *ACM SIGMETRICS Performance Evaluation Review* **49**(4), 111–117 (2022)
26. Sarhan, M., Layeghy, S., Portmann, M.: Towards a standard feature set for network intrusion detection system datasets. *Mobile networks and applications* pp. 1–14 (2022)
27. Scarselli, F., Gori, M., Tsoi, A.C., Hagenbuchner, M., Monfardini, G.: The graph neural network model. *IEEE transactions on neural networks* **20**(1), 61–80 (2008)
28. Sharafaldin, I., Lashkari, A.H., Ghorbani, A.A.: Toward generating a new intrusion detection dataset and intrusion traffic characterization. *ICISSp* **1**, 108–116 (2018)
29. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1199–1208 (2018)
30. Thein, T.T., Shiraishi, Y., Morii, M.: Few-shot learning-based malicious iot traffic detection with prototypical graph neural networks. *IEICE TRANSACTIONS on Information and Systems* **106**(9), 1480–1489 (2023)
31. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., Bengio, Y.: Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017)
32. Veličković, P., Fedus, W., Hamilton, W.L., Liò, P., Bengio, Y., Hjelm, R.D.: Deep graph infomax. *arXiv preprint arXiv:1809.10341* (2018)
33. Wang, Y., Yao, Q., Kwok, J.T., Ni, L.M.: Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)* **53**(3), 1–34 (2020)
34. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., Philip, S.Y.: A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* **32**(1), 4–24 (2020)
35. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826* (2018)
36. You, Y., Chen, T., Sui, Y., Chen, T., Wang, Z., Shen, Y.: Graph contrastive learning with augmentations. *Advances in neural information processing systems* **33**, 5812–5823 (2020)
37. Zhong, M., Lin, M., Zhang, C., Xu, Z.: A survey on graph neural networks for intrusion detection systems: Methods, trends and challenges. *Computers & Security* p. 103821 (2024)
38. Zhou, J., Xu, Z., Rush, A.M., Yu, M.: Automating botnet detection with graph neural networks. *arXiv preprint arXiv:2003.06344* (2020)