



Universiteit
Leiden
The Netherlands

Cross-domain graph level anomaly detection

Li, Z.; Liang, S.; Shi, J.; Leeuwen, M. van

Citation

Li, Z., Liang, S., Shi, J., & Leeuwen, M. van. (2024). Cross-domain graph level anomaly detection. *Ieee Transactions On Knowledge And Data Engineering*, 36(12), 7839-7850. doi:10.1109/TKDE.2024.3462442

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/4170186>

Note: To cite this publication please use the final published version (if applicable).

Cross-Domain Graph Level Anomaly Detection

Zhong Li , Sheng Liang, Jiayang Shi , and Matthijs van Leeuwen 

Abstract—Existing graph level anomaly detection methods are predominantly unsupervised due to high costs for obtaining labels, yielding sub-optimal detection accuracy when compared to supervised methods. Moreover, they heavily rely on the assumption that the training data exclusively consists of normal graphs. Hence, even the presence of a few anomalous graphs can lead to substantial performance degradation. To alleviate these problems, we propose a *cross-domain graph level anomaly detection method*, aiming to identify anomalous graphs from a set of unlabeled graphs (*target domain*) by using easily accessible normal graphs from a different but related domain (*source domain*). Our method consists of four components: a feature extractor that preserves semantic and topological information of individual graphs while incorporating the distance between different graphs; an adversarial domain classifier to make graph level representations domain-invariant; a one-class classifier to exploit label information in the source domain; and a class aligner to align classes from both domains based on pseudolabels. Experiments on seven benchmark datasets show that the proposed method largely outperforms state-of-the-art methods.

Index Terms—Graph anomaly detection, graph neural networks, graph transfer learning.

I. INTRODUCTION

GRAPH-STRUCTURED data is ubiquitous, as it can represent relations between objects using edges and semantic characteristics of objects using node attributes. As a result, graph level anomaly detection has a wide range of potential applications, such as criminal detection in financial network [1], error detection in system logs [2], identifying specific molecules in drug discovery [3], and detecting unhealthy brain structures [4].

Graph neural networks (GNNs) are capable of learning discriminative feature representations for graphs and have significantly advanced benchmark results for graph level anomaly detection [5]. Similar to other types of neural networks, the impressive performance of graph neural networks (GNNs) is often attained by using a substantial amount of labeled data. However, the process of manually annotating graph data is laborious and thus often impractical. To circumvent this challenge,

recent studies have turned to unsupervised (or semi-supervised) learning instead. These methods, however, strongly rely on the assumption that the training data exclusively consists of normal graphs. Our experiments demonstrate that even a minor presence of anomalous graphs in the training data can lead to substantial performance degradation for these methods (c.f. Table III).

In practice, labeled data may be accessible or relatively cheap to obtain in some domains. Hence, in situations where a certain ‘target’ domain of interest suffers from a dearth of labeled data (or its purity cannot be guaranteed), there is a strong motivation to construct learners that can exploit abundant labeled data from a different but related domain.

Recent work on graph level anomaly detection [1], [5], [6], [7] is mostly unsupervised (or semi-supervised), and has been limited to detecting anomalies within a single domain. That is, the potential benefits of incorporating labeled information from a related domain has not yet been researched. In this paper we investigate how to transfer ‘anomaly knowledge’ from a source to a target graph database.

Unsupervised domain adaptation (UDA) is an attractive approach to achieve this: it adapts models learned from a source domain with plenty labeled data to a target domain without labels, and has demonstrated remarkable performance in computer vision and natural language processing [8], [9]. Although a few studies have explored UDA for cross-domain node classification, there has been no prior research on cross-domain graph level anomaly detection. Two challenges need to be overcome. First, most existing UDA methods are developed for vector-based data, such as image and text data, for which a distance in a euclidean space can be defined, while for graph-structured data distance is typically defined in a non-euclidean space due to graph isomorphism. This makes directly applying off-the-shelf UDA methods to graphs impractical. Second, graph level anomaly detection is inherently more challenging than node level anomaly detection, as anomalies at the graph level may involve global patterns and interactions that cannot be easily discerned by examining individual nodes.

To fill this gap, being motivated and supported by domain adaptation theory [10], we propose an unsupervised domain adaptation based graph level anomaly detection method called ARMET. It addresses the following cross-domain graph level anomaly detection problem: given a target graph database with fully unlabeled graphs and a *different but related* source graph database that contains only normal graphs, learn a one-class classifier that identifies anomalous graphs from the target graph database.

To achieve this, ARMET leverages an adversarial learning approach consisting of four main components. First, to learn

Received 10 March 2024; revised 20 July 2024; accepted 10 September 2024. Date of publication 20 September 2024; date of current version 13 November 2024. Zhong Li and Matthijs van Leeuwen: This publication is part of the project Digital Twin with project number P18-03 of the research programme TTW Perspective, which is (partly) financed by the Dutch Research Council (NWO). The work of Sheng Liang was supported in part by European Research Council under Grant 740516. The work of Jiayang Shi was supported in part by European Union H2020-MSCA-ITN-2020 under Grant 956172 (xCTing). Recommended for acceptance by Y. Gao. (*Corresponding author: Zhong Li.*)

Zhong Li, Jiayang Shi, and Matthijs van Leeuwen are with LIACS, Leiden University, 2311 EZ Leiden, The Netherlands (e-mail: z.li@liacs.leidenuniv.nl). Sheng Liang is with CIS, LMU Munich, 80539 München, Germany. Digital Object Identifier 10.1109/TKDE.2024.3462442

graph level representations, it utilizes a two-part feature extractor: a semantic feature extractor to jointly preserve the semantic and topological information of each graph, and a structure feature extractor to extract the structure of each graph domain. Second, a domain classifier is learned to make graph level representations domain-invariant, thereby reducing the domain discrepancy. Third, a one-class classifier is trained using normal source graphs, aiming to make the learned graph level representations label-discriminative. Finally, a class aligner is trained to align normal graphs in both domains while separating anomalous graphs and normal graphs in the target domain. As a result, in an end-to-end manner, ARMET can learn both domain-invariant and label-discriminative graph level representations, and thus effectively identify anomalous graphs from the target domain.

Our contributions can be summarized as follows:

- We introduce the cross-domain graph level anomaly detection problem, and develop ARMET, an effective approach to address this problem;
- ARMET is the first attempt to combine graph neural networks and adversarial domain adaptation techniques for performing graph level anomaly detection tasks;
- Experiments demonstrate its improved performance for graph level anomaly detection when compared to state-of-the-art unsupervised graph level anomaly detectors.

The rest of this work is organized as follows. Section II reviews related literature. Section III introduces relevant notations and formulates the research problem. Section IV presents the framework of our method ARMET. Sections V and VI describe the design details of ARMET. Section VII provides experimental setup, and Section VIII presents results and corresponding analysis. Section X concludes this work.

II. RELATED WORK

We review work that is most closely related; readers are referred to the following surveys for more on transfer learning [11], [12], graph representation learning [13], and graph anomaly detection [14], [15].

Graph Level Anomaly Detection Unsupervised/semi-supervised methods include OCGIN [6], GLAM [5], GLocalKD [7], OCGTL [1] and CODEtect [16]. Unlike ARMET, CODEtect can only handle unattributed graphs. Meanwhile, OCGIN, GLAM, GLocalKD, and OCGTL can handle attributed graphs, but they heavily depend on the assumption that the training data exclusively contains normal graphs. This assumption is impractical or expensive in real-world applications. As we will show later, the presence of anomalies in the training data may largely decrease the performance of these methods (c.f. Table III). Moreover, the supervised method iGAD requires fully labelled training data, which is expensive or sometimes impossible to obtain. In contrast, ARMET only requires that the source domain data exclusively contains normal graphs, while imposing no additional assumptions on the target domain data. Further, it is the first graph level anomaly detection method that leverages labeled data from a different but related domain.

Traditional Unsupervised Domain Adaptation Traditional UDA techniques, such as DeepCoral [17], DANN [18], ADDA [19], and CDAN [20], are primarily designed for addressing multi-class classification problems in computer vision and natural language processing. Consequently, these methods fail to account for the unique characteristics of graph-structured data, as well as the highly imbalanced nature of anomaly detection problems. Moreover, these methods assume that the source domain contains all classes and is fully labeled. As a result, their effectiveness diminishes when the source domain contains only partial classes (e.g., only the normal class in anomaly detection), and this will be shown later in Table VI.

Domain Adaptation on Graphs Most existing domain adaptation methods on graphs, such as SDA-DAGL [21], DANE [22], UDA-GCN [23], DASGA [24], ACDNE [25], CDNE [26], DANE [27], COMMANDER [28], AdaGCN [29], DGL [30], AdaGIn [31] and CD-GAD [32], only consider domain adaptation from an individual graph to another graph.

Only a few works, including DGDA [33], CDA [34], and [35], consider domain adaptation from a set of graphs to another set of graphs. However, they focus on the multi-class graph classification problem, while ARMET is the first to consider cross-domain graph level anomaly detection, which is more challenging due to the extreme class imbalance. Moreover, unlike ARMET, these methods require the source domain data to contain all classes and be fully labeled.

III. PROBLEM STATEMENT

Following the notation commonly used in transfer learning [11], a *domain* $\mathcal{D}$ consists of two components, namely a feature space $\mathcal{X}$ and its marginal probability distribution $\mathbb{P}(\mathcal{X})$. Meanwhile, a *task* contains two components, that is, a label space $\mathcal{Y}$ and a predictive function $h(\cdot)$ that can be expressed as $\mathbb{P}(\mathcal{Y}|\mathcal{X})$. Moreover, we consider an attributed and undirected graph $G = (\mathcal{V}, \mathcal{E})$, where the node set $\mathcal{V}$ is associated with a node feature matrix $\mathbf{X} \in \mathcal{R}^{|\mathcal{V}| \times C}$ and the edge set $\mathcal{E}$ is associated with the adjacency matrix $\mathbf{A} \in \mathcal{R}^{|\mathcal{V}| \times |\mathcal{V}|}$. In the following sections, we use the superscripts $(\cdot)^s$ and $(\cdot)^t$ to denote concepts in the source and target domains (also tasks), respectively.

Traditional Unsupervised Domain Adaptation on Graphs Traditional unsupervised domain adaptation (UDA) assumes that there is a set of labeled source graphs $\mathcal{G}^s = \{G_n^s, Y_n^s\}_{n=1}^{N_s}$ that are i.i.d drawn from $\mathbb{P}(\mathcal{X}^s, \mathcal{Y}^s)$, and another set of unlabeled target graphs $\mathcal{G}^t = \{G_n^t\}_{n=1}^{N_t}$ that are i.i.d drawn from $\mathbb{P}(\mathcal{X}^t)$. Moreover, UDA usually imposes the *covariate shift assumption*, i.e., $\mathbb{P}^s(\mathcal{X}^s) \neq \mathbb{P}^t(\mathcal{X}^t)$ but $\mathbb{P}^s(\mathcal{Y}^s|\mathcal{X}^s) = \mathbb{P}^t(\mathcal{Y}^t|\mathcal{X}^t)$. In other words, it assumes that the source and the target are different (in the sense that $\mathbb{P}^s(\mathcal{X}^s) \neq \mathbb{P}^t(\mathcal{X}^t)$) but related (in the sense that $\mathcal{X}^s = \mathcal{X}^t$ and $\mathbb{P}^s(\mathcal{Y}^s|\mathcal{X}^s) = \mathbb{P}^t(\mathcal{Y}^t|\mathcal{X}^t)$). On this basis, UDA aims to learn a classifier $h : \mathcal{X}^t \rightarrow \mathcal{Y}^t$ by using fully labelled source graphs and unlabelled target graphs.

For simplicity, we assume that the node feature matrices of the source (i.e., $\mathbf{X}^s$) and target graphs (i.e., $\mathbf{X}^t$) have the same dimensionality and their columns share the same semantic meanings. Otherwise, we can construct a unified node feature set $\mathcal{X} = \mathcal{X}^s \cup \mathcal{X}^t$ by following the practice in [29], [31]. This




Fig. 1. An overview of ARMET. Graphs and their learned representations are framed in oval rectangles, where the target domain is highlighted in green and the source domain in orange. Meanwhile, the semantic feature extractor, structure feature extractor, and other learners including one-class classifier, domain classifier, and class aligner are depicted in blue rectangles.

assumption is important for utilising parameters-shared graph embedding models and fulfilling the covariate shift assumption in UDA.

Cross-Domain Graph Level Anomaly Detection The anomaly detection problem can be considered as a binary classification problem. Hence, we have $\mathcal{Y}^s = \mathcal{Y}^t = \{0, 1\}$, where 0 and 1 represent normal and abnormal graphs, respectively. To further relax the dependency on fully labeled source data, we assume that $\mathcal{G}^s$ only contains normal graphs. Specifically, given $\mathcal{G}^s = \{G_n^s, Y_n^s\}_{n=1}^{N_s} \in \mathcal{X}^s \times \mathcal{Y}^s$ with $Y_n^s = 0$ for $n = 1, \dots, N_s$, and $\mathcal{G}^t = \{G_n^t\}_{n=1}^{N_t} \in \mathcal{X}^t$, *Cross-Domain Graph Level Anomaly Detection* (CD-GLAD) aims to learn a binary classifier $g: \mathcal{X}^t \rightarrow \mathcal{Y}^t$ that accurately predicts anomaly labels for graphs in the target domain, with the assistance of both normal graphs from the source domain and unlabeled graphs from the target domain. Hence, the CD-GLAD problem is more practical but entails greater challenges than traditional UDA.

IV. PROPOSED METHOD: ARMET

Motivation According to domain adaptation theory [10], given the source domain $\mathcal{D}^s = \{\mathcal{X}^s, \mathbb{P}(\mathcal{X}^s)\}$ and the target domain $\mathcal{D}^t = \{\mathcal{X}^t, \mathbb{P}(\mathcal{X}^t)\}$, given any binary classifier h drawn from a hypothesis class $\mathcal{H}$, for any $\delta \in (0, 1)$, with the probability at least of $1 - \delta$, we have

$$\epsilon^t(h) \leq \epsilon^s(h) + \frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^s, \mathcal{D}^t) + [\epsilon^s(h^*) + \epsilon^t(h^*)] + \omega, \quad (1)$$

where $\epsilon^t(h)$ and $\epsilon^s(h)$ indicate the expected error of classifier h on the target domain and source domain, respectively. Moreover, $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}^s, \mathcal{D}^t)$ represents the domain discrepancy. Importantly, $[\epsilon^s(h^*) + \epsilon^t(h^*)]$ means the combined error of the ideal joint

classifier h^* on both domains, where $h^* =: \arg \min_{h \in \mathcal{H}} [\epsilon^s(h) + \epsilon^t(h)]$. Besides, ω is the item associated with the model complexity and the sample sizes.

In this work, we aim to learn a classifier h that has the minimal expected error on the target domain. Therefore, the sum of terms on the right side of (1) should be minimized. This sheds key insights into the design of our algorithm, which considers the following overall objective:

$$\mathcal{L}(\mathcal{X}^s, \mathcal{Y}^s, \mathcal{X}^t) = \lambda_1 \mathcal{L}_{SC}(\mathcal{X}^s, \mathcal{Y}^s) - \lambda_2 \mathcal{L}_{DA}(\mathcal{X}^s, \mathcal{X}^t) + \lambda_3 \mathcal{L}_{CA}(\mathcal{X}^s, \mathcal{Y}^s, \mathcal{X}^t), \quad (2)$$

where $\mathcal{L}(\mathcal{X}^s, \mathcal{Y}^s, \mathcal{X}^t)$ corresponds to the upper bound of $\epsilon^t(h)$, $\mathcal{L}_{SC}(\mathcal{X}^s, \mathcal{Y}^s)$ is the source classifier loss corresponding to $\epsilon^s(h)$, $\mathcal{L}_{DA}(\mathcal{X}^s, \mathcal{X}^t)$ is the domain classifier loss that approximates $d(\mathcal{D}^s, \mathcal{D}^t)$ (larger loss indicates smaller discrepancy), and $\mathcal{L}_{CA}(\mathcal{X}^s, \mathcal{Y}^s, \mathcal{X}^t)$ is the class alignment loss that corresponds to $[\epsilon^s(h^*) + \epsilon^t(h^*)]$. Further, the balance parameters $\lambda_1 > 0$, $\lambda_2 > 0$ and $\lambda_3 > 0$.

Approach Fig. 1 depicts the architecture of our proposed method, dubbed ARMET (adversarial cross domain graph level anomaly detection). Specifically, ARMET adapts an adversarial learning framework to perform cross-domain graph level anomaly detection. It involves the four following main components:

- **Parameters-Shared Feature Extractor h^{FE} :** takes a source graph database $\mathcal{G}^s$ consisting of exclusively normal graphs and an unlabeled target graph database $\mathcal{G}^t$ as inputs, and aims to learn a representation vector $h^{FE}(G_i)$ for each graph $G_i \in \mathcal{G}^s \cup \mathcal{G}^t$ such that similar graphs (in terms of semantic and structure properties) from both domains have similar embeddings;

- **One-Class Classifier h^s** : takes the representation of each graph $h^{FE}(G_i)$ from the source domain as input, and learns a hypersphere to include embeddings of normal graphs while excluding those of anomalous graphs. This classifier can be directly used to predict labels for graphs in the target domain and corresponds to $\mathcal{L}_{SC}$;
- **Domain Classifier h^d** : takes the representation of each graph $h^{FE}(G_i)$ as input, and attempts to discriminate whether it is drawn from the source domain or the target domain such that the distributions of embeddings of both domains are aligned. And it corresponds to $\mathcal{L}_{DA}$;
- **Class Aligner**: takes $\{G_n^s, Y_n^s\}_{n=1}^{N_s}$ and $\{G_n^t, \hat{Y}_n^t\}_{n=1}^{N_t}$ as inputs, further making normal graphs from both domains have close embeddings, while normal graphs and anomalous graphs in the target domain have distant embeddings. Particularly, we apply the One-Class Classifier h^s to obtain pseudo-labels $\hat{Y}_n^t$ for graphs in the target domain while using the true labels Y_n^s for graphs in the source domain; This component corresponds to $\mathcal{L}_{CA}$.

For better organization, we divide these components into two modules, as shown in Fig. 1: the *graph feature extraction module* that contains the feature extractor (steps ①, ②, ③), and the *cross-domain anomaly detection module* that includes the domain classifier (step ④), the one-class classifier (step ⑤), and the class aligner (step ⑥). Importantly, the two modules are trained jointly in an end-to-end manner, although they are introduced separately in the following.

V. MODULE 1: GRAPH FEATURE EXTRACTION

The graph feature extraction module consists of a feature extractor dubbed $h^{FE}(\cdot)$ with trainable parameters Θ^{FE} , aiming to extract graph level representations for graphs from source domain and target domain. This component is further decomposed to three sub-components: a semantic feature extractor (step ①) and a structure feature extractor (step ②), followed by a feature concatenation operator (step ③).

Semantic Feature Extractor We denote the semantic feature extractor as $h^{SE}(\cdot)$ with trainable parameters Θ^{SE} , which learns a graph level representation for each input graph. Specifically, we adapt a K -layer GIN model [36] followed by a READOUT function to obtain graph level representations due to the superior performance of GIN compared to other competing methods [37]. Concretely, GIN updates node representations as

$$f_v^{(k+1)} = \text{MLP} \left(\left(1 + \alpha^{(k)} \right) f_v^{(k)} + \sum_{u \in \mathcal{N}(v)} f_u^{(k)} \right), \quad (3)$$

where $f_v^{(k)}$ denotes the representation of node v learned at the k -th layer, $\mathcal{N}(v)$ indicates the neighbour set for node v , $\alpha^{(k)}$ is a trainable parameter while MLP represents a multilayer perceptron. Next, we obtain the graph level representation by leveraging READOUT($f_v^{(k)} | k = 1, \dots, K$), which can be a simple permutation-invariant function such as the maximum, sum or mean. Using this, we can preserve the semantic and topological information of each graph.

Structure Feature Extractor We denote the structure extractor as $h^{ST}(\cdot)$ with parameters Θ^{ST} . We assume that $\mathcal{G}^s =$

$\{G_1^s, \dots, G_j^s, \dots, G_N^s\}$ and $\mathcal{G}^t = \{G_1^t, \dots, G_j^t, \dots, G_M^t\}$ exhibit some inherent data structures such as clusters in $\mathcal{X}^s$ and $\mathcal{X}^t$, respectively. Inspired by [38], for each graph database, we construct a k -nearest neighbours graph (*KNN graph*) in the latent space to model its data structure, aiming to capture the neighbourhood information between different graphs. Without loss of generality, we use $\mathcal{G}^s$ to demonstrate the construction of a *KNN graph*.

The *KNN graph* of $\mathcal{G}^s$ contains N nodes, with each node representing a source graph and its node attribute indicating the corresponding graph level representation vector learned by $h^{ST}(\cdot)$. Next, to generate edges, we can construct the adjacency matrix as

$$A_{ij} = \begin{cases} 1, & \text{if } G_i \in \mathcal{N}_k(G_j) \text{ or } G_j \in \mathcal{N}_k(G_i) \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

where $\mathcal{N}_k(G_j)$ represents the k -nearest neighbours set of graph G_j based on the euclidean distance between graph level embeddings. After constructing the *KNN graph*, we can leverage another GIN model (without READOUT function) to learn the node representations, wherein a node represents a source graph instance.

Feature Concatenation Operator We utilize a concatenation operator that directly appends the structure feature for each graph to its semantic feature. Hence, this operator has no parameters. Overall, given a graph G_i , its final extracted feature can be expressed as $h^{FE}(G_i) = h^{SE}(G_i) \parallel h^{ST}(G_i)$. As a result, the corresponding trainable parameter $\Theta^{FE} = (\Theta^{SE}, \Theta^{ST})$. Importantly, these parameters are shared when learning representations for graphs from the source domain and the target domain, respectively.

VI. MODULE 2: CROSS-DOMAIN GRAPH LEVEL ANOMALY DETECTION

This module contains three components: the one-class classifier, the domain classifier, and the class aligner. We first elucidate the rationale and design of each component, followed by introducing the theory of adversarial training.

One-Class Classifier: We use $h^s(\cdot)$ to represent the source classifier, which has no additional parameters beyond the feature extraction parameters Θ^{FE} . We train a one-class classifier on the source domain by minimizing the following One-Class Deep SVDD objective [39]:

$$\mathcal{L}_{SC}(\mathcal{X}^s, \mathcal{Y}^s; \Theta^{FE}) =: \frac{1}{N_s} \sum_{m=1}^{N_s} \|h^{FE}(G_m) - \mathbf{o}\|_2^2, \quad (5)$$

where $h^{FE}(G_m)$ is the final extracted feature of graph G_m (from the source domain), and $\mathbf{o}$ is the learned center that represents the normality defined in the source domain. Minimizing this SVDD loss ensures the model learns the label information from the source domain.

Domain Classifier: We denote the domain classifier as $h^d(\cdot)$, with trainable parameters Θ^d in addition to parameters Θ^{FE} .

We maximize the binary cross-entropy loss to learn domain-invariant features:

$$\mathcal{L}_{DA}(\mathcal{X}^s, \mathcal{X}^t; \Theta^d, \Theta^{FE}) =:$$

$$\left[\frac{1}{N_s + N_t} \sum_{i=1}^{N_s + N_t} \left[d_i \log \left(\frac{1}{\hat{d}_i} \right) + (1 - d_i) \log \left(\frac{1}{1 - \hat{d}_i} \right) \right] \right], \quad (6)$$

where d_i denotes the binary ground-truth domain label for graph G_i . Specifically, d_i is 0 for graphs from the source domain and 1 for graphs from the target domain. Additionally, $\hat{d}_i$ represents the predicted probability that G_i belongs to the target domain, as determined by $h^d(\cdot)$. More precisely, the prediction $\hat{d}_i$ is given by $\hat{d}_i = h^d(h^{FE}(G_i))$, where $h^{FE}(\cdot)$ is the final feature extractor and $h^d(\cdot)$ is the domain classifier. In loss term (6), we consider the negative log-probability, expressed as $-\log(\hat{d}_i)$. This can be rewritten as $\log(\frac{1}{h^d(h^{FE}(G_i))})$. The goal of this loss function is to maximize it, which encourages the feature extractor $h^{FE}(\cdot)$ to create similar representations for graphs from both the source and target domains. This helps align the domains, making the feature distributions of the source and target domains indistinguishable and reducing discrepancies between them.

Class Aligner: Structure consistency between domains (via parameters-shared feature extractor), discriminability in source domain (via source classifier), and domain-invariant features (via domain classifier) do not necessarily lead to discriminability in the target domain. That is, although we manage to align features cross domains, the learned features may be distorted in the sense that they are not representative of the underlying patterns in the target domain. As a result, features of the normal and abnormal classes in the target domain may exhibit close proximity or even overlap, leading to a high value of $\epsilon^t(h)$ in the target domain. This is known as excessive alignment in [40] and collapse of target neighborhood structure in [41].

To alleviate this problem, we should consider the third term in (1), namely $[\epsilon^s(h^*) + \epsilon^t(h^*)]$, that considers labels from both domains. Although the true label information in the target domain cannot be obtained, we can generate pseudolabels and minimize the class centroid alignment loss:

$$\mathcal{L}_{CA}(\mathcal{X}^s, \mathcal{Y}^s, \mathcal{X}^t, \hat{\mathcal{Y}}^t; \Theta^{FE}) =: [\psi(C_n^s, C_n^t) - \psi(C_a^t, C_n^t)], \quad (7)$$

where $\hat{\mathcal{Y}}^t$ are the *pseudolabels* obtained by directly applying $h^s(\cdot)$ to the target graphs, and $h^s(\cdot)$ is obtained by optimizing loss term (5) in previous iteration. Moreover, C_n^s , C_n^t , and C_a^t represent the centroid of normal source class, normal target class, and anomalous target class, respectively. The function $\psi(\cdot, \cdot)$ is a distance metric such as the euclidean distance. Minimizing this loss can reduce the inter-domain distance between centroids of normal classes, while simultaneously maximizing the intra-domain distances between centroids of normal and abnormal classes within the target domain. Particularly, with an increase of training epochs, we expect that the discrepancy between source domain and target domain is reduced, the data structures are better aligned, the $h^s(\cdot)$ is further improved, and thus the pseudolabels are gradually updated to approach the ground-truth labels, progressively improving the class centroid alignment.

Algorithm 1: Mini-Batch Algorithm of ARMET.

Input: Labelled source graphs $\mathcal{G}^s = \{G_n^s, Y_n^s\}_{n=1}^{N_s}$; Unlabelled target graphs $\mathcal{G}^t = \{G_n^t\}_{n=1}^{N_t}$; Balance parameters λ_1, λ_2 and λ_3 ; Batch size N_b ; Maximal training epochs N_e ; Maximal iteration per epoch N_i

Output: Predicted labels $\hat{\mathcal{Y}}^t$ of target graphs

```

1: Initialise parameters  $\Theta^{FE}, \Theta^d$ 
2: for epoch  $< N_e$  and not converge do
3:   for iteration  $< N_i$  do
4:     Sample a batch  $\mathcal{B}^s$  and  $\mathcal{B}^t$ 
5:     Learn representations using  $h^{FE}$  for  $\mathcal{B}^s$  and  $\mathcal{B}^t$ 
6:     Compute source classifier loss  $\mathcal{L}_{SC}$  using (5)
7:     Compute domain classifier loss  $\mathcal{L}_{DC}$  using (6)
8:     Backpropagate  $\mathcal{L}_{DC}$  and update  $\Theta^d$  using (9)
9:     Compute class aligner loss  $\mathcal{L}_{CA}$  using (7)
10:    Compute total loss  $\mathcal{L}$  using (2)
11:    Backpropagate  $\mathcal{L}$  and update  $\Theta^{FE}$  using (9)
12:  end for
13: end for
14:  $\hat{h}^{FE}(\mathcal{G}^t) \leftarrow$  Use feature extractor with optimised
    parameters  $\hat{\Theta}^{FE}$  to extract features for  $\mathcal{G}^t$ 
15:  $\hat{\mathcal{Y}}^t \leftarrow$  Use optimised source classifier  $\hat{h}^s(\cdot)$  to predict
    labels for  $\hat{h}^{FE}(\mathcal{G}^t)$ 

```

Adversarial Training: It can be seen that the optimization of (5) and (7) involves a minimization w.r.t. parameters Θ^{FE} , while the optimization of (6) concerns a maximization w.r.t. parameters Θ^d and Θ^{FE} . In other words, objective (6) competes against objectives (5) and (7) during training over the overall objective (2). To obtain a good trade-off, adversarial training techniques have been explored, with impressive results [18].

For simplicity, we rewrite $\mathcal{L}(\mathcal{X}^s, \mathcal{Y}^s, \mathcal{X}^t; \Theta^{FE}, \Theta^d)$ as $\mathcal{L}(\Theta^{FE}, \Theta^d)$, which contains two sets of parameters. Adversarial training attempts to find a saddle point $(\hat{\Theta}^{FE}, \hat{\Theta}^d)$ such that $\hat{\Theta}^{FE} = \operatorname{argmin}_{\Theta^{FE}} \mathcal{L}(\Theta^{FE}, \hat{\Theta}^d)$ and $\hat{\Theta}^d = \operatorname{argmax}_{\Theta^d} \mathcal{L}(\hat{\Theta}^{FE}, \Theta^d)$. In other words, we perform a minmax optimization over (2):

$$\min_{\Theta^{FE}} \max_{\Theta^d} [\lambda_1 \mathcal{L}_{SC}(\Theta^{FE}) - \lambda_2 \mathcal{L}_{DA}(\Theta^{FE}, \Theta^d) + \lambda_3 \mathcal{L}_{CA}(\Theta^{FE})]. \quad (8)$$

Hence, the trainable parameters can be optimized alternatively as follows:

$$\begin{aligned} \Theta^d &\leftarrow \Theta^d - \mu \frac{\partial \mathcal{L}_{DA}}{\partial \Theta^d}, \\ \Theta^{FE} &\leftarrow \Theta^{FE} - \mu \left(\lambda_1 \frac{\partial \mathcal{L}_{SC}}{\partial \Theta^{FE}} - \lambda_2 \frac{\partial \mathcal{L}_{DA}}{\partial \Theta^{FE}} + \lambda_3 \frac{\partial \mathcal{L}_{CA}}{\partial \Theta^{FE}} \right), \end{aligned} \quad (9)$$

where μ is the learning rate. We perform mini-batch training with gradient descent and the corresponding pseudo-code is provided in Algorithm 1.

VII. EXPERIMENT SETUP

We aim to answer the following research questions (RQ) via experiments:

TABLE I
DATASETS. #NODES, #EDGES, DEGREE, #ATTR DENOTE THE AVERAGE
NUMBER OF NODES AND EDGES, THE AVERAGE DEGREE, AND THE
DIMENSIONALITY OF NODE ATTRIBUTES, RESPECTIVELY

Data	#Graphs	#Nodes	#Edges	Degree	#Attr
BG	5000(5%)	10	30	3	200
HD	5000(5%)	7	20	2.86	200
SP	5000(5%)	6	24	4	200
TB	5000(5%)	16	52	3.25	200
LL	500(10%)	4.7	4.5	0.96	2
LM	500(10%)	4.7	3.1	0.66	2
LH	500(10%)	4.7	3.3	0.70	2

RQ1 How does ARMET perform when compared to state-of-the-art graph level anomaly detection methods?

RQ2 How does each component of ARMET affect the performance? (Ablation study)

RQ3 How does the performance of ARMET change with different hyperparameter values? (Sensitivity analysis)

In addition, to get a better understanding of ARMET, we utilize t-SNE [42] to visualize the learned graph representations from both domains.

A. Benchmark Datasets

As summarized in Table I, we study the following datasets collected from various application fields of graph mining:

System Logs We use four benchmark datasets for log anomaly detection, as converting logs into graphs and then leveraging GNNs to detect anomalies can achieve superior performance [2], [43]. Hence, following [2], we construct four graph datasets: HDFS (HD) [44], BGL (BG), SPIRIT (SP), and THUNDER-BIRD (TH) [45], where each dataset contains 5000 graphs with 5% anomalous graphs. Particularly, HD consists of Hadoop Distributed File System logs while BG, SP, and TH containing system logs collected from three different supercomputing systems. For this group of datasets, we create 12 transfer tasks.

Letter Drawings We use three benchmark graph datasets with letter drawings of varying levels of distortion: low (LL), medium (LM), and high (LH) [46]. Following the practice of downsampling classification datasets for anomaly detection [47], for each dataset the letters N, M, and W are selected as the normal class, comprising a total of 450 instances (150 instances per letter), while the letter F is chosen as the anomalous class (downsampled to 50 instances). For each graph, a node denotes the end point of a line, an edge represents a line, and a node attribute represents its two-dimensional coordinate. For these datasets, we create 6 transfer tasks.

Discussion on Dataset Selection: Some commonly used benchmark graph datasets, including BZR, DHFR and COX2 (small molecules), as well as IMDB-Binary and REDDIT (social networks), have been excluded from our analysis for the following reasons: 1) the UDA assumptions do not hold for them, as these datasets have different definitions of classification problem, namely they have different label semantics; 2) transferability cannot be guaranteed due to huge domain discrepancy between datasets. For instance, node attributes of graphs in these datasets have different dimensionalities and/or meanings, and these datasets have very different graph statistics; and 3) even

supervised classification methods can only achieve very limited in-domain classification accuracy (i.e., with an accuracy lower than 0.7) on these datasets.

B. Baselines

We compare ARMET to the following baselines:

- *Unsupervised graph level anomaly detection methods:* We directly apply OCGIN [6], GLAM [5], and GLocalKD [7] on unlabeled target graphs.
- *Traditional Domain Adaptation Methods:* ADDA [19], CDAN [20], and DeepCoral [17] are state-of-the-art UDA methods for image classification. As they are not designed for graph-structured data, we modify these models for cross-domain graph level anomaly detection by following [32].

To explore the ceiling performance of GLAD methods, we additionally present the outcomes of iGAD [48], a supervised method that is not considered a direct competitor due to its reliance on labeled target data.

C. Evaluation

We employ the widely used Area Under the Curve of the Receiver Operating Characteristics curve (AUC ROC), Area Under the Curve of Precision Recall (AUC PR), and F1-Score to evaluate and compare the different methods, where a higher value (closer to 1) represents better anomaly detection accuracy. Particularly, we report the average values of AUC ROC, AUC PR, F1-Score and the corresponding standard deviations across 10 independent runs.

D. Implementation and Model Configuration

We use the publicly available implementations of OCGIN¹, GLAM,¹ GLocalKD² and iGAD³ with their recommended configurations. Besides, ADDA,⁴ CDAN⁵ and DeepCoral⁵ are adapted from their publicly available implementations. As they are not designed for graph-structured data, we modify these models for cross-domain graph level anomaly detection by following the practice in [32]:

- CDAN: We replace the AlexNet encoder with a GIN plus a mean readout function;
- ADDA: Similarly, we replace the CNN encoder with a GIN plus a mean readout function;
- DeepCoral: We replace the CNN encoder (CaffeNet) with a GIN plus a mean readout function and apply CORAL loss to the last classification layer.

For ARMET, by following the practice in [17], the values of hyperparameters λ_1 , λ_2 and λ_3 can be set such that, after the training process (e.g., 100 epochs), the losses associated with one-class classification ($\mathcal{L}_{SC}$), domain discrimination

¹[Online]. Available: <https://github.com/lingxiaoshawn/glam>

²[Online]. Available: <https://github.com/RongrongMa/GLocalKD>

³[Online]. Available: <https://github.com/graph-level-anomalies/iGAD>

⁴[Online]. Available: <https://github.com/yuhui-zh15/pytorch-adda>

⁵[Online]. Available: <https://github.com/agrija9/deep-unsupervised-domain-adaptation>

TABLE II
DESCRIPTION OF HYPERPARAMETERS AND THEIR RECOMMENDED VALUES

Symbol	Meaning	Range
d	embedding dimension	{16, 32, 64 , 96, 128, 160, 192, 224, 256, 300}
k	number of neighbours in KDTree	{2, 5 , 8, 11, 14, 17, 20}
L	number of layers	{1, 2 , 3, 4, 5}
Bs	batch size	{16, 32, 64, 128 , 256, 512 , 1024}
λ	weight decay parameter	{ 0.0001 , 0.001, 0.01, 0.1}
η	learning rate	{0.0001, 0.001, 0.01 }
Re	readout function	mean , sum, max
Ep	Epochs for training	{ 100 , 200, 300, 400}

Range indicates the values that we have tested in the sensitivity analysis, and boldfaced values represent the values suggested to use in the experiments.

($\mathcal{L}_{DA}$) and class-alignment ($\mathcal{L}_{CA}$) are approximately at the same magnitude (after being multiplied by their corresponding weights). The rational behind it is that we aim to learn feature representations that are both label-discriminative and domain-invariant [17]. Particularly, we found that there is usually not a single set of optimal weights for each transfer task, as pointed out in multi-task learning by [49]. Although this hyperparameter tuning method works well, it is time-consuming and labor-intensive. For simplicity, we can adapt a *de-facto* strategy for hyperparameter tuning in UDA, namely splitting the target dataset according to a ratio of 20% : 80%, where the 20% data with labels is used as the validation dataset to select these three hyperparameters and the remaining 80% data without labels is used as the test data.

For fair comparisons, all GIN models used in different domain adaptation methods are configured with the same backbone architecture, namely a backbone of two layers (64 hidden units) with each layer followed by a ReLU activation. Besides, all neural network based methods are trained using mini-batch gradient descent, with a batchsize of 512 on log anomaly detection datasets and a batchsize of 128 on other datasets respectively, initial learning rate of 0.01, weight decay rate of 0.0001, and a maximum of 200 training epochs. The settings for these hyperparameters for ARMET are summarised in Table II. Other algorithm-specific hyperparameters are set in accordance with their respective references given by the original authors.

E. Training Hardware and Reproducibility

We implemented and ran all algorithms in Python 3.8, using PyTorch [50] and PyTorch Geometric [51] when applicable, on a workstation equipped with an Intel i7-11700KF CPU and Nvidia RTX3070 GPU. For reproducibility, all code and datasets are made available on GitHub.⁶

VIII. EXPERIMENT RESULTS AND ANALYSIS

We answer the three research questions as follows.

TABLE III
ANOMALY DETECTION ACCURACY (AVERAGE AUC ROC AND CORRESPONDING STANDARD DEVIATIONS ACROSS 10 RUNS) OF UNSUPERVISED METHODS (OCCIN, GLAM, AND GLOCALKD) UNDER TWO SCENARIOS: 1) WHEN THE TRAINING DATASET IS CLEAN, NAMELY CONTAINING EXCLUSIVELY NORMAL GRAPHS (SHOWN ON THE LEFT SIDE OF ‘→’), AND 2) WHEN THE TRAINING DATASET IS CONTAMINATED, NAMELY CONTAINING BOTH NORMAL AND ABNORMAL INSTANCES (SHOWN ON THE RIGHT SIDE OF ‘→’)

Dataset	OCGIN	GLAM	GLocalKD
BG	0.93±0.04 → 0.71±0.11(↓)	0.64±0.09 → 0.74±0.06(↑)	0.93±0.02 → 0.91±0.01(↓)
SP	0.62±0.16 → 0.59±0.16(↓)	0.70±0.09 → 0.71±0.09(↑)	0.81±0.08 → 0.66±0.15(↓)
HD	0.98±0.01 → 0.88±0.03(↓)	0.87±0.06 → 0.77±0.09(↓)	0.45±0.05 → 0.45±0.05(−)
TH	0.81±0.20 → 0.43±0.28(↓)	0.93±0.05 → 0.96±0.09(↑)	0.93±0.05 → 0.76±0.11(↓)
LL	0.99±0.00 → 0.70±0.10(↓)	0.82±0.18 → 0.72±0.18(↓)	0.56±0.11 → 0.53±0.10(↓)
LM	0.91±0.02 → 0.64±0.05(↓)	0.57±0.10 → 0.55±0.07(↓)	0.56±0.11 → 0.53±0.10(↓)
LH	0.65±0.10 → 0.52±0.08(↓)	0.66±0.08 → 0.54±0.10(↓)	0.56±0.11 → 0.53±0.10(↓)

A. Detection Accuracy (RQ1)

We perform the following analysis based on the experiment results in Tables III, IV and VI. Particularly, the results in terms of AUC PR and F1-Score are consistent with those in terms of AUC ROC. Therefore, these results are either deferred to Table VIII or omitted.

1) *Accuracy of Single-Domain GLAD*: Table III demonstrates that unsupervised/semi-supervised methods OCGIN, GLAM,⁷ and GLocalKD generally suffer from large performance degradation when the training data is contaminated with anomalies. Moreover, these unsupervised/semi-supervised methods deliver very unstable results across different datasets even when the training data is clean. For example, GLocalKD yields high performance on BGL (ROC = 0.91), but very poor performance on HDFS (ROC = 0.45). In contrast, Table IV shows that supervised method iGAD achieves perfect results, with an ROC of 1.0, on all cases. This indicates that these anomaly detection problems are solvable when sufficient labeled data is available in the target domain. However, this is unrealistic in many real-world scenarios as fully labeled data is usually expensive and often even impossible to obtain in practice.

2) *Accuracy of Traditional UDA Methods*: Table IV shows that when there is only one class in the source domain, ADDA, CDAN, and DeepCoral behave like random guessing (e.g., with ROC ≈ 0.50), performing much worse than ARMET for most transfer tasks. The potential factors contributing to their subpar performance are as follows. First, CDAN considers the conditional distribution that captures the cross-variance between feature representations and classifier predictions, but the conditional distribution degenerates to the marginal distribution when there is only one class in the source domain. Second, DeepCoral aligns the second-order statistics of layer activations in the source encoder and the target encoder, leading to random-guessing results since the source training data contains only normal graphs while the target training data includes both normal and abnormal graphs (and thus the statistics of layer activations should be different by nature). Third, ADDA performs poorly as it utilises different encoders for the source and target domains, implying the importance of parameters-shared models

⁶[Online]. Available: <https://github.com/ZhongLIFR/ARMET/>

⁷With a few exceptions for GLAM, where the performance is slightly increased or stable.

TABLE IV
ANOMALY DETECTION ACCURACY (AVERAGE AUC ROC AND CORRESPONDING STANDARD DEVIATIONS ACROSS 10 RUNS)

Dataset	Supervised	Unsupervised			Traditional UDA			Ours
	iGAD	OCGIN	GLAM	GLocalKD	ADDA	CDAN	DeepCoral	ARMET
HD → BG					0.50 $\pm$ 0.00	0.50 $\pm$ 0.02	0.50 $\pm$ 0.00	<u>0.77</u> $\pm$ 0.06
SP → BG	1.0 $\pm$ 0.00	0.71 $\pm$ 0.11	0.74 $\pm$ 0.06	0.91 $\pm$ 0.01	0.50 $\pm$ 0.00	0.49 $\pm$ 0.03	0.50 $\pm$ 0.00	<u>0.85</u> $\pm$ 0.07
TH → BG					0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	<u>0.83</u> $\pm$ 0.05
BG → SP					0.50 $\pm$ 0.00	0.50 $\pm$ 0.01	0.50 $\pm$ 0.00	0.89 $\pm$ 0.07
HD → SP	1.0 $\pm$ 0.00	0.59 $\pm$ 0.16	0.71 $\pm$ 0.09	0.66 $\pm$ 0.15	0.50 $\pm$ 0.00	0.51 $\pm$ 0.04	0.50 $\pm$ 0.00	0.91 $\pm$ 0.04
TH → SP					0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.88 $\pm$ 0.06
BG → HD					0.50 $\pm$ 0.00	0.50 $\pm$ 0.01	0.50 $\pm$ 0.00	<u>0.79</u> $\pm$ 0.04
SP → HD	1.0 $\pm$ 0.00	0.88 $\pm$ 0.03	0.77 $\pm$ 0.09	0.45 $\pm$ 0.05	0.50 $\pm$ 0.00	0.51 $\pm$ 0.01	0.50 $\pm$ 0.00	0.90 $\pm$ 0.03
TH → HD					0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	<u>0.81</u> $\pm$ 0.06
BG → TH					0.50 $\pm$ 0.00	0.60 $\pm$ 0.20	0.50 $\pm$ 0.00	1.0 $\pm$ 0.00
SP → TH	1.0 $\pm$ 0.00	0.43 $\pm$ 0.28	0.96 $\pm$ 0.09	0.76 $\pm$ 0.11	0.50 $\pm$ 0.00	0.69 $\pm$ 0.24	0.50 $\pm$ 0.00	1.0 $\pm$ 0.00
HD → TH					0.50 $\pm$ 0.00	0.72 $\pm$ 0.24	0.50 $\pm$ 0.00	1.0 $\pm$ 0.00
LM → LL	1.0 $\pm$ 0.00	0.70 $\pm$ 0.10	0.72 $\pm$ 0.18	0.53 $\pm$ 0.10	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.98 $\pm$ 0.02
LH → LL					0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.96 $\pm$ 0.00
LL → LM	1.0 $\pm$ 0.00	0.64 $\pm$ 0.05	0.55 $\pm$ 0.07	0.53 $\pm$ 0.10	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.88 $\pm$ 0.02
LH → LM					0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.78 $\pm$ 0.03
LL → LH	1.0 $\pm$ 0.00	0.52 $\pm$ 0.08	0.54 $\pm$ 0.10	0.53 $\pm$ 0.10	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.79 $\pm$ 0.02
LM → LH					0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.50 $\pm$ 0.00	0.80 $\pm$ 0.05

The best and runner-up results are highlighted in BOLD and underlined, respectively. For unsupervised methods, we report the ROC values when the training dataset contains both normal and abnormal instances. For cross-domain methods, we report the ROC values when the source dataset contains only normal graphs. The results of supervised method are shown only to be able to put the performance of graph level anomaly detection methods in perspective (i.e., iGAD should not be considered as baseline). Moreover, the poor performance of traditional UDA methods (namely ADDA, CDAN and DeepCoral) are not due to intentionally choosing weak baselines or poor hyper-parameters. The underlying reasons for their subpar performance are given in Section VIII-A2.

TABLE V
ABLATION STUDIES WITH THE FOLLOWING STRIPPED-DOWN VARIATIONS, WHERE ‘✓’ AND ‘✗’ MEAN THE CORRESPONDING COMPONENT IS INCLUDED OR EXCLUDED, RESPECTIVELY

Components	SD	/SF	/SC	/DC	/CA	ARMET
Structure Feature Extractor	✓	✗	✓	✓	✓	✓
One-Class Classifier	✓	✓	✗	✓	✓	✓
Domain Classifier	✗	✓	✓	✗	✓	✓
Class Aligner	✗	✓	✓	✓	✗	✓

when encoding graphs from two domains. The efficacy of the source encoder degrades to random-guessing when the source domain contains only a single class, as the source classifier fails to acquire any informative knowledge. Furthermore, these traditional UDA methods are primarily designed for balanced multiple classification problems, making them ill-suited for anomaly detection, which involves an extremely imbalanced binary classification task. As a reference, we report the performance of ADDA, CDAN, and DeepCoral when the source domain contains both labeled anomalous and normal instances in Table VI. One can see that the performance of ADDA, CDAN, and DeepCoral is largely boosted with auxiliary label information in most transfer tasks. However, even with auxiliary labels, they are still outperformed by ARMET in most cases.

3) *Accuracy of ARMET*: Table IV shows that ARMET achieves the best performance on 13 out of 18 transfer tasks, and the second-best performance on the remaining tasks. Further, it

TABLE VI
ANOMALY DETECTION ACCURACY (AVERAGE AUC ROC AND CORRESPONDING STANDARD DEVIATIONS ACROSS 10 RUNS) OF TRADITIONAL UDA METHODS (INCLUDING ADDA, CDAN, AND DEEPCORAL) WHEN THE SOURCE DOMAIN CONTAINS BOTH LABELED ANOMALOUS AND NORMAL INSTANCES. (‘↑’ INDICATES THE PERFORMANCE IS BOOSTED COMPARED TO THE CASE WHERE THE SOURCE DOMAIN CONTAINS ONLY NORMAL INSTANCES, WHILE ‘↓’ MEANS IT IS DEGRADED.)

Dataset	ADDA*	CDAN*	DeepCoral*	ARMET
HD → BG	0.50 $\pm$ 0.03(−)	0.64 $\pm$ 0.11(↑)	0.51 $\pm$ 0.11(↑)	0.77 $\pm$ 0.02
SP → BG	0.47 $\pm$ 0.00(↓)	0.48 $\pm$ 0.06(↓)	0.49 $\pm$ 0.01(↓)	0.85 $\pm$ 0.07
TH → BG	0.49 $\pm$ 0.02(↓)	0.50 $\pm$ 0.00(−)	0.50 $\pm$ 0.00(−)	0.83 $\pm$ 0.05
BG → SP	0.56 $\pm$ 0.17(↑)	0.50 $\pm$ 0.08(−)	0.49 $\pm$ 0.19(↓)	0.89 $\pm$ 0.07
HD → SP	0.63 $\pm$ 0.24(↑)	0.50 $\pm$ 0.03(↓)	0.60 $\pm$ 0.16(↑)	0.91 $\pm$ 0.04
TH → SP	0.54 $\pm$ 0.13(↑)	0.50 $\pm$ 0.00(−)	0.50 $\pm$ 0.01(−)	0.88 $\pm$ 0.06
BG → HD	0.62 $\pm$ 0.13(↑)	0.50 $\pm$ 0.08(−)	0.48 $\pm$ 0.11(↓)	0.79 $\pm$ 0.04
SP → HD	0.57 $\pm$ 0.14(↑)	0.51 $\pm$ 0.03(−)	0.51 $\pm$ 0.01(↑)	0.90 $\pm$ 0.03
TH → HD	0.57 $\pm$ 0.06(↑)	0.51 $\pm$ 0.02(↑)	0.50 $\pm$ 0.00(−)	0.81 $\pm$ 0.06
BG → TH	0.48 $\pm$ 0.01(↓)	0.48 $\pm$ 0.20(↓)	0.54 $\pm$ 0.24(↑)	1.00 $\pm$ 0.00
SP → TH	0.53 $\pm$ 0.17(↑)	0.77 $\pm$ 0.24(↑)	0.73 $\pm$ 0.27(↑)	1.00 $\pm$ 0.00
HD → TH	0.58 $\pm$ 0.22(↑)	0.50 $\pm$ 0.17(↓)	0.84 $\pm$ 0.21(↑)	1.00 $\pm$ 0.00
LM → LL	0.93 $\pm$ 0.15(↑)	1.0 $\pm$ 0.00(↑)	1.0 $\pm$ 0.00(↑)	0.98 $\pm$ 0.02
LH → LL	0.78 $\pm$ 0.23(↑)	0.99 $\pm$ 0.00(↑)	0.95 $\pm$ 0.05(↑)	0.96 $\pm$ 0.00
LL → LM	0.53 $\pm$ 0.04(↑)	0.74 $\pm$ 0.00(↑)	0.63 $\pm$ 0.01(↑)	0.88 $\pm$ 0.02
LH → LM	0.65 $\pm$ 0.05(↑)	0.89 $\pm$ 0.00(↑)	0.76 $\pm$ 0.03(↑)	0.78 $\pm$ 0.03
LL → LH	0.50 $\pm$ 0.01(−)	0.78 $\pm$ 0.00(↑)	0.59 $\pm$ 0.02(↑)	0.79 $\pm$ 0.02
LM → LH	0.66 $\pm$ 0.08(↑)	0.87 $\pm$ 0.00(↑)	0.80 $\pm$ 0.01(↑)	0.80 $\pm$ 0.05

largely outperforms unsupervised graph level anomaly detection methods on certain target datasets, demonstrating the benefit of leveraging label information from a different but related domain.

TABLE VII
ABLATION STUDY (AVERAGE AUC ROC WITH 10 RUNS)

Dataset	ARMET	SD	/SF	/SC	/DC	/CA
HD $\rightarrow$ BG	0.77	0.55	0.74	0.48	0.55	0.57
SP $\rightarrow$ BG	0.85	0.66	0.81	0.55	0.60	0.58
TH $\rightarrow$ BG	0.83	0.60	0.74	0.64	0.64	0.77
BG $\rightarrow$ SP	0.89	0.71	0.55	0.67	0.62	0.59
HD $\rightarrow$ SP	0.91	0.74	0.62	0.73	0.72	0.64
TH $\rightarrow$ SP	0.88	0.61	0.52	0.73	0.77	0.58
BG $\rightarrow$ HD	0.79	0.60	0.58	0.59	0.59	0.60
SP $\rightarrow$ HD	0.90	0.61	0.74	0.67	0.53	0.67
TH $\rightarrow$ HD	0.81	0.68	0.66	0.70	0.64	0.68
BG $\rightarrow$ TH	1.0	0.77	0.61	1.0	0.90	0.73
SP $\rightarrow$ TH	1.0	0.56	0.06	1.0	0.77	0.68
HD $\rightarrow$ TH	1.0	0.88	0.79	1.0	0.96	0.88
LM $\rightarrow$ LL	0.98	0.98	0.71	0.70	0.99	0.98
LH $\rightarrow$ LL	0.96	0.92	0.65	0.73	0.87	0.90
LL $\rightarrow$ LM	0.88	0.74	0.75	0.62	0.77	0.85
LH $\rightarrow$ LM	0.78	0.76	0.55	0.42	0.71	0.73
LL $\rightarrow$ LH	0.79	0.74	0.60	0.59	0.72	0.71
LM $\rightarrow$ LH	0.80	0.71	0.56	0.63	0.67	0.70

SD: trained on source domain and then directly applied on target domain; /SF: ARMET without structure feature extractor; /SC: ARMET without source one-class classifier; /DC: ARMET without domain classifier; /CA: ARMET without class aligner.

For example, transferring knowledge from **HD** to **SP** leads to 54%, 28%, and 38% performance gains compared to OCGIN, GLAM, and GLocalKD, respectively. Finally, ARMET achieves impressive results under the setting that the source domain contains only normal graphs, while other traditional unsupervised domain adaptation methods provide “random-guessing” results, making ARMET more applicable to real-world scenarios.

B. Ablation Study (RQ2)

As summarised in Table V, we here compare ARMET to five of its stripped-down variations: SD is trained on the source domain and then directly applied on the target domain; /SF is ARMET trained without the structure feature extractor; /SC is ARMET trained without the source one-class classifier; /DC is ARMET trained without the domain classifier; and /CA is ARMET trained without the class aligner. We have the following important observations from Table VII:

- *SD versus ARMET.* ARMET is always superior to the case where we train the model on the source domain and then directly apply it to the target domain. This exemplifies the benefits and necessity of performing transfer learning.
- */SF versus ARMET.* ARMET consistently outperforms its counterpart without the structural feature extractor. This underlines the importance of considering the neighbourhood information in a set of graphs for CD-GLAD.
- */SC versus ARMET.* It shows that explicitly incorporating the source label information via a source classifier is typically beneficial.




Fig. 2. Sensitivity analysis of hyperparameters on four representative transfer tasks (HD $\rightarrow$ SP, TH $\rightarrow$ BG, HD $\rightarrow$ TH, LM $\rightarrow$ LL). Average results over five runs.

- */DC versus ARMET.* In most cases, explicitly aligning the embeddings of both domains via a domain classifier is favorable. In certain cases, the presence of a domain classifier may have a small adverse impact on performance. One possible reason is that the domain classifier overly distorts the embedding space by aligning the embedding space in a brute-force manner. Another possible reason is that the embeddings are also implicitly aligned via the parameters-shared feature extractor and the class aligner.
- */CA versus ARMET.* The removal of the class aligner always leads to a performance degradation. This corroborates the importance of considering labels (or pseudo-labels) in the target domain when performing CD-GLAD.

C. Sensitivity Analysis (RQ3)

We examine the effects of the following hyperparameters on the detection performance, and Fig. 2 depicts the selected results on four representative transfer tasks:

- The number of embedding dimensions d in parameters-shared feature extractor: most transfer tasks can achieve the best performance with an embedding dimension of 64. Small fluctuations can be observed with varying number of dimensions. Moreover, an overly small value of d may lead to suboptimal performance, while a large value introduces a considerable computational burden.
- The number of hidden layers L in the GIN model in the parameters-shared feature extractor: optimal performances are obtained when $L = 1$ or $L = 2$ on most transfer tasks. A further increase of its value usually results in performance degradation, which is widely known as over-smoothing [52].
- The number of neighbours k in the KNN graph: most transfer tasks can obtain the best performance on a wide range of k 's values (namely $1 \leq k \leq 10$). However, further increasing the value of k may cause large performance fluctuations.

D. Visualisation With T-SNE

Fig. 3 visualizes the domain alignment and anomaly separation process on the transfer task **LM** $\rightarrow$ **LL**. The target normal graphs (green dots) are progressively adapted to the source normal graphs (blue dots), while the target anomalous graphs (red crosses) become increasingly separable in the embedding space.

TABLE VIII
ANOMALY DETECTION ACCURACY: AVERAGE **AUC PR** AND CORRESPONDING STANDARD DEVIATIONS ACROSS 10 RUNS (TOP), AND AVERAGE **F1-SCORE**(BINARY) AND CORRESPONDING STANDARD DEVIATIONS ACROSS 10 RUNS (BOTTOM)

Dataset	Supervised	Unsupervised			Traditional UDA			Ours
	iGAD	OCGIN	GLAM	GLocalKD	ADDA	CDAN	DeepCoral	ARMET
HD → BG	1.0±0.00	0.19±0.09	0.14±0.03	0.46 ±0.10	0.05±0.00	0.05±0.01	0.05±0.00	0.14±0.09
SP → BG					0.05±0.00	0.05±0.00	0.05±0.00	<u>0.25</u> ±0.13
TH → BG					0.05±0.00	0.05±0.00	0.05±0.00	<u>0.29</u> ±0.11
BG → SP	1.0±0.00	0.08±0.03	0.09±0.02	0.06±0.01	0.05±0.00	0.05±0.00	0.05±0.00	0.76 ±0.12
HD → SP					0.05±0.00	0.09±0.08	0.05±0.00	0.83 ±0.15
TH → SP					0.05±0.00	0.05±0.00	0.05±0.00	0.90 ±0.05
BG → HD	1.0±0.00	0.43±0.06	0.62 ±0.07	0.40±0.04	0.05±0.00	0.05±0.00	0.05±0.00	<u>0.48</u> ±0.15
SP → HD					0.05±0.00	0.06±0.02	0.05±0.00	0.29±0.10
TH → HD					0.05±0.00	0.05±0.00	0.05±0.00	<u>0.49</u> ±0.10
BG → TH	1.0±0.00	0.12±0.15	0.51±0.14	0.08±0.02	0.06±0.00	0.06±0.00	0.06±0.00	0.93 ±0.13
SP → TH					0.06±0.00	0.19±0.29	0.06±0.00	1.0 ±0.00
HD → TH					0.06±0.00	0.41±0.48	0.06±0.00	0.98 ±0.01
LM → LL	1.0±0.00	0.19±0.06	0.42±0.20	0.13±0.03	0.03±0.00	0.03±0.00	0.03±0.00	0.92 ±0.06
LH → LL					0.03±0.00	0.03±0.00	0.03±0.00	0.51 ±0.15
LL → LM	1.0±0.00	0.36±0.04	0.30±0.07	0.13±0.03	0.10±0.00	0.10±0.00	0.10±0.00	0.62 ±0.11
LH → LM					0.10±0.00	0.10±0.00	0.10±0.00	0.37 ±0.05
LL → LH	1.0±0.00	0.25±0.04	0.39±0.09	0.13±0.03	0.10±0.00	0.10±0.00	0.10±0.00	0.43 ±0.08
LM → LH					0.10±0.00	0.10±0.00	0.10±0.00	0.47 ±0.12
HD → BG	1.0±0.00	0.15±0.08	0.18±0.09	0.34 ±0.06	0.00±0.00	0.03±0.02	0.00±0.00	0.08±0.15
SP → BG					0.00±0.00	0.00±0.00	0.00±0.00	0.15±0.18
TH → BG					0.00±0.00	0.00±0.00	0.00±0.00	<u>0.29</u> ±0.11
BG → SP	1.0±0.00	0.05±0.08	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.76 ±0.18
HD → SP					0.00±0.00	0.09±0.18	0.00±0.00	0.79 ±0.39
TH → SP					0.00±0.00	0.00±0.00	0.00±0.00	0.85 ±0.06
BG → HD	1.0±0.00	0.47±0.07	0.56 ±0.10	0.26±0.02	0.00±0.00	0.02±0.04	0.00±0.00	<u>0.50</u> ±0.13
SP → HD					0.00±0.00	0.04±0.06	0.00±0.00	0.09±0.09
TH → HD					0.00±0.00	0.00±0.00	0.00±0.00	<u>0.50</u> ±0.10
BG → TH	1.0±0.00	0.07±0.22	0.55±0.32	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.97 ±0.05
SP → TH					0.00±0.00	0.17±0.37	0.00±0.00	1.0 ±0.00
HD → TH					0.00±0.00	0.39±0.53	0.00±0.00	0.99 ±0.00
LM → LL	0.99±0.01	0.14±0.07	0.39±0.21	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.91 ±0.01
LH → LL					0.00±0.00	0.00±0.00	0.00±0.00	0.57 ±0.10
LL → LM	0.99±0.01	0.40±0.03	0.29±0.09	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.62 ±0.11
LH → LM					0.00±0.00	0.00±0.00	0.00±0.00	0.41 ±0.07
LL → LH	0.99±0.01	0.25±0.08	0.37±0.09	0.00±0.00	0.00±0.00	0.00±0.00	0.00±0.00	0.49 ±0.06
LM → LH					0.00±0.00	0.00±0.00	0.00±0.00	0.48 ±0.12

IX. DISCUSSION ON TRANSFERABILITY

As shown in Table VII, when comparing ARMET with its counterpart without explicit transfer learning (namely the column ‘SD’ that represents the scenarios where we train the model on the source domain and then directly apply it to the target domain), the performance gains from transfer learning are consistently non-negative. Moreover, when comparing ARMET with the best performing single-domain graph level anomaly detectors that are trained directly on the target domain, the

performance gains are often positive (in 13 out of 18 cases, see Table IV). This further demonstrates the benefits of transfer learning.

The underlying reason why transfer learning is beneficial in this context is that the extent of “relatedness” among System Logs datasets (i.e., BG, HD, SP and TH) is large enough, and so is the extent of “relatedness” among Letter Drawings datasets (i.e., LL, LM, and LH). In cross domain graph-level anomaly detection, it is crucial to ensure that the semantic meanings of




Fig. 3. T-SNE visualization of the graph representation space during domain adaptation on task **LM** $\rightarrow$ **LL**. From left to right: the number of epochs correspond to zero, two, five and one hundred, respectively.

normal patterns in the source and target domains are approximately the same. For instance, in System Logs data, these are normal patterns of system operations, while in Letter Drawings data, they represent the same set of letters.

As in other typical unsupervised domain adaptation settings, we assume that the source and target domains are different yet related. However, the extent of “relatedness” in this study is ensured based on our domain knowledge rather than a quantifiable transferability metric. To our knowledge, how to quantify the extent of this “relatedness” (namely *transferability* cross datasets) remains a long-standing problem [53]. Notably, when this “relatedness” is low, it may introduce negative transfer [54]. Moreover, it is critical to note that the transferability is usually *asymmetrical*, e.g., the performance gain for transfer task SP $\rightarrow$ TH is 0.44 (namely 1.0 minus 0.56 in terms of ROC AUC), which is different from the performance gain for TH $\rightarrow$ SP (namely 0.23 in terms of ROC AUC). Therefore, we should exercise caution (especially in safety-critical fields such as healthcare) when the transferability from the source domain to the target domain cannot be guaranteed, whether from a quantifiable metric or domain knowledge perspective.

X. CONCLUSION

This paper studies the problem of cross-domain graph level anomaly detection, wherein a set of unlabelled graphs from the target domain and a set of normal graphs from a different but related domain are given. We propose ARMET, a theoretically motivated, novel method to solve this widely encountered but largely understudied problem. Specifically, ARMET consists of four components: a feature extractor, an adversarial domain classifier, a one-class classifier, and a class aligner. It is the first attempt to combine graph neural networks and adversarial domain adaptation techniques for performing graph level anomaly detection tasks. Extensive experiments demonstrate the efficacy of ARMET and its superiority to single-domain graph level anomaly detection methods and traditional unsupervised domain adaptation methods. Moreover, extensive ablation studies validate the benefits of incorporating each component into ARMET. Understanding and quantifying the transferability across different graph domains should be addressed in future research.

REFERENCES

- [1] C. Qiu, M. Kloft, S. Mandt, and M. Rudolph, “Raising the bar in graph-level anomaly detection,” 2022, *arXiv:2205.13845*.
- [2] Z. Li, J. Shi, and M. Van Leeuwen, “Graph neural networks based log anomaly detection and explanation,” in *Proc. IEEE/ACM 46th Int. Conf. Softw. Eng.: Companion Proc.*, 2024, pp. 306–307.
- [3] H. G. Vogel, W. H. Vogel, H. G. Vogel, G. Müller, J. Sandow, and B. A. Schölkens, *Drug Discovery and Evaluation: Pharmacological Assays*, vol. 2. Berlin, Germany: Springer, 1997.
- [4] T. Lanciano, F. Bonchi, and A. Gionis, “Explainable classification of brain networks via contrast subgraphs,” in *Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2020, pp. 3308–3318.
- [5] L. Zhao, S. Sawlani, A. Srinivasan, and L. Akoglu, “Graph anomaly detection with unsupervised GNNs,” 2022, *arXiv:2210.09535*.
- [6] L. Zhao and L. Akoglu, “On using classification datasets to evaluate graph outlier detection: Peculiar observations and new insights,” *Big Data*, vol. 11, no. 3, pp. 151–180, 2021.
- [7] R. Ma, G. Pang, L. Chen, and A. van denHengel, “Deep graph-level anomaly detection by glocal knowledge distillation,” in *Proc. 15th ACM Int. Conf. Web Search Data Mining*, 2022, pp. 704–714.
- [8] G. Wilson and D. J. Cook, “A survey of unsupervised deep domain adaptation,” *ACM Trans. Intell. Syst. Technol.*, vol. 11, no. 5, pp. 1–46, 2020.
- [9] E. Nie, S. Liang, H. Schmid, and H. Schütze, “Cross-lingual retrieval augmented prompt for low-resource languages,” in *Proc. Findings Assoc. Computat. Linguistics*, 2023, pp. 8320–8340.
- [10] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan, “A theory of learning from different domains,” *Mach. Learn.*, vol. 79, pp. 151–175, 2010.
- [11] S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010.
- [12] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, “A survey on deep transfer learning,” in *Proc. 27th Int. Conf. Artif. Neural Netw.*, Rhodes, Greece, Springer, 2018, pp. 270–279.
- [13] F. Chen, Y.-C. Wang, B. Wang, and C.-C. J. Kuo, “Graph representation learning: A survey,” *APSIPA Trans. Signal Inf. Process.*, vol. 9, 2020, Art. no. e15.
- [14] L. Akoglu, H. Tong, and D. Koutra, “Graph based anomaly detection and description: A survey,” *Data Mining Knowl. Discov.*, vol. 29, pp. 626–688, 2015.
- [15] X. Ma et al., “A comprehensive survey on graph anomaly detection with deep learning,” *IEEE Trans. Knowl. Data Eng.*, vol. 35, no. 12, pp. 12012–12038, Dec. 2023.
- [16] H. T. Nguyen, P. J. Liang, and L. Akoglu, “Anomaly detection in large labeled multi-graph databases,” 2020, *arXiv: 2010.03600*.
- [17] B. Sun and K. Saenko, “Deep coral: Correlation alignment for deep domain adaptation,” in *Proc. Eur. Conf. Comput. Vis.*, Amsterdam, The Netherlands, Springer, 2016, pp. 443–450.
- [18] Y. Ganin et al., “Domain-adversarial training of neural networks,” *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [19] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, “Adversarial discriminative domain adaptation,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 7167–7176.
- [20] M. Long, Z. Cao, J. Wang, and M. I. Jordan, “Conditional adversarial domain adaptation,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 1647–1657.
- [21] E. Vural, “Domain adaptation on graphs by learning graph topologies: Theoretical analysis and an algorithm,” *Turkish J. Elect. Eng. Comput. Sci.*, vol. 27, no. 3, pp. 1619–1635, 2019.
- [22] Y. Zhang, G. Song, L. Du, S. Yang, and Y. Jin, “DANE: Domain adaptive network embedding,” 2019, *arXiv: 1906.00684*.
- [23] M. Wu, S. Pan, C. Zhou, X. Chang, and X. Zhu, “Unsupervised domain adaptive graph convolutional networks,” in *Proc. Web Conf.*, 2020, pp. 1457–1467.

- [24] M. Pilanci and E. Vural, "Domain adaptation on graphs by learning aligned graph bases," *IEEE Trans. Knowl. Data Eng.*, vol. 34, no. 2, pp. 587–600, Feb. 2022.
- [25] X. Shen, Q. Dai, F.-L. Chung, W. Lu, and K.-S. Choi, "Adversarial deep network embedding for cross-network node classification," in *Proc. AAAI Conf. Artif. Intell.*, 2020, pp. 2991–2999.
- [26] X. Shen, Q. Dai, S. Mao, F.-L. Chung, and K.-S. Choi, "Network together: Node classification via cross-network deep network embedding," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 5, pp. 1935–1948, May 2021.
- [27] G. Song, Y. Zhang, L. Xu, and H. Lu, "Domain adaptive network embedding," *IEEE Trans. Big Data*, vol. 8, no. 5, pp. 1220–1232, Oct. 2020.
- [28] K. Ding, K. Shu, X. Shan, J. Li, and H. Liu, "Cross-domain graph anomaly detection," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 6, pp. 2406–2415, Jun. 2022.
- [29] Q. Dai, X.-M. Wu, J. Xiao, X. Shen, and D. Wang, "Graph transfer learning via adversarial domain adaptation with graph convolution," *IEEE Trans. Knowl. Data Eng.*, vol. 35, no. 5, pp. 4908–4922, May 2023.
- [30] J. Li, W. Liu, Y. Zhou, J. Yu, D. Tao, and C. Xu, "Domain-invariant graph for adaptive semi-supervised domain adaptation," *ACM Trans. Multimedia Comput., Commun., Appl.*, vol. 18, no. 3, pp. 1–18, 2022.
- [31] J. Xiao, Q. Dai, X. Xie, Q. Dou, K.-W. Kwok, and J. Lam, "Domain adaptive graph infomax via conditional adversarial networks," *IEEE Trans. Netw. Sci. Eng.*, vol. 10, no. 1, pp. 35–52, Jan./Feb. 2022.
- [32] Q. Wang, G. Pang, M. Salehi, W. Buntine, and C. Leckie, "Cross-domain graph anomaly detection via anomaly-aware contrastive alignment," 2022, *arXiv:2212.01096*.
- [33] R. Cai, F. Wu, Z. Li, P. Wei, L. Yi, and K. Zhang, "Graph domain adaptation: A generative view," 2021, *arXiv:2106.07482*.
- [34] M. Wu and M. Rostami, "Unsupervised domain adaptation for graph-structured data using class-conditional distribution alignment," 2023, *arXiv:2301.12361*.
- [35] Y. You, T. Chen, Z. Wang, and Y. Shen, "Graph domain adaptation via theory-grounded spectral regularization," in *Proc. 11th Int. Conf. Learn. Representations*, 2023.
- [36] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, "How powerful are graph neural networks?," 2018, *arXiv:1810.00826*.
- [37] F. Errica, M. Podda, D. Bacciu, and A. Micheli, "A fair comparison of graph neural networks for graph classification," 2019, *arXiv:1912.09893*.
- [38] X. Ma, T. Zhang, and C. Xu, "GCAN: Graph convolutional adversarial network for unsupervised domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 8266–8276.
- [39] L. Ruff et al., "Deep one-class classification," in *Proc. Int. Conf. Mach. Learn.*, PMLR, 2018, pp. 4393–4402.
- [40] N. Xiao and L. Zhang, "Dynamic weighted learning for unsupervised domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 15242–15251.
- [41] K. Saito, D. Kim, P. Teterwak, S. Sclaroff, T. Darrell, and K. Saenko, "Tune it the right way: Unsupervised validation of domain adaptation via soft neighborhood density," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 9184–9193.
- [42] L. Van der Maaten and G. Hinton, "Visualizing data using t-SNE," *J. Mach. Learn. Res.*, vol. 9, no. 11, pp. 2579–2605, 2008.
- [43] Y. Xie, H. Zhang, and M. A. Babar, "LogGD: Detecting anomalies from system logs by graph neural networks," 2022, *arXiv:2209.07869*.
- [44] W. Xu, L. Huang, A. Fox, D. Patterson, and M. I. Jordan, "Detecting large-scale system problems by mining console logs," in *Proc. ACM SIGOPS 22nd Symp. Operating Syst. Princ.*, 2009, pp. 117–132.
- [45] A. Oliner and J. Stearley, "What supercomputers say: A study of five system logs," in *Proc. 37th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw.*, 2007, pp. 575–584.
- [46] K. Riesen et al., "IAM graph database repository for graph based pattern recognition and machine learning," in *Proc. Int. Workshop Structural, Syntactic, Statist. Pattern Recognit.*, 2008, pp. 287–297.
- [47] G. O. Campos et al., "On the evaluation of unsupervised outlier detection: Measures, datasets, and an empirical study," *Data Mining Knowl. Discov.*, vol. 30, pp. 891–927, 2016.
- [48] G. Zhang et al., "Dual-discriminative graph neural network for imbalanced graph-level anomaly detection," in *Proc. Adv. Neural Inf. Process. Syst.*, 2022, pp. 24144–24157.
- [49] A. Kendall, Y. Gal, and R. Cipolla, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7482–7491.
- [50] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 8026–8037.
- [51] M. Fey and J. E. Lenssen, "Fast graph representation learning with pytorch geometric," 2019, *arXiv:1903.02428*.
- [52] D. Chen, Y. Lin, W. Li, P. Li, J. Zhou, and X. Sun, "Measuring and relieving the over-smoothing problem for graph neural networks from the topological view," in *Proc. AAAI Conf. Artif. Intell.*, vol. 2020, pp. 3438–3445.
- [53] S. Ibrahim, N. Ponomareva, and R. Mazumder, "Newer is not always better: Rethinking transferability metrics, their peculiarities, stability and performance," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discov. Databases*, Springer, 2022, pp. 693–709.
- [54] W. Zhang, L. Deng, L. Zhang, and D. Wu, "A survey on negative transfer," *IEEE/CAA J. Automatica Sinica*, vol. 10, no. 2, pp. 305–329, Feb. 2023.

Zhong Li received the MSc degree in mathematics from Tongji University, China, and the Diplôme d'Ingénieur degree in data science from ENSAI, France. He is currently working toward the PhD degree in computer science with Leiden University, The Netherlands. His research focuses on trustworthy anomaly detection, particularly in complex data such as event sequences and graph-structured data. He has published papers in leading data mining journals such as *Data Mining and Knowledge Discovery*, *ACM Transactions on Knowledge Discovery from Data*, and *SIGKDD Explorations*. He also serves as reviewer for leading conferences/journals such as KDD, *Data Mining and Knowledge Discovery*, and *IEEE Transactions on Knowledge and Data Engineering*.

Sheng Liang received the MSc degree in computer science from the University of Sheffield in U.K. He is currently working toward the PhD degree in computer science with the Ludwig Maximilian University of Munich, Germany. His current research focuses on multilingual modelling and retrieval augmented language model.

Jiayang Shi received the MSc degree in electrical and computer engineering from the Karlsruhe Institute of Technology, Germany. He is currently working toward the PhD degree in computer science with Leiden University, The Netherlands. He focuses on integrating machine learning with denoising and artifact reduction techniques for computed tomography (CT) imaging data. This encompasses tomographic reconstruction and addressing inverse problems in general.

Matthijs van Leeuwen is associate professor and group leader with Explanatory Data Analysis Group, LIACS, Leiden University, The Netherlands. His primary research interest is exploratory data mining: how can we enable domain experts to explore and analyse their data, to discover structure and ultimately novel knowledge? He was awarded several grants and best paper awards, co-organised international conferences and workshops, is action editor for *Data Mining and Knowledge Discovery* and on the guest editorial board of the ECML PKDD Journal Track. He was guest editor of a *ACM Transactions on Knowledge Discovery from Data* special issue on Interactive Data Exploration and Analytics.