

Can Modifying Data Address Graph Domain Adaptation?

Renhong Huang*
Zhejiang University
Hangzhou, China
Fudan University
Shanghai, China
renh2@zju.edu.cn

Jiarong Xu†
Fudan University
Shanghai, China
jiarongxu@fudan.edu.cn

Xin Jiang
Lehigh University
Bethlehem, United States
xjiang@lehigh.edu

Ruichuan An
Xi'an Jiaotong University
Xi'an, China
arctanx@stu.xjtu.edu.cn

Yang Yang
Zhejiang University
Hangzhou, China
yangya@zju.edu.cn

ABSTRACT

Graph neural networks (GNNs) have demonstrated remarkable success in numerous graph analytical tasks. Yet, their effectiveness is often compromised in real-world scenarios due to distribution shifts, limiting their capacity for knowledge transfer across changing environments or domains. Recently, Unsupervised Graph Domain Adaptation (UGDA) has been introduced to resolve this issue. UGDA aims to facilitate knowledge transfer from a labeled source graph to an unlabeled target graph. Current UGDA efforts primarily focus on model-centric methods, such as employing domain invariant learning strategies and designing model architectures. However, our critical examination reveals the limitations inherent to these model-centric methods, while a data-centric method allowed to modify the source graph provably demonstrates considerable potential. This insight motivates us to explore UGDA from a data-centric perspective. By revisiting the theoretical generalization bound for UGDA, we identify two data-centric principles for UGDA: alignment principle and rescaling principle. Guided by these principles, we propose GraphAlign, a novel UGDA method that generates a small yet transferable graph. By exclusively training a GNN on this new graph with classic Empirical Risk Minimization (ERM), GraphAlign attains exceptional performance on the target graph. Extensive experiments under various transfer scenarios demonstrate the GraphAlign outperforms the best baselines by an average of 2.16%, training on the generated graph as small as 0.25~1% of the original training graph.

CCS CONCEPTS

• Networks → Network algorithms.

*This work was done when the author was a visiting student at Fudan University.

†Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '24, August 25–29, 2024, Barcelona, Spain

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0490-1/24/08...\$15.00

<https://doi.org/10.1145/3637528.3672058>

KEYWORDS


Graph Neural Network; Domain Adaptation; Data Centric

ACM Reference Format:

Renhong Huang, Jiarong Xu, Xin Jiang, Ruichuan An, and Yang Yang. 2024. Can Modifying Data Address Graph Domain Adaptation?. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24)*, August 25–29, 2024, Barcelona, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3637528.3672058>

1 INTRODUCTION

(a) Existing UGDA methods: model-centric



(b) GraphAlign (ours): data-centric




Figure 1: Comparison between existing UGDA methods (which are all model-centric) and our data-centric method GraphAlign. Guided by the rescaling and alignment principles, GraphAlign generates a small yet transferable graph, on which a simple GNN is trained with classic ERM. GraphAlign deviates from conventional approaches that employ sophisticated model design, and achieves outstanding practical performance.

Graph is a ubiquitous data structure that models complex dependencies among entities. Typical examples include social networks [11, 14], biological networks [33, 66] and web networks [27]. Graph Neural Networks (GNNs) have demonstrated considerable potential in a variety of tasks related to graph data [15, 16, 21, 28, 34, 49, 59]. However, they face notable challenges in real-world scenarios when distribution shift exists, e.g., GNNs are trained in one

environment and then deployed in a different environment. This frequently occurs, for example, in social networks where interaction patterns among nodes change over time [53], and in molecular networks with diverse species [8].

To address distribution shift in GNNs, Unsupervised Graph Domain Adaptation (UGDA) has emerged as a crucial solution. The goal of UGDA is to take advantage of the labeled source graph to facilitate the transfer of knowledge to an unlabeled target graph. Most existing efforts in UGDA have been made from the model-centric perspective, *e.g.*, employing domain invariant learning strategies and designing model architectures, as depicted in Figure 1 (a). Specifically, they aim to learn consistent representations across different domains by minimizing domain discrepancy metrics [47, 62, 69] or by incorporating adversarial training with a domain discriminator [10, 56, 57].

In this work, we investigate the inherent limitations of existing model-centric UGDA methods. We identify scenarios where, regardless of how the model parameters are modified, these methods consistently fail in node classification tasks. On the contrary, we find that a data-centric approach that is allowed to modify source graph can theoretically achieve an arbitrarily low classification error with classic empirical risk minimization (ERM) setup. This highlights the potential of data-centric methods in addressing UGDA.

Inspired by these findings, we aim to tackle UGDA from a data-centric perspective. By revisiting the theoretical generalization bound for UGDA [45], we propose two data-centric principles: (1) *Alignment principle* suggests reducing the discrepancy between the modified source graph and the target graph. (2) *Rescaling principle* states that a smaller graph modified from the source graph can achieve a generalization error comparable to that of a larger graph.

With the guidance of these principles, we present a data-centric UGDA method, aiming to generate a new graph that is significantly smaller than the original source graph, yet retains enough information from the source graph and effectively aligns with the target graph. The purpose of such a data-centric UGDA method is to achieve outstanding performance on the target graph with a GNN trained on the generated graph with standard ERM but without sophisticated model design, as shown in Figure 1 (b). In essence, we encounter three main difficulties: (1) how to measure and compute the distribution discrepancy of non-Euclidean structure of graphs when aligning the generated graph with the target graph; (2) how to ensure that the generated graph (with a much smaller size) retains sufficient information from the source graph; and (3) how to efficiently optimize the generated graph, which involves multiple interdependent variables related to graph structure, node features, and node labels.

To address the above difficulties, we propose a novel UGDA method, named GraphAlign. We first use a surrogate GNN model to map graphs into Euclidean representation spaces, and utilize a computationally more efficient Maximum Mean Discrepancy (MMD) distance as the discrepancy metric. The generated graph is expected to have a much smaller size (inspired by rescaling principle) and at the same time, to be so informative that a GNN model trained on the generated graph behaves similarly to that trained on the source graph. To achieve this, we match the gradients of the GNN parameters w.r.t. the generated graph and the source graph. Then, during the optimization process, we model the graph structure

as a function of node features, so the main decision variables are only the node features. We also introduce a novel initialization approach for the generated graph, inspired by the theoretical connection between GNN transferability and the spectral distance to the target graph. This initialization is shown to enhance practical performance and accelerate the optimization process.

Our contributions are summarized as follows:

- **New perspective:** For the first time, UGDA is addressed from a data-centric perspective.
- **New principles:** Inspired by a theoretical generalization bound for UGDA, we propose two data-centric principles that serve as the guidelines for modifying graph data: the alignment principle and the rescaling principle.
- **New method:** We propose GraphAlign, a novel UGDA method. GraphAlign generates a small yet transferable new graph, on which a simple GNN model is trained using classic ERM. In particular, GraphAlign does not need sophisticated GNN design and achieves outstanding performance on target graphs.
- **Extensive experiments:** Experiments on four scenarios and twelve transfer setups demonstrate the effectiveness and efficiency of our method in tackling UGDA. In particular, our method beats the best baselines by an average of +2.16% and trains only on a smaller generated graph, with size up to 1% of the original training graph.

The rest of the paper is organized as follows. We first present the paradigm of UGDA and basic definitions in §2. Then, we present the limitations of existing UGDA methods and propose data-centric principles for UGDA in §3. Guided by these principles, we show in §4 that the limitation of existing UGDA methods can be mitigated by using our proposed model GraphAlign, which employs a data-centric approach that generates a small yet transferable graph for GNN training. Finally, we evaluate the effectiveness and efficiency of our method in §5.

2 PRELIMINARIES

In this section, we present the basic paradigm of UGDA, and introduce the contextual stochastic block model (CSBM), which will be used in §3 to build a motivating example.

Unsupervised Graph Domain Adaptation (UGDA). We focus on UGDA for node classification tasks, where we have a labeled source domain graph and an unlabeled target domain graph. We denote the labeled source graph as $G^S = (A^S, X^S, Y^S)$ with n^S nodes, where A^S , X^S and Y^S represent the adjacency matrix, node features, and node labels of the source graph, respectively. The unlabeled target graph is denoted by $G^T = (A^T, X^T)$ with n^T nodes, where A^T and X^T are the adjacency matrix and node features, respectively.

Given the labeled source graph G^S and unlabeled target graph G^T , UGDA aims to train a GNN h that predicts accurately the node labels Y^T of target graph. The GNN $h = g \circ f$ typically consists of a feature extractor $f : \mathcal{G} \rightarrow \mathcal{Z}$ and a classifier $g : \mathcal{Z} \rightarrow \mathcal{Y}$, where \mathcal{G} , \mathcal{Z} and \mathcal{Y} represent input space, representation space, and label space, respectively. A common approach of UGDA is to learn invariant representations by ensuring that the feature extractor f outputs representations whose distribution remains consistent across both source and target graphs [10, 56, 69].

Contextual Stochastic Block Model. The contextual stochastic block model (CSBM) is an integration of the stochastic block model (SBM) with node features for random graph generation [12]. CSBM generates a graph based on a prescribed edge connection probability matrix, and the distribution for node features is determined by distinct Gaussian mixture models for each class. In the context of UGDA, two distinct CSBMs can be used to model the source and target graphs, facilitating the examination of the domain shift. Before we formally define CSBM, we remark that CSBM is only used to build the motivating example in §3 and is *not* needed in the design of our model. For simplicity, we consider the CSBM with two classes.

Definition 1 (Contextual Stochastic Block Model). CSBM is a generative model that builds a labeled graph $G = (A, X, Y)$ (with node size n) as follows. The node labels are random variables drawn from a Bernoulli distribution (e.g., $Y_i \sim \text{Bern}(0.5)$). The entries of the adjacency matrix follow a Bernoulli distribution $a_{ij} \sim \text{Bern}(C_{pq})$ if node i belongs to class p and j belongs to class q , where the matrix $C \in [0, 1]^{2 \times 2}$ is a prescribed probability matrix that is used to model edge connections. Node features are drawn independently from normal distributions $X_i \sim \mathcal{N}(\mu, I)$ if $Y_i = 0$ and $X_i \sim \mathcal{N}(\nu, I)$ if $Y_i = 1$, where I is the identity matrix. Such a CSBM is denoted by $\text{CSBM}(n, C, \mu, \nu)$.

3 DATA-CENTRIC PRINCIPLES

In §3.1, we present a constructive example to demonstrate the inherent limitation in current model-centric UGDA methods that only focus on sophisticated GNN model design. This example further motivates our exploration of data-centric principles for UGDA in §3.2.

3.1 Motivating Example

To understand the potential issues in existing UGDA models, we investigate the performance of GNN models on a constructive graph pair (G^S, G^T) . Example 1 describes the source and target graphs constructed via CSBMs. Proposition 1 identifies the limitation of existing model-centric-based UGDA approaches, that are typically designed with the goal of learning domain-invariant representations. Proposition 2 shows the potential benefits of adopting a data-centric approach in UGDA.

Example 1. Consider the source and target graphs generated by two CSBMs: $\text{CSBM}(n, C^S, \mu, \nu)$ and $\text{CSBM}(n, C^T, \tilde{\mu}, \tilde{\nu})$, respectively. In both CSBMs, each class consists of $n/2$ nodes, and their edge connection probability matrices are

$$C^S = \begin{bmatrix} a & a \\ a & a - \Delta \end{bmatrix}, \quad C^T = \begin{bmatrix} a - \Delta & a \\ a & a \end{bmatrix},$$

where a and Δ are constants with $0 < \Delta < a < 1$.

The following proposition shows that existing UGDA model-centric methods focusing solely on sophisticated GNN model design would fail even for the simple case in Example 1.

Proposition 1. Assuming the feature extractor f is a single-layer GNN, and it is trained with the domain-invariant constraint $\mathbb{P}(f(G^S)) = \mathbb{P}(f(G^T))$, and then used for inference on the target graph. When such a GNN f is applied to Example 1, the classification error in

the target domain is always larger than a strictly positive constant, regardless of the parameters of the GNN.

The proofs of Proposition 1 can be found in Appendix A.2. Proposition 1 suggests that model-centric UGDA models would fail the node classification task for the graphs in Example 1. In comparison, if we are allowed to “modify” the source graph, it could yield a GNN model with an arbitrarily small classification error, as shown in the following proposition.

Proposition 2. Suppose that the feature extractor f is a single-layer GNN. Also, suppose that a data-centric approach is employed to construct a new graph G' by modifying G^S with the constraint $\mathbb{P}(G') = \mathbb{P}(G^T)$. The GNN f is trained with standard ERM on G' , which minimizes the classification error on G' , and then used for inference on the target graph. There exist examples of graphs generated in Example 1 such that the classification error in the target domain is arbitrarily small.

The proofs of Proposition 2 can be found in Appendix A.2. Proposition 2 highlights that, in certain cases, adapting the data can be more beneficial than adapting the model. An intuitive reason is that data-centric approaches that modify the source data could mitigate the inherent difference of the data distribution between the source and target domains (i.e., $\mathbb{P}(G') = \mathbb{P}(G^T)$). Thus, the GNN model trained on the modified graph G' can be seamlessly applied to the target graph, resulting in a reduction in the error in the target domain. Overall, Propositions 1 and 2 reveal the potential benefits of a data-centric approach for UGDA.

3.2 Data-Centric Principles for UGDA

As demonstrated previously, our objective is to address UGDA through a data-centric approach. Before we delve into the detailed methods, we first discuss two principles that serve as the guidelines for modifying source data. Specifically, we present a generalization bound for UGDA that lays the theoretical foundation for these principles. In our case, the generalization error is defined as the classification error in the target domain [45]:

$$\epsilon^T(g, f) = \mathbb{E}_{\mathbb{P}(G^T)} (\|g \circ f(G^T) - \psi^T(G^T)\|),$$

where $\psi^T: \mathcal{G}^T \rightarrow \mathcal{Y}^T$ is the true labeling function on the target graph. Based on [45], we can derive the generalization bound in the following theorem.

Theorem 1 (Generalization bound for UGDA [45]). Denote by L_{GNN} the Lipschitz constant of the GNN model $g \circ f$. Let the hypothesis set be $\mathcal{H} = \{h = g \circ f: \mathcal{G} \rightarrow \mathcal{Y}\}$, and let the pseudo-dimension be $\text{Pdim}(\mathcal{H}) = d$. The following inequality holds with a probability of at least $1 - \delta$:

$$\epsilon^T(g, f) \leq \hat{\epsilon}^S(g, f) + \underbrace{\eta + 2L_{\text{GNN}}W_1(\mathbb{P}(G^S), \mathbb{P}(G^T))}_{\text{alignment term}} + \underbrace{\sqrt{\frac{4d}{n^S} \log\left(\frac{en^S}{d}\right) + \frac{1}{n^S} \log\left(\frac{1}{\delta}\right)}}_{\text{rescaling term}}, \quad (1)$$

where $\hat{\epsilon}^S(g, f) = (1/n^S) \|g \circ f(G^S) - \psi^S(G^S)\|$ is the empirical classification error in source domain with ψ^S the true labeling function on the source domain, $\eta = \min_{h \in \mathcal{H}} \{\epsilon^S(g^*, f^*) + \epsilon^T(g^*, f^*)\}$




Figure 2: The figure illustrates how the rescaling term varies with the scale of the source graph. We specify $\delta = 0.01$ to ensure that the (1) holds with a probability of at least 99%. The pseudo-dimension d is set to 1000, which is a reasonable assumption based on [13] (note that the trend of the rescaling term’s variation is consistent, regardless of the value of d). The horizontal axis is presented on a logarithmic scale.

denotes the optimal combined error that can be achieved on both source and target graphs by the optimal hypothesis g^* and f^* , $\mathbb{P}(G^S)$ and $\mathbb{P}(G^T)$ are the graph distribution of source and target domain respectively, the probability distribution $\mathbb{P}(G)$ of a graph G is defined as the distribution of all the ego-graphs of G , and $W_1(\cdot, \cdot)$ is the Wasserstein distance.

Proof of Theorem 1 can be found in Appendix A.2. The last two terms on the right-hand side of (1), labeled as the alignment term and rescaling term, inspire the following two principles that will serve as critical guidelines when modifying source data.

Alignment principle: *Modifying the source graph to minimize the distribution discrepancy between the modified source and target graphs can reduce the generalization error.*

This principle is inspired by the alignment term in the bound (1). The smaller the divergence of distribution between the source and target graphs (i.e., $W_1(\mathbb{P}(G^S), \mathbb{P}(G^T))$), the smaller the generalization bound.

Rescaling principle: *Modifying the source graph to reduce its scale could achieve a generalization error comparable to that of a larger-scale source graph.*

This principle encourages us to decrease graph size to improve efficiency while achieving comparable accuracy. It is drawn from the behavior of the rescaling term in the bound (1). Figure 2 shows that this term increases at first, and then decreases as the node size further grows. Notably, a smaller-scale source graph is capable of achieving an accuracy comparable to that of a larger-scale source graph (indicated by the red points in the figure). Yet, as expected, an overly small node size is not advisable, as it loses too much information.

We further discuss the first two terms in the bound (1). The first term is related to the performance of the source domain, and it is a common practice to enhance source domain performance in order to obtain a good performance on the target domain [10, 42, 45, 56]. Regarding the second term, it is often overlooked by other unsupervised domain adaptation methods [10, 42, 45, 56, 62, 69]. Given that

previous research [29] have demonstrated that generating transferable examples by confusing domain discriminator can effectively bridge the domain divergence and reduce this term, similar effects can be achieved by modifying the source graph to better align with the target graph by our approach.

4 PROPOSED METHOD: GRAPHALIGN

In this section, we propose a novel UGDA method GraphAlign. GraphAlign strictly adheres to the alignment and rescaling principles and generates a new graph to replace the source graph for training. These two principles guide us in the generation of a new graph that (1) is much smaller than the original source graph, (2) aligns well with the target graph, and (3) retains enough information from the source graph.

Definition 2 (Data-Centric UGDA). Given the labeled source graph $G^S = (A^S, X^S, Y^S)$ with $A^S \in \mathbb{R}^{n^S \times n^S}$, $X^S \in \mathbb{R}^{n^S \times d}$, $Y^S \in \{0, \dots, c-1\}^{n^S}$, and the unlabeled target graph $G^T = (A^T, X^T)$, Data-Centric UGDA generates a new graph $G' = (A', X', Y')$ with $A' \in \mathbb{R}^{n' \times n'}$, $X' \in \mathbb{R}^{n' \times d}$, $Y' \in \{0, \dots, c-1\}^{n'}$ and $n' \ll n^S$. The graph G' is designed to (1) align with the target graph G^T , and (2) incorporate sufficient information from the source graph G^S , such that the GNN model trained with standard ERM on G' rather than G^S yields enhanced performance on G^T .

Next, we will introduce the problem for optimizing G' in §4.1. Following this, we describe our approach to relaxing the optimization problem and modeling the generated graph G' in §4.2. We finally provide the complexity analysis of our method in §4.3.

4.1 Optimization Problem

We here formulate the optimization problem that guides the alignment of graph G' with the target graph while incorporating sufficient information from the source graph.

Enhancing generalization based on alignment principle. The alignment principle tells us to achieve a lower generalization bound, and it is better to align the distribution of the generated graph G' more closely with that of the target graph G^T . Referring to the alignment term in the generalization bound, this can be achieved by optimizing G' so as to minimize the Wasserstein distance:

$$W_1(\mathbb{P}(G'), \mathbb{P}(G^T)) = \inf_{\gamma \in \Gamma(\mathbb{P}(G'), \mathbb{P}(G^T))} \mathbb{E}_{(u,v) \sim \gamma} c(u,v), \quad (2)$$

where u and v are ego-graphs sampled from $\mathbb{P}(G')$ and $\mathbb{P}(G^T)$ respectively, $c(u,v)$ is the distance function between the ego-graphs u and v , Γ is the set of all joint distribution of $\gamma \in \Gamma(\mathbb{P}(G'), \mathbb{P}(G^T))$, and the marginals for γ are $\mathbb{P}(G')$ and $\mathbb{P}(G^T)$ on the first and second factors respectively.

Although the Wasserstein distance is a natural objective in the optimization problem, its minimization remains a great challenge for the following reasons. First, calculating the Wasserstein distance for a given pair of (u,v) involves solving a large-scale linear program, which is itself computationally expensive, let alone the minimization of the Wasserstein distance. What’s worse, the computation and minimization of the Wasserstein distance in the non-Euclidean space, such as graph data, are even more difficult. We present our solutions as follows.

First, the computational complexity of the Wasserstein distance grows cubically in the problem dimension [39], which is unacceptably

expensive in our case. To resolve this, a common alternative to the Wasserstein distance is the MMD distance, which is computationally cheaper and more efficient [3, 4].

However, replacing Wasserstein distance with the MMD distance does not resolve the second issue caused by non-Euclidean structure of graphs. To handle this issue, existing efforts often use graph kernels that map graphs into (Euclidean) representation spaces [9, 37, 50]. Yet, such mapping process is typically non-differentiable, which complicates the optimization process, and computing graph kernels is still unacceptably costly in our case. In this work, we employ a surrogate GNN model to represent the mapping process. A good example of the GNN model is GIN [60], owing to its discriminative power akin to the Weisfeiler-Lehman test [54].

Consequently, we can update generated graph G' by minimizing the following objective derived from the alignment principle as

$$\mathcal{L}_{\text{alignment}} = \text{MMD} \left(\hat{\mathbb{P}}(\text{GNN}(A', X')), \hat{\mathbb{P}}(\text{GNN}(A^{\mathcal{T}}, X^{\mathcal{T}})) \right), \quad (3)$$

where $\hat{\mathbb{P}}$ is the empirical distribution computed via random sampling. Details regarding the implementation details for MMD distance are provided in Appendix A.1.

Incorporating sufficient information from source graph. The generated graph G' should retain enough information from the source graph, which can be guaranteed by the following two strategies.

First, if G' is adequately informative as the source graph $G^{\mathcal{S}}$, a GNN model trained on G' would behave similarly to that trained on $G^{\mathcal{S}}$. Inspired by [67], we aim to match the gradients of the GNN parameters w.r.t. G' and $G^{\mathcal{S}}$. This is crucial for preserving the essential information from the source graph in G' . To achieve this, we focus on minimizing the following objective function:

$$\mathcal{L}_{\text{mimic}} = \text{Cos} \left(\nabla \mathcal{L}_{\text{CE}}(\text{GNN}(A', X'), Y'), \nabla \mathcal{L}_{\text{CE}}(\text{GNN}(A^{\mathcal{S}}, X^{\mathcal{S}}), Y^{\mathcal{S}}) \right),$$

where \mathcal{L}_{CE} denotes the cross entropy loss, GNN is the surrogate GNN model and Cos is the cosine similarity function.

On the other hand, the generated graph G' needs to reflect generally observed properties in real-world networks. A typical property of real-world networks is feature smoothness, where connected nodes often share similar features [1, 36]. Moreover, real-world graphs are usually sparse [68]. Therefore, to ensure that G' accurately represents these real-world characteristics, we focus on minimizing the following objective function:

$$\mathcal{L}_{\text{prop}} = \text{tr}(X'^{\mathcal{T}} L X') + \|A'\|_F^2, \quad (4)$$

where $L = I - D^{-\frac{1}{2}} A' D^{-\frac{1}{2}}$ is the normalized Laplacian matrix and D is the diagonal degree matrix for A' . The first term in $\mathcal{L}_{\text{prop}}$ captures feature smoothness while the second one characterizes sparsity.

Optimization problem. We here outline the construction of the generated new graph G' . Specifically, our goal is to build $G' = (A', X', Y') \in \mathbb{R}^{n' \times n'} \times \mathbb{R}^{n' \times d} \times \mathbb{R}^{n'}$ so that it aligns with the target graph and retains enough information from the source graph. Integrating the aforementioned objectives, the construction of G' can be formulated as the following optimization problem

$$\min_{A', X', Y'} \mathcal{L}_{\text{mimic}} + \alpha_1 \mathcal{L}_{\text{alignment}} + \alpha_2 \mathcal{L}_{\text{prop}}, \quad (5)$$

where $\alpha_1, \alpha_2 > 0$ are hyper-parameters.

4.2 Modeling the Generated Graph

Note that the decision variables in (5) are A', X' , and Y' . Optimizing the three variables is extremely difficult due to their interdependence. To this end, the node size n' is pre-chosen and proportional to $n^{\mathcal{S}}$, i.e., $n' = rn^{\mathcal{S}}$ for some prescribed $0 < r \ll 1$. We also require that the node labels Y' have the same distribution as $Y^{\mathcal{S}}$ when randomly choose n' nodes from the source graph.

Even if n' is pre-fixed and much smaller than n , the number of parameters in A' is still quadratic in n' and prohibitively large to optimize. To further reduce the number of parameters in A' , we propose to model the graph structure A' as a function of X' . This is motivated by the observation in real-world networks that the graph structure and the node features are implicitly correlated [40]. So, in our implementation, A' is modeled as

$$A' = \rho_{\phi}(X'), \text{ with } A'_{ij} \sim \text{Bernoulli} \left(\text{Sigmoid} \left(\text{MLP}_{\phi} \left(X'_i, X'_j \right) \right) \right),$$

where ρ_{ϕ} , parameterized by ϕ , is the function that transforms node features to graph structure, and MLP_{ϕ} is a multi-layer neural network. As is common in the literature, the non-differentiability of Bernoulli sampling can be handled by the Gumbel-Max reparametrization technique [23].

In summary, the decision variables in (5) are effectively reduced to X' and ϕ , and then the original problem (5) can be rewritten as

$$\min_{X', \phi} \mathcal{L}_{\text{mimic}} + \alpha_1 \mathcal{L}_{\text{alignment}} + \alpha_2 \mathcal{L}_{\text{prop}}. \quad (6)$$

Initialization of the generated graph. As one may expect, the initial value of X' is crucial for solving (6). A good initialization not only helps improve practical performance, but can also accelerate the optimization process.

To further improve the performance of proposed method, we propose an initialization strategy for X' . Intuitively, we hope to select from the source graph those nodes and features that already present transferability. This intuition can be further supported by the following theorem, which builds theoretical connections between the property of the newly generated G' and the transferability of GNNs.

Theorem 2 (GNN transferability). *Let G' and $G^{\mathcal{T}}$ be the newly generated graph and the target graph. Given a GNN graph encoder f , the transferability of the GNN f satisfies*

$$\|f(G') - f(G^{\mathcal{T}})\|_2 \leq \xi_1 \Delta_{\text{spectral}}(G', G^{\mathcal{T}}) + \xi_2, \quad (7)$$

where ξ_1 and ξ_2 are two positive constants, and $\Delta_{\text{spectral}}(G', G^{\mathcal{T}}) = \frac{1}{n'n^{\mathcal{T}}} \sum_{i=1}^{n'} \sum_{j=1}^{n^{\mathcal{T}}} \|L_{G'_i} - L_{G^{\mathcal{T}}_j}\|_2$ measures the spectral distance between G' and $G^{\mathcal{T}}$. Here G'_i is the ego-graph of node i in G' , and $L_{G'_i}$ is its normalized graph Laplacian. The graph Laplacian $L_{G^{\mathcal{T}}}$ is defined in a similar manner.

Theorem 2 suggests that a smaller spectral distance between G' and $G^{\mathcal{T}}$ indicates better transferability. Based on this interpretation, we propose to select n' nodes from G' whose features are used as the initial values of X' . Such selection guarantees that the graph G' constructed by these nodes and features has a small spectral distance $\Delta_{\text{spectral}}(G', G^{\mathcal{T}})$.

4.3 Complexity Analysis

The traditional domain adaptation methods typically involve GNN training and domain-invariant learning, e.g., using MMD loss for optimization. The time complexity of these methods is typically $O(d^2n + dn^2)$. In comparison, the time complexity of GraphAlign primarily depends on the optimization process of G' . Assume the dimension of representations is denoted as d and the number of nodes in G' is n' . Constructing G' involves calculating n'^2 edges, which is equivalently $O(r^2dn^2)$. Regarding the GNN forward pass, it requires a time complexity of $O(d^2n)$. For the computation of $\mathcal{L}_{\text{alignment}}$, the time complexity is $O(rdn^2)$ because of the computation of kernel function. For the $\mathcal{L}_{\text{mimic}}$ loss, we need to inference GNN on both G^S and G' , resulting in a time complexity of $O(d^2n + rd^2n)$. Besides, the time complexity of computing $\mathcal{L}_{\text{prop}}$ primarily focuses on the calculations of $\text{trace}(X'^T LX')$ an is $O(r^2n^2)$, with the trace operation requires $O(rn)$. Overall, the time complexity of GraphAlign is determined by $O((r^2d + rd + r^2)n^2 + (1 + r)d^2n)$. In our experiments, the hyper-parameter r is chosen as $r = 0.01$, so the complexity of GraphAlign is much cheaper compared with classic model-centric UGDA methods.

5 EXPERIMENTS

In this section, we evaluate the performance of GraphAlign under various transfer scenarios. We first generate a new graph G' using the proposed method, and then train a GNN on G' with classic ERM. The GNN is then tested on the target graph. The experiments span a range of transfer scenarios, and we also include ablation studies, hyper-parameter analysis, and runtime comparison to demonstrate the effectiveness of GraphAlign.

5.1 Experimental Setup

Datasets. We conduct experiments on node classification in the transfer setting across six scenarios. In each scenario, we train the GNN on one graph and evaluate it on the others.

- **ACMv9 (A), DBLPv7 (D), Citationv1 (C)** [10]: These datasets are citation networks from different sources, where each node represents a research article and an edge indicates citation relationship between two articles. The data are collected from ACM (prior to 2008), DBLP (between 2004 and 2008), and Microsoft Academic Graph (after 2010), respectively. We include six transfer settings: $C \rightarrow D$, $A \rightarrow D$, $D \rightarrow C$, $A \rightarrow C$, $D \rightarrow A$ and $C \rightarrow A$.
- **ACM_{small} (Ā), DBLP_{small} (Ď)** [56]: These two are also citation networks, with articles collected between the years 2000 and 2010, and after year 2010. We include two transfer settings: $\hat{A} \rightarrow \hat{D}$, $\hat{D} \rightarrow \hat{A}$.
- **Cora-degree, Cora-word** [18]: They are two transfer settings for citation networks provided by [5], derived from the full Cora dataset [5]. The data pre-process involves partitioning the original Cora dataset into two graphs based on node degrees and selected word count in publications, respectively. Each setting evaluates the transferability from one graph to the other.
- **Arxiv-degree, Arxiv-time** [18]: They are two transfer settings for citation networks provided by [5], adapted from the Arxiv dataset that comprises computer science arXiv papers [5]. The partitioning of Arxiv into two graphs is based on node degrees and time, respectively.

- **USA, Brazil, Europe** [43]: They are collected from transportation statistics and primarily comprise airline activity data, where each node represents to an airport. We include six transfer settings: $USA \rightarrow Brazil$, $USA \rightarrow Europe$, $Brazil \rightarrow USA$, $Brazil \rightarrow Europe$, $Europe \rightarrow USA$, $Europe \rightarrow Brazil$.
- **Blog1, Blog2** [46]: They are collected from the BlogCatalog dataset, where node represents a blogger, and edge indicates friendship between bloggers. The node attributes comprise keywords extracted from self-descriptions of blogger, and the task is to predict their corresponding group affiliations. We include two transfer settings: $Blog1 \rightarrow Blog2$, $Blog2 \rightarrow Blog1$.

Baselines. We compare our method with the following baselines, which can be categorized into three classes: (1) **Vanilla ERM**, including GCN [26], GraphSAGE [19] and GIN [60]. They are trained on the source graph with ERM and then directly evaluated on the target graph. (2) **Non-graph domain adaptation methods**, including MMD [31], CMD [64], DANN [17], CDAN [32] are considered. To adapt them to the UGDA setting, we replace the encoder with GCN [26]. (3) **UGDA methods**, including UDAGCN [56], AdaGCN [10], MIXUP [20], EERM [57], MFRReg [62], SSReg [62], GRADE [55], JHGDA [48] and STRURW [30].

Implementation details. We evaluate our proposed method by training a GCN [26] on the generated graph G' with ERM and test the GCN on G^T . When computing $\mathcal{L}_{\text{mimic}}$, we train a surrogate GCN on G' for the supervised node classification task, with a cross-entropy loss. When computing $\mathcal{L}_{\text{alignment}}$, we train GIN on G' under the infomax principle following [51]. All the GNN models adopt a two-layer structure with 256 hidden units, while the other hyper-parameters are set to default. After training all the surrogate models, we freeze all the parameters when optimizing G' . We set the reduction rate $r = 0.25\%$ for Arxiv (due to its large scale) and $r = 1\%$ for the remaining datasets. The values of α_1 , α_2 are set to 1 and 30. For the optimizer, we use Adam [25] with a learning rate of 1×10^{-3} and weight decay of 5×10^{-3} . We use mini-batch training with batch size 32. The total iterations of training is 300.

When evaluating the baselines, for Vanilla ERM and non-graph domain adaptation baselines, we employ two-layer GCN with 256 hidden units and the remaining hyper-parameters are set to default values. We follow the setting of [56] to perform a grid search on the trade-off between classification loss and the loss function designed to address domain adaptation, exploring values within [0.01, 0.1, 1.0, 10.0] and reporting the best performance. Adam is employed for optimization with a learning rate of 1×10^{-3} and weight decay of 5×10^{-3} . We use mini-batches of size 32 over 300 training iterations. For UGDA methods, we adopt their default hyper-parameters.

The reported numbers in all experiments are the mean and standard deviation over 10 trials. More details can be found in Appendix A.1. Our codes are available at <https://github.com/zjunet/GraphAlign>.

5.2 Experimental Results

Main results. Table 1 and Table 2 presents the experiment results across various settings. For all the datasets, GraphAlign surpasses all baselines and shows an average improvement of +2.16% over the best baseline. Notably, this superior performance is attained with our generated small graphs. In comparison, the suboptimal

Table 1: Micro F1 scores across various transfer settings on citation networks. The bold numbers denote the best result. “OOM” denotes the instances where the method ran out of memory.

Methods	ACMv9 (A), DBLPv7 (D), Citationv1 (C)						ACM _{small} (Ā), DBLP _{small} (Ď)		Cora		Arxiv		Avg.rank
	C→D	A→D	D→C	A→C	D→A	C→A	Ā → Ď	Ď → Ā	Cora-word	Cora-degree	Arxiv-degree	Arxiv-time	
ERM (GIN)	43.32(2.40)	39.32(2.10)	38.86(1.57)	37.27(2.92)	37.14(1.36)	35.40(2.71)	50.12(3.06)	63.24(1.53)	30.64(1.63)	19.91(1.76)	22.69(1.51)	32.25(2.42)	14.9
ERM (SAGE)	64.22(0.89)	61.73(0.88)	60.92(1.25)	61.90(1.19)	55.66(0.92)	57.66(1.04)	77.51(3.17)	36.59(4.05)	61.63(0.18)	53.56(0.30)	47.84(0.52)	44.84(0.56)	8.3
ERM (GCN)	67.80(3.49)	61.05(0.37)	62.36(5.01)	61.78(4.33)	53.78(4.13)	62.93(5.32)	63.28(2.05)	68.48(0.89)	63.26(0.35)	54.42(0.51)	43.86(1.12)	12.86(5.16)	7.8
DANN	66.02(1.89)	61.44(2.89)	54.68(3.66)	59.61(4.88)	49.01(3.59)	55.02(2.45)	63.38(2.21)	68.53(0.90)	63.24(0.41)	54.44(0.54)	43.86(1.12)	12.87(5.17)	9.8
CDAN	53.69(3.90)	61.53(1.64)	61.13(2.78)	60.48(2.61)	53.69(3.90)	58.22(0.95)	63.53(2.09)	70.73(0.86)	63.42(0.27)	54.48(0.42)	43.91(1.00)	12.86(5.18)	8.2
MMD	63.06(0.61)	59.72(0.35)	59.46(0.49)	62.98(0.50)	53.57(0.33)	59.34(0.47)	61.20(0.69)	69.22(0.87)	63.27(0.38)	54.21(0.60)	OOM	OOM	10.25
CMD	47.89(12.49)	46.67(11.03)	48.49(5.13)	53.02(10.79)	44.8(2.83)	49.31(8.70)	50.56(6.17)	64.86(3.73)	54.95(0.64)	49.61(0.48)	39.13(1.40)	13.58(0.85)	13.3
UDAGCN	70.70(2.64)	64.64(3.12)	56.34(8.55)	62.40(5.81)	49.57(4.95)	55.92(5.85)	69.97(4.10)	70.43(1.36)	63.40(0.21)	53.98(0.41)	37.82(6.70)	47.44(0.77)	7.3
AdaGCN	69.72(2.05)	67.67(0.92)	66.38(2.86)	69.34(1.44)	56.78(2.53)	63.34(1.24)	69.28(2.34)	69.33(1.57)	62.91(0.49)	53.24(0.47)	38.93(1.62)	12.38(5.50)	5.8
MIXUP	67.60(1.88)	63.08(2.68)	60.75(5.95)	66.04(3.36)	54.01(5.00)	62.06(2.35)	51.10(1.40)	68.43(1.46)	65.44(5.95)	62.79(4.01)	52.13(0.40)	23.63(0.24)	6.3
EERM	43.60(1.59)	46.27(5.36)	43.97(2.74)	46.37(3.25)	38.39(1.89)	43.87(2.10)	67.05(1.33)	47.11(4.82)	13.10(0.77)	7.42(0.86)	OOM	OOM	14.9
GRADE	64.68(1.30)	54.46(1.13)	59.61(0.09)	66.05(0.35)	57.30(0.18)	61.19(0.46)	64.92(0.09)	55.47(0.21)	48.37(2.45)	42.84(1.87)	41.36(1.64)	11.32(0.78)	9.8
JHGDA	65.09(4.98)	58.20(1.33)	51.31(3.56)	64.51(2.65)	46.46(4.72)	59.25(3.44)	71.51(2.40)	62.26(3.28)	OOM	OOM	OOM	OOM	12.3
MFRReg	66.99(7.77)	61.39(8.45)	65.71(6.33)	70.72(8.71)	56.65(7.21)	59.81(8.53)	65.80(10.42)	71.13(0.95)	59.23(1.81)	53.04(1.28)	OOM	OOM	8.3
SSReg	63.36(17.73)	61.78(9.98)	66.53(4.47)	61.91(11.87)	56.05(9.41)	60.41(6.75)	71.10(8.31)	70.00(1.41)	59.61(2.26)	52.19(1.29)	OOM	OOM	9.1
STRURW	64.10(0.35)	59.45(0.60)	58.91(1.02)	63.42(0.38)	55.83(1.11)	62.41(1.27)	77.47(1.01)	72.11(2.21)	62.41(0.72)	67.76(0.39)	57.45(0.15)	49.98(0.12)	5.9
GraphAlign	72.56(0.61)	69.65(0.26)	68.08(0.32)	75.61(0.24)	62.06(0.68)	67.36(0.40)	79.51(3.75)	72.63(2.21)	66.37(1.46)	69.83(1.53)	57.51(1.42)	51.17(1.37)	1.0

Table 2: Micro F1 scores across various transfer settings on airport networks and social networks.

Methods	Airport						Social			Avg.rank
	USA→Brazil	USA→Europe	Brazil→USA	Brazil→Europe	Europe→USA	Europe→Brazil	Blog1→Blog2	Blog2→Blog1		
ERM (GIN)	35.88(4.71)	32.33(1.78)	41.43(9.32)	33.18(5.72)	44.52(5.72)	42.90(7.39)	18.36(0.25)	19.47(1.14)	15.5	
ERM (SAGE)	49.62(1.37)	43.96(0.61)	53.29(0.40)	53.78(1.01)	52.79(0.68)	67.63(1.04)	45.40(0.61)	47.00(1.05)	6.9	
ERM (GCN)	43.36(4.81)	37.99(4.39)	49.95(0.82)	42.21(1.87)	56.82(0.34)	71.76(0.84)	44.57(1.24)	41.25(2.34)	8.9	
DANN	59.85(8.34)	52.48(2.09)	53.38(0.33)	57.74(1.61)	57.48(0.48)	70.99(0.68)	42.30(1.23)	41.30(2.67)	4.8	
CDAN	46.87(3.82)	42.61(1.79)	52.29(0.99)	45.01(1.51)	56.76(0.39)	72.06(1.57)	42.56(1.68)	41.21(3.32)	8.0	
MMD	52.90(0.78)	55.64(0.69)	52.49(0.33)	56.41(0.31)	56.77(0.14)	72.98(0.37)	43.98(1.83)	40.98(2.26)	4.6	
CMD	60.84(0.57)	54.99(0.88)	48.99(0.70)	58.50(0.75)	55.21(0.68)	72.98(1.50)	28.65(5.31)	26.55(6.42)	7.3	
UDAGCN	34.42(3.14)	51.78(1.06)	24.96(6.12)	44.81(1.93)	55.45(0.44)	44.43(2.60)	36.47(7.45)	36.89(5.48)	12.8	
AdaGCN	52.37(3.79)	48.67(0.58)	45.63(4.44)	51.48(2.61)	48.97(4.23)	66.11(3.82)	38.25(1.86)	36.82(3.65)	11.5	
MIXUP	43.18(4.13)	41.63(0.66)	50.34(4.73)	42.74(0.98)	49.32(1.62)	68.76(1.48)	40.31(1.84)	39.63(2.82)	11.6	
EERM	39.12(5.87)	43.42(0.90)	42.98(6.40)	55.72(1.91)	48.92(2.20)	48.92(6.37)	41.89(2.67)	40.21(1.96)	11.5	
GRADE	24.43(2.44)	24.81(4.33)	34.29(1.42)	27.32(1.90)	34.03(2.35)	29.77(3.53)	17.33(1.17)	16.39(1.42)	16.9	
JHGDA	61.41(4.73)	52.09(3.73)	51.42(5.71)	58.04(9.05)	51.88(2.94)	72.77(7.43)	17.86(2.56)	17.93(2.60)	8.4	
MFRReg	56.99(6.16)	53.19(6.45)	50.71(5.07)	50.72(8.71)	56.65(7.21)	69.18(8.53)	40.84(3.12)	46.34(6.72)	7.1	
SSReg	53.36(7.73)	53.78(1.72)	49.34(4.88)	52.43(6.23)	54.28(6.22)	55.71(6.25)	40.93(4.29)	45.37(8.11)	8.3	
STRURW	60.73(0.34)	53.77(0.98)	52.19(2.01)	53.48(0.23)	49.67(2.88)	63.40(1.27)	46.02(0.95)	38.64(1.76)	7.5	
GraphAlign	62.90(0.78)	54.32(0.94)	54.38(0.10)	58.80(0.86)	57.34(2.02)	73.12(0.90)	47.14(1.72)	45.83(5.01)	1.6	

performance of Vanilla ERM and non-graph domain adaptation methods highlights the critical need for a carefully tailored UGDA strategy to address domain discrepancy on graphs. Compared with other UGDA methods, our superior performance emphasizes the advantage of adopting a data-centric approach over a model-centric approach in UGDA.

Ablation studies. To validate the effectiveness of components, ablation studies are conducted on: (1) GraphAlign-init, employing random initialization instead of our proposed initialization; (2) GraphAlign- \mathcal{L}_{prop} , removing the loss \mathcal{L}_{prop} . (3) GraphAlign- $\mathcal{L}_{alignment}$, removing the loss $\mathcal{L}_{alignment}$. (4) GraphAlign- \mathcal{L}_{mimic} ,

removing the loss \mathcal{L}_{mimic} . The results, depicted in Figure 3, clearly demonstrate the contribution of each component to enhancing performance of GraphAlign. Particularly, the superiority of GraphAlign over GraphAlign- $\mathcal{L}_{alignment}$, GraphAlign- \mathcal{L}_{prop} and GraphAlign- \mathcal{L}_{mimic} highlights indispensable roles of the alignment principle, the gradient matching strategy, and the preservation of graph properties when generating the new graph, respectively. The notable performance drop with GraphAlign- \mathcal{L}_{mimic} underscores the critical importance of the loss \mathcal{L}_{mimic} . In addition, the observed decrease in performance with GraphAlign-init validates the effectiveness of our initialization strategy in improving graph transferability.




Figure 3: Ablation studies on A \rightarrow D and C \rightarrow A tasks.




Figure 4: Our results on D \rightarrow A task w.r.t varying r , α_1 and α_2 . The dashed line represents the performance of the best baseline.

Hyper-parameters analysis. Figure 4 shows the effects of the reduction rate r , the coefficients α_1 for $\mathcal{L}_{alignment}$ and α_2 for \mathcal{L}_{prop} , respectively. Regarding the reduction rate r , we observe that our model consistently outperforms the most competitive baseline model. This indicates our model’s ability to scale down the graph while enhancing the transferability to the target graph. We also note that an excessively low reduction rate could hamper performance due to the difficulty in obtaining an informative yet overly small graph. Regarding α_1 and α_2 , we find that both too small and large values could deteriorate performance, as a value too small would weaken the influence of $\mathcal{L}_{alignment}$ and α_2 for \mathcal{L}_{prop} , while a value too large would negatively affect the balance with other loss components in (5). Despite variations, our model consistently surpasses the best baseline, demonstrating that GraphAlign exhibits low sensitivity to hyper-parameter changes.

Transferability under different levels of domain shifts. We further evaluate the transferability of our approach and baselines across different levels of domain shifts. We here focus on structural shifts and feature shifts between graphs. Specifically, we conduct experiments on a series of synthetic graphs constructed based on CSBM. The source graph is constructed as CSBM($n = 128, C^S = [[0.6, 0.3], [0.3, 0.6]]$, $\mu = 1, \nu = -1$) and remains fixed. For structural shift, the target graphs are defined by a sequence of CSBMs with parameters $n = 128, \mu = 1, \nu = -1$, and C^T randomly generated. The level of structural shift is quantified as $\|C^T - C^S\|_F$ and ranges from 0.05 to 0.45. As for feature shift, the target graphs are defined by a sequence of CSBM($n = 128, C^T = [[0.6, 0.3], [0.3, 0.6]]$, $\mu = 1 + \Delta, \nu = -1 + \Delta$), where Δ varies in $[0.1, 0.9]$ and quantifies the level of feature shift.

Table 3: Micro F1 scores under different levels of domain shift (structural shift and feature shift).

Structural shift	0.05	0.15	0.25	0.35	0.45
ERM (GCN)	77.31(7.54)	68.72(5.83)	65.47(6.67)	65.50(6.97)	56.38(5.10)
CDAN	77.31(7.66)	67.27(5.68)	67.56(8.00)	68.68(7.08)	55.50(6.92)
AdaGCN	82.21(6.37)	77.62(7.18)	72.58(7.63)	70.73(7.87)	55.78(0.62)
GraphAlign	83.34(4.26)	78.32(9.45)	73.25(6.63)	72.54(8.86)	58.64(5.10)
Gain (%)	+1.37	+0.90	+0.92	+2.56	+4.01
Feature shift	0.1	0.3	0.5	0.7	0.9
ERM (GCN)	80.61(6.32)	73.94(6.63)	70.07(6.60)	66.02(5.44)	65.38(6.71)
CDAN	73.08(13.87)	75.35(7.05)	73.56(7.08)	63.83(11.03)	66.72(6.97)
AdaGCN	82.68(5.50)	77.73(6.40)	76.67(8.05)	66.33(7.35)	67.52(8.66)
GraphAlign	82.47(6.22)	78.21(6.28)	77.34(3.75)	66.98(7.70)	68.46(5.90)
Gain (%)	-0.25	+0.61	+0.87	+0.98	+1.39

The results of our model and the most competitive methods in the three categories (vanilla ERM, non-graph domain adaptation methods, and UGDA methods) are shown in Table 3. We observe a decrease in the performance of all methods as domain shift increases, while GraphAlign performs the best under various levels of domain shifts in most cases. Besides, we find that GraphAlign is more effective under large domain shift, achieving more substantial gains. When there is a small domain divergence, the effectiveness of our method may not be as significant, and existing domain adaptation methods may achieve comparable performance.

Adopting GraphAlign with different GNNs with ERM and model-centric UGDAs. After obtaining the generated graph G' , it can be utilized to train diverse GNN architectures with ERM or fed into model-centric UGDA methods. The results are shown in Table 4. We find that (1) using the generated graph G' to train different GNNs with ERM still exhibits superior performance compared to existing UGDA methods. This underscores the versatility and superior efficacy of our approach in bridging the gap between graphs from different domains; (2) by combining our method with established model-centric UGDA methods, further enhancements in performance are observed in some cases, notably in cross-domain scenarios like D \rightarrow C and C \rightarrow A.

Runtime comparison. Table 5 presents the runtime comparisons between GraphAlign and the most competitive baselines of three categories: vanilla ERM, Non-graph domain adaptation methods, and UGDA methods. We observe that the runtime of GraphAlign is comparable to vanilla ERM methods, yet it delivers superior performance. This efficiency stems from the rescaled graph, which significantly reduces the time required for GNN training with ERM. We also note that the efficiency of our method facilitates the possibility of further applications, such as neural architecture search of GNNs and hyper-parameter optimization.

6 RELATED WORK

Unsupervised domain adaptation. Unsupervised domain adaptation (UDA) aims to transfer knowledge from a labeled source domain to an unlabeled target domain and has demonstrated success in computer vision and natural language processing [2, 52]. Most existing work on UDA attempts to learn invariant representations across different domains [17, 31, 63].

Recently, domain adaptation has been adapted to graph, and numerous studies have been conducted to address UGDA. These

Table 4: Micro F1 scores when adopting GraphAlign with different GNNs with ERM and model-centric UGDAs. The bold numbers denote the best result. The asterisk (*) indicates that the result surpasses the most competitive baseline in Table 1.

Methods	C→D	A→D	D→C	A→C	D→A	C→A
GraphAlign+ERM (GCN)	72.56(0.61)*	69.65(0.26)*	68.08(0.32)*	75.61(0.24)*	62.06(0.68)*	67.36(0.40)*
GraphAlign+ERM (GraphSAGE)	70.73(0.41)*	68.97(1.28)*	67.70(0.48)*	73.21(0.50)*	61.08(0.52)*	66.79(0.22)*
GraphAlign+ERM (SGC)	69.93(0.02)	66.23(0.21)	64.12(0.01)	70.56(0.28)	58.65(0.01)*	65.42(0.03)*
GraphAlign+ERM (APPNP)	68.22(0.05)	65.59(0.01)	61.89(1.70)	67.89(0.73)	56.84(0.02)*	63.49(0.04)*
GraphAlign+MMD	71.13(0.32)*	67.01(0.35)	68.05(0.42)*	73.69(0.46)*	60.57(0.61)*	66.53(0.32)*
GraphAlign+UDAGCN	70.83(0.17)*	69.35(0.23)*	68.31(0.32)*	75.61(0.24)*	60.23(0.80)*	67.40(0.47)*

Table 5: Runtime (sec) and Memory (Mb) comparison on Arxiv-degree. All the models are trained within 300 iterations and the reduction rate for GraphAlign is 0.25%.

	Runtime			Memory
	generate	G'	GNN training total	total
ERM (GCN)	-	110.49	110.49	2868
CDAN	-	148.57	148.57	2868
AdaGCN	-	577.91	577.91	2956
GraphAlign	257.53	18.57	276.10	2050

studies can be generally divided into two categories: one focuses on minimizing domain discrepancy metrics, while the other utilizes adversarial training techniques. The first class of methods aims to learn domain-invariant representations by minimizing pre-defined domain discrepancy metrics, such as class-conditional MMD [47], central moment discrepancy [64], spectral regularization [62], graph subtree discrepancy [55], tree mover’s distance [9]. In comparison, the adversarial training-based methods typically incorporate a domain classifier that adversarially predicts the domain of the representation. For instance, Dai et al. [10] and Wu et al. [56] utilize GNN models as feature extractors and train them in an adversarial manner to align the cross-domain distributions. Qiao et al. [41] further introduces shift parameters to enhance transferability to target graph during adversarial training. However, these methods consider UGDA from the model-centric perspective. Different from aforementioned works on UGDA, this paper addresses UGDA from a novel data-centric perspective that is allowed to modify the source graph.

To the best of our knowledge, the only UGDA method that modifies the source graph is a recent one [30]. But, their modification relies on the accurate estimation of target labels, and the process of estimating target labels is highly dependent on model-centric UGDA methods. In addition, their modification to the source graph is limited to edge weights. In contrast, our method is purely data-centric and allows the modified graph to be directly utilized for training a GNN by ERM, achieving competitive performance without the need for sophisticated design. Besides, we enable modifications to graph structure, node features and node labels. We further investigate the impact of graph scale on the generalization bound and suggest that a smaller-scale modified graph can be effective.

Another line of research assume that source graph is not accessible during the adaptation process [24, 35]. Specifically, [35] focuses on enhancing the discriminative capability of the source model

through structure consistency and information maximization. [24] explores modifying graph data during testing to improve model generalization and robustness. However, this setting is different from ours, and the absence of source graph poses difficulties in addressing shifts caused by differences in data distributions.

Data-centric AI. This recently introduced concept emphasizes the potential of data modification over model design [7, 61, 65]. Subsequent works leverage the data-centric approach to address various questions in machine learning. Some works [24, 58] take into account the impact of graph data on model generalization. Xu et al. [58] considers the co-evolution of data and model to enhance data quality for pre-training without considering information in a downstream graph. Jin et al. [24] modifies test graph to address distribution shift at test time and then, the pretrained model is provided for inferences on the modified test graph. However, both of them overlook the rich information inherent in either the target graph or the source graph itself and, and thus fail to address UGDA.

Another line of research focuses on reducing graph scale. Several works have proposed methods for graph condensation or sparsification while preserving information of the original graph [6, 22]. However, these methods don’t consider the transferability of the generated graphs, making them unsuitable for direct adoption in UGDA. In contrast, our method focuses on generating a new graph that can transfer effectively to the target domain, rather than merely retaining information. Additionally, the process of reducing the graph scale is often sensitive to initialization, but there is currently no existing work that discusses this aspect.

7 CONCLUSION

This work investigates UGDA through a data-centric lens. Our analysis pinpoints the limitations in model-centric UGDA methods and shows the potential of data-centric methods for UGDA. We therefore introduce two data-centric principles for UGDA: the alignment principle and the rescaling principle, rooted in the generalization bound for UGDA. Guided by these principles, we propose GraphAlign, a novel data-centric UGDA method. GraphAlign first generates a small yet transferable graph in replacement of the original training graph and trains a GNN on the newly generated graph with classic ERM setting. Numerical experiments demonstrate that GraphAlign achieves remarkable transferability to the target graph.

ACKNOWLEDGEMENTS

This work is supported by NSFC (62206056, 92270121), Zhejiang NSF (LR22F020005) and SMP-IDATA Open Youth Fund.

REFERENCES

- [1] Mikhail Belkin and Partha Niyogi. 2001. Laplacian eigenmaps and spectral techniques for embedding and clustering. *NeurIPS* 14 (2001).
- [2] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. 2010. A theory of learning from different domains. *Machine learning* 79 (2010), 151–175.
- [3] Tolga Birdal, Michael Arbel, Umüt Simsekli, and Leonidas J Guibas. 2020. Synchronizing probability measures on rotations via optimal transport. In *CVPR*.
- [4] Mikolaj Bińkowski, Dougal J. Sutherland, Michael Arbel, and Arthur Gretton. 2018. Demystifying MMD GANs. In *ICLR*.
- [5] Aleksandar Bojchevski and Stephan Günnemann. 2018. Deep Gaussian Embedding of Graphs: Unsupervised Inductive Learning via Ranking.
- [6] Chen Cai, Dingkang Wang, and Yusu Wang. 2021. Graph coarsening with neural networks. *ICLR* (2021).
- [7] Yuxuan Cao, Jiarong Xu, Carl Yang, Jiaan Wang, Yunchao Zhang, Chunping Wang, Lei Chen, and Yang Yang. 2023. When to Pre-Train Graph Neural Networks? From Data Generation Perspective!. In *SIGKDD*. 142–153.
- [8] Hyunghoon Cho, Bonnie Berger, and Jian Peng. 2016. Compact integration of multi-network topology for functional analysis of genes. *Cell systems* 3, 6 (2016).
- [9] Ching-Yao Chuang and Stefanie Jegelka. 2022. Tree Mover’s Distance: Bridging Graph Metrics and Stability of Graph Neural Networks. *NeurIPS* 35 (2022).
- [10] Quanyu Dai, Xiao-Ming Wu, Jiaren Xiao, Xiao Shen, and Dan Wang. 2022. Graph Transfer Learning via Adversarial Domain Adaptation with Graph Convolution. *TKDE* (2022), 1–1.
- [11] Alex Davies and Nirav Ajmeri. 2022. Realistic Synthetic Social Networks with Graph Neural Networks. arXiv:2212.07843 [cs.SI]
- [12] Yash Deshpande, Subhabrata Sen, Andrea Montanari, and Elchanan Mossel. 2018. Contextual stochastic block models. *NeurIPS* 31 (2018).
- [13] Luc Devroye, László Györfi, Gábor Lugosi, Luc Devroye, László Györfi, and Gábor Lugosi. 1996. Vapnik-Chervonenkis Theory. *A probabilistic theory of pattern recognition* (1996), 187–213.
- [14] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph Neural Networks for Social Recommendation.
- [15] Taoran Fang, Yunchao Zhang, Yang Yang, Chunping Wang, and Lei Chen. 2024. Universal prompt tuning for graph neural networks. *Advances in Neural Information Processing Systems* 36 (2024).
- [16] Taoran Fang, Wei Zhou, Yifei Sun, Kaiqiao Han, Lvbin Ma, and Yang Yang. 2024. Exploring Correlations of Self-supervised Tasks for Graphs. *arXiv preprint arXiv:2405.04245* (2024).
- [17] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *JMLR* 17, 59 (2016), 1–35.
- [18] Shurui Gui, Xiner Li, Limei Wang, and Shuiwang Ji. 2022. GOOD: A Graph Out-of-Distribution Benchmark. In *NeurIPS*.
- [19] Will Hamilton, Zhitaoying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *NeurIPS*.
- [20] Xiaotian Han, Zhimeng Jiang, Ninghao Liu, and Xia Hu. 2022. G-Mixup: Graph Data Augmentation for Graph Classification.
- [21] Renhong Huang, Jiarong Xu, Xin Jiang, Chenglu Pan, Zhiming Yang, Chunping Wang, and Yang Yang. 2024. Measuring Task Similarity and Its Implication in Fine-Tuning Graph Neural Networks. 38, 11 (2024), 12617–12625.
- [22] Zengfeng Huang, Shengzhong Zhang, Chong Xi, Tang Liu, and Min Zhou. 2021. Scaling up graph neural networks via graph coarsening. In *SIGKDD*. 675–684.
- [23] Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical Reparameterization with Gumbel-Softmax.
- [24] Wei Jin, Tong Zhao, Jiayuan Ding, Yozen Liu, Jiliang Tang, and Neil Shah. 2022. Empowering graph representation learning with test-time graph transformation. *ICLR* (2022).
- [25] Diederik P. Kingma and Jimmy Ba. 2017. Adam: A Method for Stochastic Optimization. arXiv:1412.6980 [cs.LG]
- [26] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks.
- [27] Jon M Kleinberg, Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, and Andrew S Tomkins. 1999. The web as a graph: Measurements, models, and methods. In *COCOON*. 1–17.
- [28] Haoyu Kuang, Jiarong Xu, Haozhe Zhang, Zuyu Zhao, Qi Zhang, Xuan-Jing Huang, and Zhongyu Wei. 2023. Unleashing the Power of Language Models in Text-Attributed Graph. In *EMNLP*. 8429–8441.
- [29] Hong Liu, Mingsheng Long, Jianmin Wang, and Michael Jordan. 2019. Transferable adversarial training: A general approach to adapting deep classifiers. In *ICML*.
- [30] Shikun Liu, Tianchun Li, Yongbin Feng, Nhan Tran, Han Zhao, Qiang Qiu, and Pan Li. 2023. Structural Re-weighting Improves Graph Domain Adaptation. In *ICML*. 21778–21793.
- [31] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I. Jordan. 2015. Learning Transferable Features with Deep Adaptation Networks.
- [32] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. 2018. Conditional adversarial domain adaptation. In *NeurIPS*. 1645–1655.
- [33] Hehuan Ma, Yatao Bian, Yu Rong, Wenbing Huang, Tingyang Xu, Weiyang Xie, Geyan Ye, and Junzhou Huang. 2020. Multi-View Graph Neural Networks for Molecular Property Prediction.
- [34] Ruoxue Ma, Jiarong Xu, Xinnong Zhang, Haozhe Zhang, Zuyu Zhao, Qi Zhang, Xuan-Jing Huang, and Zhongyu Wei. 2023. One-Model-Connects-All: A Unified Graph Pre-Training Model for Online Community Modeling. In *EMNLP*.
- [35] Haitao Mao, Lun Du, Yujia Zheng, Qiang Fu, Zelin Li, Xu Chen, Shi Han, and Dongmei Zhang. 2024. Source Free Graph Unsupervised Domain Adaptation. In *WSDM*. 520–528.
- [36] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annu. Rev. Sociol.* 27, 1 (2001), 415–444.
- [37] Giannis Nikolentzos, Polykarpos Meladinos, and Michalis Vazirgiannis. 2017. Matching node embeddings for graph similarity. In *AAAI*, Vol. 31.
- [38] Sinno Jialin Pan, James T Kwok, Qiang Yang, et al. 2008. Transfer learning via dimensionality reduction. In *AAAI*, Vol. 8. 677–682.
- [39] Ofir Pele and Michael Werman. 2009. Fast and robust earth mover’s distances. In *ICCV*. 460–467.
- [40] Joseph J Pfeiffer III, Sebastian Moreno, Timothy La Fond, Jennifer Neville, and Brian Gallagher. 2014. Attributed graph models: Modeling network structure with correlated attributes. In *WWW*. 831–842.
- [41] Ziyue Qiao, Xiao Luo, Meng Xiao, Hao Dong, Yuanchun Zhou, and Hui Xiong. 2023. Semi-supervised Domain Adaptation in Graph Transfer Learning. *IJCAI* (2023).
- [42] Ievgen Redko, Amaury Habrard, and Marc Sebban. 2017. Theoretical analysis of domain adaptation with optimal transport. In *ECML*. 737–753.
- [43] Leonardo FR Ribeiro, Pedro HP Saverese, and Daniel R Figueiredo. 2017. struc2vec: Learning node representations from structural identity. In *SIGKDD*.
- [44] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE transactions on neural networks* 20, 1 (2008), 61–80.
- [45] Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. 2018. Wasserstein distance guided representation learning for domain adaptation. In *AAAI*, Vol. 32.
- [46] Xiao Shen, Quanyu Dai, Fu-lai Chung, Wei Lu, and Kup-Sze Choi. 2020. Adversarial deep network embedding for cross-network node classification. In *AAAI*, Vol. 34. 2991–2999.
- [47] Xiao Shen, Quanyu Dai, Sitong Mao, Fu-Lai Chung, and Kup-Sze Choi. 2021. Network Together: Node Classification via Cross-Network Deep Network Embedding. *TNNLS* 32, 5 (2021), 1935–1948.
- [48] Boshen Shi, Yongqing Wang, Fangda Guo, Jiangli Shao, Huawei Shen, and Xueqi Cheng. 2023. Improving graph domain adaptation with network hierarchy. In *CIKM*. 2249–2258.
- [49] Yifei Sun, Haoran Deng, Yang Yang, Chunping Wang, Jiarong Xu, Renhong Huang, Linfeng Cao, Yang Wang, and Lei Chen. 2022. Beyond Homophily: Structure-aware Path Aggregation Graph Neural Network. In *IJCAI*. 2233–2240.
- [50] Yaver Titouan, Nicolas Courty, Romain Tavenard, and Rémi Flamary. 2019. Optimal transport for structured data with application on graphs. In *ICML*.
- [51] Petar Velicković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. 2018. Deep graph infomax. *ICLR* (2018).
- [52] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. 2017. Deep hashing network for unsupervised domain adaptation. In *CVPR*. 5018–5027.
- [53] Yanbang Wang, Yen-Yu Chang, Yunyu Liu, Jure Leskovec, and Pan Li. 2021. Inductive representation learning in temporal networks via causal anonymous walks. *ICLR* (2021).
- [54] Boris Weisfeiler and Andrei Leman. 1968. The reduction of a graph to canonical form and the algebra which appears therein. *nti, Series 2*, 9 (1968), 12–16.
- [55] Jun Wu, Jingrui He, and Elizabeth Ainsworth. 2023. Non-iid transfer learning on graphs. In *AAAI*, Vol. 37. 10342–10350.
- [56] Man Wu, Shirui Pan, Chuan Zhou, Xiaojun Chang, and Xingquan Zhu. 2020. Unsupervised Domain Adaptive Graph Convolutional Networks. In *WWW*.
- [57] Qitian Wu, Hengrui Zhang, Junchi Yan, and David Wipf. 2022. Handling Distribution Shifts on Graphs: An Invariance Perspective. In *ICLR*.
- [58] Jiarong Xu, Renhong Huang, Xin Jiang, Yuxuan Cao, Carl Yang, Chunping Wang, and Yang Yang. 2023. Better with Less: A Data-Active Perspective on Pre-Training Graph Neural Networks. *NeurIPS* (2023).
- [59] Jiarong Xu, Yang Yang, Junru Chen, Xin Jiang, Chunping Wang, Jiangang Lu, and Yizhou Sun. 2022. Unsupervised adversarially robust representation learning on graphs. In *AAAI*, Vol. 36. 4290–4298.
- [60] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How Powerful are Graph Neural Networks?. In *ICLR*.
- [61] Cheng Yang, Deyu Bo, Jixi Liu, Yufei Peng, Boyu Chen, Haoran Dai, Ao Sun, Yue Yu, Yixin Xiao, Qi Zhang, et al. 2023. Data-centric Graph Learning: A Survey. *arXiv preprint arXiv:2310.04987* (2023).
- [62] Yuning You, Tianlong Chen, Zhenyang Wang, and Yang Shen. 2023. Graph Domain Adaptation via Theory-Grounded Spectral Regularization. In *ICML*.
- [63] Werner Zellinger, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, and Susanne Saminger-Platz. 2017. Central moment discrepancy (cmd) for domain-invariant representation learning. *ICLR* (2017).

- [64] Werner Zellinger, Bernhard A. Moser, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, and Susanne Saminger-Platz. 2019. Robust unsupervised domain adaptation for neural networks via moment alignment. *Information Sciences* 483 (2019), 174–191.
- [65] Daochen Zha, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, Zhimeng Jiang, Shaochen Zhong, and Xia Hu. 2023. Data-centric artificial intelligence: A survey. *arXiv preprint arXiv:2303.10158* (2023).
- [66] Shuke Zhang, Yanzhao Jin, Tianmeng Liu, Qi Wang, Zhaohui Zhang, Shuliang Zhao, and Bo Shan. 2023. SS-GNN: a simple-structured graph neural network for affinity prediction. *ACS omega* 8, 25 (2023), 22496–22507.
- [67] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. 2020. Dataset condensation with gradient matching. *ICLR* (2020).
- [68] Ke Zhou, Hongyuan Zha, and Le Song. 2013. Learning social infectivity in sparse low-rank networks using multi-dimensional hawkes processes. In *AISTATS*.
- [69] Qi Zhu, Natalia Ponomareva, Jiawei Han, and Bryan Perozzi. 2021. Shift-robust gnn: Overcoming the limitations of localized graph training data. *NeurIPS* 34 (2021).

A APPENDIX

A.1 Addition Experimental Setup

Additional implementation details. We further elaborate on the hyper-parameters used and the running environment. In the implementation of GraphAlign, we set $\tau = 0.07$, $\kappa = 0.52$. As for the running environment, our model is implemented under the following software setting: Pytorch version 1.12.0+cu113, CUDA version 11.3, networkx version 2.3, torch-geometric version 2.3.0, sklearn version 1.0.2, numpy version 1.21.5, Python version 3.7.10. We conduct all experiments on the Linux system with an Intel Xeon Gold 5118 (128G memory) and a GeForce GTX Tesla P4 (8GB memory).

Estimation of MMD distance. We follow the standard procedure outlined by Pan et al. [38] to estimate MMD:

$$\text{MMD}^2(\mathbb{P}(f(G')), \mathbb{P}(f(G^{\mathcal{T}}))) \approx \frac{1}{n'(n'-1)} \sum_i \sum_{j \neq i}^{n'} k(x_i, x_j) + \frac{1}{n'(n'-1)} \sum_i \sum_{j \neq i}^{n'} k(y_i, y_j) - \frac{2}{n'n'} \sum_i \sum_j^{n'} k(x_i, y_j),$$

where k is the gaussian kernel defined by $k(x, x') = e^{-\frac{\|x-x'\|^2}{2\sigma^2}}$, and x and y are node representation matrix on the G' and $G^{\mathcal{T}}$, provided by the surrogate model. In practice, we sample an equivalent number of nodes n' from the target graph to compute the MMD to enhance the efficiency of estimation of the distribution of the target graph.

A.2 Proofs

Here, we begin by presenting the proof of Proposition 1&2. Furthermore, we provide the proof for Theorem 1 and Theorem 2.

Proof for Proposition 1. We start the proof as follows: Denote the GNN encoder as f and classifier g . Most of the previous methods primarily focus on learning domain-invariant representation, that is, the distribution of learnt representation $\mathbb{P}(f(G^{\mathcal{S}})) = \mathbb{P}(f(G^{\mathcal{T}}))$.

Let's recall that the CSBM models for the source and target graph come with specific structures:

$$C^{\mathcal{S}} = \begin{bmatrix} a & a \\ a & a - \Delta \end{bmatrix}, \quad C^{\mathcal{T}} = \begin{bmatrix} a - \Delta & a \\ a & a \end{bmatrix},$$

In this context, f is the most basic single-layer GNN, which typically employs message-passing mechanism for aggregation as follows:

$$h_v^{(k)} = \text{COMBINE}(h_v^{(k-1)}, \text{AGGREGATE}(\{h_u^{(k-1)} \mid u \in \mathcal{N}(v)\}))$$

where k is the layer number. From the above equation, we are aware that representations are impacted by neighbors as well as

the corresponding features (without loss of the generality, we let $\mu = \tilde{\mu} = \nu = \tilde{\nu}$ in CSBM). Therefore, the representation h is dominated by c_0 and c_1 , where c_0 or c_1 is the number of nodes in label 0 or 1. For simplicity, we denote the representation space for label 0 as $\zeta_0(g, f) = \{h \mid g \circ f(c_0, c_1) = 0, h = f(c_0, c_1)\}$ and $\zeta_1(g, f) = \{h \mid g \circ f(c_0, c_1) = 1, h = f(c_0, c_1)\}$ for label 1.

From the CSBM model, we can observe that the edge probability among the node labeled with $Y = 0$ in source domain is equal to node labelled with $Y = 1$ in the target domain, leading to the same neighbourhood structure. Therefore, no matter what model f, g are chosen, $P[\zeta_0(g, f) \mid Y^{\mathcal{S}} = 0] = P[\zeta_1(g, f) \mid Y^{\mathcal{T}} = 1]$. Therefore, denote the classification error for graph encoder f and classifier g in domain \mathcal{D} as $\epsilon_{\mathcal{D}}(g, f)$ as [45]. We have:

$$\begin{aligned} 1 &= P[\zeta_0(g, f) \mid Y^{\mathcal{S}} = 0] + P[\zeta_1(g, f) \mid Y^{\mathcal{S}} = 0] \\ &= P[\zeta_0(g, f) \mid Y^{\mathcal{T}} = 1] + P[\zeta_1(g, f) \mid Y^{\mathcal{S}} = 0] \leq 2(\epsilon^{\mathcal{T}}(g, f) + \epsilon^{\mathcal{S}}(g, f)). \end{aligned}$$

The last inequality is because

$$\begin{aligned} \epsilon^{\mathcal{S}}(g, f) &= P[\zeta_0(g, f) \mid Y^{\mathcal{S}} = 1]P[Y^{\mathcal{S}} = 1] + P[\zeta_1(g, f) \mid Y^{\mathcal{S}} = 0]P[Y^{\mathcal{S}} = 0] \\ &\geq \frac{1}{2} \max\{P[\zeta_0(g, f) \mid Y^{\mathcal{S}} = 1], P[\zeta_1(g, f) \mid Y^{\mathcal{S}} = 0]\}. \end{aligned}$$

The above deduction is the same with \mathcal{T} . In practical scenarios, models trained on seen data tend to perform well on seen data but poorly on unseen data. Therefore, based on practical considerations, we assume $\epsilon^{\mathcal{T}}(g, f) \geq \epsilon^{\mathcal{S}}(g, f)$. Under this assumption, we can prove $\epsilon^{\mathcal{T}}(g, f) \geq 0.25$.

Proof for Proposition 2. For simplicity but without losing generality, let's assume that the GNN encoder possesses the following characteristics: (1) Linearity. This property has already been demonstrated in some prior research [44]; (2) During message propagation, the normalization step is applied. Based on these properties, we can represent the encoder as follows: $f(c_0, c_1) = (c_0 + c_1)/n$.

Let's explore the constraint $\mathbb{P}(G') = \mathbb{P}(G^{\mathcal{T}})$ by the divergence of node representations. Intuitively, when $\mathbb{P}(G') = \mathbb{P}(G^{\mathcal{T}})$ is satisfied, it implies that for nodes with the same label, their nodes representation distributions have the same expectation. The expectation of the representation distribution can be expressed as follows (let's consider nodes of category 0):

$$\mathbb{E}(h') = a, \mathbb{E}(h^{\mathcal{T}}) = a - \frac{\Delta}{2}.$$

Therefore, the divergence of node representations can be calculated by $L_{\text{distance}} = \text{dis}(\mathbb{E}(h'), \mathbb{E}(h^{\mathcal{T}}))$, where $\text{dis}(\cdot, \cdot)$ is the distance metrics. Here, we choose the most common Frobenius norm function as our distance function, $L_{\text{distance}} = \text{dis}(\mathbb{E}(h') - \mathbb{E}(h^{\mathcal{T}})) = \|a - (a - \frac{\Delta}{2})\|_F = \frac{\Delta}{2}$. It's important to recall that we have two distinct cases as follows (where $B(\cdot)$ and $\text{Bern}(\cdot)$ represent the binomial distribution and Bernoulli distribution):

- If v is from class 0 in the source domain, $c_0 \sim B(n/2, a)$, $c_1 \sim B(n/2, a)$.
- If v is from class 0 in the target domain, $c_0 \sim B(n/2, a - \Delta)$, $c_1 \sim B(n/2, a)$.

As c_1 and c_0 are always independent, if v is from class 0 in the target domain, the node e $h = \frac{1}{n} (\sum_{i=1}^{n/2} Z_i - \sum_{i=1}^{n/2} Z'_i)$, where $Z_i \sim \text{Bern}(a)$ and $Z'_i \sim \text{Bern}(a)$, and all Z_i 's and Z'_i 's are independent. Therefore, using Hoeffding's inequality, we have

$$P(h - \mathbb{E}[h] > t) \leq \exp\left(-\frac{nt^2}{2}\right).$$

Defining the classifier as $g(h) = 0$ when $h < x$ or $g(h) = 1$ when $h > x$. Here, x is a constant. Under this classifier, we can find that the classification error for node with label 0 in source domain is $\exp\left(-\frac{nx^2}{2}\right)$. By setting $t = x$ and $\mathbb{E}[h] = 0$ for node with label 0 in source domain, $P(h > x) \leq \exp\left(-\frac{nx^2}{2}\right)$. Similarly, the classification error for node with label 1 in source domain is $\exp\left(-\frac{n(x-\frac{\Delta}{2})^2}{2}\right)$. In the context of ERM, the error on the source is minimized. Therefore, we optimize the optimal classifier based on $\frac{1}{2}\left(\exp\left(-\frac{nx^2}{2}\right) + \exp\left(-\frac{n(x-\frac{\Delta}{2})^2}{2}\right)\right)$. By analyzing the extreme points of the function, we can determine that the optimal classifier is given by setting $x = \frac{\Delta}{4}$. After using such a classifier in the target domain, we can similarly achieve a classification error of $\epsilon^{\mathcal{T}}(g, f) \leq 1 - \exp\left(-\frac{n\Delta^2}{32}\right)$. It shows that as Δ decreases, the classification error $\epsilon^{\mathcal{T}}$ becomes smaller. And, the distance function loss can serve as an upper bound to optimize and reduce the error

$$\epsilon^{\mathcal{T}}(g, f) \leq 1 - \exp\left(-\frac{n\Delta^2}{32}\right) = 1 - \exp\left(-\frac{nL_{\text{distance}}^2}{8}\right).$$

Therefore, there exist cases that classification error in the target domain can approach 0 by adopting a data-centric method.

Proof for Theorem 1. We first introduce the following inequality to be used that:

$$\begin{aligned} \epsilon^{\mathcal{T}}(g, f) &\leq \epsilon^{\mathcal{T}}(g^*, f^*) + \epsilon^{\mathcal{T}}(g, f|g^*, f^*) \\ &= \epsilon^{\mathcal{T}}(g^*, f^*) + \epsilon^{\mathcal{S}}(g, f|g^*, f^*) + \epsilon^{\mathcal{T}}(g, f|g^*, f^*) - \epsilon^{\mathcal{S}}(g, f|g^*, f^*). \end{aligned}$$

Here, we assume that the Lipschitz constant for the GNN model $g \circ f$ is denoted as L_{GNN} . We denote $\epsilon^{\mathcal{T}}(g, f|g^*, f^*)$ as $\epsilon^{\mathcal{T}}(g, f|g^*, f^*) = \mathbb{E}_{\mathbb{P}(G^{\mathcal{T}})}\left(\|g \circ f(G^{\mathcal{T}}) - g^* \circ f^*(G^{\mathcal{T}})\|\right)$ and $\epsilon^{\mathcal{S}}(g, f|g^*, f^*) = \mathbb{E}_{\mathbb{P}(G^{\mathcal{S}})}\left(\|g \circ f(G^{\mathcal{S}}) - g^* \circ f^*(G^{\mathcal{S}})\|\right)$.

According to Lemma 1 from [45], we proof the following equation:

$$\begin{aligned} \epsilon^{\mathcal{T}}(g, f) &\leq \epsilon^{\mathcal{T}}(g^*, f^*) + \epsilon^{\mathcal{S}}(g, f|g^*, f^*) + \epsilon^{\mathcal{T}}(g, f|g^*, f^*) - \epsilon^{\mathcal{S}}(g, f|g^*, f^*) \\ &\leq \epsilon^{\mathcal{T}}(g^*, f^*) + \epsilon^{\mathcal{S}}(g, f|g^*, f^*) + 2L_{\text{GNN}}W_1\left(\mathbb{P}(G^{\mathcal{S}}), \mathbb{P}(G^{\mathcal{T}})\right) \\ &\leq \epsilon^{\mathcal{T}}(g^*, f^*) + \epsilon^{\mathcal{S}}(g, f) + \epsilon^{\mathcal{S}}(g^*, f^*) + 2L_{\text{GNN}}W_1\left(\mathbb{P}(G^{\mathcal{S}}), \mathbb{P}(G^{\mathcal{T}})\right) \\ &= \epsilon^{\mathcal{S}}(g, f) + 2L_{\text{GNN}}W_1\left(\mathbb{P}(G^{\mathcal{S}}), \mathbb{P}(G^{\mathcal{T}})\right) + \eta. \end{aligned}$$

We next link the bound with the empirical risk and labeled sample size by showing, with probability at least $1 - \delta$ that:

$$\begin{aligned} \epsilon^{\mathcal{T}}(g, f) &\leq \epsilon^{\mathcal{S}}(g, f) + 2L_{\text{GNN}}W_1\left(\mathbb{P}(G^{\mathcal{S}}), \mathbb{P}(G^{\mathcal{T}})\right) + \eta \\ &\leq \hat{\epsilon}^{\mathcal{S}}(g, f) + 2L_{\text{GNN}}W_1\left(\mathbb{P}(G^{\mathcal{S}}), \mathbb{P}(G^{\mathcal{T}})\right) + \sqrt{\frac{2d}{n^{\mathcal{S}}}\log\left(\frac{en^{\mathcal{S}}}{d}\right)} + \sqrt{\frac{1}{2n^{\mathcal{S}}}\log\left(\frac{1}{\delta}\right)} + \eta. \end{aligned}$$

With the assistance of the Cauchy-Schwarz inequality, we can further derive the following inequality and provide the final conclusion.

$$\epsilon^{\mathcal{T}}(g, f) \leq \hat{\epsilon}^{\mathcal{S}}(g, f) + 2L_{\text{GNN}}W_1\left(\mathbb{P}(G^{\mathcal{S}}), \mathbb{P}(G^{\mathcal{T}})\right) + \sqrt{\frac{4d}{n^{\mathcal{S}}}\log\left(\frac{en^{\mathcal{S}}}{d}\right)} + \frac{1}{n^{\mathcal{S}}}\log\left(\frac{1}{\delta}\right) + \eta.$$

Proof of Theorem 2. Consider a graph encoder f (usually instantiated as a GNN) with k layers and 1-hop graph filter $\Lambda(L)$. Our focus lies on the central node's representation, which is acquired through a k -layer GCN utilizing a 1-hop polynomial filter $\Lambda(L) = Id - L$. This particular GNN model is widely employed in

various applications. We denote the representations of nodes i for all $i = 1, \dots, n$ in the final layer of the GCN, taking a node-wise perspective: $Z_i^{(k)} = \sigma\left(\sum_{j \in \mathcal{N}(i)} a_{ij} Z_j^{(k-1)T} \theta^{(k)}\right) \in \mathbb{R}^d$, where $a_{ij} = [\Lambda(L)]_{ij} \in \mathbb{R}$ is the weighted link between node i and j ; and $\theta^{(k)} \in \mathbb{R}^{d \times d}$ is the weight for the k -th layer sharing across nodes. Then $\theta = \left\{\theta^{(\ell)}\right\}_{\ell=1}^k$. We may denote $Z_i^{(\ell)} \in \mathbb{R}^d$ similarly for $\ell = 1, \dots, k-1$, and $Z_i^0 = X_i \in \mathbb{R}^d$ as the node feature of center node i . With the assumption of GCN in the statement, we consider that only the k -hop ego-graph G'_i centered at X_i is needed to compute $Z_i^{(k)}$ for any $i = 1, \dots, n$.

Denote $L_{G'_i}$ as the out-degree normalised graph Laplacian of G'_i , which is defined with respect to the direction from leaves to the centre node in G'_i . We write the ℓ -th layer representation as follows

$$\left[Z_i^{(\ell)}\right]_{k-\ell+1} = \sigma\left([\Lambda(L_{G'_i})]_{k-\ell+1} \left[Z_i^{(\ell-1)}\right]_{k-\ell+1} \theta^{(\ell)}\right).$$

Assume that $\forall i, \max_{\ell} \left\| \left[Z_i^{(\ell)}\right]_{k-\ell+1} \right\|_2 \leq c_z$, and $\max_{\ell} \left\| \theta^{(\ell)} \right\|_2 \leq c_{\theta}$. Suppose that the activation function σ is ρ_{σ} -Lipschitz function. Then, for $\ell = 1, \dots, k-1$, we have

$$\begin{aligned} \left\| \left[Z_i^{(\ell)}\right]_{k-\ell} - \left[Z_{i'}^{(\ell)}\right]_{k-\ell} \right\|_2 &\leq \rho_{\sigma} c_{\theta} \left\| \Lambda(L_{G'_i}) \right\|_2 \left\| \left[Z_i^{(\ell-1)}\right]_{k-\ell+1} - \left[Z_{i'}^{(\ell-1)}\right]_{k-\ell+1} \right\|_2 \\ &\quad + \rho_{\sigma} c_{\theta} c_z \left\| \Lambda(L_{G'_i}) - \Lambda(L_{G'_{i'}}) \right\|_2. \end{aligned}$$

Since $[\Lambda(L_{G'_i})]_{k-\ell+1}$ is the principle submatrix of $\Lambda(L_{G'_i})$. We equivalently write the above equation as $A_{\ell} \leq aA_{\ell-1} + b$, where a and b are the coefficient. And we have

$$\begin{aligned} A_{\ell} &\leq aA_{\ell-1} + b \leq a^2A_{\ell-2} + b(1+a) \leq \dots \\ &\leq a^{\ell}A_0 + \frac{a^{\ell}-1}{a-1}b. \end{aligned}$$

Therefore, for any $\ell = 1, \dots, k$, we have an upper bound for the hidden representation difference between G'_i and $G'_{i'}$ by substitute coefficient a and b ,

$$\begin{aligned} \left\| \left[Z_i^{(\ell)}\right]_{k-\ell} - \left[Z_{i'}^{(\ell)}\right]_{k-\ell} \right\|_2 &\leq (\rho_{\sigma} c_{\theta})^{\ell} \left\| \Lambda(L_{G'_i}) \right\|_2^{\ell} \left\| [X_i] - [X_{i'}] \right\|_2 \\ &\quad + \frac{(\rho_{\sigma} c_{\theta})^{\ell} \left\| \Lambda(L_{G'_i}) \right\|_2^{\ell-1}}{\rho_{\sigma} c_{\theta} \left\| \Lambda(L_{G'_i}) \right\|_2 - 1} \rho_{\sigma} c_{\theta} c_z \left\| \Lambda(L_{G'_i}) - \Lambda(L_{G'_{i'}}) \right\|_2. \end{aligned}$$

Specifically, for $\ell = k$, we obtain the upper bound for center node representation $\left\| \left[Z_i^{(k)}\right]_0 - \left[Z_{i'}^{(k)}\right]_0 \right\| \equiv \|Z_i - Z_{i'}\|$. Assuming that the difference in features between any two nodes does not exceed a constant, namely, $\|[X_i] - [X_{i'}]\|_2 \leq c_x$. Suppose that $\forall i, \left\| \Lambda(L_{G'_i}) \right\|_2 \leq c_L$ because that graph Laplacians are normalised. Since Λ is a linear function for L , we have

$$\begin{aligned} \|Z_i - Z_{i'}\|_2 &\leq (\rho_{\sigma} c_{\theta} c_L)^k c_x + \frac{(\rho_{\sigma} c_{\theta} c_L)^k - 1}{\rho_{\sigma} c_{\theta} c_L - 1} c_{\theta} c_z \left\| \Lambda(L_{G'_i}) - \Lambda(L_{G'_{i'}}) \right\|_2 \\ &\leq \chi_1 \left\| L_{G'_i} - L_{G'_{i'}} \right\|_2 + \chi_2, \end{aligned}$$

where $\chi_1 = \frac{(\rho_{\sigma} c_{\theta} c_L)^k - 1}{\rho_{\sigma} c_{\theta} c_L - 1} c_{\theta} c_z$ and $\chi_2 = (\rho_{\sigma} c_{\theta} c_L)^k c_x$.

Therefore, let's rephrase the following equation.

$$\left\| f(G') - f(G^{\mathcal{T}}) \right\|_2 \leq \frac{\chi_1}{n^{\mathcal{T}}} \sum_{i=1}^{n'} \sum_{j'=1}^{n'} \left\| L_{G'_i} - L_{G'_{j'}} \right\|_2 + \chi_2.$$

Finally, let $\xi_1 = \chi_1$ and $\xi_2 = \chi_2$, concluding the proof.