

Smoothness Really Matters: A Simple Yet Effective Approach for Unsupervised Graph Domain Adaptation

Wei Chen¹, Guo Ye², Yakun Wang², Zhao Zhang³, Libang Zhang²,
Daixin Wang², Zhiqiang Zhang², Fuzhen Zhuang^{1,4*}

¹Institute of Artificial Intelligence, Beihang University, Beijing, China

²Independent Researcher, Beijing, China

³Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

⁴Zhongguancun Laboratory, Beijing, China

{chenwei23,zhuangfuzhen}@buaa.edu.cn, zhangzhao2021@ict.ac.cn

Abstract

Unsupervised Graph Domain Adaptation (UGDA) seeks to bridge distribution shifts between domains by transferring knowledge from labeled source graphs to given unlabeled target graphs. Existing UGDA methods primarily focus on aligning features in the latent space learned by graph neural networks (GNNs) across domains, often overlooking structural shifts, resulting in limited effectiveness when addressing structurally complex transfer scenarios. Given the sensitivity of GNNs to local structural features, even slight discrepancies between source and target graphs could lead to significant shifts in node embeddings, thereby reducing the effectiveness of knowledge transfer. To address this issue, we introduce a novel approach for UGDA called Target-Domain Structural Smoothing (TDSS). TDSS is a simple and effective method designed to perform structural smoothing directly on the target graph, thereby mitigating structural distribution shifts and ensuring the consistency of node representations. Specifically, by integrating smoothing techniques with neighborhood sampling, TDSS maintains the structural coherence of the target graph while mitigating the risk of over-smoothing. Our theoretical analysis shows that TDSS effectively reduces target risk by improving model smoothness. Empirical results on three real-world datasets demonstrate that TDSS outperforms recent state-of-the-art baselines, achieving significant improvements across six transfer scenarios.

Code — <https://github.com/cwei01/TDSS>

Introduction

Graph Neural Networks (GNNs) have become a powerful tool for processing graph-structured data (Zhu et al. 2023; Wang et al. 2024; Zhu et al. 2024; Liao et al. 2024; Chen et al. 2024b,c; Liang et al. 2024a,b). Despite their notable success, GNNs often struggle with generalizing to distribution shifts, leading to unsatisfactory performance in the new domains (Li et al. 2022; Gao et al. 2023; Luo et al. 2024).

To tackle the challenge of distribution shifts across domains, Unsupervised Graph Domain Adaptation (UGDA) (Wilson and Cook 2020; Shi et al. 2024) provides an effective strategy by focusing on minimizing domain

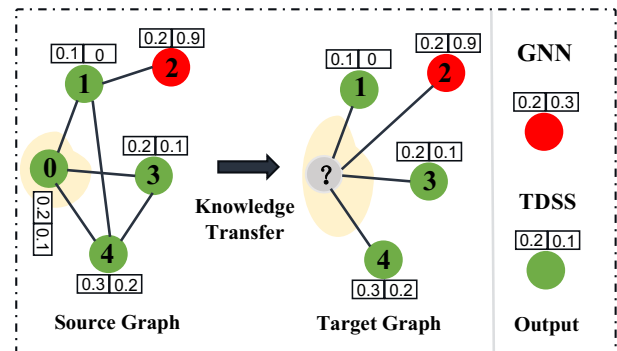


Figure 1: An example of GNNs sensitivity to similar inputs. Nodes have two features; green and red represent different classes. We assume the second feature determines node class: values ≤ 0.2 are green, and values > 0.2 are red.

discrepancies, enabling the transfer of knowledge from a well-labeled source domain to a target domain abundant in unlabeled samples. Inspired by domain adaptation success in computer vision (CV) (Xu et al. 2024b) and natural language processing (NLP) (Xu et al. 2023), most UGDA methods use GNNs to create dense node representations and apply regularization techniques like Maximum Mean Discrepancy (MMD) (Shen et al. 2020b; Liu et al. 2024b), graph subtree discrepancy (Wu, He, and Ainsworth 2023), and adversarial learning (Xiao et al. 2023; Li et al. 2024) to maintain consistency across source and target domains.

Nonetheless, these approaches often overlook the distinct impacts of distribution shifts driven by graph structures, which can yield suboptimal results. For instance, when performing domain adaptation between the *Cora* and *PubMed* citation networks¹ (Chen and Redner 2010), it’s crucial to consider the structural differences. Network motifs (Alon 2007), which are specific subgraph patterns capturing local characteristics, offer a means to quantify these differences. Specifically, in *Cora*, about 30% of the structures are triangular motifs (i.e., \triangle), indicating tight citation clusters and suggesting knowledge propagation within small, cohesive

*Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹They are widely used citation networks: *Cora* focuses on machine learning papers, while *PubMed* covers biomedical research.

communities. In contrast, the *PubMed* features star-shaped motifs (i.e., $*$), where many papers cite a few core papers, indicating reliance on a few influential works with limited direct connections among the citing papers. Although both networks involve overlapping research topics and citation patterns, the differences in their structural motifs could impact the effectiveness of domain adaptation strategies.

Conversely, in image and text domain (Chen et al. 2024a; Xu et al. 2024a), such similarities in feature distributions or content patterns can yield modest performance gains during knowledge transfer. This contrast underscores the fundamental differences in domain adaptation between graph structures and other data types. The reason lies in GNNs’ pronounced sensitivity to local structural stability, where even slight discrepancies between source and target graphs can result in significant shifts in node embeddings, thereby diminishing the effectiveness of knowledge transfer. Consider Figure 1 as an example. In the source graph, the input neighbor vectors $[0.1, 0]$, $[0.2, 0.1]$, and $[0.3, 0.2]$ result in a GNNs output vector of $[0.2, 0.1]$. However, in the target graph, despite a similar input, the GNNs produces a distinctly different output vector of $[0.2, 0.3]$ due to slight structural changes in neighbor node. This highlights the need for a method to constrain local variations and improve model robustness. In light of this, a few recent studies (Liu et al. 2023, 2024b) have primarily addressed structural shifts by reweighting edges to modify the source graph’s structure, thereby aligning it more closely with target graph. However, it may fall short when significant differences in nodes and edges exist between domains, as it might not fully capture target graph’s structural characteristics, particularly when transferring from a simpler graph to a more complex one.

Smoothing techniques (Field 1988), known for their ability to reduce inter-node discrepancies and improve the robustness of GNNs models, have been successfully applied in various tasks like node classification and link prediction. Inspired by the aforementioned analysis, this paper proposes a simple yet effective approach for UGDA called Target-Domain Structural Smoothing (TDSS) to address structural shifts problem. Our method is adaptable to various complex transfer scenarios. Specifically, by applying structural smoothing directly to the target graph, TDSS better accommodates the complex structures of the target domain. Specifically, it integrates a smoothing term with a neighborhood sampling strategy. The sampling mechanism flexibly explores the graph structure, avoiding over-smoothing while preserving key distinctions between nodes. Simultaneously, the smoothing term enhances model consistency, maintaining structural integrity in node representations and mitigating the effects of local perturbations and unstable features. Our theoretical analysis shows that target risk is primarily influenced by both model smoothness and domain discrepancies. Empirical results on three datasets demonstrate that TDSS outperforms recent state-of-the-art baselines.

Our major contributions are summarized as follows: **(Novel Perspective):** We are the first to approach UGDA through the lens of model smoothness. Our theoretical analysis reveals that target risk is predominantly driven by model smoothness and domain discrepancies, establishing a robust

theoretical foundation for our method. **(Simple yet Effective Method):** We introduce TDSS, a straightforward and effective method that seamlessly integrates into existing frameworks, providing a plug-and-play solution to enhance model generalization. **(Experimental Results):** Our experiments demonstrate TDSS’s adaptability and effectiveness across graph settings, surpassing state-of-the-art methods and consistently boosting performance in diverse transfer scenarios.

Related Work

Unsupervised Graph Domain Adaptation. UGDA aims to transfer knowledge from a well-labeled source graph to an unlabeled target graph, tackling the distribution shifts inherent in graph-structured data. Unlike traditional domain adaptation techniques (Zhuang et al. 2020), UGDA encounters greater challenges due to the non-IID nature of graph data, which complicates the alignment of node embeddings between the source and target graphs and can lead to mismatches and inconsistencies in the transferred knowledge. Consequently, many approaches have explored various techniques to achieve effective knowledge transfer, such as MMD (Shen et al. 2020b), graph subtree discrepancy (Wu, He, and Ainsworth 2023), and adversarial learning (Xiao et al. 2023; Li et al. 2024). However, these methods often miss the specific challenges posed by structural shifts in graphs, which can result in less effective solutions. Recently, a few studies have begun addressing distributional changes in graph structures by designing methods for structural alignment to reduce variations. For example, StruRW (Liu et al. 2023) reweights edges in the source graph to mitigate neighborhood shifts, while PairAlign (Liu et al. 2024b) recalibrates node influence through edge weights and adjusts classification loss with label weights to manage both structural and label shifts. Yet, when there is a significant disparity in the number of nodes and edges between the source and target graphs, direct edge reweighting may not fully capture the target graph’s intricate structure and may still result in suboptimal performance.

Model Smoothness. Model smoothness refers to designing a model that produces consistent and stable predictions across similar data points, ensuring that the prediction outputs are uniformly flat within certain neighborhoods (Rosca et al. 2020; Wang and Manchester 2023). This concept has already been both empirically and theoretically validated to improve model generalization and robustness in independent and identically distributed scenarios (Wang and Manchester 2023). In the UGDA setting, several approaches have sought to improve UGDA performance by incorporating additional model generalization constraints into their objectives. For example, SpecReg (You et al. 2023) derives a generalization bound for single-layer GNNs and enhances it by introducing spectral regularization. Similarly, A2GNN (Liu et al. 2024a) applies a Lipschitz formulation for k -layer GNNs, demonstrating that the target risk bound can be tightened by decreasing propagation layers in the source graph while increasing them in the target graph, though this approach can easily lead to over-smoothing. Different from these approaches, our work focuses on directly reducing target risk

by leveraging model smoothness, providing a novel perspective to enhance the effectiveness of UGDA.

Problem Formulation

Consider an undirected, unweighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of nodes connected by the edge set \mathcal{E} . The adjacency matrix of \mathcal{G} is a binary symmetric matrix \mathbf{A} such that $\mathbf{A}_{i,j} = 1$ if the i -th and j -th nodes are connected by an edge, and $\mathbf{A}_{i,j} = 0$ otherwise.

Given a source graph $\mathcal{G}_s = (\mathcal{V}_s, \mathcal{E}_s)$ with m labeled nodes and a target domain graph $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$ with n unlabeled nodes, we assume that both graphs share the same feature and label space but have different marginal distributions, i.e., $P(\mathcal{X}_s) \neq P(\mathcal{X}_t)$ but $P(\mathcal{Y}_s | \mathcal{X}_s) = P(\mathcal{Y}_t | \mathcal{X}_t)$. In this paper, our goal is to develop a GNN-based classifier that can predict the node labels $\{y_i^t\}_{i=1}^n$ in the target domain, where $y_i^t \in \mathcal{Y}_t$. We structure the model as a composition $g \circ \phi$, where the mapping $\phi : \mathcal{X} \rightarrow \mathcal{H}$ transforms the features into a latent space \mathcal{H} , and the function $g : \mathcal{H} \rightarrow \mathcal{Y}$ performs classification in the label space \mathcal{Y} .

Methodology

In this section, we begin by presenting target domain structural smoothing tailored for UGDA. Following this, we provide a theoretical analysis showing how target error is influenced by model smoothness, demonstrating the method’s effectiveness in enhancing knowledge transfer. Lastly, we offer a detailed analysis of the model’s complexity.

Target Domain Structural Smoothing

As previously discussed, even minor structural differences between the source and target domains can significantly alter node representations. To counteract this, we introduce a straightforward yet powerful approach called target domain structural smoothing, designed to mitigate the effects of structural shifts. The key concept is to enhance the smoothness of node representations in the target graph by constraining the variations between neighboring nodes, thereby minimizing the adverse impact of motif-related structural changes on the model’s predictive performance.

The overall process comprises two main steps: neighboring node generation and Laplacian smoothness constraint. In the first step, semantically similar nodes are sampled as neighbors for each node. The sampling method is flexible, accommodating various strategies such as random sampling, k -hop sampling, or random walk sampling. In this context, we present two specific techniques for neighbor sampling.

k-hop Sampling: In the k -hop sampling, the neighbors of each node are selected by traversing up to k edges in the graph, where k determines the size of the neighborhood. Formally, let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ represent an undirected graph, where \mathcal{V} is the set of nodes and \mathcal{E} is the set of edges. For a given node v_i , its k -hop neighbors $\mathcal{N}_k(v_i)$ are defined as:

$$\mathcal{N}_k(v_i) = \{v_j \in \mathcal{V} \mid \text{shortest_path}(v_i, v_j) \leq k\}, \quad (1)$$

where $\text{shortest_path}(v_i, v_j)$ denotes the shortest distance between nodes v_i and v_j .

Random Walk Sampling: The effectiveness of k -hop sampling is highly dependent on the parameter k . An inappropriate k can compromise the distinguishability of the learned embeddings. For example, a large k might lead to different nodes sharing identical subgraphs, causing their embeddings to collapse into a single point, which is a characteristic of over-smoothing in GNNs.

Instead, the random walk sampling enables flexible exploration of the graph structure during sampling, helping to avoid over-smoothing and maintain meaningful distinctions between nodes (Nikolentzos and Vazirgiannis 2020). Formally, for a given node $v_i \in \mathcal{V}$, the set of neighboring nodes $\mathcal{N}_r(v_i)$ is defined as:

$$\mathcal{N}_r(v_i) = \{v_j \in \mathcal{V} \mid \text{RW}_\lambda(v_i, v_j) = \text{True}\}, \quad (2)$$

where $\text{RW}_\lambda(v_i, v_j)$ indicates if node v_j is visited during a random walk of length λ starting at node v_i .

Next, let \mathbf{A} be the original adjacency matrix of graph \mathcal{G} . Here, we use random walk sampling as an example to illustrate the subsequent process. The updated adjacency matrix $\tilde{\mathbf{A}}$, reflecting the sampled neighborhood, can be defined as:

$$\tilde{\mathbf{A}}_{ij} = \begin{cases} 1, & \text{if } v_j \in \mathcal{N}_r(v_i) \text{ and } (v_i, v_j) \in \mathcal{E} \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

This updated matrix $\tilde{\mathbf{A}}$ ensures that only edges between node v_i and its sampled neighbors $\mathcal{N}_r(v_i)$ are retained, thereby emphasizing relevant connections within the graph.

Finally, we apply a smoothness constraint for each node to ensure consistency among neighboring nodes in target graph. In this context, we focus on applying this constraint solely to the target graph, as explained in **Remark 2**. Let $f(\mathbf{x}_i)$ denote the feature representation of node i , and d_i represent the degree of node i . The Laplacian smoothing regularization loss \mathcal{L}_{SR} is defined as follows:

$$\mathcal{L}_{\text{SR}} = \frac{1}{2} \sum_i \sum_j \tilde{\mathbf{A}}_{ij} \left\| \frac{f(\mathbf{x}_i)}{\sqrt{d_i}} - \frac{f(\mathbf{x}_j)}{\sqrt{d_j}} \right\|^2. \quad (4)$$

This constraint minimizes feature differences between nodes and their neighbors, normalized by the square root of their degrees. This normalization ensures that high-degree nodes do not overly dominate the smoothing process.

Target Risk Bound with Model Smoothness

In this part, we show that how the target risk is bounded above by the model’s smoothness, providing a theoretical foundation for the effectiveness of our approach. Formally, we define the model smoothness (Rosca et al. 2020) and total variation distance (TVD) (Villani et al. 2009).

Definition 1. Model Smoothness: A model f is (k, r) -cover with Φ smoothness on graph \mathcal{G} , if

$$\mathbb{E}_{\mathcal{G}} \left[\sup_{d(i,j) \leq k, \|\mathbf{x}_i - \mathbf{x}_j\|_\infty \leq r} |f(\boldsymbol{\theta}, \mathbf{x}_i) - f(\boldsymbol{\theta}, \mathbf{x}_j)| \right] \leq \Phi, \quad (5)$$

where ‘sup’ refers to the supremum, k is the maximum distance between neighboring nodes, and r denotes the maximum difference in the feature space between node features,

and θ is the model parameters. It should note that a lower Φ value indicates a higher level of model smoothness.

Definition 2. Total Variation Distance: Given two graphs \mathcal{G}_1 and \mathcal{G}_2 , $\mathcal{G}_1(v)$ and $\mathcal{G}_2(v)$ represent the attribute distributions of node v on graphs \mathcal{G}_1 and \mathcal{G}_2 respectively. The total variation distance is defined as:

$$\text{TVD}(\mathcal{G}_1, \mathcal{G}_2) = \frac{1}{2} \sum_{v \in V} |\mathcal{G}_1(v) - \mathcal{G}_2(v)|, \quad (6)$$

where it measures the maximum difference between the attribute distributions of two nodes on the graph.

As we proceed, let's assume that the graphs \mathcal{G}_1 and \mathcal{G}_2 from the source and target domains have a compact support, denoted as $\mathcal{X} \subseteq \mathbb{R}^d$, where d is the feature dimension. This implies there exists a constant $\Gamma > 0$ such that for any two nodes $\mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}$, the difference between their features, $\|\mathbf{x}_i - \mathbf{x}_j\|$, is also constrained by Γ . Consider $\mathcal{L}(f(\mathbf{x}), y)$ as the loss function, which is continuous and differentiable. Here, $f(\mathbf{x})$ is the output of the GNNs predicting the node label, and y is the ground truth of node. We define the expected risk for the model f across the distribution \mathbb{S} as $\mathcal{E}_{\mathbb{S}}(f) = \mathbb{E}_{\{\mathbf{x}, y\} \sim \mathbb{S}} [\mathcal{L}(f(\mathbf{x}), y)]$. Additionally, we assume that the loss is bounded, i.e., $0 \leq \mathcal{L}(f(\mathbf{x}), y) \leq \Upsilon$. By combining our understanding of model smoothness and the concept of TVD, we can establish the following theorem.

Theorem 1. Consider two distributions \mathbb{S} and \mathbb{T} , if a model f is $(k, 2r)$ -cover with Φ smoothness over distributions \mathbb{S} and \mathbb{T} , then with probability at least $1 - \xi$, the following holds:

$$\mathcal{E}_{\mathbb{T}}(f) \leq \mathcal{E}_{\mathbb{S}}(f) + 2\text{TVD}(\mathbb{S}, \mathbb{T}) + \Phi_{\mathbb{S}} + \Phi_{\mathbb{T}} + K, \quad (7)$$

where

$$K = \frac{\Upsilon Z}{\sqrt{m}} + \frac{\Upsilon Z}{\sqrt{n}} + \Upsilon \sqrt{\frac{\log(1/\xi)}{2m}}, \quad (8)$$

and

$$Z = \sqrt{(2d)^{\frac{2\Phi^2\Gamma}{r^2}+1} \log 2 + 2 \log(1/\xi)}. \quad (9)$$

Proof. Refer to the Appendix for the details. \square

In this theorem, K and Z are intermediate variables, Γ and Υ are constants specific to the given dataset, and m and n denote the number of nodes in the source and target domains, respectively. It is important to note that, unlike previous methods that derive upper bounds for target error (You et al. 2023; Liu et al. 2024a), we propose that the target risk $\mathcal{E}_{\mathbb{T}}(f)$ is constrained by the source risk $\mathcal{E}_{\mathbb{S}}(f)$, domain discrepancy $\text{TVD}(\mathbb{S}, \mathbb{T})$, and the model's smoothness Φ . This relationship is encapsulated in Theorem 1, which provides a more comprehensive understanding of the factors impacting the target risk. Therefore, by reducing model's smoothness Φ , we can effectively decrease the target error, offering a novel perspective for addressing the UGDA problem. It is also the core of our proposed TDSS approach.

Remark 1: (Why not directly increase GNNs depth?) Some readers may wonder why we do not simply increase the depth of GNNs to enhance model smoothness. While adding more layers can increase smoothness, it often results

in over-smoothing, which leads to the loss of critical node-specific features. In contrast, our proposed TDSS provides a better balance by enhancing smoothness while preserving local structural details, making it ideal for improving model robustness in UGDA tasks. Additionally, TDSS aligns more closely with the mathematical formulation of the target risk bound (i.e., Eq. (7)), as it directly influences Φ , the smoothness parameter we aim to optimize. This alignment ensures that our theoretical insights translate into tangible improvements in model's generalization capabilities.

Remark 2: (Target or source domain?) Since UGDA involves two distributions, \mathbb{S} (source domain) and \mathbb{T} (target domain), one might intuitively consider smoothing both the source and target graphs simultaneously. However, this approach can be computationally expensive, necessitating further optimization. In fact, the target domain \mathbb{T} generally requires greater smoothness than the source domain \mathbb{S} because the source domain benefits from labeled data, enabling more stable feature representations. Conversely, the target domain lacks supervision, making it more susceptible to noise and variability, thus requiring enhanced smoothness for effective generalization. Thus, we have $\Phi_{\mathbb{S}} < \Phi_{\mathbb{T}}$. As a consequence, Eq. (7) can be reformulated as:

$$\mathcal{E}_{\mathbb{T}}(f) \leq \mathcal{E}_{\mathbb{S}}(f) + 2\text{TVD}(\mathbb{S}, \mathbb{T}) + 2\Phi_{\mathbb{T}} + K, \quad (10)$$

where K is defined in Eq. (8). This approach reduces computational costs while minimizing the target error.

Remark 3: (Why total variation distance?) In proving target error, we opt to use TVD instead of MMD or other measures due to its robust and interpretable assessment of distributional discrepancy. TVD evaluates the maximum difference between probabilities assigned by two distributions, making it particularly suitable for scenarios where absolute differences are critical, and is favored for its convenience in theoretical derivations (Devroye, Mehrabian, and Reddad 2018). On the contrary, MMD captures higher-order moment differences but relies on kernel choices and large sample sizes for convergence, adding complexity and computational overhead. Thus, we use TVD to analyze the new target error bound $\mathcal{E}_{\mathbb{T}}(f)$. While in fact, with suitable kernel functions, MMD can approximate an upper bound for TVD:

$$\text{TVD}(\mathbb{S}, \mathbb{T}) \leq \text{MMD}_k(\mathbb{S}, \mathbb{T}) \cdot C_k, \quad (11)$$

where C_k is a constant related to the kernel function. Consequently, these measures of domain discrepancy can be regarded as equivalent under mild conditions, thereby preserving the correctness of our derivations regardless of the metric chosen. Furthermore, by combining Eq. (10) and Eq. (11), the new target error can be rewritten as:

$$\mathcal{E}_{\mathbb{T}}(f) \leq \underbrace{\mathcal{E}_{\mathbb{S}}(f)}_{\mathcal{L}_{\text{GC}}} + \underbrace{2\text{MMD}_k(\mathbb{S}, \mathbb{T}) \cdot C_k}_{\mathcal{L}_{\text{DA}}} + \underbrace{2\Phi_{\mathbb{T}}}_{\mathcal{L}_{\text{SR}}} + K, \quad (12)$$

where K is defined in Eq. (8). Interestingly, we observe that the target error is primarily constrained by three factors (i.e., source error, the MMD distance between domains, and model smoothness of target domain), each of which aligns closely with our optimization objectives in Eq. (13). This alignment indicates a strong consistency between our theoretical framework and the practical optimization goals, further reinforcing the validity of our approach.

Overall Optimization

The overall training process of our proposed TDSS consists of three main components:

$$\mathcal{L} = \mathcal{L}_{GC} + \alpha\mathcal{L}_{DA} + \beta\mathcal{L}_{SR}, \quad (13)$$

where the first term, \mathcal{L}_{GC} , represents the GNNs classifier loss. In this paper, we employ the state-of-the-art A2GNN (Liu et al. 2024a) to optimize this component. The \mathcal{L}_{DA} is the domain alignment loss, where we utilize the MMD method to align the source and target domains, aiming to minimize domain discrepancies. Finally, \mathcal{L}_{SR} is the smoothness loss, designed to ensure node smoothness by constraining feature variations between neighboring nodes. The α and β act as trade-offs to balance the contributions of the domain alignment and smoothness losses, respectively.

Complexity Analysis. In this section, we analyze the time complexity of TDSS, using the random walk sampling method as an example. Firstly, for each node performing γ random walks with a walk length of λ , the time complexity is $O(n\gamma\lambda)$, where n is the node number in target graph. The values of γ and λ are chosen to be small, thus not incurring significant computational overhead. Secondly, for each node calculating the smoothing constraint with its random walk sampled neighbors, assuming that each node samples on average ρ neighbors (i.e., $\rho \leq \gamma\lambda$), the time complexity is $O(n\rho d)$, where d is the number of features. Therefore, the overall time complexity is $O(n\gamma\lambda + n\rho d)$.

Experiment

In this section, we conduct experiments on three real-world datasets to address the following five research questions: **(RQ1)**: How does the proposed method TDSS perform when compared to state-of-the-art baseline methods? **(RQ2)**: How is the effectiveness of the proposed model evaluated across different backbone GNNs architectures? **(RQ3)**: Which part of the model primarily contributes to the effective prediction of optimal UGDA? **(RQ4)**: How do the hyper-parameters affect the performance of the proposed approach? **(RQ5)**: How about the intuitive effect of TDSS?

Experimental Setups

Datasets We conduct experiments utilizing three real-world graphs from the ArnetMiner dataset (Tang et al. 2008): **ACMv9 (A)**, **Citationv1 (C)**, and **DBLPv7 (D)**. To address the discrepancies in node attributes across these graphs, we integrate their attribute sets and standardized the attribute dimension to 6775, as elaborated in (Qiao et al. 2023). Each node signifies a paper, while each edge denotes a citation between two papers. Our objective is to categorize all the papers into five distinct research topics: Databases, Artificial Intelligence, Computer Vision, Information Security, and Networking. This study focuses on six key transfer tasks: **A** \rightarrow **C**, **A** \rightarrow **D**, **C** \rightarrow **A**, **C** \rightarrow **D**, **D** \rightarrow **A**, and **D** \rightarrow **C**. The statistics of three datasets are summarized in Table 1.

Baselines For performance evaluation, we compare TDSS with three categories of baseline methods: (1) Unsupervised graph learning: DeepWalk (Perozzi, Al-Rfou, and

Graph	#Nodes	#Edges	#Attr.	#Labels	#Density
ACMv9 (A)	9,360	15,556	6,775	5	0.00036
Citationv1(C)	8,935	15,098	6,775	5	0.00038
DBLPv7(D)	5,484	8,117	6,775	5	0.00054

Table 1: Statistics of the three real-world graphs. Note: ‘#’ denotes ‘number of’, ‘Attr.’ refers to attributes.

Skiena 2014), node2vec (Grover and Leskovec 2016), and ANRL (Zhang et al. 2018) learn node embeddings without supervision and then evaluate target graph representations using a classifier trained on the source graph. (2) Source-only GNNs: GAT (Velickovic et al. 2017), GSAGE (Hamilton, Ying, and Leskovec 2017), SGC (Wu et al. 2019), and GCN (Kipf and Welling 2016), are trained on the source graph in an end-to-end fashion, allowing direct application to the target graph. (3) Graph domain adaptation methods: CDNE (Shen et al. 2020b), AdaGCN (Dai et al. 2022), ACDNE (Shen et al. 2020a), UDAGCN (Wu et al. 2020), GRADE (Wu, He, and Ainsworth 2023), StruRW (Liu et al. 2023), PairAlign (Liu et al. 2024b) SpecReg (You et al. 2023) and A2GNN (Liu et al. 2024a) are designed specifically to handle graph domain adaptation.

Implementation Details Consistent with previous studies (Shen et al. 2020b; Liu et al. 2024a), we use Micro-F1 and Macro-F1 scores to evaluate model performance. Node features are set to a dimensionality of 128, and the learning rate was fine-tuned within the set $\{0.01, 0.03, 0.05\}$ for optimal performance. Domain alignment is achieved using the MMD method with a Gaussian kernel (Filippone et al. 2008). To prevent overfitting, a dropout rate of 0.5 is applied. We tune the α and β parameters within the range $[0, 1]$. All experiments are conducted using PyTorch, ensuring consistency with baseline models by adhering to the implementation details from their original papers.

Experiment Results

Performance Comparison (RQ1) We present the performance of TDSS compared with baselines for node classification, as shown in Table 2. From the results, we can observe:

(1) Graph domain adaptation methods outperform vanilla graph neural networks and hypothesis transfer approaches, underscoring the importance of addressing domain discrepancies. Vanilla GNNs struggle with domain shifts due to the assumption of identical training and test distributions. Hypothesis transfer methods adapt to new domains but fail to explicitly address graph structure discrepancies. In contrast, graph domain adaptation methods model domain shifts and adapt graph representations, leading to superior performance and highlighting the need for domain-aware adaptations in graph-based learning tasks.

(2) The proposed TDSS_{RW} model achieves the best performance across six cross-domain node classification datasets, outperforming all baseline methods. On average, TDSS_{RW} shows a 2.04% improvement in Macro-F1 and a 3.05% improvement in Micro-F1 over the state-of-the-art A2GNN model. These gains are due to the incorpo-

Methods	A → C		A → D		C → A		C → D		D → A		D → C		Avg.	
	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1	Micro-F1	Macro-F1
DeepWalk	0.2105	0.1772	0.2594	0.1987	0.2194	0.1933	0.2257	0.1751	0.2623	0.1983	0.2946	0.2276	0.2453	0.1956
node2vec	0.2989	0.2584	0.2454	0.1950	0.2176	0.1799	0.2895	0.2498	0.2861	0.2205	0.2116	0.1622	0.2582	0.2108
ANRL	0.3031	0.2093	0.2954	0.2333	0.3184	0.2204	0.2590	0.2271	0.2956	0.1912	0.2599	0.1825	0.2886	0.2099
GAT	0.5713	0.4364	0.5380	0.4136	0.5037	0.4214	0.5585	0.4525	0.5293	0.4395	0.5552	0.5004	0.5427	0.4436
GSAGE	0.7140	0.6914	0.6482	0.6180	0.6522	0.6469	0.6996	0.6686	0.5922	0.5731	0.6785	0.6490	0.6641	0.6379
SGC	0.7740	0.7213	0.6913	0.6254	0.7231	0.6292	0.7380	0.6693	0.6332	0.5394	0.7231	0.6292	0.7138	0.6290
GCN	0.7738	0.7478	0.6905	0.6529	0.7058	0.7039	0.7417	0.7137	0.6335	0.5942	0.7417	0.6979	0.7145	0.6824
UDAGCN	0.7215	0.6033	0.6695	0.6483	0.5816	0.5589	0.7177	0.6946	0.5816	0.5589	0.7328	0.6112	0.6675	0.5986
GRADE	0.7604	0.7252	0.6822	0.6303	0.6372	0.5935	0.7395	0.7002	0.6372	0.5935	0.7432	0.6932	0.6999	0.6551
CDNE	0.7876	0.7683	0.7158	0.6924	0.6962	0.7045	0.7436	0.7134	0.6962	0.7045	0.7888	0.7736	0.7380	0.7362
AdaGCN	0.7932	0.7651	0.7504	0.7139	0.6967	0.6947	0.7559	0.7234	0.6967	0.6947	0.7820	0.7422	0.7458	0.7254
StruRW	0.7735	0.7207	0.6910	0.6251	0.6781	0.5977	0.7381	0.6689	0.6327	0.5382	0.7241	0.6294	0.7016	0.6234
PairAlign	0.7088	0.6788	0.6591	0.6235	0.6585	0.6509	0.7104	0.6756	0.5934	0.5877	0.6707	0.6461	0.6602	0.6338
SpecReg	0.8055	0.7883	0.7593	0.7398	0.7101	0.7234	0.7574	0.7364	0.7101	0.7234	0.7904	0.7778	0.7556	0.7501
ACDNE	0.8175	0.8009	0.7624	0.7359	0.7129	0.7264	0.7721	0.7574	0.7129	0.7264	0.8014	0.7883	0.7513	0.7425
A2GNN	0.8275	0.8086	0.7784	0.7518	<u>0.7463</u>	<u>0.7631</u>	<u>0.7828</u>	<u>0.7590</u>	0.7344	0.7494	0.8109	0.7884	0.7795	0.7672
TDSS _{k-hop}	<u>0.8294</u>	<u>0.8131</u>	<u>0.7846</u>	<u>0.7611</u>	0.7433	0.7573	0.7713	0.7527	<u>0.7399</u>	<u>0.7573</u>	<u>0.8177</u>	<u>0.8011</u>	<u>0.7815</u>	<u>0.7765</u>
TDSS _{RW}	0.8320*	0.8140*	0.7954*	0.7813*	0.7694*	0.7833*	0.7896*	0.7755*	0.7489*	0.7646*	0.8229*	0.8075*	0.7954	0.7906

Table 2: Node classification performance comparisons on six cross-domain tasks. The highest scores are highlighted in **bold**, and the second-highest scores are underlined. * indicates that TDSS significantly outperforms A2GNN at the 0.05 level.

Backbone		GCN		GAT		SGC	
		Original	Ours	Original	Ours	Original	Ours
A → C	Mi-F1	0.7813	0.7930	0.7545	0.7996	0.7632	0.7765
	Ma-F1	0.7598	0.7758	0.7360	0.7748	0.6983	0.7224
A → D	Mi-F1	0.7055	0.7255	0.7213	0.7490	0.7535	0.7644
	Ma-F1	0.6750	0.6967	0.6861	0.7218	0.6852	0.7083
C → A	Mi-F1	0.7150	0.7230	0.7233	0.7425	0.6657	0.6818
	Ma-F1	0.7184	0.7298	0.7291	0.7529	0.5954	0.6394
C → D	Mi-F1	0.7583	0.7621	0.7403	0.7740	0.7485	0.7546
	Ma-F1	0.7343	0.7405	0.7153	0.7529	0.6791	0.7029
D → A	Mi-F1	0.6568	0.6709	0.6795	0.7157	0.6252	0.6467
	Ma-F1	0.6465	0.6700	0.6862	0.7136	0.5364	0.5549
D → C	Mi-F1	0.7487	0.7632	0.7652	0.7902	0.7374	0.7637
	Ma-F1	0.7187	0.7339	0.7347	0.7756	0.6380	0.6732

Table 3: Node classification performance on various GNNs for unsupervised domain adaptation.

ration of model smoothness, which mitigates the impact of local structural changes. The random walk-based Laplacian smoothing constraint in TDSS_{RW} enables better capture of the intrinsic graph structure and enhances robustness against noise, resulting in more accurate node classification. Moreover, the TDSS_{k-hop} model underperforms compared to TDSS_{RW}, likely due to over-smoothing during the optimization process. In k-hop sampling, nodes may share similar subgraphs as the hop count increases, leading to indistinguishable embeddings and a loss of critical node-specific features, which reduces the model’s effectiveness.

(3) We observe that the improvement in the C → A setting is the most significant. This enhancement may be attributed to the similarity in the number of nodes and edges between the Citationv1 and ACMv9 datasets. Such similarity enables the random walk-based smoothing constraint method to more effectively leverage the structural and attribute correspondences, resulting in smoother adaptation and higher classification accuracy in this specific transfer setting.

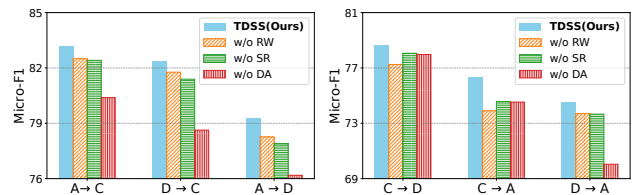


Figure 2: Classification Micro-F1 comparisons between TDSS variants on six cross-domain tasks.

Results for Various GNNs (RQ2). Since TDSS² is designed to be model-agnostic and seamlessly integrable into existing GNNs frameworks, it can be effectively integrated not only into state-of-the-art models like A2GNN but also into standard GNNs. To validate its broad applicability, we conduct experiments with GCN, GAT, and SGC. As shown in Table 3, the results demonstrate significant performance gains with our method compared to the original models without the smoothing constraint. Overall, these findings clearly indicate that our proposed TDSS approach enhances the generalization capabilities of various GNNs models. The consistent improvements across different tasks and backbones underscore the effectiveness and versatility of TDSS in advancing graph domain adaptation techniques.

Ablation Study (RQ3). To investigate the specific contributions of different components, we conduct an in-depth analysis and ablation studies on TDSS. The results are shown in Figure 2, and the variants of TDSS are: (1) **w/o DA**: which eliminates the alignment loss to assess the impact of domain alignment (i.e., setting $\alpha = 0$). (2) **w/o SR**: which eliminates the smooth regularization loss to evaluate the impact of model smoothness (i.e., $\beta = 0$). (3) **w/o RW**:

²To simplify notation, we use TDSS to refer to TDSS_{RW} in the following experiments.

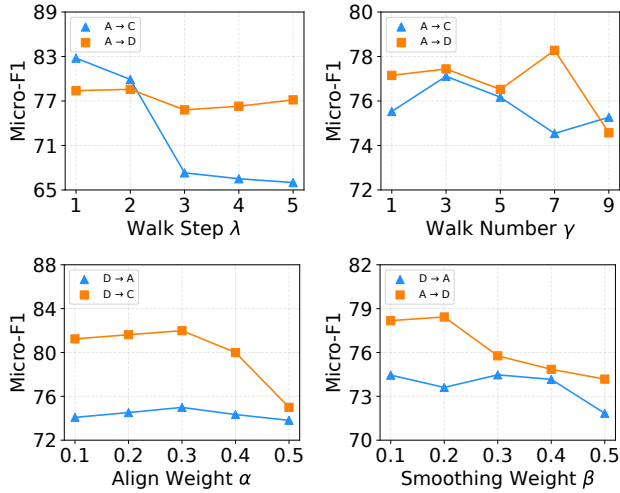


Figure 3: The Micro-F1 performances of our TDSS w.r.t varying λ , γ , α , and β on different transfer tasks.

which removes the random walk sampling mechanism and reverts to 1-hop neighborhood sampling.

First, we observe that removing specific modules consistently leads to a decrease in F1 scores, highlighting the essential role that various model components play in determining overall performance. Notably, in the majority of transfer tasks, the removal of MMD alignment (w/o DA) causes the most significant performance drop, emphasizing the critical importance of explicit domain alignment in minimizing distributional differences between two domains.

Further analysis of the comparison between TDSS and its variants on $C \rightarrow D$ and $C \rightarrow A$ reveals a key insight: the “w/o RW” variant shows inconsistent performance compared to other tasks, likely due to insufficient smoothness when using 1-hop neighbors for Laplacian constraints.

Sensitivity Analysis (RQ4). In this section, we investigate the sensitivity of TDSS to three key components: the random walk mechanism (walk number γ and walk length λ), domain alignment (weight α), and the smoothing constraint (weight β). Due to limited space, we present the experimental results of two transfer tasks per parameter in Figure 3.

• **Impact of walk length λ and walk number γ .** Firstly, we investigate the influence of λ and γ . When analyzing the walk step λ , distinct trends appear in the performance metrics. As λ increases, the Micro-F1 for $A \rightarrow C$ decreases, suggesting that longer walk steps introduce noise and over-smoothing. For $A \rightarrow D$, the Micro-F1 rises initially, peaking at $\lambda = 2$, before declining, indicating that a moderate walk step captures useful information without noise, emphasizing the need for domain-specific tuning. For the walk number γ , performance slightly drops at higher values, suggesting that a moderate number of walks (e.g., $\gamma = 3$ to $\gamma = 5$) balances information capture and redundancy. Excessive walks result in overlap, reducing performance. In summary, fine-tuning the walk step and number is essential for optimal UGDA performance, avoiding over-smoothing and redundancy.

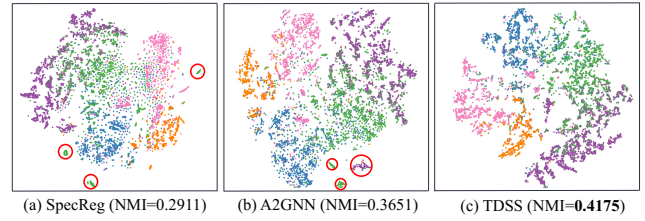


Figure 4: Visualization of learned node embedding on $D \rightarrow A$ for different models. We use Normalized Mutual Information (NMI) (Zhang 2015) to evaluate the clustering effect, with higher values indicating better performance.

• **Impact of domain aligning weight α .** Secondly, we study the impact of α , a parameter that regulates the intensity of domain aligning. In this analysis, we manipulate the value of α , selecting from the set $\{0.1, 0.2, 0.3, 0.4, 0.5\}$. Notably, we observe that optimal Micro-F1 performance is attained when $\alpha = 0.3$ for both of these transfer tasks. As α continues to increase beyond this point, we observe a decline in Micro-F1 score, suggesting the significance of selecting an appropriate value for α in our model.

• **Impact of smoothing weight β .** Finally, we analyze the impact of the weight β , which controls the strength of the smoothing constraint. We vary its values within the set $\{0.1, 0.2, 0.3, 0.4, 0.5\}$. In the $A \rightarrow D$ setting, optimal performance is achieved at $\beta = 0.2$, after which the Micro-F1 score gradually decreases. Similarly, in the $D \rightarrow A$ transfer scenario, the model reaches its best trade-off at $\beta = 0.3$. These results suggest that a smaller β is more effective for optimizing the smoothing constraint loss in TDSS.

Visualization Analysis (RQ5). To provide an intuitive understanding of TDSS, we visualize the node representations learned in the target domain for the $D \rightarrow A$ task, comparing it with three models. Node embeddings are visualized using Stochastic Neighbor Embedding (SNE) (Hinton and Roweis 2002), as shown in Fig 4. TDSS achieves the best clustering performance with an NMI of 0.4175, indicating clear separation between clusters and tighter groupings of nodes. In contrast, SpecReg and A2GNN exhibit more overlap and dispersion, with NMIs of 0.2911 and 0.3651, respectively. The red circles highlight areas where TDSS reduces inter-cluster mixing, demonstrating superior discriminative power and feature representation in complex domains.

Conclusion

In this paper, we present target-domain structural smoothing, a simple yet effective method for enhancing model transferability in UGDA tasks. By integrating Laplacian smoothing with a neighborhood sampling mechanism, TDSS preserves consistent node representations and strengthens model robustness against structural distribution shifts. Our theoretical analysis underscores the importance of addressing model smoothness for successful domain adaptation. Empirical results on three real-world datasets validate the superiority of TDSS over state-of-the-art baselines, and demonstrate significant improvements across a range of transfer scenarios.

Acknowledgments

We sincerely thank all the anonymous reviewers for their valuable comments to improve this paper. The research work is supported by the National Key Research and Development Program of China under Grant Nos. 2021ZD0113602, the National Natural Science Foundation of China under Grant No. 62176014 and No. 62206266, the Fundamental Research Funds for the Central Universities.

References

- Alon, U. 2007. Network motifs: theory and experimental approaches. *Nature Reviews Genetics*, 8(6): 450–461.
- Chen, P.; and Redner, S. 2010. Community structure of the physical review citation network. *Journal of Informetrics*, 4(3): 278–290.
- Chen, W.; Liu, Y.; Zhang, Z.; Zhuang, F.; and Zhong, J. 2024a. Modeling Adaptive Inter-Task Feature Interactions via Sentiment-Aware Contrastive Learning for Joint Aspect-Sentiment Prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 17781–17789.
- Chen, W.; Wu, Y.; Zhang, Z.; Zhuang, F.; He, Z.; Xie, R.; and Xia, F. 2024b. FairGap: Fairness-aware Recommendation via Generating Counterfactual Graph. *ACM Transactions on Information Systems*, 42(4): 1–25.
- Chen, W.; Yuan, M.; Zhang, Z.; Xie, R.; Zhuang, F.; Wang, D.; and Liu, R. 2024c. FairDgcl: Fairness-aware Recommendation with Dynamic Graph Contrastive Learning. *arXiv preprint arXiv:2410.17555*.
- Dai, Q.; Wu, X.-M.; Xiao, J.; Shen, X.; and Wang, D. 2022. Graph transfer learning via adversarial domain adaptation with graph convolution. *IEEE Transactions on Knowledge and Data Engineering*, 35(5): 4908–4922.
- Devroye, L.; Mehrabian, A.; and Reddad, T. 2018. The total variation distance between high-dimensional Gaussians with the same mean. *arXiv preprint arXiv:1810.08693*.
- Field, D. A. 1988. Laplacian smoothing and Delaunay triangulations. *Communications in applied numerical methods*, 4(6): 709–712.
- Filippone, M.; Camastra, F.; Masulli, F.; and Rovetta, S. 2008. A survey of kernel and spectral methods for clustering. *Pattern recognition*, 41(1): 176–190.
- Gao, C.; Zheng, Y.; Li, N.; Li, Y.; Qin, Y.; Piao, J.; Quan, Y.; Chang, J.; Jin, D.; He, X.; et al. 2023. A survey of graph neural networks for recommender systems: Challenges, methods, and directions. *ACM Transactions on Recommender Systems*, 1(1): 1–51.
- Grover, A.; and Leskovec, J. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 855–864.
- Hamilton, W.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30.
- Hinton, G. E.; and Roweis, S. 2002. Stochastic neighbor embedding. *Advances in neural information processing systems*, 15.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Li, H.; Wang, X.; Zhang, Z.; and Zhu, W. 2022. Ood-gnn: Out-of-distribution generalized graph neural network. *IEEE Transactions on Knowledge and Data Engineering*, 35(7): 7328–7340.
- Li, J.; Yu, Z.; Du, Z.; Zhu, L.; and Shen, H. T. 2024. A comprehensive survey on source-free domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Liang, K.; Meng, L.; Liu, M.; Liu, Y.; Tu, W.; Wang, S.; Zhou, S.; Liu, X.; Sun, F.; and He, K. 2024a. A survey of knowledge graph reasoning on graph types: Static, dynamic, and multi-modal. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Liang, K.; Meng, L.; Zhou, S.; Tu, W.; Wang, S.; Liu, Y.; Liu, M.; Zhao, L.; Dong, X.; and Liu, X. 2024b. MINES: Message Intercommunication for Inductive Relation Reasoning over Neighbor-Enhanced Subgraphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 10645–10653.
- Liao, W.; Zhu, Y.; Li, Y.; Zhang, Q.; Ou, Z.; and Li, X. 2024. RevGNN: Negative Sampling Enhanced Contrastive Graph Learning for Academic Reviewer Recommendation. *ACM Transactions on Information Systems*, 43(1): 1–26.
- Liu, M.; Fang, Z.; Zhang, Z.; Gu, M.; Zhou, S.; Wang, X.; and Bu, J. 2024a. Rethinking propagation for unsupervised graph domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 13963–13971.
- Liu, S.; Li, T.; Feng, Y.; Tran, N.; Zhao, H.; Qiu, Q.; and Li, P. 2023. Structural re-weighting improves graph domain adaptation. In *International Conference on Machine Learning*, 21778–21793. PMLR.
- Liu, S.; Zou, D.; Zhao, H.; and Li, P. 2024b. Pairwise Alignment Improves Graph Domain Adaptation. *arXiv preprint arXiv:2403.01092*.
- Luo, H.; Zhuang, F.; Xie, R.; Zhu, H.; Wang, D.; An, Z.; and Xu, Y. 2024. A survey on causal inference for recommendation. *The Innovation*, 5(2): 100590.
- Nikolentzos, G.; and Vazirgiannis, M. 2020. Random walk graph neural networks. *Advances in Neural Information Processing Systems*, 33: 16211–16222.
- Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 701–710.
- Qiao, Z.; Luo, X.; Xiao, M.; Dong, H.; Zhou, Y.; and Xiong, H. 2023. Semi-supervised domain adaptation in graph transfer learning. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 2279–2287.

- Rosca, M.; Weber, T.; Gretton, A.; and Mohamed, S. 2020. A case for new neural network smoothness constraints. In Zosa Forde, J.; Ruiz, F.; Pradier, M. F.; and Schein, A., eds., *Proceedings on "I Can't Believe It's Not Better!" at NeurIPS Workshops*, volume 137 of *Proceedings of Machine Learning Research*, 21–32. PMLR.
- Shen, X.; Dai, Q.; Chung, F.-l.; Lu, W.; and Choi, K.-S. 2020a. Adversarial deep network embedding for cross-network node classification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 2991–2999.
- Shen, X.; Dai, Q.; Mao, S.; Chung, F.-l.; and Choi, K.-S. 2020b. Network together: Node classification via cross-network deep network embedding. *IEEE Transactions on Neural Networks and Learning Systems*, 32(5): 1935–1948.
- Shi, B.; Wang, Y.; Guo, F.; Xu, B.; Shen, H.; and Cheng, X. 2024. Graph Domain Adaptation: Challenges, Progress and Prospects. *arXiv preprint arXiv:2402.00904*.
- Tang, J.; Zhang, J.; Yao, L.; Li, J.; Zhang, L.; and Su, Z. 2008. Arnetminer: extraction and mining of academic social networks. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 990–998.
- Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Lio, P.; Bengio, Y.; et al. 2017. Graph attention networks. *stat*, 1050(20): 10–48550.
- Villani, C.; et al. 2009. *Optimal transport: old and new*, volume 338. Springer.
- Wang, R.; and Manchester, I. 2023. Direct parameterization of lipschitz-bounded deep networks. In *International Conference on Machine Learning*, 36093–36110. PMLR.
- Wang, Y.; Wang, D.; Liu, H.; Hu, B.; Yan, Y.; Zhang, Q.; and Zhang, Z. 2024. Optimizing Long-tailed Link Prediction in Graph Neural Networks through Structure Representation Enhancement. *arXiv preprint arXiv:2407.20499*.
- Wilson, G.; and Cook, D. J. 2020. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(5): 1–46.
- Wu, F.; Souza, A.; Zhang, T.; Fifty, C.; Yu, T.; and Weinberger, K. 2019. Simplifying graph convolutional networks. In *International conference on machine learning*, 6861–6871. PMLR.
- Wu, J.; He, J.; and Ainsworth, E. 2023. Non-iid transfer learning on graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 10342–10350.
- Wu, M.; Pan, S.; Zhou, C.; Chang, X.; and Zhu, X. 2020. Unsupervised domain adaptive graph convolutional networks. In *Proceedings of the web conference 2020*, 1457–1467.
- Xiao, Z.; Wang, H.; Jin, Y.; Feng, L.; Chen, G.; Huang, F.; and Zhao, J. 2023. SPA: a graph spectral alignment perspective for domain adaptation. *Advances in Neural Information Processing Systems*, 36.
- Xu, B.; Zhao, C.; Jiang, W.; Zhu, P.; Dai, S.; Pang, C.; Sun, Z.; Wang, S.; and Sun, Y. 2023. Retrieval-augmented domain adaptation of language models. In *Proceedings of the 8th Workshop on Representation Learning for NLP (RepL4NLP 2023)*, 54–64.
- Xu, K.; Ma, Z.; Xu, L.; He, G.; Li, Y.; Yu, W.; Han, T.; and Yang, C. 2024a. An End-to-End Real-World Camera Imaging Pipeline. In *ACM Multimedia 2024*.
- Xu, K.; Xu, L.; He, G.; Yu, W.; and Li, Y. 2024b. Beyond Alignment: Blind Video Face Restoration via Parsing-Guided Temporal-Coherent Transformer. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, {IJCAI-24}*, 1489–1497.
- You, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2023. Graph domain adaptation via theory-grounded spectral regularization. In *The eleventh international conference on learning representations*.
- Zhang, P. 2015. Evaluating accuracy of community detection using the relative normalized mutual information. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(11): P11006.
- Zhang, Z.; Yang, H.; Bu, J.; Zhou, S.; Yu, P.; Zhang, J.; Ester, M.; and Wang, C. 2018. ANRL: attributed network representation learning via deep neural networks. In *Ijcai*, volume 18, 3155–3161.
- Zhu, H.; Xiong, F.; Chen, H.; Xiong, X.; and Wang, L. 2024. Incorporating a triple graph neural network with multiple implicit feedback for social recommendation. *ACM Transactions on the Web*, 18(2): 1–26.
- Zhu, Y.; Cong, F.; Zhang, D.; Gong, W.; Lin, Q.; Feng, W.; Dong, Y.; and Tang, J. 2023. Wingnn: Dynamic graph neural networks with random gradient aggregation window. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3650–3662.
- Zhuang, F.; Qi, Z.; Duan, K.; Xi, D.; Zhu, Y.; Zhu, H.; Xiong, H.; and He, Q. 2020. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1): 43–76.