

LEARNING STRUCTURE-SEMANTIC EVOLUTION TRAJECTORIES FOR GRAPH DOMAIN ADAPTATION

Wei Chen¹, Xingyu Guo¹, Shuang Li¹, Yan Zhong², Zhao Zhang³,

Fuzhen Zhuang^{1,4,†}, Hongrui Liu⁵, Libang Zhang⁵, Guo Ye⁵, Huimei He⁵

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

²School of Mathematical Sciences, Peking University, Beijing, China

³School of Computer Science and Engineering, Beihang University, Beijing, China

⁴Zhongguancun Laboratory, Beijing, China

⁵Independent Researcher, Beijing, China

{chenwei23, shuangliai, zhuangfuzhen}@buaa.edu.cn

ABSTRACT

Graph Domain Adaptation (GDA) aims to bridge distribution shifts between domains by transferring knowledge from well-labeled source graphs to given unlabeled target graphs. One promising recent approach addresses graph transfer by discretizing the adaptation process, typically through the construction of intermediate graphs or stepwise alignment procedures. However, such discrete strategies often fail in real-world scenarios, where graph structures evolve continuously and nonlinearly, making it difficult for fixed-step alignment to approximate the actual transformation process. To address these limitations, we propose **DiffGDA**, a **Diffusion-based GDA** method that models the domain adaptation process as a continuous-time generative process. We formulate the evolution from source to target graphs using stochastic differential equations (SDEs), enabling the joint modeling of structural and semantic transitions. To guide this evolution, a domain-aware network is introduced to steer the generative process toward the target domain, encouraging the diffusion trajectory to follow an optimal adaptation path. We theoretically show that the diffusion process converges to the optimal solution bridging the source and target domains in the latent space. Extensive experiments on 14 graph transfer tasks across 8 real-world datasets demonstrate DiffGDA consistently outperforms state-of-the-art baselines. Code is available at [DiffGDA](#).

1 INTRODUCTION

Graph Neural Networks (GNNs) (Zhu et al., 2021; Wu et al., 2022) have demonstrated remarkable success in various real-world applications (Zhang et al., 2022a; Therrien et al., 2025; Chen et al., 2024; 2025b). However, their performance often degrades under distribution shifts across training and testing graphs (Liu et al., 2020; Qin et al., 2022). To overcome this limitation, Graph Domain Adaptation (GDA) aims to transfer knowledge from a labeled source graph to an unlabeled target graph by mitigating both feature distribution discrepancy and structural divergence (Dai et al., 2022; Liu et al., 2024a; Chen et al., 2025a). This work focuses on node-level GDA (Ding et al., 2018; Wang et al., 2024), tackling critical challenges posed by feature and topological shifts across graphs.

In recent years, a wide range of methods have emerged to enhance GDA, which can be broadly categorized into two paradigms (Huang et al., 2024): (1) **Model-oriented** approaches (Xiao et al., 2023; Pilancı & Vural, 2020; Dan et al., 2024; Fang et al., 2025b) aim to learn domain-invariant representations by aligning both semantic and structural information through objectives such as minimizing subgraph structural distances (Liu et al., 2023a; You et al., 2023), reducing Maximum Mean Discrepancy (MMD) (Chen et al., 2019; 2025a), and using adversarial discriminators (Dai et al., 2022). While these methods promote representation sharing, they assume limited structural variation between domains, which restricts their effectiveness when substantial structural differences exist (Huang et al., 2024). (2) **Data-oriented** methods (Na et al., 2021; Huang et al., 2024; Lei

[†]Corresponding author.

et al., 2025) address these challenges by constructing intermediate graphs to bridge structural gaps, projecting both domains into a shared semantic space. This facilitates more stable training by aligning structures and capturing cross-graph correspondences, offering a more flexible and interpretable framework for GDA, and showing promising potential in tackling complex graph domain shifts.

However, most data-oriented methods rely on the assumption that **the graph transfer process is discrete**, meaning the source graph can be transformed into the target graph through a limited number of alignment steps. Unfortunately, this assumption often fails in real-world scenarios, particularly when dealing with unlabeled graphs, due to the key limitations: graphs are governed by different underlying processes (Deng et al., 2019; Sharma et al., 2024), such as social dynamics, citation growth, or knowledge diffusion, which evolve nonlinearly and in domain-specific ways. Their evolution is driven by complex, context-dependent dynamics, rather than fixed or proportional structural patterns. Additionally, semantic information arises from the interplay between node features and structural context, but it often lacks explicit anchors for alignment (Heimann et al., 2018; Zhao et al., 2023). These factors make it difficult to align graphs with a fixed number of steps, as structural and semantic differences cannot be effectively bridged through simple alignment processes.

Along this research line, **continuous evolution** (Zeng et al., 2024; Zhang et al., 2021) provides a more natural modeling paradigm for graph transfer. Rather than manually constructing discrete intermediate graphs, this approach directly models the underlying evolution process, where the source graph gradually transforms into the target graph over continuous time. This paradigm not only bypasses the limitations of step-by-step alignment but also offers several advantages for GDA: (1) Structural changes are represented as smooth temporal trajectories, enabling flexible adaptation to nonlinear and heterogeneous topologies without relying on rigid graph (or subgraph) level alignment. (2) Semantic information evolves continuously along the transformation path, allowing the model to automatically learn optimal alignment trajectories and establish accurate cross-domain correspondences in an end-to-end manner. This perspective motivates a generative view of GDA, where the transfer of structure and semantics is unified as a time-driven evolutionary process. Inspired by the success of diffusion models in capturing complex distributional transformations (Croitoru et al., 2023; Zeng et al., 2024), we adopt a stochastic differential equation (SDE) (Protter, 2012; Li et al., 2020) framework to represent cross-graph transfer as a continuous probabilistic flow.

In this work, we propose **DiffGDA**, a **Diffusion-based GDA** that models both structural and semantic adaptation as a continuous-time evolution process. To guide this evolution, we introduce a domain-aware network that learns to steer the generative trajectory toward the target domain by uncovering the optimal adaptation path. This adaptive guidance allows the model to overcome the limitations of conventional discrete alignment methods by continuously directing the evolution toward the target domain, resulting in more coherent structural alignment and semantic adaptation. Furthermore, we provide a theoretical guarantee that the guided diffusion process converges to an optimal adaptation trajectory, effectively bridging the source and target domains.

Overall, our contributions can be summarized as:

(i) **Continuous Evolution Modeling.** We formulate GDA as a continuous-time generative process, jointly capturing structural and semantic transitions via stochastic differential equations. To the best of our knowledge, DiffGDA is the first work that brings diffusion into GDA. (ii) **Domain-Aware Guidance Network.** We propose a density-ratio-based guidance mechanism that adaptively steers the diffusion trajectory, enabling precise, target-aware evolution. (iii) **Empirical Superiority.** Extensive experiments on 14 real-world transferring tasks demonstrate the effectiveness of the proposed DiffGDA, consistently outperforming state-of-the-art baselines.

2 RELATED WORK

Graph Domain Adaptation (GDA) has attracted significant research attention due to its crucial role in enabling knowledge transfer between graphs with distributional shifts (Ding et al., 2018; Zhang et al., 2022b; Liu et al., 2023a; 2024a; Chen et al., 2025a; Guo et al., 2025; Kou et al., 2025b;a). Early approaches (Zhang et al., 2019; Wu et al., 2020; Dai et al., 2022; Pilanci & Vural, 2020; Dan et al., 2024) primarily adopted a model-driven paradigm, focusing on learning domain-invariant representations through techniques like adversarial training and MMD minimization (Liu et al., 2024a; You et al., 2023). While effective, these methods often struggle due to their reliance on rigid parametric

frameworks, which impose inherent limitations on their performance. Recently, data-driven methods (Huang et al., 2024; Lei et al., 2025) have emerged as a promising alternative, addressing domain gaps through intermediate graph construction and transformation. However, these approaches typically rely on heuristic alignment rules, which often lack generality and cannot adaptively discover or generate intermediate graphs that best suit the structural and semantic characteristics of the graph data, thereby limiting their effectiveness in handling complex graph domain transitions.

Graph Diffusion Models. Diffusion models (Cao et al., 2024) have achieved significant success across a variety of domains (Kong et al., 2021; Saharia et al., 2022; Luo et al., 2024b; Zhang et al., 2025; Chen et al., 2025c). In graph domain, these models are increasingly applied to model complex graph-structured data, particularly for tasks such as node and edge generation, molecule creation, and protein design (Niu et al., 2020; Jo et al., 2022; Konstantin Haefeli et al., 2022). By carefully designing the noising process and model architectures, symmetry constraints on the transition kernel and prior distribution can be explicitly incorporated, ensuring that the generated data follows roto-translationally invariant distributions. Recent advancements have expanded these frameworks beyond score-based formulations, introducing techniques like flow matching (Lipman et al., 2023; Liu et al., 2023b) and optimal transport-based diffusion (Peyré et al., 2019). These approaches provide alternative mechanisms to enforce geometric and permutation symmetry constraints in graph generation tasks. However, most existing methods have focused on symmetric diffusion processes. This presents a key challenge for diffusion-based domain adaptation models: how to leverage asymmetric diffusion mechanisms to enable transfer learning across heterogeneous graphs.

3 PRELIMINARIES

Notation. Given a labeled source graph $\mathbf{G}^S = (\mathcal{V}^S, \mathcal{E}^S)$, where \mathcal{V}^S is the set of nodes, $\mathcal{E}^S \subseteq \mathcal{V}^S \times \mathcal{V}^S$ is the set of edges, and the nodes \mathcal{V}^S are associated with a feature matrix $\mathbf{X}^S \in \mathbb{R}^{N_S \times F}$ and a label matrix $\mathbf{Y}^S \in \{0, 1\}^{N_S \times C}$. The adjacency matrix of the source graph is denoted as $\mathbf{A}^S \in \{0, 1\}^{N_S \times N_S}$. Similarly, we have an unlabeled target graph $\mathbf{G}^T = (\mathcal{V}^T, \mathcal{E}^T)$ with node features $\mathbf{X}^T \in \mathbb{R}^{N_T \times F}$ and adjacency matrix $\mathbf{A}^T \in \{0, 1\}^{N_T \times N_T}$. Both graphs share the same feature space (F -dimensional) and label space (C classes), but they exhibit domain shift. In this work, our goal is to train a GNN that accurately predicts the node labels of the target graph \mathbf{G}^T .

Diffusion Models. Diffusion models (Croitoru et al., 2023; Cao et al., 2024) have emerged as leading generative techniques in various fields (Luo et al., 2024b; Kong et al., 2021; Liu et al., 2025). These models are built upon the concept of a diffusion process that gradually transforms data into noise and then trains models to reverse this process using a well-defined objective function. Formally, the forward process incrementally corrupts data $\mathbf{X}_0 \sim p(\mathbf{X}_0)$ into noise through a continuous-time stochastic differential equation (SDE) (Protter, 2012):

$$d\mathbf{X}_t = \mathbf{f}(\mathbf{X}_t, t)dt + g(t)d\mathbf{w}_t, \quad t \in [0, T], \quad (1)$$

where \mathbf{w}_t denotes the standard Brownian motion (also called the Wiener process) (Luo et al., 2024a), $\mathbf{f}(\mathbf{X}_t, t) : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$ is a vector-valued drift coefficient determining the deterministic evolution of the system, and $g(t) : [0, T] \rightarrow \mathbb{R}$ is a scalar diffusion coefficient controlling the magnitude of Gaussian noise added at each time step. A canonical example is the Ornstein-Uhlenbeck process (Maller et al., 2009), where $\mathbf{f}(\mathbf{X}_t, t) = -\frac{1}{2}\beta(t)\mathbf{X}_t$ and $g(t) = \sqrt{\beta(t)}$, with $\beta(t)$ being a noise schedule that determines the rate of data corruption. As $t \rightarrow T$, the distribution $p_t(\mathbf{X}_t)$ converges to a simple prior distribution (typically standard Gaussian), regardless of the initial data distribution $p_0(\mathbf{X}_0)$. The reverse process aims to invert this corruption using the reverse-time SDE:

$$d\mathbf{X}_t = [\mathbf{f}(\mathbf{X}_t, t) - g^2(t)\nabla_{\mathbf{X}_t} \log p_t(\mathbf{X}_t)] d\bar{t} + g(t)d\bar{\mathbf{w}}_t, \quad (2)$$

where $\nabla_{\mathbf{X}_t} \log p_t(\mathbf{X}_t)$ denotes score function, $d\bar{t}$ is an infinitesimal negative time step, and $\bar{\mathbf{w}}_t$ is the reverse-time Brownian motion (Freedman, 2012). The score function is approximated by a neural network trained via score matching (Kane et al., 2020), which minimizes the Fisher divergence between the model and the true data distribution to ensure accurate and data-consistent generation.

4 METHODOLOGY

In this section, we present the **DiffGDA** framework (Figure 1). We first introduce §4.1 *Domain-Guided Graph Diffusion*, which formulates adaptation as a continuous-time evolution. We then de-

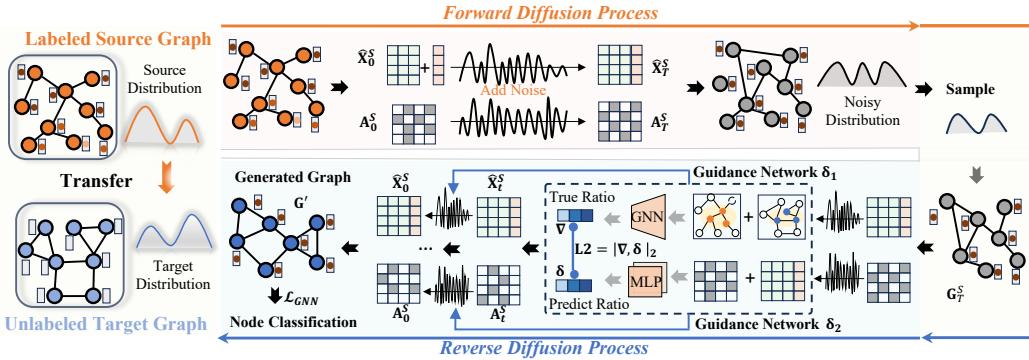


Figure 1: Overall architecture of our proposed DiffGDA framework: (1) In the forward diffusion process, labeled source graphs \mathbf{G}^S are progressively perturbed with noise, approximating a Gaussian distribution. (2) During the reverse diffusion process, a guidance network aligns predicted and true density ratios, reconstructing intermediate graphs to match the target distribution. The generated labeled graphs are then utilized for node classification training with a GNN.

scribe §4.2 *Practical Implementation*, covering the design of score and guidance networks. Finally, we detail the model optimization that integrates diffusion with GNN for end-to-end training.

4.1 DOMAIN-GUIDED GRAPH DIFFUSION PROCESS

Forward Diffusion Process. DiffGDA formulates GDA as a continuous-time stochastic process, governed by a system of SDEs. For a source graph denoted as $\mathbf{G}^S = (\mathbf{X}^S, \mathbf{A}^S, \mathbf{Y}^S)$, we design a forward diffusion process that systematically perturbs both node features and adjacency matrices toward a tractable noise distribution. To integrate the source domain label knowledge into the diffusion process, we propose a structured feature concatenation to explicitly establish the semantic association between node features and labels. Specifically, the original node feature matrix \mathbf{X}^S and the label matrix \mathbf{Y}^S are concatenated along the channel dimension to construct augmented features: $\tilde{\mathbf{X}}^S = [\mathbf{X}^S || \mathbf{Y}^S] \in \mathbb{R}^{N_S \times (F+C)}$, where $||$ denotes the concatenation operation.

Formally, this process is defined as a continuous-time stochastic trajectory $\{\mathbf{G}_t^S = (\tilde{\mathbf{X}}_t^S, \mathbf{A}_t^S)\}_{t \in [0, T]}$ over a fixed temporal domain $[0, T]$, where \mathbf{G}_0^S represents the initial source graph. The complete dynamics of the forward process are governed by the following SDE:

$$d\mathbf{G}_t^S = \mathbf{f}_t(\mathbf{G}_t^S)dt + g_t(\mathbf{G}_t^S)d\mathbf{w}, \quad (3)$$

where \mathbf{f}_t governs the deterministic drift, g_t controls the noise injection rate. The Wiener process \mathbf{w} drives the stochastic component. Overall, this process transforms the initial graph structure $\mathbf{G}_0^S = (\tilde{\mathbf{X}}_0^S, \mathbf{A}_0^S)$ into a simple noise distribution through designed drift term \mathbf{f}_t and noise term $g_t d\mathbf{w}$.

Reverse Diffusion Process. The reverse diffusion process aims to generate graphs that adhere to the target data distribution. It begins by sampling from the prior distribution and traversing the diffusion process of Eq. (3) backward in time. This reverse process is formulated as:

$$d\mathbf{G}_t^S = [\mathbf{f}_t(\mathbf{G}_t^S) - g_t^2 \nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S)] d\bar{t} + g_t d\bar{\mathbf{w}}, \quad (4)$$

where p_t denotes the marginal distribution under the forward diffusion process at time t , $\bar{\mathbf{w}}$ is a reverse-time standard Wiener process, and $d\bar{t}$ represents an infinitesimal negative time step.

The key in reverse process is estimating the score function $\nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S)$. Typically, we can train a neural network $\mathbb{P}(\ell)$ via score matching to approximate this score function (Ouyang et al., 2024), thereby facilitating discrete-time generation via Eq. (4). However, this unconditional generation process only reconstructs the source graph distribution. In our scenario, the objective is to generate graphs conforming to a predefined target graph distribution \mathbf{G}^T , which differs from the source distribution \mathbf{G}^S . To address this issue, we propose an additional auxiliary guidance network $\mathbb{Q}(\delta)$, which steers generation process from the source distribution toward target distribution. Therefore,

the optimization objective in reverse process is expressed as:

$$\mathcal{L}_{\text{SDE}} = \mathcal{F}\left(\mathbb{P}(\ell), \nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S)\right) + \mathcal{F}\left(\mathbb{Q}(\delta), \nabla_{\mathbf{G}_t^S} \Psi(\mathbf{G}_t^S \rightarrow \mathbf{G}_t^T)\right), \quad (5)$$

where $\mathcal{F}(\cdot, \cdot)$ denotes a generic supervised loss (i.e., mean squared error), with the first argument as the trainable predictor and the second argument as its supervision signal. Crucially, $\Psi(\mathbf{G}_t^S \rightarrow \mathbf{G}_t^T)$ acts as a guidance potential, whose gradient $\nabla_{\mathbf{G}_t^S} \Psi$ explicitly encourages alignment with the target domain. Building on the insights from Theorem 1, we can deduce the specific $\Psi(\mathbf{G}_t^S \rightarrow \mathbf{G}_t^T)$:

Theorem 1. *Let \mathbf{G}^S and \mathbf{G}^T denote the source and target graphs, respectively. Suppose the data distributions p and q define the forward diffusion processes on these two domains, respectively. Following (Ouyang et al., 2024), the optimal diffusion network $\mathbb{P}(\ell^*)$ for the target graph \mathbf{G}^T satisfies:*

$$\mathbb{P}(\ell^*) = \nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S) + \nabla_{\mathbf{G}_t^S} \log \mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} q(\mathbf{G}_0^T) / p(\mathbf{G}_0^S), \quad (6)$$

where the first term is the score function of the source graph, and the second term represents the gradient of the density ratio between the target and source graph distributions. The density ratio is defined as the likelihood ratio between the target and source domains in the latent space.

Proof: For a detailed derivation, please refer to the Appendix A.1.

Consequently, the optimization objective in Eq. (5) can be rewritten as:

$$\mathcal{L}_{\text{SDE}} = \mathcal{F}\left(\mathbb{P}(\ell), \nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S)\right) + \mathcal{F}\left(\mathbb{Q}(\delta), \nabla_{\mathbf{G}_t^S} \log \mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} q(\mathbf{G}_0^T) / p(\mathbf{G}_0^S)\right). \quad (7)$$

The theorem demonstrates that our guidance network adaptively learns the optimal transformation path, directing the generative trajectory toward the target domain. This adaptive guidance enables the model to overcome the limitations of traditional discrete alignment methods by progressively steering the evolution toward the target domain, resulting in more coherent graph adaptation.

4.2 PRACTICAL IMPLEMENTATION

Optimizing Score Network $\mathbb{P}(\ell)$. To fit the score network, we first need to estimate the score function $\nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S) \in \mathbb{R}^{N_S \times (F+C)} \times \mathbb{R}^{N_S \times N_S}$, where N_S represents the node number in source graph, F denotes the feature dimensionality, and C corresponds to the class number. However, directly estimating this score function is computationally prohibitive due to the high dimensionality of the problem. Thus, we adopt the reverse-time diffusion process (Jo et al., 2022), which is mathematically equivalent to the formulation in Eq. (4) and is modeled by the following SDEs:

$$\begin{aligned} d\tilde{\mathbf{X}}_t^S &= [\mathbf{f}_{1,t}(\tilde{\mathbf{X}}_t^S) - g_{1,t}^2 \nabla_{\tilde{\mathbf{X}}_t^S} \log p_t(\tilde{\mathbf{X}}_t^S, \mathbf{A}_t^S)] d\bar{t} + g_{1,t} d\bar{\mathbf{w}}_1, \\ d\mathbf{A}_t^S &= [\mathbf{f}_{2,t}(\mathbf{A}_t^S) - g_{2,t}^2 \nabla_{\mathbf{A}_t^S} \log p_t(\tilde{\mathbf{X}}_t^S, \mathbf{A}_t^S)] d\bar{t} + g_{2,t} d\bar{\mathbf{w}}_2, \end{aligned} \quad (8)$$

where $\mathbf{f}_{1,t}$ and $\mathbf{f}_{2,t}$ are linear drift coefficients satisfying $\mathbf{f}_t(\tilde{\mathbf{X}}, \mathbf{A}) = (\mathbf{f}_{1,t}(\tilde{\mathbf{X}}), \mathbf{f}_{2,t}(\mathbf{A}))$, $g_{1,t}$ and $g_{2,t}$ are scalar diffusion coefficients, and $\bar{\mathbf{w}}_1, \bar{\mathbf{w}}_2$ represent reverse-time standard Wiener processes. This formulation effectively decomposes the original high-dimensional score function into two partial score functions: $\nabla_{\tilde{\mathbf{X}}_t^S} \log p_t(\tilde{\mathbf{X}}_t^S, \mathbf{A}_t^S)$ and $\nabla_{\mathbf{A}_t^S} \log p_t(\tilde{\mathbf{X}}_t^S, \mathbf{A}_t^S)$. By splitting the computation, the problem becomes more tractable, as it reduces the complexity of estimating the high-dimensional score function into manageable components, which can be estimated using the Variance Exploding (VE) (Yang et al., 2023) technique to compute these two gradients. Consequently, the score network $\mathbb{P}(\ell)$ is also divided into two parts: $\mathbb{P}(\ell) \triangleq (\mathbb{P}(\ell_1), \mathbb{P}(\ell_2))$, where $\mathbb{P}(\ell_1)$ is responsible for computing the score function associated with the node features, and $\mathbb{P}(\ell_2)$ is tasked with computing the score function for the graph adjacency structure, “ \triangleq ” denotes a decomposition. To model these two score functions, we utilize two independent multilayer perceptron (MLP) networks:

$$\mathbb{P}(\ell_1) = \text{MLP}\left(\left[\{\mathbf{H}_i\}_{i=0}^L\right]\right), \quad \mathbb{P}(\ell_2) = \text{MLP}\left(\left[\{\text{GMH}(\mathbf{H}_i, \mathbf{A}_t^\kappa)\}_{i=0, \kappa=1}^{L, K}\right]\right), \quad (9)$$

where \mathbf{A}_t^κ denotes the higher-order adjacency matrices, and $\mathbf{H}_{i+1} = \text{GNN}(\mathbf{H}_i, \mathbf{A}_t)$ with $\mathbf{H}_0 = \tilde{\mathbf{X}}_0$ as the input. The GMH refers to the graph multi-head attention block, K is the number of GMH

layers, and L is the number of GNN layers. To approximate these score functions using score network $\mathbb{P}(\ell)$, we define the following training objectives for nodes and edges, respectively:

$$\begin{aligned}\ell_1^* &= \arg \min_{\ell_1} \mathbb{E}_t \left\{ \lambda_{1,t} \mathbb{E}_{\mathbf{G}_0^S} \mathbb{E}_{\mathbf{G}_t^S | \mathbf{G}_0^S} \left\| \mathbb{P}(\ell_1) - \nabla_{\mathbf{X}_t^S} \log p_t(\tilde{\mathbf{X}}_t^S | \tilde{\mathbf{X}}_0^S) \right\|_2^2 \right\}, \\ \ell_2^* &= \arg \min_{\ell_2} \mathbb{E}_t \left\{ \lambda_{2,t} \mathbb{E}_{\mathbf{G}_0^S} \mathbb{E}_{\mathbf{G}_t^S | \mathbf{G}_0^S} \left\| \mathbb{P}(\ell_2) - \nabla_{\mathbf{A}_t^S} \log p_t(\mathbf{A}_t | \mathbf{A}_0^S) \right\|_2^2 \right\},\end{aligned}\quad (10)$$

where $\lambda_{1,t}$ and $\lambda_{2,t}$ are weighting functions and t is uniformly sampled from $[0, T]$, the \mathbb{E} can be efficiently computed using the Monte Carlo estimate (McCool, 1999) with the samples $(t, \mathbf{G}_0, \mathbf{G}_t)$.

Optimizing Guidance Network $\mathbb{Q}(\delta)$. We begin by formulating the gradient estimation objective $\nabla_{\mathbf{G}_t^S} \log \mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} q(\mathbf{G}_0^T) / p(\mathbf{G}_0^S)$. This objective can be decomposed into two key components: (1) the density ratio $q(\mathbf{G}_0^T) / p(\mathbf{G}_0^S)$ between target and source domains, and (2) the expectation operator over the posterior distribution $\mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)}$. To address the intractability of direct computation, we estimate these two items respectively through the following methods inspired by (Ouyang et al., 2024): (1) We approximate the density ratio through domain discrimination. A graph neural classifier $\mathcal{C}_{\text{gnn}} : \mathbf{G}^S \cup \mathbf{G}^T \rightarrow \mathbf{y}$ is trained to distinguish between source and target domain nodes, using a subset of nodes from each domain. Let $\mathbf{y}(\mathbf{x}) \in [0, 1]$ denote the classifier's output indicating source membership. The density ratio can then be estimated as: $q(\mathbf{G}_0^T) / p(\mathbf{G}_0^S) \approx (1 - \mathbf{y}(\mathbf{x})) / \mathbf{y}(\mathbf{x})$. (2) We compute the term $\mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} q(\mathbf{G}_0^T) / p(\mathbf{G}_0^S)$ using Monte Carlo sampling. This process involves generating multiple perturbed graph instances from the joint distribution $p(\mathbf{G}_0^S, \mathbf{G}_t^S)$ and then averaging the computed gradients across these samples (Lu et al., 2023). We approximate the true gradient by sampling from this distribution. The sampling details are provided in Appendix B.

Similarly to the score network $\mathbb{P}(\ell)$, we approximate the guidance function using a parameterized network $\mathbb{Q}(\delta)$, where $\mathbb{Q}(\delta) \triangleq (\mathbb{Q}(\delta_1), \mathbb{Q}(\delta_2))$. Here, $\mathbb{Q}(\delta_1)$ acts as the domain estimator for node-level features, while $\mathbb{Q}(\delta_2)$ serves as the domain estimator for the adjacency matrix. The two components are formally expressed as follows:

$$\mathbb{Q}(\delta_1) = \text{MLP}(\tilde{\mathbf{X}}_t^S, t), \quad \mathbb{Q}(\delta_2) = \text{MLP}(\mathbf{A}_t^S, t). \quad (11)$$

Thus, the final learning objective for the guidance network $\mathbb{Q}(\delta)$ can be described as follows:

$$\begin{aligned}\delta_1^* &= \arg \min_{\delta_1} \mathbb{E}_{p(\mathbf{G}_0^S, \mathbf{G}_t^S)} \left\| \mathbb{Q}(\delta_1) - q(\mathbf{X}_0^T) / p(\mathbf{X}_0^S) \right\|_2^2, \\ \delta_2^* &= \arg \min_{\delta_2} \mathbb{E}_{p(\mathbf{G}_0^S, \mathbf{G}_t^S)} \left\| \mathbb{Q}(\delta_2) - q(\mathbf{A}_0^T) / p(\mathbf{A}_0^S) \right\|_2^2.\end{aligned}\quad (12)$$

In this manner, we can achieve an accurate estimation of the guiding function in Eq. (7) as:

$$\mathbb{Q}(\delta^*) = (\mathbb{Q}(\delta_1^*), \mathbb{Q}(\delta_2^*)) = \log \mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} q(\mathbf{G}_0^T) / p(\mathbf{G}_0^S). \quad (13)$$

The detailed proof can be found in Appendix A.2.

Applying diffusion indiscriminately to all nodes is computationally costly and may weaken informative signals. To mitigate this, we adopt a **targeted stochastic diffusion** strategy that selectively perturbs a subset of nodes while keeping the rest intact, i.e., $p(\mathbf{G}_0^T) \sim p(\mathbf{G}^T)$. A hyperparameter α controls the diffusion ratio, balancing efficiency and information preservation.

Training on Generated Graph \mathbf{G}' . After training the diffusion model, we obtain $\mathbf{G}' = (\mathbf{X}', \mathbf{A}', \mathbf{Y}')$ with node features, adjacency matrix, and labels. Since \mathbf{G}' may still differ from the target graph, we adopt an additional MMD alignment (Huang et al., 2024). The training loss is:

$$\mathcal{L}_{\text{GNN}} = \mathcal{L}_{\text{CE}}(\text{GNN}(\mathbf{X}', \mathbf{A}'), \mathbf{Y}') + \eta \mathcal{L}_{\text{MMD}}(\text{GNN}(\mathbf{X}', \mathbf{A}'), \text{GNN}(\mathbf{X}^T, \mathbf{A}^T)), \quad (14)$$

where \mathcal{L}_{CE} is the cross-entropy loss, \mathcal{L}_{MMD} aligns \mathbf{G}' with \mathbf{G}^T , and η is a weighting factor.

The design of \mathbf{G}' generation offers two main advantages in GDA. First, it allows the diffusion model to learn class-aware dynamics directly in the joint feature-label space, rather than aligning feature-level distributions alone, which can be ambiguous without labels. Second, by reconstructing labels with features, the model can better capture the relationship between node attributes, structure, and class membership, leading to improved node classification performance on the target domain.

Methods	ACMv9 (A), Citationv1 (C), DBLPv7 (D)												Avg.
	A → C		A → D		C → A		C → D		D → A		D → C		
	Mi-F1	Ma-F1	Mi-F1	Ma-F1	Mi-F1	Ma-F1	Mi-F1	Ma-F1	Mi-F1	Ma-F1	Mi-F1	Ma-F1	
GAT (ICLR'18)	62.77 _{±2.24}	59.60 _{±2.82}	60.29 _{±1.33}	55.09 _{±1.16}	58.35 _{±1.51}	58.06 _{±1.88}	67.07 _{±2.05}	63.82 _{±1.85}	54.33 _{±2.17}	52.99 _{±2.36}	63.24 _{±2.21}	59.96 _{±2.86}	59.63
GIN (ICLR'19)	69.95 _{±0.90}	62.34 _{±1.08}	64.69 _{±1.43}	53.08 _{±2.06}	62.63 _{±0.23}	61.34 _{±0.62}	68.54 _{±0.83}	64.83 _{±0.44}	58.18 _{±0.60}	51.36 _{±2.36}	69.91 _{±1.83}	63.09 _{±1.92}	62.50
GCN (ICLR'17)	70.82 _{±1.26}	66.49 _{±2.21}	65.05 _{±2.15}	59.53 _{±0.44}	65.44 _{±1.14}	65.06 _{±1.38}	69.46 _{±0.83}	65.80 _{±1.32}	59.92 _{±0.72}	58.95 _{±0.57}	66.83 _{±0.94}	64.66 _{±1.01}	64.83
DANE (ICAF'19)	69.77 _{±2.14}	66.89 _{±2.90}	62.41 _{±2.15}	59.09 _{±2.60}	63.93 _{±1.32}	64.04 _{±2.01}	65.05 _{±2.07}	60.34 _{±3.01}	58.21 _{±1.17}	57.98 _{±1.47}	67.41 _{±2.29}	62.37 _{±2.47}	63.12
UDAGCN (WWW'20)	80.68 _{±0.31}	78.74 _{±0.55}	74.66 _{±0.93}	72.59 _{±1.05}	73.46 _{±0.40}	74.37 _{±0.35}	76.97 _{±0.31}	75.56 _{±0.42}	69.36 _{±0.31}	70.10 _{±0.60}	77.81 _{±0.50}	76.09 _{±0.78}	75.03
AdaGCN (TKDE'22)	68.07 _{±0.86}	64.15 _{±1.88}	66.72 _{±1.07}	61.97 _{±1.25}	63.25 _{±1.15}	63.22 _{±1.28}	70.97 _{±0.48}	66.43 _{±0.43}	60.68 _{±0.67}	58.99 _{±1.42}	65.83 _{±2.33}	58.27 _{±1.49}	64.05
StruRW (ICML'23)	60.48 _{±1.41}	54.49 _{±1.41}	59.63 _{±0.34}	52.54 _{±0.57}	56.18 _{±0.86}	51.77 _{±2.01}	62.50 _{±1.40}	56.74 _{±2.01}	52.25 _{±1.10}	44.68 _{±1.25}	57.55 _{±0.84}	50.47 _{±1.25}	54.94
GRADE (AAAI'23)	75.02 _{±0.48}	71.66 _{±0.50}	68.17 _{±0.25}	63.05 _{±0.41}	68.96 _{±0.10}	68.43 _{±0.14}	73.48 _{±0.80}	69.76 _{±0.98}	61.72 _{±0.30}	56.45 _{±0.09}	71.69 _{±0.90}	66.54 _{±0.75}	67.91
PairAlign (ICML'24)	60.25 _{±0.89}	55.11 _{±1.02}	59.58 _{±0.58}	53.80 _{±0.85}	56.02 _{±0.87}	51.57 _{±1.36}	63.49 _{±0.70}	59.02 _{±0.82}	51.83 _{±0.69}	46.47 _{±1.13}	58.60 _{±0.42}	53.66 _{±0.47}	55.78
GraphAlign (KDD'24)	75.18 _{±0.71}	71.09 _{±1.21}	68.81 _{±0.78}	65.51 _{±1.21}	65.21 _{±0.55}	61.37 _{±0.28}	72.21 _{±0.45}	71.18 _{±1.20}	61.66 _{±1.23}	62.33 _{±0.45}	68.85 _{±0.56}	65.56 _{±0.21}	67.41
A2GNN (AAAI'23)	80.93 _{±0.52}	78.06 _{±0.73}	75.94 _{±0.33}	71.31 _{±0.34}	75.09 _{±0.43}	76.41 _{±0.54}	77.16 _{±0.23}	73.28 _{±0.32}	73.21 _{±0.44}	74.48 _{±0.68}	79.72 _{±0.63}	76.01 _{±0.79}	75.97
GGDA (Arxiv'25)	79.33 _{±1.00}	76.02 _{±1.93}	73.58 _{±1.73}	68.97 _{±3.80}	73.50 _{±0.60}	73.91 _{±0.92}	76.62 _{±0.59}	72.94 _{±0.64}	70.94 _{±0.11}	72.27 _{±0.32}	77.95 _{±0.54}	75.24 _{±1.52}	74.27
TDSS (AAAI'25)	80.41 _{±0.73}	79.10 _{±1.64}	74.04 _{±3.64}	71.59 _{±0.78}	73.50 _{±0.60}	72.88 _{±0.83}	76.09 _{±1.06}	77.23 _{±2.71}	74.96 _{±0.75}	72.38 _{±1.09}	73.23 _{±3.09}	79.04 _{±3.61}	77.70 _{±2.68}
DGSDA (ICML'25)	78.59 _{±0.55}	77.03 _{±0.52}	73.71 _{±0.24}	72.63 _{±0.25}	73.96 _{±0.70}	71.61 _{±0.80}	76.56 _{±0.29}	75.52 _{±0.34}	72.69 _{±0.34}	71.82 _{±0.45}	77.41 _{±0.11}	76.13 _{±0.23}	74.75
GAA (ICLR'25)	80.03 _{±0.37}	75.85 _{±0.65}	73.32 _{±0.67}	68.07 _{±0.09}	73.15 _{±0.32}	73.57 _{±0.55}	76.04 _{±0.39}	71.40 _{±0.16}	68.32 _{±0.43}	62.66 _{±0.93}	78.27 _{±0.31}	72.74 _{±0.76}	72.65
DiffGDA (Ours)	82.28_{±0.46}	81.07_{±0.37}	76.70_{±0.92}	73.20_{±1.68}	75.75_{±0.30}	77.04_{±0.47}	78.11_{±0.40}	77.09_{±1.09}	74.55_{±0.20}	75.96_{±0.06}	80.71_{±0.20}	78.53_{±0.58}	77.58
Methods	USA (U), Brazil (B), Europe (E)												
	U → B	U → E	B → U	B → E	E → U	E → B	E → B	E → B	E → B	E → B	E → B	E → B	Avg.
	Mi-F1	Ma-F1	Mi-F1	Ma-F1	Mi-F1	Ma-F1	Mi-F1	Ma-F1	Mi-F1	Ma-F1	Mi-F1	Ma-F1	
GAT (ICLR'18)	50.53 _{±3.25}	48.01 _{±4.03}	40.25 _{±3.82}	35.57 _{±3.78}	44.25 _{±1.30}	41.25 _{±1.57}	46.67 _{±1.99}	46.18 _{±1.23}	43.90 _{±1.69}	42.13 _{±1.72}	50.84 _{±3.22}	49.02 _{±4.50}	44.88
GIN (ICLR'19)	33.13 _{±2.96}	23.29 _{±4.55}	32.28 _{±4.55}	22.50 _{±4.12}	35.87 _{±1.96}	28.06 _{±1.93}	32.98 _{±2.54}	26.03 _{±4.30}	36.07 _{±2.44}	28.03 _{±2.24}	33.89 _{±4.03}	26.21 _{±4.23}	29.86
GCN (ICLR'17)	56.34 _{±3.75}	55.12 _{±5.22}	44.71 _{±0.52}	42.63 _{±1.04}	40.22 _{±2.11}	31.21 _{±2.51}	48.92 _{±1.88}	46.52 _{±1.77}	44.67 _{±0.98}	37.39 _{±4.21}	58.47 _{±2.39}	57.27 _{±3.13}	46.96
DANE (ICAF'19)	44.12 _{±1.56}	41.23 _{±1.63}	39.90 _{±3.25}	34.24 _{±4.76}	42.20 _{±4.43}	35.73 _{±2.85}	38.45 _{±3.29}	33.57 _{±3.90}	36.13 _{±2.75}	28.26 _{±2.44}	46.56 _{±0.48}	39.94 _{±0.80}	38.36
UDAGCN (WWW'20)	58.17 _{±2.07}	56.82 _{±2.59}	44.51 _{±0.41}	41.85 _{±0.86}	42.17 _{±1.14}	34.69 _{±1.26}	47.97 _{±1.70}	47.60 _{±1.75}	42.79 _{±1.47}	41.83 _{±1.07}	62.29 _{±3.79}	61.97 _{±1.15}	48.55
AdaGCN (TKDE'22)	65.65 _{±1.93}	66.15 _{±1.38}	50.63 _{±0.55}	51.45 _{±0.93}	46.87 _{±1.00}	43.89 _{±1.07}	54.44 _{±0.93}	54.92 _{±0.91}	48.62 _{±0.53}	44.37 _{±0.64}	73.74 _{±2.24}	73.34 _{±2.46}	56.17
StruRW (ICML'23)	62.44 _{±1.96}	62.25 _{±2.61}	39.70 _{±1.71}	37.26 _{±2.15}	41.45 _{±1.30}	40.02 _{±1.56}	38.30 _{±1.52}	36.84 _{±1.54}	41.21 _{±1.50}	38.10 _{±0.78}	55.11 _{±0.75}	52.15 _{±0.57}	45.40
GRADE (AAAI'23)	57.71 _{±0.78}	56.16 _{±0.96}	48.97 _{±0.20}	46.86 _{±0.22}	43.31 _{±0.58}	38.68 _{±0.41}	55.94 _{±0.19}	53.93 _{±0.33}	46.64 _{±0.80}	42.47 _{±0.65}	63.05 _{±0.37}	60.13 _{±0.36}	51.15
PairAlign (ICML'24)	66.26 _{±1.77}	65.35 _{±2.65}	39.50 _{±0.95}	35.61 _{±0.45}	41.92 _{±2.30}	40.89 _{±2.33}	38.10 _{±2.66}	37.57 _{±3.00}	39.63 _{±0.90}	37.71 _{±0.48}	55.57 _{±0.75}	53.26 _{±1.27}	45.95
GraphAlign (KDD'24)	62.54 _{±0.80}	61.33 _{±1.21}	52.18 _{±0.17}	50.09 _{±0.23}	50.33 _{±0.77}	47.11 _{±1.17}	55.23 _{±2.33}	53.19 _{±0.75}	54.35 _{±0.30}	49.31 _{±0.37}	71.02 _{±0.48}	69.97 _{±1.22}	56.39
A2GNN (AAAI'24)	64.24 _{±1.37}	61.08 _{±2.57}	46.62 _{±1.61}	43.51 _{±3.94}	47.66 _{±3.51}	38.60 _{±1.73}	47.67 _{±1.85}	42.80 _{±3.35}	52.72 _{±1.16}	48.39 _{±1.07}	54.81 _{±1.63}	46.76 _{±2.47}	49.57
GGDA (Arxiv'25)	62.44 _{±2.36}	60.05 _{±1.19}	48.17 _{±1.72}	43.48 _{±3.08}	47.03 _{±1.29}	38.80 _{±0.74}	48.67 _{±0.99}	38.88 _{±2.28}	49.94 _{±1.52}	47.09 _{±2.20}	65.80 _{±1.70}	62.46 _{±3.17}	51.07
TDSS (AAAI'25)	67.43 _{±1.62}	63.36 _{±1.33}	52.05 _{±1.01}	48.17 _{±1.50}	47.62 _{±0.77}	36.46 _{±2.63}	51.80 _{±0.12}	40.38 _{±1.15}	46.08 _{±0.51}	36.55 _{±2.14}	55.73 _{±1.08}	45.27 _{±1.12}	49.24
DGSDA (ICML'25)	61.22 _{±1.63}	60.28 _{±1.55}	49.35 _{±2.76}	48.20 _{±1.46}	54.91 _{±0.51}	53.70 _{±0.78}	42.56 _{±1.64}	42.58 _{±1.73}	49.76 _{±0.82}	47.59 _{±0.76}	60.15 _{±0.57}	59.30 _{±0.63}	52.28
GAA (ICLR'25)	64.89 _{±2.36}	59.28 _{±2.67}	51.70 _{±1.97}	50.16 _{±1.01}	52.50 _{±1.67}	50.18 _{±1.03}	48.22 _{±1.07}	47.76 _{±0.64}	48.70 _{±2.34}	47.17 _{±1.30}	56.07 _{±1.51}	55.34 _{±2.73}	51.98
DiffGDA (Ours)	71.76_{±1.25}	72.21_{±0.98}	54.18_{±0.58}	52.58_{±1.00}	54.37_{±0.88}	57.31_{±0.85}	57.56_{±0.12}	56.20_{±0.00}	52.93_{±0.34}	74.81_{±1.63}	74.84_{±0.85}	60.75	

Table 1: Node classification performance (%) $\pm \sigma$ on the Citation and Airport domains. The highest scores are highlighted in **bold**, while the second-highest scores are underlined.

Model Optimization. Overall, the objective consists of optimizing both GNN and diffusion model in an end-to-end manner. Specifically, the optimization problem can be formulated as:

$$\min_{\ell_1, \ell_2, \delta_1, \delta_2, \theta} \mathcal{L}_{\text{SDE}} + \mathcal{L}_{\text{GNN}}, \quad (15)$$

where $\ell_1, \ell_2, \delta_1, \delta_2$ represent the diffusion parameters, and θ denotes the parameters of the GNN.

Complexity Analysis. The overall computational complexity of DiffGDA is governed by two main components: the diffusion model and the graph neural network (GNN). For the diffusion model, which is implemented as a multilayer perceptron (MLP), the complexity is $\mathcal{O}(\mathcal{T} \cdot n^2)$, where \mathcal{T} denotes the number of diffusion steps and n is diffusion node number. Importantly, the targeted diffusion strategy controls $n = \alpha |\mathcal{V}^S|$, making α the key knob to balance efficiency. The GNN introduces a complexity of $\mathcal{O}(L \cdot (|\mathcal{V}^S| + |\mathcal{E}^S|))$, where L is the number of layers, $|\mathcal{V}^S|$ is the number of source nodes, and $|\mathcal{E}^S|$ is the number of edges in source graph. Thus, the total complexity is $\mathcal{O}(\mathcal{T} \cdot n^2 + L \cdot (|\mathcal{V}^S| + |\mathcal{E}^S|))$. The **error bound** analysis is provided in Appendix C.

Remark. (Discussion on Computational Cost) Firstly, despite being diffusion-based, our framework introduces a stochastic diffusion strategy that processes subsets of nodes while preserving critical source information. With an adjustable sampling ratio, it balances computational cost and performance, achieving faster convergence and higher accuracy compared to other methods. Secondly, the model complexity is $\mathcal{O}(\mathcal{T} \cdot n^2 + L \cdot (|\mathcal{V}^S| + |\mathcal{E}^S|))$, where \mathcal{T} is the diffusion steps, n the sampled nodes, and L the GNN layers. In practice, $\mathcal{T} \leq 100$ and n is controlled via stochastic sampling, keeping n^2 manageable. The shallow GNNs further enhance efficiency. Thus, DiffGDA achieves competitive performance with modest computational resources, making it practical and efficient.

Domain	# Datasets	# Nodes	# Edges	# Feat Dims	# Labels
Citation	ACMv9 (A)	9,360	31,112	6,775	5
	Citationv1 (C)	8,935	30,196		
	DBLPv7 (D)	5,484	16,234		
Airport	USA (U)	1,190	27,198	241	4
	Brazil (B)	131	2,148		
	Europe (E)	399	11,990		
Social	Blog1 (B1)	2,300	66,942	8,189	6
	Blog2 (B2)	2,896	107,672		

Table 2: Statistics of the eight real-world datasets.

5 EXPERIMENTS

In this section, we empirically evaluate the DiffGDA framework through experiments addressing five key research questions: **RQ1**: How does DiffGDA compare to state-of-the-art methods? **RQ2**: Which part of the DiffGDA model majorly contributed to the effective prediction? **RQ3**: How do key parameters impact model performance? **RQ4**: What are the computational efficiency and resource overhead of DiffGDA? **RQ5**: How can we intuitively understand DiffGDA’s strengths?

5.1 EXPERIMENTAL SETUP

(1) Datasets. We employ three domains, as detailed in Table 2, which include a total of eight commonly used datasets: Citation (ACMv9, Citationv1, DBLPv7) (Dai et al., 2022), Airport (USA, Brazil, Europe) (Ribeiro et al., 2017), and Social (Blog1, Blog2) (Liu et al., 2024a). **(2) Baselines.** To validate the effectiveness of our model, we compare against two categories of baseline methods: (1) source-only graph neural networks including GAT, GIN, and GCN, which are trained exclusively on the source domain and directly transferred to the target domain; and (2) graph domain adaptation methods: DANE, UDAGCN, AdaGCN, StruRW, GRADE, PairAlign, GraphAlign, A2GNN, GGDA, TDSS, DGSDA and GAA. The complete experimental setup is provided in Appendix D.

5.2 MAIN EXPERIMENTAL RESULTS (RQ1)

We evaluate the performance of DiffGDA framework on node classification tasks across three domains, as shown in Table 1 (with additional Social results in Appendix E.1). The results consistently demonstrate that DiffGDA outperforms state-of-the-art baselines across a variety of transfer scenarios. On the first three datasets, DiffGDA achieves the highest average score, showcasing its superiority in modeling complex cross-domain transitions via continuous-time diffusion processes. Unlike traditional discrete intermediate approaches (i.e., GGDA, GraphAlign), DiffGDA effectively captures nonlinear and multi-scale structural discrepancies. Likewise, in the USA, Brazil, and Europe transfer settings, DiffGDA maintains strong generalization capability, significantly outperforming competitive methods such as GraphAlign and TDSS. These findings validate the efficacy of our generative continuous evolution perspective, which reduces information distortion and enhances structural alignment in GDA. The significance tests ($p < 0.05$) are provided in Appendix E.2.

5.3 ABLATION STUDIES (RQ2)

To evaluate the contribution of each module in DiffGDA framework, we conduct ablation studies on four cross-domain tasks, as shown in Figure 2. Specifically, we investigate the impact of three key components: (i) the domain-aware guidance network, which dynamically steers the diffusion trajectory; (ii) MMD alignment, which matches marginal distributions between domains; and (iii) adjacency constraints, which preserve local structural information during transfer. The results reveal that the full model achieves the best performance across all tasks,

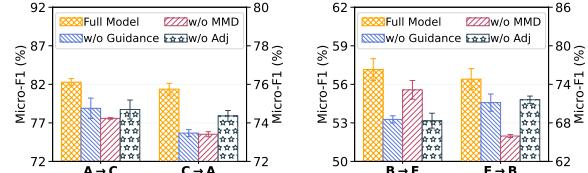


Figure 2: Classification Mi-F1 comparisons between DiffGDA variants on four cross-domain tasks.

confirming the effectiveness of integrating all three components. Removing the guidance network leads to substantial drops in performance, especially on the more challenging tasks (e.g., $\mathbf{B} \rightarrow \mathbf{E}$), highlighting its role in maintaining task-relevant diffusion dynamics. MMD alignment has a more pronounced effect on simpler shifts such as in $\mathbf{A} \rightarrow \mathbf{C}$ transfer task, where distribution-level matching is crucial. The removal of adjacency constraints consistently degrades performance, indicating their utility in preserving structural dependencies. Overall, each module plays a complementary role: the guidance network stabilizes the diffusion path, MMD facilitates cross-domain alignment, and adjacency constraints reinforce relational structure.

Additional ablation results are provided in Appendix E.3.

5.4 HYPER-PARAMETER ANALYSIS (RQ3)

In this section, we analyze the impact of three key parameters on the proposed DiffGDA framework, as illustrated in Figure 3: the sample ratio α , the loss weight η , and the diffusion step T .

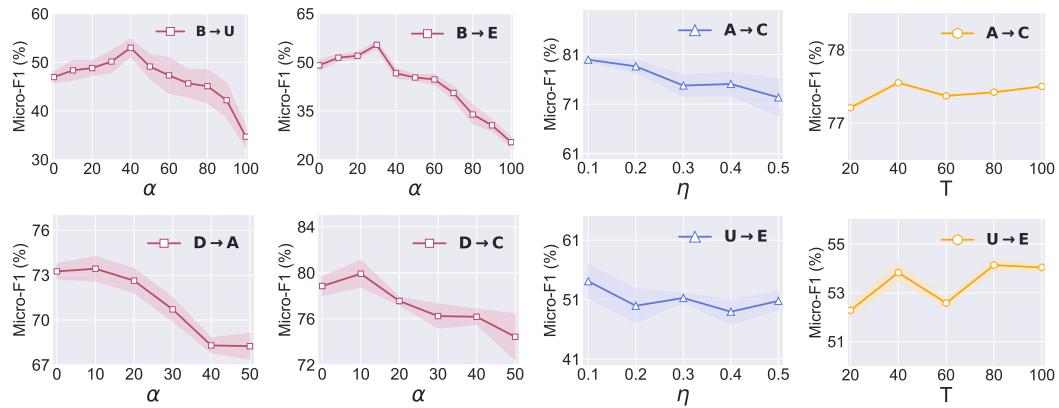


Figure 3: The performances of our DiffGDA w.r.t varying α , η and T on different transfer tasks.

Study on Sample Ratio α . The results demonstrate the impact of varying subgraph sampling sizes on the Mi-F1 performance across different transfer tasks. For the first two tasks ($\mathbf{B} \rightarrow \mathbf{U}$ and $\mathbf{B} \rightarrow \mathbf{E}$), the model was evaluated across the full range of sampling sizes ($0 \sim 100\%$), revealing a trend where performance initially improves with larger subgraphs, as critical structural information is preserved, but may plateau or decline beyond a certain point due to computational overhead or over-smoothing. In contrast, for the latter tasks, the sampling range was truncated to $0 \sim 50\%$ due to out-of-memory issues, indicating that these tasks involve larger or more complex graphs where aggressive subgraph sampling is necessary to maintain feasibility. This aligns with the model’s design goal of leveraging diffusion-based generation while managing scalability. Further investigation is needed to determine the relationship between graph properties and the ideal sampling size.

Study on Loss Weight η . We analyze the impact of coefficient η in Eq. (14). As shown in Figure 3, both tasks exhibit a consistent decline in Mi-F1 as η increases. This suggests that stronger alignment (i.e., larger η) may induce excessive regularization, hindering the model’s ability to preserve task-specific discriminative features. Additional results are provided in Appendix E.4.

Study on Diffusion Step T . As shown in Figure 3, the Mi-F1 score for $\mathbf{A} \rightarrow \mathbf{C}$ improves with increasing T and stabilizes beyond $T = 40$, suggesting that a sufficient number of diffusion steps is necessary to effectively capture structural dependencies. A similar trend is observed in $\mathbf{U} \rightarrow \mathbf{E}$, where performance peaks at $T = 80$ and fluctuates slightly thereafter. These results indicate that while adequate diffusion depth is important, excessively large T yields diminishing returns. Additional results across datasets are provided in Appendix E.4.

5.5 MODEL EFFICIENCY ANALYSIS (RQ4)

Table 3 compares the runtime performance of four approaches: UDAGCN, DANE, GraphAlign and DiffGDA on two tasks, $\mathbf{A} \rightarrow \mathbf{D}$ and $\mathbf{A} \rightarrow \mathbf{C}$, with all models trained for 150 epochs. Addition-

Model	A → D				A → C			
	Generating	Training	Total	Mi-F1	Generating	Training	Total	Mi-F1
UDAGCN	-	83.19	83.19	67.52	-	104.91	104.91	72.64
DANE	-	134.92	134.92	62.41	-	167.00	167.00	69.77
GraphAlign	250.13	19.52	269.65	70.14	274.86	22.29	297.15	76.62
DiffGDA	103.98	22.75	126.73	73.41	102.31	23.10	125.41	80.23

Table 3: Runtime (seconds) and Mi-F1 (%) comparison on two transfer tasks. All models are trained using a single RTX 4090 GPU.

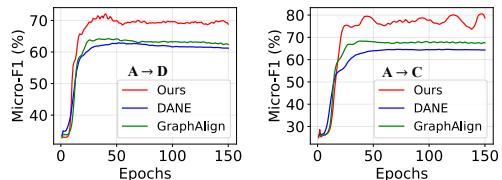


Figure 5: The training curves of Mi-F1 scores on three baselines across two cross-domain transfer scenarios over 150 epochs.

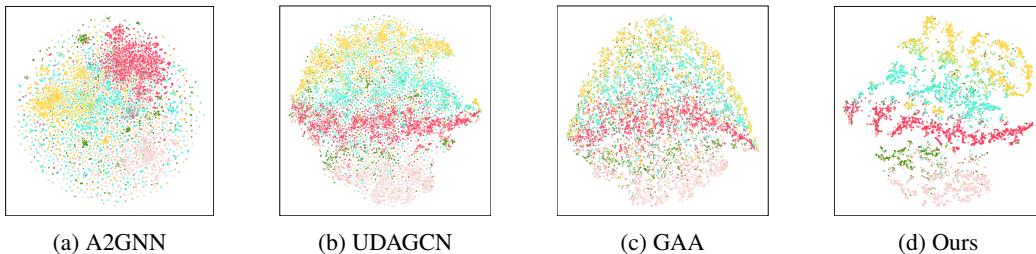


Figure 6: Visualization of node embeddings inferred by four models for the $\mathbf{A} \rightarrow \mathbf{C}$ task.

ally, the training progress is visualized in Figure 5. While our method does not have the lowest absolute runtime, it achieves a relatively lower total runtime compared to GraphAlign, a state-of-the-art approach for GDA based on graph generation, which exhibits significant time overhead due to intermediate graph generation. As shown in Table 3, our method effectively strikes a balance between graph generation and training time, reducing the total runtime by over 50% compared to GraphAlign, with a smaller training time and a manageable graph generation cost. At the same time, it achieves superior Mi-F1 scores on both tasks. Furthermore, Figure 5 highlights that our approach not only converges faster but also consistently outperforms the baselines in terms of Mi-F1 scores across epochs. These results collectively demonstrate that our diffusion-based method offers a compelling trade-off between computational efficiency and performance, making it a suitable choice for cross-domain tasks that demand rigorous experimentation and reliable outcomes.

5.6 VISUALIZATION FOR REPRESENTATION SPACE (RQ5)

To qualitatively assess the learned node representations, we visualize their distributions using t-SNE (Maaten & Hinton, 2008; Yuan et al., 2025), as shown in Figure 6. In the plot, each color represents a distinct class, enabling an intuitive evaluation of the clustering quality. Our method, DiffGDA, produces more compact and well-separated clusters compared to other models, especially the strong data-driven baseline, GAA. This improved clustering suggests a more effective separation of the underlying class distributions within the target domain. These results demonstrate that our diffusion-based adaptation process effectively mitigates domain-irrelevant noise, enhancing inter-class discriminability and leading to a cleaner, more structured representation of the target graph.

6 CONCLUSION

In this work, we revisit the problem of Graph Domain Adaptation (GDA) from a continuous generative perspective. By framing cross-domain adaptation as a stochastic diffusion process, our proposed method, DiffGDA, effectively models the joint evolution of structure and semantics in continuous time. This continuous formulation naturally captures smooth structural transitions and avoids the rigid assumptions of discrete alignment frameworks. The incorporation of a guidance network enables the model to automatically learn the optimal adaptation path. Extensive empirical results across 14 transfer tasks validate the effectiveness and superiority of our approach.

ACKNOWLEDGEMENTS

This work was supported by the National Key Research and Development Program of China under Grant Nos. 2024YFF0729003, the National Natural Science Foundation of China under Grant Nos. 62176014, 62276015, 62206266, the Fundamental Research Funds for the Central Universities.

REFERENCES

Hanqun Cao, Cheng Tan, Zhangyang Gao, Yilun Xu, Guangyong Chen, Pheng-Ann Heng, and Stan Z Li. A survey on generative diffusion models. *TKDE*, 2024.

Wei Chen, Yiqing Wu, Zhao Zhang, Fuzhen Zhuang, Zhongshi He, Ruobing Xie, and Feng Xia. Fairgap: Fairness-aware recommendation via generating counterfactual graph. *TOIS*, 2024.

Wei Chen, Guo Ye, Yakun Wang, Zhao Zhang, Libang Zhang, Daixin Wang, Zhiqiang Zhang, and Fuzhen Zhuang. Smoothness really matters: A simple yet effective approach for unsupervised graph domain adaptation. In *AAAI*, 2025a.

Wei Chen, Meng Yuan, Zhao Zhang, Ruobing Xie, Fuzhen Zhuang, Deqing Wang, and Rui Liu. Fairdgc: Fairness-aware recommendation with dynamic graph contrastive learning. *TKDE*, 2025b.

Wei Chen, Zhao Zhang, Meng Yuan, Kepeng Xu, and Fuzhen Zhuang. Fckt: Fine-grained cross-task knowledge transfer with semantic contrastive learning for targeted sentiment analysis. In *IJCAI*, 2025c.

Yiming Chen, Shiji Song, Shuang Li, and Cheng Wu. A graph embedding framework for maximum mean discrepancy-based domain adaptation algorithms. *TIP*, 2019.

Florinel-Alin Croitoru, Vlad Hondu, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *TPAMI*, 2023.

Quanyu Dai, Xiao-Ming Wu, Jiaren Xiao, Xiao Shen, and Dan Wang. Graph transfer learning via adversarial domain adaptation with graph convolution. *TKDE*, 2022.

Jun Dan, Weiming Liu, Xie Xie, Hua Yu, Shunjie Dong, and Yanchao Tan. Tfgda: Exploring topology and feature alignment in semi-supervised graph domain adaptation through robust clustering. *NeurIPS*, 2024.

Songgaojun Deng, Huzeifa Rangwala, and Yue Ning. Learning dynamic context graphs for predicting social events. In *KDD*, 2019.

Zhengming Ding, Sheng Li, Ming Shao, and Yun Fu. Graph adaptive knowledge transfer for unsupervised domain adaptation. In *ECCV*, 2018.

Ruiyi Fang, Bingheng Li, Zhao Kang, Qiuhan Zeng, Nima Hosseini Dashtbayaz, Ruizhi Pu, Boyu Wang, and Charles Ling. On the benefits of attribute-driven graph domain adaptation. *ICLR*, 2025a.

Ruiyi Fang, Bingheng Li, Jingyu Zhao, Ruizhi Pu, QIUHAO Zeng, Gezheng Xu, Charles Ling, and Boyu Wang. Homophily enhanced graph domain adaptation. In *ICML*, 2025b.

David Freedman. *Brownian motion and diffusion*. Springer Science & Business Media, 2012.

Wei Guo, Yiqi Tong, Yiyang Duan, Fuzhen Zhuang, Xiao Zhang, Zhaojun Hu, and Jin Dong. Feze: Alignment-flexible zero-shot vertical federated learning. In *KDD*, 2025.

Mark Heimann, Haoming Shen, Tara Safavi, and Danai Koutra. Regal: Representation learning-based graph alignment. In *CIKM*, 2018.

Renhong Huang, Jiarong Xu, Xin Jiang, Ruichuan An, and Yang Yang. Can modifying data address graph domain adaptation? In *KDD*, 2024.

Jaehyeong Jo, Seul Lee, and Sung Ju Hwang. Score-based generative modeling of graphs via the system of stochastic differential equations. In *ICML*, 2022.

Liam T Kane, Taolin Fang, Matthew S Galetta, Dhruv KC Goyal, Kristen J Nicholson, Christopher K Kepler, Alexander R Vaccaro, and Gregory D Schroeder. Propensity score matching: a statistical method. *Clinical spine surgery*, 2020.

S Sathiya Keerthi and Chih-Jen Lin. Asymptotic behaviors of support vector machines with gaussian kernel. *Neural Computation*, 2003.

Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *ICLR*, 2017.

Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. *ICLR*, 2021.

Kilian Konstantin Haefeli, Karolis Martinkus, Nathanaël Perraudin, and Roger Wattenhofer. Diffusion models for graphs benefit from discrete state spaces. *arXiv e-prints*, 2022.

Zhiqiang Kou, Jing Wang, Yuheng Jia, Biao Liu, and Xin Geng. Instance-dependent inaccurate label distribution learning. *TNNLS*, 2025a.

Zhiqiang Kou, Yucheng Xie, Hailin Wang, Jing Wang, Mingkun Xie, Shuo Chen, Yuheng Jia, Tongliang Liu, and Xin Geng. Rankmatch: A novel approach to semi-supervised label distribution learning leveraging rank correlation between labels. In *NeurIPS*, 2025b.

Pui Ieng Lei, Ximing Chen, Yijun Sheng, Yanyan Liu, Jingzhi Guo, and Zhiguo Gong. Gradual domain adaptation for graph learning. *arXiv preprint arXiv:2501.17443*, 2025.

Xuechen Li, Ting-Kam Leonard Wong, Ricky TQ Chen, and David Duvenaud. Scalable gradients for stochastic differential equations. In *ICAIIS*, 2020.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *ICLR*, 2023.

Chengyi Liu, Jiahao Zhang, Shijie Wang, Wenqi Fan, and Qing Li. Score-based generative diffusion models for social recommendations. *TKDE*, 2025.

Meihan Liu, Zeyu Fang, Zhen Zhang, Ming Gu, Sheng Zhou, Xin Wang, and Jiajun Bu. Rethinking propagation for unsupervised graph domain adaptation. In *AAAI*, 2024a.

Meng Liu, Hongyang Gao, and Shuiwang Ji. Towards deeper graph neural networks. In *KDD*, 2020.

Shikun Liu, Tianchun Li, Yongbin Feng, Nhan Tran, Han Zhao, Qiang Qiu, and Pan Li. Structural re-weighting improves graph domain adaptation. In *ICML*, 2023a.

Shikun Liu, Deyu Zou, Han Zhao, and Pan Li. Pairwise alignment improves graph domain adaptation. *ICML*, 2024b.

Xingchao Liu, Chengyue Gong, et al. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *ICLR*, 2023b.

Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. Contrastive energy prediction for exact energy-guided diffusion sampling in offline reinforcement learning. In *ICML*, 2023.

Junyu Luo, Yiyang Gu, Xiao Luo, Wei Ju, Zhiping Xiao, Yusheng Zhao, Jingyang Yuan, and Ming Zhang. Gala: Graph diffusion-based alignment with jigsaw for source-free domain adaptation. *TPAMI*, 2024a.

Shengjie Luo, Yixian Xu, Di He, Shuxin Zheng, Tie-Yan Liu, and Liwei Wang. Bridging geometric states via geometric diffusion bridge. In *NeurIPS*, 2024b.

Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 2008.

Ross A Maller, Gernot Müller, and Alex Szimayer. Ornstein–uhlenbeck processes and extensions. *Handbook of financial time series*, 2009.

Michael D McCool. Anisotropic diffusion for monte carlo noise reduction. *TOG*, 1999.

Jaemin Na, Heechul Jung, Hyung Jin Chang, and Wonjun Hwang. Fixbi: Bridging domain spaces for unsupervised domain adaptation. In *CVPR*, 2021.

Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. Permutation invariant graph generation via score-based generative modeling. In *IJCAI*, 2020.

Yidong Ouyang, Liyan Xie, Hongyuan Zha, and Guang Cheng. Transfer learning for diffusion models. In *NeurIPS*, 2024.

Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 2019.

Mehmet Pilancı and Elif Vural. Domain adaptation on graphs by learning aligned graph bases. *TKDE*, 2020.

Philip E Protter. Stochastic differential equations. In *Stochastic integration and differential equations*. 2012.

Yijian Qin, Xin Wang, Ziwei Zhang, Pengtao Xie, and Wenwu Zhu. Graph neural architecture search under distribution shifts. In *ICML*, 2022.

Leonardo FR Ribeiro, Pedro HP Saverese, and Daniel R Figueiredo. struc2vec: Learning node representations from structural identity. In *KDD*, 2017.

Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*, 2022.

Kartik Sharma, Yeon-Chang Lee, Sivagami Nambi, Aditya Salian, Shlok Shah, Sang-Wook Kim, and Srijan Kumar. A survey of graph neural networks for social recommender systems. *ACM Computing Surveys*, 2024.

Félix Therrien, Edward H Sargent, and Oleksandr Voznyy. Using gnn property predictors as molecule generators. *Nature Communications*, 2025.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *ICLR*, 2018.

Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 2011.

Yu Wang, Ronghang Zhu, Pengsheng Ji, and Sheng Li. Open-set graph domain adaptation via separate domain alignment. In *AAAI*, 2024.

Jun Wu, Jingrui He, and Elizabeth Ainsworth. Non-iid transfer learning on graphs. In *AAAI*, 2023.

Lingfei Wu, Peng Cui, Jian Pei, Liang Zhao, and Xiaojie Guo. Graph neural networks: foundation, frontiers and applications. In *KDD*, 2022.

Man Wu, Shirui Pan, Chuan Zhou, Xiaojun Chang, and Xingquan Zhu. Unsupervised domain adaptive graph convolutional networks. In *WWW*, 2020.

Zhiqing Xiao, Haobo Wang, Ying Jin, Lei Feng, Gang Chen, Fei Huang, and Junbo Zhao. Spa: A graph spectral alignment perspective for domain adaptation. *NeurIPS*, 2023.

Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *ICLR*, 2019.

Liang Yang, Xin Chen, Jiaming Zhuo, Di Jin, Chuan Wang, Xiaochun Cao, Zhen Wang, and Yuanfang Guo. Disentangled graph spectral domain adaptation. In *ICML*, 2025.

Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *CSUR*, 2023.

Yuning You, Tianlong Chen, Zhangyang Wang, and Yang Shen. Graph domain adaptation via theory-grounded spectral regularization. In *ICLR*, 2023.

Meng Yuan, Zhao Zhang, Wei Chen, Chu Zhao, Tong Cai, Deqing Wang, Rui Liu, and Fuzhen Zhuang. Hek-cl: Hierarchical enhanced knowledge-aware contrastive learning for recommendation. *TOIS*, 2025.

Qiuhan Zeng, Changjian Shui, Long-Kai Huang, Xi Chen, Charles Ling, and Boyu Wang. Latent trajectory learning for limited timestamps under distribution shift over time. In *ICLR*, 2024.

Fuwei Zhang, Zhao Zhang, Xiang Ao, Fuzhen Zhuang, Yongjun Xu, and Qing He. Along the time: Timeline-traced embedding for temporal knowledge graph completion. In *CIKM*, 2022a.

Fuwei Zhang, Xiaoyu Liu, Xinyu Jia, Yingfei Zhang, Shuai Zhang, Xiang Li, Fuzhen Zhuang, Wei Lin, and Zhao Zhang. Multi-level relevance document identifier learning for generative retrieval. In *ACL*, 2025.

Yao Zhang, Yun Xiong, Dongsheng Li, Caihua Shan, Kan Ren, and Yangyong Zhu. Cope: modeling continuous propagation and evolution on interaction graph. In *CIKM*, 2021.

Yizhou Zhang, Guojie Song, Lun Du, Shuwen Yang, and Yilun Jin. Dane: Domain adaptive network embedding. *IJCAI*, 2019.

Zhiqiang Zhang, Jinjia Zhou, Wenxin Yu, and Ning Jiang. Text-to-image synthesis: Starting composite from the foreground content. *Inf. Sci.*, 2022b.

Tianxiang Zhao, Dongsheng Luo, Xiang Zhang, and Suhang Wang. Faithful and consistent graph neural network explanations with rationale alignment. *TIST*, 2023.

Qi Zhu, Natalia Ponomareva, Jiawei Han, and Bryan Perozzi. Shift-robust gnns: Overcoming the limitations of localized graph training data. *NeurIPS*, 2021.

Appendix

Table of Contents

A Theoretical Proof	15
A.1 Proof of Theorem 1	15
A.2 Proof of Guidance Network \mathbb{Q}	16
B Implementation details	17
C Theoretical Analysis of DiffGDA’s Error Bound	17
C.1 Error Bound Expression	17
C.2 Analysis and Insights	17
C.3 Practical Implications	18
D Experiment Setup	18
D.1 Dataset Description	18
D.2 Baselines	18
D.3 Implementation Details	19
E Experiment Results	20
E.1 Main Experiment Results	20
E.2 Significance test (t-test) of DiffGDA	20
E.3 Ablation Study	20
E.4 Hyper-parameter Analysis	21

A THEORETICAL PROOF

A.1 PROOF OF THEOREM 1

Let \mathbf{G}^S and \mathbf{G}^T denote the source graph and target graph, respectively. Assume that the data distributions p and q define the forward diffusion processes on the source and target domains, respectively. For consistency of modeling, we consider the forward process in the target domain to share the same functional form as that in the source domain, i.e., $p(\mathbf{G}_0^S | \mathbf{G}_t^S) = q(\mathbf{G}_0^T | \mathbf{G}_t^T)$. Under this formulation, the optimal diffusion network $\mathbb{P}(\ell^*)$ for the target graph \mathbf{G}^T satisfies:

$$\mathbb{P}(\ell^*) = \nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S) + \nabla_{\mathbf{G}_t^S} \log \mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)}, \quad (16)$$

where the first term is score function of the source graph, and the second term represents the gradient of the density ratio between the target and source graph distributions. The density ratio is defined as the likelihood ratio between the target and source domains in latent space. It reflects how domain-specific a given representation is, and provides a directional signal to guide diffusion process.

Remark on Consistent Diffusion. In this theorem, we assume that clean graph data are corrupted in the same way in the source and target domains. Although the data distributions in the two domains ($p(\mathbf{G}_0^S)$ and $q(\mathbf{G}_0^T)$) are different, we assume they follow the same physical diffusion model (e.g., the same variance-exploding (VE) SDE). This is a common simplifying assumption that allows us to embed both domains into the same diffusion latent space for operations and comparisons. Based on these shared diffusion dynamics, inter-domain differences are entirely reflected in the differences of their initial-state ($t = 0$) data distributions, and the core task of our guidance network \mathbb{Q} is to estimate and compensate for this discrepancy via the density ratio $q(\mathbf{G}_0^T)/p(\mathbf{G}_0^S)$.

Proof: For clarity in the proof, we denote the left-hand side and right-hand side of Eq. (16) as **LE** and **RE**, respectively. Our goal is to prove the equivalence by transforming **RE** into **LE**. In the first step, we establish a connection between score matching on the target graph and importance-weighted denoising score matching on the source graph. Drawing inspiration from (Vincent, 2011; Ouyang et al., 2024), this relationship can be formalized as:

$$\begin{aligned} \ell^* &= \arg \min_{\ell} \mathbb{E}_t \left\{ \lambda_t \mathbb{E}_{q_t(\mathbf{G}_t^T)} \left[\left\| \mathbb{P}(\ell) - \nabla_{\mathbf{G}_t^T} \log q_t(\mathbf{G}_t^T) \right\|_2^2 \right] \right\} \\ &\implies \arg \min_{\ell} \mathbb{E}_t \left\{ \lambda_t \mathbb{E}_{p(\mathbf{G}_0^S)} \mathbb{E}_{p(\mathbf{G}_t^S | \mathbf{G}_0^S)} \left[\left\| \mathbb{P}(\ell) - \nabla_{\mathbf{G}_0^S} \log p(\mathbf{G}_t^S | \mathbf{G}_0^S) \right\|_2^2 \frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \right] \right\}. \end{aligned} \quad (17)$$

The first term corresponds to score matching on source graph, while the second term incorporates importance weighting, bridging the connection to denoising score matching on the source graph. This adjustment is modulated by $\frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)}$, which accounts for the distributional differences between the source and target graphs, thereby ensuring a consistent and unbiased score estimation.

Leveraging importance-weighted denoising score matching on the source graph, the optimal score can be formulated as:

$$\mathbb{P}(\ell^*) = \frac{\mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \left[\frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \nabla_{\mathbf{G}_t^S} \log p(\mathbf{G}_t^S | \mathbf{G}_0^S) \right]}{\mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \left[\frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \right]} = \mathbf{LE}. \quad (18)$$

The final optimal solution $\mathbb{P}(\ell^*)$ is computed by normalizing the weighted gradient of the source graph's log-probability. The numerator captures the weighted score evaluation, while the denominator ensures proper scaling by accounting for the overall importance weights. This formulation guarantees an unbiased and consistent estimation, as the importance weights $q(\mathbf{G}_0^T)/p(\mathbf{G}_0^S)$ correct for any domain-level discrepancies.

In the second step, without loss of generality, we consider an arbitrary graph \mathbf{G} within the framework of the diffusion process p and proceed with the following derivation:

$$\nabla_{\mathbf{G}_t} \log p(\mathbf{G}_0 | \mathbf{G}_t) = \nabla_{\mathbf{G}_t} \log p(\mathbf{G}_t | \mathbf{G}_0) - \nabla_{\mathbf{G}_t} \log p_t(\mathbf{G}_t). \quad (19)$$

This expression establishes the relationship between the conditional score $\nabla_{\mathbf{G}_t} \log p(\mathbf{G}_0 | \mathbf{G}_t)$ and the components of the forward diffusion process, specifically the reverse transition probability and the marginal distribution at time t .

Finally, based on Eq. (19), we can rewrite Eq. (16) as:

$$\begin{aligned}
\mathbf{RE} &= \nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S) + \nabla_{\mathbf{G}_t^S} \log \mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \left[\frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \right] \\
&= \nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S) + \frac{\mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \left[\frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \nabla_{\mathbf{G}_t^S} \log p(\mathbf{G}_t^S | \mathbf{G}_0^S) \right]}{\mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \left[\frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \right]} \\
&= \nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S) + \frac{\mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \left[\frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \nabla_{\mathbf{G}_t^S} \log p(\mathbf{G}_t^S | \mathbf{G}_0^S) \right]}{\mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \left[\frac{q(\mathbf{G}_0^S)}{p(\mathbf{G}_0^S)} \right]} - \nabla_{\mathbf{G}_t^S} \log p_t(\mathbf{G}_t^S) \\
&= \frac{\mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \left[\frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \nabla_{\mathbf{G}_t^S} \log p(\mathbf{G}_t^S | \mathbf{G}_0^S) \right]}{\mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \left[\frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \right]} = \mathbf{LE}.
\end{aligned} \tag{20}$$

Therefore, we have proven that Theorem 1 holds.

A.2 PROOF OF GUIDANCE NETWORK \mathbb{Q}

In this section, we demonstrate that Eq. (13) holds, namely:

$$\mathbb{Q}(\boldsymbol{\delta}^*) = (\mathbb{Q}(\boldsymbol{\delta}_1^*), \mathbb{Q}(\boldsymbol{\delta}_2^*)) = \log \mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)}, \tag{21}$$

where $\mathbb{Q}(\boldsymbol{\delta})$ denotes a learned representation parameterized by $\boldsymbol{\delta}$, and the logarithm is applied element-wise. To proceed, let us define $p(\mathbf{G}_0^S) = p(\mathbf{X}_0^S, \mathbf{A}_0^S)$ and $q(\mathbf{G}_0^T) = q(\mathbf{X}_0^T, \mathbf{A}_0^T)$ as the source and target joint distributions over node features \mathbf{X} and adjacency matrices \mathbf{A} , respectively. Note that, for simplicity of representation, we use \mathbf{G} as a shorthand for (\mathbf{X}, \mathbf{A}) throughout this proof. Thus, we have:

$$\begin{aligned}
\boldsymbol{\delta}^* &= \arg \min_{\boldsymbol{\delta}} \int_{\mathbf{G}_t^S} \left\| \mathbb{Q}(\boldsymbol{\delta}) - \mathbb{E}_{p(\mathbf{G}_0^S | \mathbf{G}_t^S)} \left[\frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \right] \right\|_2^2 p(\mathbf{G}_t^S) d\mathbf{G}_t^S \\
&= \arg \min_{\boldsymbol{\delta}} \int_{\mathbf{G}_t^S} \left\| \mathbb{Q}(\boldsymbol{\delta}) - \int_{\mathbf{G}_0^S} p(\mathbf{G}_0^S | \mathbf{G}_t^S) \cdot \frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} d\mathbf{G}_0^S \right\|_2^2 p(\mathbf{G}_t^S) d\mathbf{G}_t^S \\
&= \arg \min_{\boldsymbol{\delta}} \int_{\mathbf{G}_t^S} \left\{ \|\mathbb{Q}(\boldsymbol{\delta})\|_2^2 - 2 \left\langle \mathbb{Q}(\boldsymbol{\delta}), \int_{\mathbf{G}_0^S} p(\mathbf{G}_0^S | \mathbf{G}_t^S) \cdot \frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} d\mathbf{G}_0^S \right\rangle \right\} p(\mathbf{G}_t^S) d\mathbf{G}_t^S + C \\
&= \arg \min_{\boldsymbol{\delta}} \iint \left\| \mathbb{Q}(\boldsymbol{\delta}) - \frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \right\|_2^2 \cdot p(\mathbf{G}_0^S | \mathbf{G}_t^S) \cdot p(\mathbf{G}_t^S) d\mathbf{G}_0^S d\mathbf{G}_t^S \\
&= \arg \min_{\boldsymbol{\delta}} \iint \left\| \mathbb{Q}(\boldsymbol{\delta}) - \frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \right\|_2^2 \cdot p(\mathbf{G}_0^S, \mathbf{G}_t^S) d\mathbf{G}_0^S d\mathbf{G}_t^S \\
&= \arg \min_{\boldsymbol{\delta}} \mathbb{E}_{p(\mathbf{G}_0^S, \mathbf{G}_t^S)} \left\| \mathbb{Q}(\boldsymbol{\delta}) - \frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \right\|_2^2,
\end{aligned} \tag{22}$$

where C is a constant independent of $\boldsymbol{\delta}$.

According to Eq. (12):

$$\boldsymbol{\delta}^* = (\boldsymbol{\delta}_1^*, \boldsymbol{\delta}_2^*) = \arg \min_{\boldsymbol{\delta}} \mathbb{E}_{p(\mathbf{G}_0^S, \mathbf{G}_t^S)} \left\| \mathbb{Q}(\boldsymbol{\delta}) - \frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)} \right\|_2^2. \tag{23}$$

Thereby we complete the proof.

B IMPLEMENTATION DETAILS

Algorithm 1 Density Ratio Estimation

Require: Source graph distribution $p(\mathbf{G}_0^S)$, target graph distribution $q(\mathbf{G}_0^T)$, pre-trained domain classifier \mathcal{C}_{gnn} , number of Monte Carlo samples S_{mc}

Ensure: Estimated density ratio: $\frac{r_1}{S_{mc}}$

- 1: Train the domain classifier \mathcal{C}_{gnn} to distinguish \mathbf{G}_0^S and \mathbf{G}_0^T
- 2: Initialize accumulator: $r_1 \leftarrow 0$
- 3: **for** $i = 1$ to S_{mc} **do**
- 4: Sample $\mathbf{G}_0^S(i) \sim p(\mathbf{G}_0^S)$
- 5: Sample $\mathbf{G}_0^T(i) \sim q(\mathbf{G}_0^T)$
- 6: Compute $\mathbf{y}^{(i)} = \mathcal{C}_{gnn}(\mathbf{G}_0^S(i), \mathbf{G}_0^T(i))$
- 7: Calculate ratio: $r^{(i)} = \frac{1 - \mathbf{y}^{(i)}}{\mathbf{y}^{(i)}}$
- 8: Accumulate: $r_1 \leftarrow r_1 + r^{(i)}$
- 9: **end for**
- 10: **return** $\frac{r_1}{S_{mc}}$

C THEORETICAL ANALYSIS OF DIFFGDA'S ERROR BOUND

In this section, we analyze the theoretical error bound of DiffGDA framework, focusing on the factors that influence its performance under GDA settings. By extending the traditional error bound framework and incorporating components specific to diffusion-based methods, we derive an expression that characterizes the generalization error on the target domain as a function of multiple factors, including source domain performance, domain discrepancy, sample complexity, and noise stability.

C.1 ERROR BOUND EXPRESSION

Inspired by (Huang et al., 2024), the generalization error on target domain \mathcal{E}_T can be bounded as:

$$\mathcal{E}_T \leq \mathcal{E}_S + \xi \cdot \mathcal{D}(\mathcal{P}_S, \mathcal{P}_T) + \frac{C}{\sqrt{N}} + \mu \cdot \mathcal{R}_{\text{diff}}, \quad (24)$$

where \mathcal{E}_T is the generalization error on the target domain, \mathcal{E}_S is the empirical error on the source domain, and $\mathcal{D}(\mathcal{P}_S, \mathcal{P}_T)$ represents the distributional discrepancy between the source and target domains, which can be measured using metrics such as MMD or Wasserstein distance. The term ξ is a weight factor that scales the impact of domain discrepancy, influenced by the model's capacity and alignment mechanism. The term $\frac{C}{\sqrt{N}}$ reflects a sample complexity term, where N represents the number of source domain samples, and C denotes model complexity (e.g., Lipschitz constant or parameter count). Finally, $\mathcal{R}_{\text{diff}}$ captures the residual error introduced by the diffusion process, which can be further decomposed into three controllable sources:

$$\mathcal{R}_{\text{diff}} = \mu_1 \mathcal{B}_{\text{sched}}(\sigma_{\min}, \sigma_{\max}) + \mu_2 \mathcal{B}_{\text{disc}}(T) + \mu_3 \varepsilon_{\text{ratio}}, \quad (25)$$

where μ_1, μ_2, μ_3 are non-negative coefficients weighting each term; $\mathcal{B}_{\text{sched}}(\sigma_{\min}, \sigma_{\max})$ denotes the error from the variance-exploding (VE) noise schedule, controlled by the minimum and maximum noise levels σ_{\min} and σ_{\max} . $\mathcal{B}_{\text{disc}}(T)$ is the discretization error of the reverse diffusion process, which decreases as the number of reverse steps T increases. And $\varepsilon_{\text{ratio}}$ is the approximation error of the density-ratio guidance network, capturing the gap between estimated and true density ratios.

C.2 ANALYSIS AND INSIGHTS

The above error bound provides a comprehensive understanding of the factors that contribute to the generalization ability of DiffGDA in the target domain:

1. Source Domain Error \mathcal{E}_S : This error serves as the foundation of the target domain performance. A lower \mathcal{E}_S indicates that the diffusion model effectively captures the structural and semantic information in the source domain, which can be transferred to the target domain. If \mathcal{E}_S is large, it becomes challenging to bridge the domain gap, even with a well-designed alignment strategy.

2. Domain Discrepancy $\mathcal{D}(\mathcal{P}_S, \mathcal{P}_T)$: This term reflects the divergence between the source and target distributions. By leveraging the domain-aware guidance network in DiffGDA, the model minimizes $\mathcal{D}(\mathcal{P}_S, \mathcal{P}_T)$ through adaptive alignment in the latent space. The density ratio term $\log \frac{q(\mathbf{G}_0^T)}{p(\mathbf{G}_0^S)}$ plays a pivotal role in reducing this discrepancy, ensuring that the generated graphs align closely with the target distribution.

3. Sample Complexity $\frac{C}{\sqrt{N}}$: This term highlights the importance of having sufficient labeled samples in the source domain. As N increases, the error bound tightens, indicating better generalization. However, larger models (with higher C) may require more samples to avoid overfitting. DiffGDA mitigates this by efficiently generating intermediate graphs that maximize the utility of source data.

4. Diffusion Residual $\mathcal{R}_{\text{diff}}$: The residual of the diffusion process is governed by three sources: (i) the noise schedule $\mathcal{B}_{\text{sched}}(\sigma_{\min}, \sigma_{\max})$, which controls the stability of forward corruption and reverse denoising; (ii) the discretization error $\mathcal{B}_{\text{disc}}(T)$, which decreases as more reverse steps T are used; and (iii) the density-ratio approximation error $\varepsilon_{\text{ratio}}$, which depends on the accuracy of the guidance network. Each of these terms provides a direct knob for improving error control.

C.3 PRACTICAL IMPLICATIONS

The error bound highlights three key aspects that influence the success of DiffGDA:

- **Domain Alignment:** Reducing $\mathcal{D}(\mathcal{P}_S, \mathcal{P}_T)$ through adaptive guidance and density ratio estimation is critical for effective knowledge transfer.
- **Sample Efficiency:** The ability to leverage limited source domain samples is vital, especially in low-resource scenarios.
- **Diffusion Stability:** Properly tuning the noise schedule, the number of reverse steps, and the accuracy of density-ratio guidance jointly ensures robustness and lowers the residual error in GDA.

In summary, this error bound clarifies that the diffusion residual is not a monolithic term but decomposes into three controllable sources, directly linked to design choices in DiffGDA. Optimizing these factors further improves the theoretical and practical guarantees of cross-domain performance.

D EXPERIMENT SETUP

D.1 DATASET DESCRIPTION

We perform node classification experiments under three migration settings in Table 2. In each scenario, we train our model on one graph and evaluate it on other graphs.

- ACMv9 (**A**), Citationv1 (**C**), DBLPv7 (**D**): These datasets are citation networks from different sources, where each node represents a research paper, and each edge represents a citation relationship between two papers. The task is to predict the category of each research paper. The data are collected from ACM (before 2008), DBLP (2004-2008), and Microsoft Academic Graph (after 2010). We set six transfer scenarios: **A** → **C**, **A** → **D**, **C** → **A**, **C** → **D**, **D** → **A**, and **D** → **C**.
- USA (**U**), Brazil (**B**), Europe (**E**): These datasets are collected from transportation statistics and primarily contain aviation activity data, where each node represents an airport, and each edge represents a flight connection between two airports. The task is to predict the category of each airport. We set six transfer scenarios: **U** → **B**, **U** → **E**, **B** → **U**, **B** → **E**, **E** → **U**, and **E** → **B**.
- Blog1 (**B1**), Blog2 (**B2**): These datasets are taken from the BlogCatalog dataset, where nodes represent bloggers, and edges represent the friendship relationships between bloggers. The node attributes consist of keywords extracted from the bloggers' self-descriptions. The task is to predict the corresponding group to which they belong. We set two scenarios: **B1** → **B2** and **B2** → **B1**.

D.2 BASELINES

We compare our method with the following baseline methods, which can be divided into two categories: (1) Graph Neural Networks (GNNs) on the source graph only: GAT (Veličković et al., 2018), GIN (Xu et al., 2019), and GCN (Kipf & Welling, 2017) are trained end-to-end on the source graph,

allowing direct application to the target graph; (2) Graph Domain Adaptation Methods: DANE (Zhang et al., 2019), UDAGCN (Wu et al., 2020), AdaGCN (Dai et al., 2022), StruRW (Liu et al., 2023a), GRADE (Wu et al., 2023), PairAlign (Liu et al., 2024b), GraphAlign (Huang et al., 2024), A2GNN (Liu et al., 2024a), GGDA (Lei et al., 2025), TDSS (Chen et al., 2025a), DGSDA (Yang et al., 2025) and GAA (Fang et al., 2025a) are specifically designed to address the GDA problem.

Below, we provide descriptions of the baselines used in the experiments.

- **GAT, GIN, GCN**, train using GAT, GIN and GCN under standard ERM.
- **DANE** uses shared weight GCNs to get node representations and then handles distribution shift via least square generative adversarial network.
- **UDAGCN** is a model-centric method that first combines adversarial learning with a GNN model.
- **AdaGCN** is a model-centric method for GDA. It utilizes the techniques of adversarial domain adaptation with graph convolution, and employs the empirical Wasserstein distance between the source and target distributions of node representation as the regularization.
- **STRURW** modifies graph data by adjusting edge weights, while the modification relies on the training of data-centric GDA methods.
- **GRADE** introduces graph subtree discrepancy as a metric to measure the graph distribution shift by connecting GNNs with WL subtree kernel.
- **PairAlign** not only recalibrates the influence between neighboring nodes using edge weights to address conditional structure shifts but also adjusts the classification loss with label weights to account for label shifts.
- **GraphAlign** generates a small, transferable graph from the source graph, which is used for GNN training with Empirical Risk Minimization (ERM).
- **A2GNN** proposes a simple GNN which stacks more propagation layers on target graph.
- **GGDA** using the Fused Gromov-Wasserstein (FGW) metric and adaptive vertex selection, ensures effective knowledge transfer between source and target domains.
- **TDSS** applies structural smoothing directly to the target graph to reduce local variations and enhance model robustness against distribution shifts. TDSS improves knowledge transfer by preserving node representation consistency while mitigating the impact of structural discrepancies.
- **DGSDA** addresses GDA by disentangling node attribute and graph topology alignment. It employs Bernstein polynomial approximation for efficient spectral filter alignment, avoiding costly eigenvalue decomposition, and enables flexible GNNs for improved cross-domain adaptability.
- **GAA** tackles both graph topology and node attribute shifts using an attention-based cross-channel module for effective alignment, enhancing the model’s focus on critical features.

D.3 IMPLEMENTATION DETAILS

We evaluate the model performance using Mi-F1 and Ma-F1 scores. The hidden layer dimension is set to 64, and the learning rate is fine-tuned within the set $\{0.0001, 0.001, 0.01\}$ to achieve optimal performance. Domain alignment is implemented using the MMD method with a Gaussian kernel (Keerthi & Lin, 2003). To prevent overfitting, a dropout rate of 0.2 is applied. We adjust the parameters α within the range $[0, 1]$, η within the range $[0, 0.5]$, and T within the range $[0, 150]$. All experiments are conducted using PyTorch on an NVIDIA RTX 4090 GPU, running for 5 rounds with 150 epochs per round. For each round, the maximum Mi-F1 and Ma-F1 scores are recorded, and their mean and variance are calculated.

In our diffusion models, we adopt the Variance Exploding (VE) SDE, defined by the following stochastic differential equation:

$$dx = \sigma_{\min} \left(\frac{\sigma_{\max}}{\sigma_{\min}} \right)^t \sqrt{2 \log \left(\frac{\sigma_{\max}}{\sigma_{\min}} \right)} dw, \quad (26)$$

where $\sigma_{\min} = 0.001$, $\sigma_{\max} = 0.01$, and $t \in (0, 1]$. The transition probability density function $p_{t, t-\delta t}(x_{t-\delta t} | x_t)$ is a Gaussian distribution with mean x_t and covariance Σ_t :

$$p_{t, t-\delta t}(x_{t-\delta t} | x_t) = \mathcal{N}(x_{t-\delta t} | x_t, \Sigma_t), \quad (27)$$

Methods	Blog1 (B1), Blog2 (B2)				Avg.	
	B1 → B2		B2 → B1			
	Mi-F1	Ma-F1	Mi-F1	Ma-F1		
GAT (ICLR'18)	21.15 \pm 2.72	11.18 \pm 3.84	19.46 \pm 1.41	9.43 \pm 3.04	15.30 \pm 2.75	
GIN (ICLR'19)	20.51 \pm 2.16	11.44 \pm 2.68	18.60 \pm 0.38	9.57 \pm 0.28	15.03 \pm 1.38	
GCN (ICLR'17)	23.17 \pm 1.10	13.80 \pm 1.15	23.46 \pm 0.48	13.76 \pm 1.33	18.59 \pm 1.02	
DANE (IJCAI'19)	28.15 \pm 1.77	25.20 \pm 2.13	30.32 \pm 0.69	28.37 \pm 0.89	28.01 \pm 1.37	
UDAGCN (WWW'22)	33.87 \pm 2.09	28.83 \pm 1.15	31.86 \pm 1.08	28.95 \pm 0.36	30.88 \pm 1.17	
AdaGCN (TKDE'22)	30.75 \pm 0.64	20.69 \pm 1.40	27.17 \pm 2.34	18.73 \pm 2.91	24.34 \pm 1.82	
StruRW (ICML'23)	40.37 \pm 0.44	39.62 \pm 0.77	42.01 \pm 0.44	41.36 \pm 0.47	40.84 \pm 0.53	
GRADE (AAAI'23)	48.44 \pm 3.12	44.46 \pm 4.01	<u>46.78</u> \pm 1.31	42.00 \pm 1.52	45.42 \pm 2.49	
PairAlign (ICML'24)	40.01 \pm 0.78	39.57 \pm 0.85	42.64 \pm 0.69	41.98 \pm 0.85	41.05 \pm 0.79	
GraphAlign (KDD'24)	45.58 \pm 1.15	43.14 \pm 0.55	43.21 \pm 0.62	42.29 \pm 3.28	43.55 \pm 1.40	
A2GNN (AAAI'24)	47.10 \pm 3.40	45.42 \pm 4.75	44.01 \pm 2.01	43.14 \pm 1.77	44.92 \pm 2.98	
GGDA (Arxiv'25)	48.25 \pm 2.45	45.09 \pm 3.60	44.85 \pm 3.13	43.45 \pm 4.74	45.41 \pm 3.48	
TDSS (AAAI'25)	<u>49.53</u> \pm 1.08	<u>48.31</u> \pm 2.00	44.20 \pm 0.29	<u>44.31</u> \pm 1.89	<u>46.59</u> \pm 1.31	
DGSDA (ICML'25)	48.61 \pm 1.73	48.01 \pm 1.71	44.12 \pm 1.26	42.06 \pm 3.66	45.70 \pm 1.68	
GAA (ICLR'25)	48.80 \pm 2.45	45.66 \pm 2.12	43.74 \pm 1.25	40.59 \pm 2.04	44.70 \pm 1.97	
DiffGDA (Ours)	53.52 \pm 2.93	51.18 \pm 0.97	48.95 \pm 2.48	47.17 \pm 1.26	50.20 \pm 1.91	

Table 4: Node classification performance ($\% \pm \sigma$) on the Social domains. The highest scores are highlighted in **bold**, while the second-highest scores are underlined.

Dataset	A → C	C → A	A → D
Baseline – $\log_2(0.05)$	4.32	4.32	4.32
DiffGDA vs TDSS (AAAI'25)	16.29	13.57	7.43
DiffGDA vs AdaGCN (TKDE'22)	30.84	26.46	19.57
DiffGDA vs A2GNN (AAAI'24)	31.98	18.96	23.67
DiffGDA vs GraphAlign (KDD'24)	45.39	43.36	40.63

Table 5: Significance test (t-test) of DiffGDA

where Σ_t is given by:

$$\Sigma_t = \sigma_{\min}^2 \left(\frac{\sigma_{\max}}{\sigma_{\min}} \right)^{2t} - \sigma_{\min}^2 \left(\frac{\sigma_{\max}}{\sigma_{\min}} \right)^{2t-2\delta t}. \quad (28)$$

After generating samples by simulating the reverse diffusion process, we quantize the elements of the adjacency matrix to $\{0, 1\}$ by clipping values in the ranges $(-\infty, 3)$ to 0 and $[3, +\infty)$ to 1.

The guidance network is a 3-layer MLP with 512 hidden units and ReLU activation functions. We train the guidance network for 100 epochs.

E EXPERIMENT RESULTS

E.1 MAIN EXPERIMENT RESULTS

The experimental results on Blog domains are presented in Table 4.

E.2 SIGNIFICANCE TEST (T-TEST) OF DIFFGDA

To ensure that our method achieves statistically significant improvements, we conducted paired t-tests to compare the performance of DiffGDA against various baseline models across multiple transfer tasks, as presented in Table 5. This rigorous statistical analysis was carried out at the 0.05 significance level, allowing us to objectively assess the superiority of our approach. We leveraged $-\log_2(p)$ values to quantify statistical evidence, where higher $-\log_2(p)$ values indicate stronger support for the observed differences. By employing this approach, we aim to demonstrate its consistent and reliable performance across diverse tasks.

E.3 ABLATION STUDY

The more ablation results are presented in Figure 7.

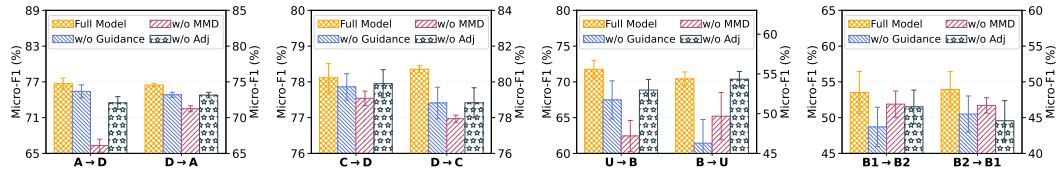
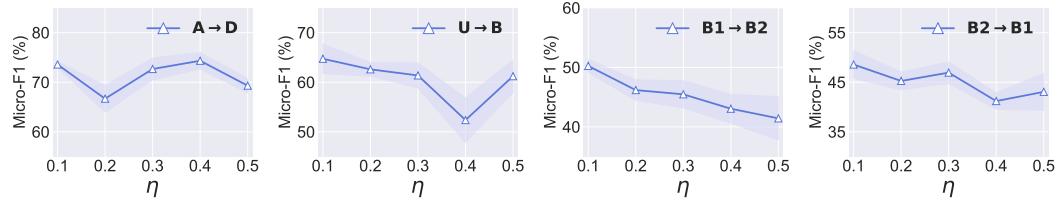
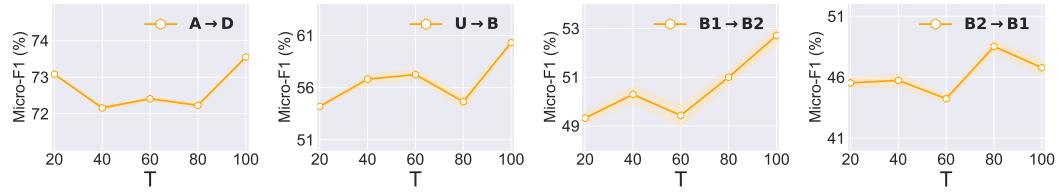


Figure 7: Classification Mi-F1 comparisons between DiffGDA variants on other cross-domain tasks.

E.4 HYPER-PARAMETER ANALYSIS

The more hyper-parameter results are presented in Figure 8 and Figure 9.

Figure 8: The performances of our DiffGDA w.r.t varying weight η on different transfer tasksFigure 9: The performances of our DiffGDA w.r.t varying step T on different transfer tasks