# Laplacian-Guided Denoising Graph Diffusion for Graph Learning with an Adaptive Prior

#### **Anonymous Author(s)**

Affiliation Address email

#### **Abstract**

Graph representation learning often relies on manually engineered, task-specific inductive biases, which limit model flexibility and generalization across diverse tasks. While diffusion models have shown promising ability in capturing arbitrary distributions, they frequently lack a deep integration of graph structure. To address this, we propose the LapDiff, a novel diffusion-based framework that learns adaptive priors to dynamically align its inductive bias with the intrinsic characteristics of graph-structured data and their tasks. The novelty of LapDiff is its use of Laplacian smoothing as a structure-aware noise mechanism in the forward process, complemented by topological perturbations. This design enables the denoising network to effectively capture the underlying data-generating factors tied to a graph's unique structure and features. By capturing priors from a task and data, LapDiff mitigates the limitations of static biases and enhances task-agnostic generalization. Extensive experiments on large-scale OGB benchmarks demonstrate that LapDiff is universally effective for both link prediction and node classification, achieving state-of-the-art performance and offering a new perspective into graph representation learning.

# 7 1 Introduction

2

3

5

9

10

11

12

13

15

16

- Large-scale network graphs are now foundational data structures across science and industry, from biological systems [1] and drug discovery [2] to recommender systems [3] and fraud detection [4]. The central challenge in this field is learning effective representations that capture the complex,
- 21 non-Euclidean relationships inherent in such data. Graph Neural Networks (GNNs) have emerged as 22 the dominant paradigm, achieving remarkable success by leveraging strong inductive biases.
- However, this reliance on inductive biases is a double-edged sword. The vast majority of GNNs are built on a fixed assumption, such as homophily—the tendency of connected nodes to be similar. While effective in some scenarios, this rigid prior leads to significant performance degradation on graphs with more complex topologies or across different downstream tasks. To compensate, a fragmented
- landscape of specialized models has emerged, each with manually engineered biases tailored to specific tasks, such as subgraph-based features for link prediction (e.g., SEAL [5]) or generalized
- 29 graph heuristics (e.g., Neo-GNNs [6]).
- This approach merely trades one fixed assumption for another, inherently limiting generalization and requiring separate models for different tasks, which is impractical for real-world deployment [7].
- The lack of a single, dominant architecture on comprehensive benchmarks like the Open Graph
- 33 Benchmark (OGB) [8] highlights this fundamental limitation. For instance, while GCN [9] achieves
- high accuracy on the OGB-PPA node classification task, its performance plummets in OGB-Collab
- link prediction, where heuristic-based methods like SEAL [5] excel. This disparity underscores that

fixed, manually-specified priors are a bottleneck for creating truly universal graph representationlearning methods.

This predicament raises a fundamental question: Can we learn powerful graph representations

without manually-designed fixed inductive bias? We argue for a paradigm shift from static, hand-

crafted biases to *adaptive* ones that emerge as data-driven priors, conditioned directly on the given task and intrinsic properties of the given graph.

To establish a formal basis for this approach, we first draw intuition from the No-Free-Lunch (NFL) theorem [10], which suggests that no single algorithm is universally optimal across all tasks and data. This implies that a model's effectiveness is tied to the alignment between its internal prior and the specific data distribution. To extend beyond the intuition behind the NFL theorem, we use the PAC-Bayes framework. This theory provides a rigorous foundation for using adaptive priors by

connecting them to a model's generalization ability. A standard PAC-Bayes generalization bound for a posterior distribution P over a hypothesis class (e.g., model parameters  $\theta$ ) and a prior distribution

49 Q is given by:

38

39

63

64

65

66

67 68

69

70

71

72

73

74

75

76

78

$$\mathbb{E}_{\theta \sim P}[\mathcal{R}_{true}(\theta)] \leq \mathbb{E}_{\theta \sim P}[\mathcal{R}_{emp}(\theta)] + \sqrt{\frac{\text{KL}(P||Q) + \ln(m/\delta)}{2m}}$$
(1)

where, with probability at least  $1-\delta$  over the draw of a training set of size m, the inequality holds for all priors Q. Here,  $\mathcal{R}_{true}(\theta)$  is a generalization error and  $\mathcal{R}_{emp}(\theta)$  is the training error for a given parameter set  $\theta$ .

The crucial insight from this bound is that the generalization error is controlled by two terms: 53 the training error and a complexity term involving the Kullback-Leibler divergence between the 54 posterior P and the prior Q. To guarantee that a low training error translates to a low generalization 55 error, the KL(P||Q) term must be minimized. Since the posterior P is learned from the data, a 56 fixed, data-independent prior Q may be far from P, resulting in a large KL divergence and a loose 57 bound. However, the PAC-Bayes framework allows for the use of an adaptive prior Q(G) [11, 12]. 58 By choosing a prior that adapts to a downstream task  $\mathcal{T}$  and an input graph  $\mathcal{G}$ , we can select a 59  $Q(\mathcal{T},\mathcal{G})$  that is already close to the expected posterior P, thereby minimizing the KL divergence and 60 tightening the generalization bound. This provides a principle for learning an adaptive prior in graph 61 representation learning methods. 62

This theoretical perspective provides a principled foundation for our goal: to learn graph representation with an adaptive prior  $P_{\theta}(z \mid \mathcal{T}, \mathcal{G})$ . We propose **LapDiff**, a novel diffusion-based framework that realizes this objective. LapDiff learns this adaptive prior by parameterizing latent representations as the outcome of a denoising process grounded in the graph's intrinsic geometry. Its core innovation is the use of \*\*Laplacian smoothing as a structure-aware noise mechanism\*\*, a principled way to infuse structural information into the diffusion process.

Our contributions are summarized as follows:

- We propose a generative framework that learns an adaptive, graph-conditioned prior,  $p_{\theta}(z|G)$ , to capture implicit structural and feature-level patterns without handcrafted biases.
- We introduce the novel use of Laplacian smoothing as a structure-aware noise mechanism
  within a diffusion model, providing a new perspective on encoding graph geometry into
  generative processes.
- We empirically demonstrate that LapDiff achieves state-of-the-art performance on diverse benchmarks for link prediction and node classification, validating the effectiveness of our adaptive, structure-aware learning approach.

#### 2 Related work

Graph Representation Learning. GNNs like GCN [13] and GAT [14] have achieved strong performance in node classification by exploiting local aggregation schemes. However, these often encode fixed assumptions (e.g., homophily), limiting generalization to diverse graph types. While GRAND [15] models feature diffusion as a continuous process, structural signals remain underutilized in most node-centric models. In contrast, SEAL [5], Neo-GNNs [6], and NBFNet [16] introduce inductive biases tailored for link prediction via subgraph extraction or path-based heuristics. However, these methods are narrowly focused on specific tasks.

**DDPMs on Graph domain.** DDPMs have recently been extended to graphs. Prior generative approaches [17–19] modeled node-edge distributions or joint graph likelihoods. Molecular graph 87 generation works (e.g., [20, 21]) employed categorical and discrete formulations, while [22] proposed 88 scalable structure perturbations. DDM citeyang2023directional introduced diffusion models for graph 89 representation learning with anisotropic node feature noise, yet ignored graph structures. Recently, 90 SGDiff [23] proposed a subgraph-based diffusion model tailored for link prediction, incorporating 91 structural context into diffusion.

#### 3 Method

93

94

95

97

98 99

100

101 102

103

105

106

108

109

110

111

112

114

115

116

117

118

119

120

121

131

Our model, **LapDiff**, implicitly learns the adaptive graph prior  $p_{\theta}(z \mid G)$  posited in our theoretical motivation via a denoising diffusion process. LapDiff consists of two complementary processes: a forward process that incrementally injects structure-aware noise into the input graph  $G_0 = (A_0, X_0)$ , and a reverse process that learns to denoise the corrupted graph to reconstruct an informative latent representation. This section details both processes and the resulting objective function.

### 3.1 LapDiff Generative Process

The forward process defines a Markov chain  $\{G_t\}_{t=1}^T$  that gradually transforms the initial graph data  $G_0$  into oversmoothed and noisy graph data at step 104 T. Unlike standard diffusion models, our forward process is designed to be structure-aware, corrupting both node 107 features and graph topology in a manner that adaptively respects the prior underlying in graph data. The reverse process, parameterized by a neural network  $f_{\theta}$ , learns to reverse this corruption stepby-step, effectively learning the datagenerating distribution. We define two

Markov chains for feature and structural

diffusion,  $q(X_{1:T}|X_0)$  and  $q(A_{1:T}|A_0)$ , respectively.

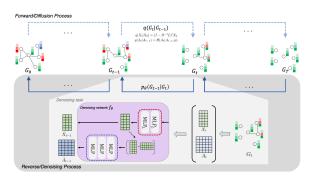


Figure 1: Graphical Model of LapDiff. LapDiff leverages the Laplacian smoothing as noise source of the forward process, inducing signal diffusion reflecting the important aspect of a graph structure. This assures noise in graph signals in the feature space.

**Feature Diffusion via Laplacian Smoothing.** To corrupt node features while retaining structural information, we define the forward process noise source using the graph Laplacian. Specifically, we use iterative Laplacian smoothing, which acts as a low-pass filter on node features, causing them to dissipate and converge towards an over-smoothed state. The feature matrix  $X_t$  at timestep t is obtained as:

$$X_t = (\mathbf{I} - \alpha \mathbf{D}^{-1} \mathbf{L}) X_{t-1} = (\mathbf{I} - \alpha \mathbf{D}^{-1} \mathbf{L})^t X_0,$$
(2)

where L is the graph Laplacian, D is the degree matrix, and  $\alpha$  is a scaling factor. For simplicity, we 122 set  $\alpha=1$ . This process defines a Markov chain  $q(X_{1:T}|X_0)=\prod_{t=1}^T q(X_t|X_{t-1})$  where the state at any step t can be computed in closed form from  $X_0$ . 123 124

Structural Diffusion via Edge Perturbation. To diffuse the graph topology, we employ stochastic 125 edge removal. At each step t, we sample a subgraph structure by randomly dropping edges from the previous state  $A_{t-1}$  with a removal probability 1-p. This process gradually sparsifies the graph, 127 destroying its topological information. The per-edge transition is defined as: 128

$$A_t[i,j] \sim \text{Bernoulli}(A_{t-1}[i,j],p)$$
 (3)

This defines the structural diffusion process  $q(A_{1:T}|A_0) = \prod_{t=1}^T q(A_t|A_{t-1})$ , which ensures that 129 structural information diffuses over time. 130

**Theoretical Grounding of Laplacian Smoothing.** Our choice of Laplacian smoothing as a noise source is theoretically grounded in spectral graph theory. Given the eigendecomposition of the Laplacian  $\mathbf{L} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^{\top}$ , the feature diffusion in Eq. (2) can be expressed in the spectral domain as:

$$X_t = \mathbf{U}(\mathbf{I} - \alpha \mathbf{\Lambda})^t \mathbf{U}^\top X_0. \tag{4}$$

The term  $(\mathbf{I} - \alpha \mathbf{\Lambda})^t$  acts as a low-pass filter, attenuating high-frequency components of the feature signal (associated with large eigenvalues  $\lambda_i$ ) more rapidly than low-frequency components (small  $\lambda_i$ ).

This selective, non-uniform information destruction preserves global structural patterns for longer, forcing the reverse process to learn representations that are deeply tied to the graph's macro-structure,

unlike isotropic Gaussian noise which degrades all signal components uniformly.

In other words, Laplacian smoothing increases the entropy of the graph signal selectively across the spectral spectrum. In the spectral domain, the smoothed signal becomes:

$$\hat{x}' = [(1 - \alpha \lambda_1)\hat{x}_1, (1 - \alpha \lambda_2)\hat{x}_2, \dots, (1 - \alpha \lambda_N)\hat{x}_N].$$
 (5)

which indicates high-frequency components associated with larger  $\lambda_i$  of x are reduced more than low-frequency components.

This selective reduction effectively smooths out rapid variations in the signal while preserving the slow variations, which correspond to the underlying structure of the graph. This ensures a stronger entropy increase of high-frequency components, leading to non-uniform entropy increase. A high-frequency variation is suppressed, while low-frequency structure is preserved. This behavior enables LapDiff enables to adaptively capture priors effectively.

# 3.2 Reverse Process and Objective Function

148

Feature Denoising Objective. The feature denoising term aims to predict the features at the previous step,  $X_{t-1}$ , which can be simplified to a mean squared error loss:

$$\mathcal{L}_{\text{feat}}(t) = \|f_{\theta}^{(X)}(X_t, A_t, t) - X_{t-1}\|_{2}^{2}. \tag{6}$$

Structural Denoising Objective. Directly optimizing the KL divergence for discrete edges is challenging. Instead, following a second-order Taylor expansion of the per-edge Bernoulli KL divergence, we can approximate the objective with a mean squared error in the Laplacian space. This provides a more stable, continuous optimization target:

$$\mathcal{L}_{\text{struct}}(t) \approx \|f_{\theta}^{(L)}(X_t, A_t, t) - (\mathbf{L}_0 - \mathbf{L}_{t-1})\|_2^2. \tag{7}$$

Here, the network predicts the change in the Laplacian, which corresponds to the structure that was removed between step 0 and t-1.

Final Loss. Combining the denoising terms for  $t \in [2, T]$  and the explicit reconstruction losses for t = 1 (using MSE for features and binary cross-entropy for the adjacency matrix), the final training loss for LapDiff is:

$$\mathcal{L}_{\text{LapDiff}} := \beta_t \sum_{t=2}^{T} \|f_{\theta}^{(X)}(X_t, A_t) - X_{t-1}\|_2^2 + \gamma \sum_{t=2}^{T} \|f_{\theta}^{(L)}(X_t, A_t) - (\mathbf{L}_0 - \mathbf{L}_{t-1})\|_2^2 + \beta_0 \|f_{\theta}^{(X)}(X_1, A_1) - X_0\|_2^2 + \lambda \mathcal{L}_{\text{BCE}}(f_{\theta}^{(A)}(X_1, A_1), A_0) + \mathcal{L}_{\text{task}},$$
(8)

where  $\beta_0, \gamma, \beta_t, \lambda$  are weighting hyperparameters. The total loss is  $\mathcal{L} = \mathcal{L}_{LapDiff} + \mathcal{L}_{task}$ , where  $\mathcal{L}_{task}$  is a downstream task-specific loss. Derivation of the objective functions is in the Appendix.

#### 4 Experiments

162

Our goal is to demonstrate that LapDiff's adaptive prior mechanism leads to universally strong perfor-163 mance across diverse graph learning tasks. We then analyze the contribution of its core components. 164 We evaluate LapDiff on seven benchmark datasets including Open Graph Benchmark (OGB), cover-165 ing two diverse network graph learning tasks: link prediction and node classification. For evaluation, we follow the standard OGB protocols, using Hits@K and MRR for link prediction, and accuracy for 167 node classification. For node classification, we adopt a challenging few-shot setting with a fixed num-168 ber of labeled nodes  $k \in \{1, 5, 10\}$  per class to test representation power under extreme label scarcity. 169 We compare LapDiff against a comprehensive set of baselines, including heuristic methods, classic GNNs (GCN, GAT, GraphSAGE), and recent specialized architectures (SEAL, Neo-GNNs, C&S, 171 GraphMAE2). Experimental details are provided in the Appendix. An anonymized implementation is available at https://anonymous.4open.science/r/NPGMLworkshop2025CDF2

Table 1: Link prediction performances on Open Graph Benchmark (OGB) datasets. OOM denotes 'out of memory'. **Bold underline** indicates the best performance and **bold** indicates the second best performance.

	Model	OGB-PPA	OGB-Collab	OGB-DDI	OGB-Citation2
Neighborhoood Heuristics	Common Neighbors   Adamic Adar   Resource Allocation	$ \begin{vmatrix} 27.65 \pm 0.00 \\ 32.45 \pm 0.00 \\ \mathbf{\underline{49.33}} \pm 0.00 \end{vmatrix} $	$ \begin{vmatrix} 50.06 \pm 0.00 \\ 53.00 \pm 0.00 \\ 52.89 \pm 0.00 \end{vmatrix} $	$ \begin{vmatrix} 17.73 \pm 0.00 \\ 18.61 \pm 0.00 \\ 6.23 \pm 0.00 \end{vmatrix} $	$ \begin{array}{c c} 76.20 \pm 0.0 \\ 76.12 \pm 0.0 \\ 76.20 \pm 0.0 \end{array} $
Shallow Methods	Matrix Factorization MLP	$\begin{vmatrix} 23.78 \pm 1.82 \\ 0.99 \pm 0.15 \end{vmatrix}$	$\begin{vmatrix} 34.87 \pm 0.23 \\ 16.05 \pm 0.48 \end{vmatrix}$	$ \begin{array}{c} 13.29 \pm 2.32 \\ \text{N/A} \end{array}$	$\begin{array}{c c} 50.48 \pm 3.09 \\ 25.13 \pm 0.28 \end{array}$
GRL Methods	GCN GAT SAGE JKNet SEAL Neo-GNN DDM LapDiff(ours)	$\begin{array}{c} 15.37 \pm 1.25 \\ \text{OOM} \\ 12.51 \pm 2.02 \\ 11.73 \pm 1.98 \\ 47.18 \pm 3.60 \\ 47.53 \pm 0.63 \\ 17.93 \pm 1.91 \\ \textbf{48.32} \pm 0.68 \end{array}$	$\begin{array}{c} \textbf{44.57} \pm 0.64 \\ \textbf{41.73} \pm 1.01 \\ \textbf{47.86} \pm 0.64 \\ \textbf{47.52} \pm 0.73 \\ \textbf{54.27} \pm 0.46 \\ 53.95 \pm 0.52 \\ \textbf{49.56} \pm 1.79 \\ \textbf{54.33} \pm 0.35 \end{array}$	$40.87 \pm 6.08$ $32.57 \pm 3.48$ $47.06 \pm 5.21$ $57.95 \pm 7.69$ $29.86 \pm 4.37$ $60.02 \pm 3.86$ $47.73 \pm 3.10$ $60.56 \pm 2.32$	$\begin{array}{c} 82.57 \pm 0.26 \\ \text{OOM} \\ 80.18 \pm 0.15 \\ \text{OOM} \\ \hline \textbf{86.72} \pm 0.31 \\ \hline 85.96 \pm 0.94 \\ 83.51 \pm 0.20 \\ \hline \textbf{86.70} \pm 0.27 \\ \end{array}$

Table 2: Node classification performance on OGB-Arxiv, OGB-Products, and PubMed dataset. OOM denotes 'out of memory'. **Bold** indicates the best performance.

Model		OGB-Arxiv			OGB-Products			PubMed	
Fixed k nodes	k = 1	k = 5	k = 10	k = 1	k = 5	k = 10	k = 1	k = 5	k = 10
GCN GAT APPNP GCNII C&S	$\begin{array}{c} 31.69 \pm 2.74 \\ 25.60 \pm 2.95 \\ 29.36 \pm 2.19 \\ 30.94 \pm 2.30 \\ 30.63 \pm 1.88 \end{array}$	$\begin{array}{c} 52.97 \pm 0.94 \\ 50.87 \pm 1.78 \\ 52.47 \pm 1.26 \\ 51.94 \pm 1.18 \\ 51.73 \pm 1.30 \end{array}$	$\begin{array}{c} 58.39 \pm 0.50 \\ 57.23 \pm 0.75 \\ 56.42 \pm 0.83 \\ 57.65 \pm 0.94 \\ 56.57 \pm 1.43 \end{array}$	$\begin{array}{c} 38.93 \pm 2.09 \\ 35.81 \pm 2.42 \\ 36.35 \pm 2.20 \\ 33.64 \pm 2.32 \\ 40.47 \pm 1.97 \end{array}$	$\begin{array}{c} 62.69 \pm 1.27 \\ 60.72 \pm 1.93 \\ 63.01 \pm 2.10 \\ 61.43 \pm 2.36 \\ 62.18 \pm 1.57 \end{array}$	$\begin{array}{c} 66.23 \pm 0.91 \\ 64.80 \pm 1.21 \\ 66.85 \pm 0.84 \\ 64.90 \pm 1.39 \\ 67.53 \pm 1.40 \end{array}$	$\begin{array}{c} 45.87 \pm 2.44 \\ 43.57 \pm 2.71 \\ 43.04 \pm 1.72 \\ 43.29 \pm 2.53 \\ 44.91 \pm 1.24 \end{array}$	$\begin{array}{c} 60.56 \pm 1.44 \\ 58.38 \pm 2.06 \\ 56.94 \pm 1.90 \\ 56.18 \pm 1.84 \\ 57.44 \pm 1.36 \end{array}$	$\begin{array}{c} 69.50 \pm 0.68 \\ 68.40 \pm 1.49 \\ 69.99 \pm 0.73 \\ 70.60 \pm 0.93 \\ 68.78 \pm 1.07 \end{array}$
CCA-SSG GraphMAE2 DDM LapDiff(ours)	$29.21 \pm 3.01$ $33.83 \pm 3.23$ $36.08 \pm 0.93$ $39.04 \pm 1.52$	$51.67 \pm 2.31$ $54.67 \pm 2.32$ $55.69 \pm 1.41$ $57.83 \pm 0.83$	$57.40 \pm 0.89$ $60.16 \pm 0.75$ $60.71 \pm 0.31$ $61.23 \pm 0.37$	$39.95 \pm 2.67$ $41.65 \pm 2.01$ $45.57 \pm 1.28$ $49.10 \pm 1.75$	$63.10 \pm 2.13$ $63.69 \pm 2.30$ $64.90 \pm 1.03$ <b>67.47</b> $\pm 1.03$	$67.62 \pm 1.56$ $68.08 \pm 1.62$ $69.62 \pm 0.15$ <b>70.69</b> $\pm 0.61$	$ \begin{vmatrix} 47.56 \pm 3.15 \\ 50.76 \pm 5.10 \\ 50.95 \pm 2.42 \\ \textbf{53.82} \pm 1.46 \end{vmatrix} $	$60.44 \pm 2.60$ $64.29 \pm 0.11$ $65.13 \pm 0.10$ $66.93 \pm 0.88$	$69.24 \pm 0.68 70.51 \pm 0.11 70.12 \pm 0.56 72.77 \pm 0.62$

#### 4.1 Performance on Downstream tasks

Across all link prediction and node classification benchmarks, LapDiff consistently outperforms task-specific baselines, indicating it captures both structural and feature-level priors without hand-crafted biases. This robustness across diverse tasks and datasets confirms its ability to learn broadly generalizable graph representations, emphasizing LapDiff's utility for diverse graph-learning applications.

Link prediction. Table 1 reports the results of OGB link prediction benchmarks. Our LapDiff generally shows significantly improved performance than other baselines on OGB-Collab and OGB-DDI datasets and second-best performance on OGB-PPA and OGB-CItation2 with low discrepancy to the best models. This indicates our LapDiff is capable of capturing latent graph priors that are crucial in the context of link prediction tasks and data. Additionally, LapDiff demonstrates robust performance across various datasets. LapDiff has the ability to handle both underlying prior distribution and internal biases in graph structures and node features. For instance, LapDiff achieves the best performance on OGB-DDI, whereas SEAL, which is designed to generalize higher-order structural heuristics shows relatively poor performance. This result illustrates that existing GNNs often show degradation under conditions that are distinct from their intrinsic objectives. Thus, it highlights that our model effectively learn latent representations with hidden, complex aspects that are captured by latent priors in a graph naturally regarding certain objectives.

**Node classification.** We conduct experiments on semi-supervised node classification benchmark datasets to validate the effectiveness of LapDiff on learning node embeddings. We constrained the training index by the fixed k nodes per label. The number k is set to 1, 5, and 10. Table 2 shows the performance of a semi-supervised node classification task that is extremely limited to label scarcity. LapDiff outperforms other baselines on all datasets and settings. According to the results, LapDiff effectively captures the latent distribution of nodes, even under very constrained conditions.

#### 4.2 Analysis on Components

We empirically validate the efficacy of each component in LapDiff through ablation experiments in Table 3. First, we demonstrate the capability of LapDiff to capture universal and comprehensive representation by training without downstream task loss. It is remarkable that it still shows a relatively

Table 3: Ablation study analyzing the efficacy of each component of LapDiff.

Dataset	OGB-Collab	OGB-PPA
LapDiff	$54.33 \pm 0.37$	$48.87 \pm 0.64$
LapDiff (w/o downstream loss)	$51.78 \pm 0.40$	$44.92 \pm 1.52$
LapDiff (w/o feature process)	$45.13 \pm 2.33$	$39.77 \pm 1.13$
LapDiff (w/o structure process)	$45.47 \pm 3.02$	$26.42 \pm 2.30$

Table 4: Analysis on the efficacy of noise source. As other baselines only consider node features, the structure process in LapDiff is excluded from evaluation.

Dataset Metric	OGB-Collab Hits@50	$\begin{array}{c} \text{OGB-Arxiv} \ (k=5) \\ \text{Accuracy} \end{array}$
LapDiff (Laplacian smoothing)	$52.30 \pm 0.97$	$57.02 \pm 0.79$
DDPM (Random noise)	$47.43 \pm 1.58$	$53.38 \pm 1.02$
DDM (Anisotropic noise)	$49.56 \pm 1.79$	$55.69 \pm 1.41$

high performance than LapDiff without feature process or structure process. Also, it is interesting that it still outperforms conventional GNNs including GCN, GAT, SAGE, and JKNet. This result implies that LapDiff is capable of learning critical latent representations within a graph that are significant by aligning hidden graph priors.

Then, we evaluate LapDiff without the feature diffusion process and structural diffusion process based on the average performance on link prediction datasets. LapDiff without feature diffusion process and LapDiff without structural diffusion process both show degraded performance on OGB-Collab and OGB-PPA.

#### 4.3 Analysis on Noise Source

209

210

212

213

214

215

216

217

218

219

220

221

222 223

224

225

226

227

228

Table 4 presents an analysis of the efficacy of different noise sources to capture underlying latent priors dynamically aligning to the internal structures of task and dataset. The noise sources compared to Laplacian smoothing are random Gaussian noise from DDPM [24] and Anisotropic Gaussian noise in DDM [25]. The effectiveness of each noise source is evaluated for link prediction and semi-supervised node classification tasks on OGB-Collab and OGB-Arxiv datasets, respectively. On both datasets, Laplacian smoothing outperforms the other two noise sources. It achieves a score of 52.30 on the OGB-Collab dataset and 57.02 on the OGB-Arxiv dataset. Compared to random Gaussian noise and anisotropic Gaussian noise, Laplacian smoothing shows a clear advantage. For instance, it exceeds random Gaussian Noise by nearly 5 percentage points and exceeds anisotropic Gaussian noise by over 3 percentage points on OGB-Collab dataset. These results indicate that utilizing Laplacian smoothing as a noise source is not only effective but also consistent, offering a significant improvement over other random noise. It provides strong empirical evidence of the efficacy of Laplacian smoothing as a noise source in a diffusion model for graph representation learning tasks, specifically for large graph data.

#### 4.4 Underlying Distribution of Graph

The goal of LapDiff is to learn universal and generalizable representations that are aligned with the underlying distribution of large graph data. To validate the capability to learn priors of network graphs and implicitly aligned with inherent priors, Fig 2 provides the visualization

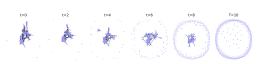


Figure 2: Visualization of graph reconstruction on PubMed at every even number time step. To improve the visibility of the figure, we sampled the subgraph of PubMed with large connectivity.

of the reconstruction of a graph on PubMed dataset at each even number of states. As shown in Fig 2, the trajectory of reconstruction is directed toward the input state, eventually, showing a similar structure at t=2 compared to the input graph t=0.

## 5 Conclusion

We introduced LAPDIFF, a diffusion-based framework for graph representation learning that instantiates adaptive inductive bias via Laplacian smoothing and stochastic structural perturbations. LAPDIFF learns a graph-specific prior instead of relying on hand-crafted assumptions. Across diverse benchmarks, it delivers strong performance and provides evidence toward more universal graph learners. Future work includes exploring richer structure-aware noise processes and theoretically analyzing the properties of the learned graph priors (e.g., does  $p_{\theta}(z|G)$ ). Our work encourage further research on adaptive inductive bias universal representation in graph machine learning.

#### References

- [1] Vladimir Gligorijevic and et al. Structure-based protein function prediction using graph convolutional networks. *Nature Communications*, 12(1):3168, 2021.
- [2] Javier Jiménez-Luna, Francesca Grisoni, and Gisbert Schneider. Drug discovery with graph neural networks: A survey. *Molecular Informatics*, 2021.
- 247 [3] Xiangnan He and et al. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference*, 2020.
- [4] Yingtong Dou and et al. Enhancing graph neural network-based fraud detection with sequential and time information. In *Proceedings of the 26th ACM SIGKDD Conference (KDD)*, 2020.
- [5] Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. In *Advances* in Neural Information Processing Systems, pages 5165–5175, 2018.
- [6] Seongjun Yun, Seoyoon Kim, Junhyun Lee, Jaewoo Kang, and Hyunwoo J. Kim. Neo-GNNs:
   Neighborhood overlap-aware graph neural networks for link prediction. In A. Beygelzimer,
   Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, Advances in Neural Information
   Processing Systems, 2021. URL https://openreview.net/forum?id=Ic9vRN3VpZ.
- [7] Vijay Prakash Dwivedi, Chaitanya K. Joshi, Thomas Laurent, Yoshua Bengio, and Xavier
   Bresson. Benchmarking graph neural networks. In *International Conference on Learning Representations (ICLR)*, 2020.
- [8] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele
   Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs.
   arXiv preprint arXiv:2005.00687, 2020.
- [9] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.
- <sup>265</sup> [10] David H Wolpert. The lack of a priori distinctions between learning algorithms. *Neural Computation*, 8(7):1341–1390, 1996.
- <sup>267</sup> [11] Olivier Catoni. Pac-bayesian supervised classification: the thermodynamics of statistical learning. *arXiv preprint arXiv:0712.0248*, 2007.
- 269 [12] Gintare Karolina Dziugaite and Daniel M. Roy. Computing nonvacuous generalization bounds 270 for deep (stochastic) neural networks with many more parameters than training data, 2017. URL 271 https://arxiv.org/abs/1703.11008.
- 272 [13] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=SJU4ayYgl.
- [14] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua
   Bengio. Graph attention networks. *International Conference on Learning Representations*,
   2018.
- [15] Benjamin Paul Chamberlain, James Rowbottom, Maria Goronova, Stefan Webb, Emanuele
   Rossi, and Michael M Bronstein. Grand: Graph neural diffusion. Proceedings of the 38th
   International Conference on Machine Learning, (ICML) 2021, 18-24 July 2021, Virtual Event,
   2021.
- Zhaocheng Zhu, Zuobai Zhang, Louis-Pascal Xhonneux, and Jian Tang. Neural bellman-ford
   networks: A general graph neural network framework for link prediction. Advances in Neural
   Information Processing Systems, 34, 2021.
- <sup>285</sup> [17] Jiaqi Ma, Weijing Tang, Ji Zhu, and Qiaozhu Mei. A flexible generative framework for graph-<sup>286</sup> based semi-supervised learning. *Advances in neural information processing systems*, 32, 2019.

- Pantelis Elinas, Edwin V Bonilla, and Louis Tiao. Variational inference for graph convolutional networks in the absence of graph data and adversarial settings. *Advances in Neural Information Processing Systems*, 33:18648–18660, 2020.
- [19] Jaehyeong Jo, Seul Lee, and Sung Ju Hwang. Score-based generative modeling of graphs
   via the system of stochastic differential equations. In *International Conference on Machine Learning*, pages 10362–10383. PMLR, 2022.
- 293 [20] Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pas-294 cal Frossard. Digress: Discrete denoising diffusion for graph generation. *arXiv preprint* 295 *arXiv*:2209.14734, 2022.
- [21] Kilian Konstantin Haefeli, Karolis Martinkus, Nathanaël Perraudin, and Roger Wattenhofer.
   Diffusion models for graphs benefit from discrete state spaces. arXiv preprint arXiv:2210.01549,
   2022.
- 299 [22] Xiaohui Chen, Jiaxing He, Xu Han, and Li-Ping Liu. Efficient and degree-guided graph generation via discrete diffusion modeling. *arXiv preprint arXiv:2305.04111*, 2023.
- [23] Hang Li, Wei Jin, Geri Skenderi, Harry Shomer, Wenzhuo Tang, Wenqi Fan, and Jiliang Tang.
   Sub-graph based diffusion model for link prediction, 2024. URL https://arxiv.org/abs/2409.08487.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- <sup>307</sup> [25] Run Yang, Yuling Yang, Fan Zhou, and Qiang Sun. Directional diffusion models for graph representation learning, 2023.
- [26] William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large
   graphs. In *Proceedings of the 31st International Conference on Neural Information Processing* Systems, NIPS'17, 2017.
- [27] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and
   Stefanie Jegelka. Representation learning on graphs with jumping knowledge networks. In
   Proceedings of the 35th International Conference on Machine Learning, volume 80, pages
   5453–5462. PMLR, 10–15 Jul 2018.
- 316 [28] Thomas N Kipf and Max Welling. Variational graph auto-encoders. *arXiv preprint* 317 *arXiv:1611.07308*, 2016.
- David Liben-Nowell and Jon Kleinberg. The link prediction problem for social networks. In Proceedings of the Twelfth International Conference on Information and Knowledge Management, CIKM '03, page 556–559. Association for Computing Machinery, 2003. doi: 10.1145/956863.956972.
- 322 [30] Lada A Adamic and Eytan Adar. Friends and neighbors on the web. *Social networks*, 25(3): 211–230, 2003.
- [31] Tao Zhou, Linyuan Lü, and Yi-Cheng Zhang. Predicting missing links via local information.
   The European Physical Journal B, 71:623–630, 2009.
- 326 [32] Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- 328 [33] Simon Haykin. Neural networks: a comprehensive foundation. Prentice Hall PTR, 1994.
- Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. Predict then propagate:
  Graph neural networks meet personalized pagerank. In *International Conference on Learning Representations (ICLR)*, 2019.
- Zhewei Wei Ming Chen, Bolin Ding Zengfeng Huang, and Yaliang Li. Simple and deep graph
   convolutional networks. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.

- Qian Huang, Horace He, Abhay Singh, Ser-Nam Lim, and Austin R Benson. Combining label propagation and simple models out-performs graph neural networks. *arXiv preprint* arXiv:2010.13993, 2020.
- [37] Hengrui Zhang, Qitian Wu, Junchi Yan, David Wipf, and S Yu Philip. From canonical correlation
   analysis to self-supervised graph neural networks. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- Image: The State of State
- 344 [39] Matthias Fey and Jan Eric Lenssen. Fast graph representation learning with pytorch geometric. 345 arXiv preprint arXiv:1903.02428, 2019.

# 346 A Derivation of Loss function of LapDiff

This section provides a derivation of the variational lower bound (ELBO) and the loss function of our proposed model, LapDiff.

#### 349 A.1 Derivation of Evidence Lower Bound (ELBO)

Let  $G_0$  represent a given observed graph data consisting of X and A, denoting node feature matrix and adjacency matrix, respectively. Taking the log-likelihood  $\log p(G_0) = \log p(X_0, A_0)$ , we obtain

$$\log p_{\theta}(X_{0}, A_{0}) = \log \int p_{\theta}(X_{0:T}, A_{0:T}) dX_{1:T} dA_{1:T}$$

$$\geq \int q(X_{1:T}, A_{1:T} | X_{0}, A_{0}) \log \frac{p_{\theta}(X_{0:T}, A_{0:T})}{q(X_{1:T}, A_{1:T} | X_{0}, A_{0})} dG_{1:T}$$
(9)
(by variational posterior and Jensen's inequality)

$$= \mathbb{E}_{q(X_{1:T}, A_{1:T}|X_0, A_0)} \left[ \log \frac{p_{\theta}(X_{0:T}, A_{0:T})}{q(X_{1:T}, A_{1:T}|X_0, A_0)} \right]$$

For simplifying the notation,  $dG_{1:T} = dX_{1:T}dA_{1:T}$  and  $\mathbb{E}_q = \mathbb{E}_{q(X_{1:T},A_{1:T}|X_0,A_0)}$ . The variational lower bound (ELBO) is obtained as follows:

ELBO: = 
$$\mathbb{E}_q \left[ \log \frac{p_{\theta}(X_{0:T}, A_{0:T})}{q(X_{1:T}, A_{1:T} | X_0, A_0)} \right]$$
 (10)

$$= \mathbb{E}_{q} \left[ \log \frac{p_{\theta}(X_{T}, A_{T}) \prod_{t=1}^{T} p_{\theta}(X_{t-1}|X_{t}, A_{t}) \cdot p_{\theta}(A_{t-1}|X_{t}, A_{t})}{\prod_{t=1}^{T} q(X_{t}|X_{t-1}) \cdot q(A_{t}|A_{t-1})} \right]$$
(11)

$$= \mathbb{E}_q \left[ \log p_{\theta}(X_T, A_T) \right]$$

$$+\sum_{t=1}^{T} \log \frac{p_{\theta}(X_{t-1}|X_{t}, A_{t}) \cdot p_{\theta}(A_{t-1}|X_{t}, A_{t})}{q(X_{t}|X_{t-1}) \cdot q(A_{t}|A_{t-1})}$$
(12)

Then, we could separate t=1 and  $t\geq 2$  because t=1 indicates the reconstruction to the given data  $G_0=(X_0,A_0)$ . Next, we flip latent variables in the variational posterior q. We reparametrize  $q(X_t|X_{t-1}),\ q(A_t|A_{t-1})$  to  $q(X_{t-1}|X_t,X_0),\ q(A_{t-1}|A_t,A_0)$ , respectively. Specifically, Markov chain properties formalize  $q(X_t|X_{t-1})=q(X_{t-1}|X_t,X_0)\frac{q(X_t|X_0)}{q(X_{t-1}|X_0)}$  and this holds in A as well.

The second term can be rewritten as

$$\sum_{t=2}^{T} \log \frac{p_{\theta}(X_{t-1}|X_{t}, A_{t}) \cdot p_{\theta}(A_{t-1}|X_{t}, A_{t})}{q(X_{t-1}|X_{t}, X_{0}) \cdot q(A_{t-1}|A_{t}, A_{0})} + \sum_{t=2}^{T} \log \frac{q(X_{t-1}|X_{0}) \cdot q(A_{t-1}|A_{0})}{q(X_{t}|X_{0}) \cdot q(A_{t}|A_{0})}$$

$$= \sum_{t=2}^{T} \left[ \log \frac{p_{\theta}(X_{t-1}|X_{t}, A_{t}) \cdot p_{\theta}(A_{t-1}|X_{t}, A_{t})}{q(X_{t-1}|X_{t}, X_{0}) \cdot q(A_{t-1}|A_{t}, A_{0})} + \log \frac{q(X_{1}|X_{0})}{q(X_{2}|X_{0})} \cdot \frac{q(X_{2}|X_{0})}{q(X_{3}|X_{0})} \cdots \frac{q(X_{T-2}|X_{0})}{q(X_{T-1}|X_{0})} \cdot \frac{q(X_{T-1}|X_{0})}{q(X_{T}|X_{0})} + \log \frac{q(A_{1}|A_{0})}{q(A_{2}|A_{0})} \cdot \frac{q(A_{2}|A_{0})}{q(A_{3}|A_{0})} \cdots \frac{q(A_{T-2}|A_{0})}{q(A_{T-1}|A_{0})} \cdot \frac{q(A_{T-1}|A_{0})}{q(A_{T}|A_{0})} \right].$$

$$(13)$$

Then, the ELBO is derived as follows:

$$ELBO = \mathbb{E}_{q} \left[ \log p_{\theta}(X_{T}, A_{T}) + \sum_{t=2}^{T} \log \frac{p_{\theta}(X_{t-1}|X_{t}, A_{t}) \cdot p_{\theta}(A_{t-1}|X_{t}, A_{t})}{q(X_{t-1}|X_{t}, X_{0}) \cdot q(A_{t-1}|A_{t}, A_{0})} + \log \frac{q(X_{T}|X_{0}) \cdot q(A_{T}|A_{0})}{q(X_{T}|X_{0}) \cdot q(A_{T}|A_{0})} + \log \frac{p_{\theta}(X_{0}|X_{1}, A_{1}) \cdot p_{\theta}(A_{0}|X_{1}, A_{1})}{q(X_{T}|X_{0}) \cdot q(A_{T}|A_{0})} \right]$$

$$= \mathbb{E}_{q} \left[ \log \frac{p_{\theta}(X_{T}, A_{T})}{q(X_{T}|X_{0}) \cdot q(A_{T}|A_{0})} + \sum_{t=2}^{T} \log \frac{p_{\theta}(X_{t-1}|X_{t}, A_{t}) \cdot p_{\theta}(A_{t-1}|X_{t}, A_{t})}{q(X_{t-1}|X_{t}, X_{0}) \cdot q(A_{t-1}|A_{t}, A_{0})} + \log p_{\theta}(X_{0}|X_{1}, A_{1}) \cdot p_{\theta}(A_{0}|X_{1}, A_{1}) \right]$$

$$(15)$$

The first term is not trainable. Since the degree of noise injection in the forward process is fixed and not optimized during training, it can be treated as a constant. The second term is equivalent to the definition of KL divergence, thus the final form of the ELBO of the model is

$$ELBO = \mathbb{E}_{q} \Big[ \log p_{\theta}(X_{0}|X_{1}, A_{1}) + \log p_{\theta}(A_{0}|X_{1}, A_{1}) \\ - \sum_{t \geq 2} KL \left[ q(X_{t-1}|X_{t}, X_{0}) \| p_{\theta}(X_{t-1}|X_{t}, A_{t}) \right] \\ - \sum_{t \geq 2} KL \left[ q(A_{t-1}|A_{t}, A_{0}) \| p_{\theta}(A_{t-1}|X_{t}, A_{t}) \right] \Big]$$

$$(17)$$

#### 363 A.2 Variational Posterior

In our proposed model, the feature-level and structural diffusion processes are designed to be Markov chains, meaning the noisy data at each timestep t depends only on the state at t-1. In the context of the forward diffusion processes implies:

$$q(X_t|X_0, X_1, \dots, X_{t-1}, A_0) = q(X_t|X_{t-1}, \cancel{A_0})$$

$$q(A_t|A_0, A_1, \dots, A_{t-1}) = q(A_t|A_{t-1}),$$
(18)

where the noise source, Laplacian smoothing, is fixed with  $A_0$ . By the definition of Markov chains,

we can remove  $A_0$ , which serves the role of a fixed noise schedule. Also, recall that we defined noise

in the feature-level diffusion process as

$$q(X_t|X_{t-1}) := (I - D^{-1}L)^t X_0 = (I - D^{-1}L)X_{t-1}$$

$$q(X_{1:T}|X_0) = \prod_{t=1}^T q(X_t|X_0) = \prod_{t=1}^T q(X_t|X_{t-1}).$$
(20)

Start from the joint conditional distribution of all  $X_t$  and  $A_t$ , given  $X_0, A_0$ :

$$q(X_{1:T}, A_{1:T}|X_0, A_0) = q(X_1, X_2, \dots, X_T, A_1, A_2, \dots, A_T|X_0, A_0).$$
(21)

By the chain rule, Eq (21) equals

$$q(X_1|X_0)q(A_1|A_0) \cdot q(X_2|X_0, X_1)q(A_2|A_0, A_1) \\ \dots q(X_T|X_0, \dots, X_{T-1})q(A_T|A_0, \dots, A_{T-1}).$$
(22)

By applying the Markov chain property that each  $X_t$ ,  $A_t$  only depends on  $X_{t-1}$ ,  $A_{t-1}$ :

$$q(X_{1:T}, A_{1:T}|X_0, A_0) = \prod_{t=1}^{T} q(X_t|X_{t-1}) \cdot q(A_t|A_{t-1})$$
(23)

#### A.3 Approximation to Laplacian Difference Prediction

For one undirected edges (i, j), let the one-edge KL be

$$p_{ij} = q(A_{t-1,ij} = 1 \mid A_t, A_0), \ \hat{p}_{ij} = p_{\theta}(A_{t-1,ij} = 1 \mid G_t).$$
(24)

The Bernoulli Kullback–Leibler term that appears in the ELBO can be rewritten as

$$D_{KL}(p_{ij} \parallel \hat{p}_{ij}) = p_{ij} \log \frac{p_{ij}}{\hat{p}_{ij}} + (1 - p_{ij}) \log \frac{1 - p_{ij}}{1 - \hat{p}_{ij}}.$$
 (25)

Fix the *target* probability  $p_{ij} \in (0,1)$  and leverage Taylor expansion,  $\varepsilon := \hat{p}_{ij} - p_{ij}$ . Because the

first derivative vanishes at  $\varepsilon = 0$  and the second derivative equals  $\frac{1}{p_{ij}(1-p_{ij})}$ ,

$$D_{KL}(p_{ij} \parallel p_{ij} + \varepsilon) = \frac{\varepsilon^2}{2 p_{ij} (1 - p_{ij})} + \mathcal{O}(\varepsilon^3)$$
 (26)

whenever  $|\varepsilon| < \min\{p_{ij}, 1-p_{ij}\}$ . Thus to second order the KL is proportional to the squared error

379  $(p_{ij} - \hat{p}_{ij})^2$ .

For a random adjacency matrix A with edge–existence probabilities q the expected combinatorial

381 Laplacian is

$$\bar{L}(q) = \operatorname{diag}\left(\sum_{k} q_{1k}, \dots\right) - q. \tag{27}$$

For off-diagonal entries  $(i \neq j)$ :  $\bar{L}(q)_{ij} = -q_{ij}$ . Therefore, we can obtain

$$p_{ij} - \hat{p}_{ij} = -(\bar{L}(p)_{ij} - \bar{L}(\hat{p})_{ij}). \tag{28}$$

Then all edges can be summed over:

$$\sum_{i < j} D_{KL} (p_{ij} \| \hat{p}_{ij}) \stackrel{(2)}{=} \frac{1}{2} \sum_{i < j} \frac{(\bar{L}(p)_{ij} - \bar{L}(\hat{p})_{ij})^2}{p_{ij} (1 - p_{ij})} + \mathcal{O}(\|\bar{L}(p) - \bar{L}(\hat{p})\|_F^3).$$
(29)

Assume every posterior edge probability is clipped away from 0,1:  $p_{ij} \in [\varepsilon,1-\varepsilon]$  for some fixed

385  $0 < \varepsilon < \frac{1}{2}$ . Then

$$p_{ij}(1 - p_{ij}) \geq \varepsilon (1 - \varepsilon)$$

$$\Longrightarrow D_{KL}(p_{ij} \parallel \hat{p}_{ij}) \leq c_{\varepsilon} \left(\bar{L}(p)_{ij} - \bar{L}(\hat{p})_{ij}\right)^{2},$$
(30)

with  $c_{\varepsilon} = \frac{1}{2\varepsilon(1-\varepsilon)}$ . Using the Frobenius norm of a symmetric matrix and (non-overlapping) edges:

$$\mathcal{L}_{\text{edge\_KL}} := \sum_{i < j} D_{KL}(p_{ij} \parallel \hat{p}_{ij}) \leq c_{\varepsilon} \left\| \bar{L}(p) - \bar{L}(\hat{p}) \right\|_{F}^{2}$$
(31)

- Hence, minimizing the Laplacian MSE upper–bounds the exact edge-wise KL. 387
- At training time we see a single Laplacian  $L_{t-1} = D(A_{t-1}) A_{t-1}$  drawn from the posterior
- $q(A_{t-1} \mid A_t, A_0)$ . Its expectation is  $\bar{L}(p)$  and its entry-wise variance is  $\mathcal{O}(\beta_t)$ , the forward step size. 389
- For small  $\beta_t$  we can drop that noise and use 390

$$\bar{L}(p) \approx L_0 - L_{t-1},$$

- because the forward diffusion from  $A_{t-1}$  to  $A_t$  removes each edge with probability  $\beta_t$ , so the 391 **expected** difference to the clean graph is  $L_0 - L_{t-1}$ . 392
- Step 6: Define the practical surrogate loss 393
- Let  $f_{\theta}(X_t, A_t)$  be the GNN output that tries to recover  $(L_0 L_{t-1})$ . Combining (5) and (6): 394

$$\mathcal{L}_{\text{Lap}} := \lambda \left\| f_{\theta}(X_t, A_t) - (L_0 - L_{t-1}) \right\|_2^2, \quad \lambda \ge c_{\varepsilon}.$$

- Because it is an **upper bound** of the true KL term in the ELBO, driving  $\mathcal{L}_{Lap}$  to zero also drives 395
- the Bernoulli KL towards zero. Empirically the continuous, degree-coupled signal contained in the 396
- Laplacian makes gradients less noisy and optimisation easier than edge-wise cross-entropy. 397
- Therefore the Laplacian-MSE loss you observe to work well is mathematically a second-order 398
- surrogate—and an upper bound—for the exact ELBO term that compares adjacency posteriors. 399
- Minimising it is guaranteed to minimise, up to a bounded factor, the information-theoretic quantity 400
- we truly care about. 401

### A.4 Training of LapDiff with denoising network

The training procedure of LapDiff is described as Algorithm 1.

#### **Algorithm 1** Training LapDiff

```
Input: Large graph G=(X,A)
```

- 1: for k=1 to T do
- $X_{k-1} \leftarrow (I D^{-1}L)^{k-1}X$

- $X_{k} \leftarrow (I D^{-1}L)^{k}X$   $A_{k} \sim \mathcal{B}(A_{k-1}; p_{k-1})$   $\theta_{k-1} \leftarrow \theta_{k} \eta \nabla_{\theta}[\mathcal{L}_{\text{feat}}(f_{\theta_{k}}(X_{k}, A_{k}), X_{k-1}) + \mathcal{L}_{\text{struct}}(f_{\theta_{k}}(X_{k}, A_{k}), A_{k-1})]$

403

404

7:  $\theta \leftarrow \eta \nabla_{\theta} \mathcal{L}_{\text{task}}(f_{\theta_0}(X, A), Y_{\text{task}})$ 

В **Hyperparameter Sensitivity Analysis** 

- We analyze LapDiff to demonstrate how hyperparameters affect the performance of LapDiff. We 405
- conduct experiments with 4 hyperparameters in LapDiff loss function,  $\mathcal{L}_{\text{LapDiff}}$ .  $\beta_t, \beta_1, \gamma$  and  $\lambda$  is 406
- weighting hyperparameters for  $\mathcal{L}_{feat}$ ,  $\mathcal{L}_{feat\text{-recon}}$ ,  $\mathcal{L}_{Lap}$ , and  $\mathcal{L}_{recon}$ , respectively. We measure Hits@50 by changing one hyperparameter while the rest of the hyperparameters are fixed to the best value. 407
- 408
- The result (Figure. 3) demonstrates that LapDiff is fairly robust to hyperparameters that weight the 409
- components of LapDiff loss  $\mathcal{L}_{LapDiff}$ . Accordingly, hyperparameters affect the performance of LapDiff 410
- slightly. Consequently, we can conclude that the performances of LapDiff are fairly consistent under 411
- various hyperparameter sets. 412

#### **Experimental Setting Details** 413

Datasets. To validate our models, we utilize Open Graph Benchmark (OGB) dataset for link prediction tasks and node classification tasks [8]. We use four OGB link property datasets for

# Hyperparameter Senstivity Analysis

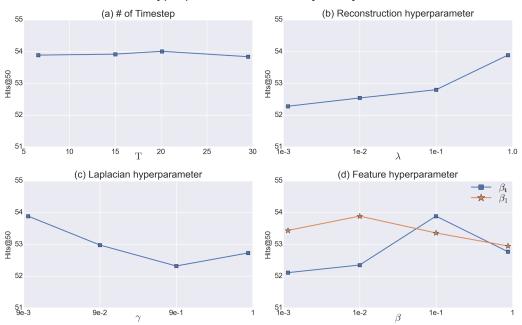


Figure 3: Visualization of hyperparameter sensitivity analysis on OGB-Collab.

link prediction tasks: OGB-PPA, OGB-Collab, OGB-DDI, and OGB-Citation2. OGB-PPA is an undirected and unweighted graph representing protein association. Nodes are proteins from different specifies and edges mean biological associations. Each node feature is a one-hot vector indicating the species to which the protein belongs. OGB-Collab is an undirected graph, which represents a collaboration network where edges denote collaborations between authors. OGB-DDI is an undirected, unweighted graph that contains drug-drug interactions, with edges indicating interactions such as combined effects. Note that this dataset lacks node features. OGB-Citation2 is a citation network graph with direction. Each node in the graph corresponds to a paper, and a directed edge indicates that one paper cites another. Both OGB-Citation2 and OGB-Collab include node features obtained from embedding models. For node classification tasks, we use three benchmark datasets: OGB-Arxiv, OGB-Products, and PubMed.

**Evaluation.** According to the evaluation protocol of OGB, we evaluate our model with Hits@K metric and Mean reciprocal rank (MRR) in link prediction. Hits@K is based on ranking positive test edges against randomly sampled negative edges. The ranking performance is measured by the ratio of positive test edges ranked at or above the K-th position. In OGB-PPA, the K-th position is set to 100, while for OGB-Collab and OGB-DDI, it is set to 50 and 20, respectively. The evaluation metric for OGB-Citation2 is MRR. It calculates the reciprocal rank of the true edges within the pool of negative candidates for each source node and then averages these values across all source nodes. To further demonstrate the ability to learn compendious underlying structures in node classification, we constrain a fixed k-nodes setting by vastly reducing the number of nodes per label in train sets. Under this setting, accuracy measures the performance on OGB-Arxiv, OGB-Products, and PubMed.

**Baselines.** For baselines on link prediction, we include prevalent GNN-based models: GCN [13], GAT [14], GraphSAGE [26], JKNet [27], Variational Graph Autoencoder [28], SEAL [5], Neo-GNNs [6], and DDM. Note that SEAL extract enclosing subgraph to utilize in link prediction. Additionally, three link prediction heuristics [29–31], Matrix factorization [32], and Multi-layer perceptron [33] are included in baselines. Baseline models for semi-supervised node classification include GCN, GAT, APPNP [34], GCNII [35], and Correct&Smooth (C&S) [36]. We also compare LapDiff with self-supervised graph learning methods, CCA-SSG [37], and generative method GraphMAE2 [38].

**Implementation Details.** We implemented link prediction heuristics, such as Common Neigh-444 bor(CN), Adamic Adar(AA), and Resource Allocation(RA), based on the paper [29–31]. For GCN, 445 GraphSAGE, GAT, JKNet, APPNP, GCNII, and MLP we used the implementation in PyTorch Geo-446 metric [39], and for SEAL and C&S, we used the implementation from the official repository. We 447 trained LapDiff with a 2-layer LapDiff encoder with latent Laplacian parameters and 3-layer MLP de-448 coder for OGB-Collab, OGB-DDI, OGB-PPA, and OGB-Citation2. For OGB-Arxiv, OGB-Products, 449 and PubMed, we used 3-layer LapDiff encoder and 3-layer MLP decoder. For the link prediction task, 450 we shared the last layer of the decoder as a predictor, and for the node classification task, we utilized 451 1 layer MLP as a classifier. Also, we set the diffusion state to 10 for OGB-Collab, OGB-DDI, and 452 OGB-PPA, and 3 for OGB-Citation2 due to the dataset's memory issue. For a fair comparison, we 453 reported performances of all baselines and LapDiff as the mean and the standard deviation obtained 454 from 10 independent runs with fixed random seed  $0, \dots, 9$ . To simulate a real-world scenario, we did 455 not use validation edges as input in OGB-Collab. The experiments are conducted on A100(40GB) 456 and A40(48GB).

# 458 D Computational Complexity

#### 459 Notation

- N = |V| : number of nodes ;  $E = |\mathcal{E}|$  : number of edges.
- d: input feature dimension; h: hidden dimension of the denoiser  $f_{\theta}$ .
- T: number of diffusion steps ( $T \ll N$  in practice).
- Diffusion Process Each Laplacian-smoothing step multiplies  $X_{t-1}$  by  $(I D^{-1}L)$ . For a sparse adjacency (E = O(N)) in large graphs), this costs

$$\mathcal{O}(E d)$$
 per step  $\implies \mathcal{O}(T E d)$  overall. (32)

Sampling edges (edge removal) is  $\mathcal{O}(E)$  per step; overall  $\mathcal{O}(TE)$ . No matrix materialization beyond the original sparse A is required. **Total (forward).** 

$$\mathcal{O}(TE(d+1)) \approx \mathcal{O}(TEd)$$

Reverse (Denoising) Process The denoiser  $f_{\theta}$  is a 2-layer MLP encoder + 3-layer MLP decoder. Given sparse adjacency, each call costs  $\mathcal{O}(E\,h+N\,h^2)$ . Invoked once per diffusion step, the reverse chain costs

$$\mathcal{O}(T(Eh + Nh^2)).$$

#### 470 D.1 Overall Time Complexity

471 Combining forward and reverse,

$$\mathcal{O}\!\!\left(T\left[E\left(d+h\right)+N\,h^2\right]\right)$$

With typical settings ( $h \approx d$ ,  $E \gg N$ ,  $T \le 10$ ), the leading term is T E d, comparable to a *single* pass of a conventional L-layer GNN when  $L \approx T$ .

#### 474 D.2 Memory Complexity

- **Parameters.** Two MLPs of width  $h: \mathcal{O}(h^2)$ , independent of T.
- Activations. We store  $X_t$  and (sparse)  $A_t$  for the current step only, so in-memory activations scale as:

$$\mathcal{O}(Nd+E) + \mathcal{O}(Nh) = \mathcal{O}(N(d+h)+E).$$

- Comparison to L-layer GNN. A standard L-layer message-passing GNN stores L intermediate node embeddings, yielding  $\mathcal{O}(L\,N\,h)$  memory. LapDiff keeps a single embedding per step and can recompute forward activations (checkpointing), requiring at most  $\mathcal{O}(N\,h)$ —often smaller than a deep GNN when L>T.
- Scalability. With T chosen  $\leq 10$ , LapDiff's runtime is on par with—or lower than—deep GNNs that rely on  $L \geq 10$  layers. The memory footprint remains modest due to sparse storage and step-wise recomputation.

## E License of the assets

Our source code is based on PyTorch which was released under Berkeley Software Distribution (BSD)
License. We implement GNN-based baselines using PyTorch Geometric, a deep learning framework
licensed under MIT. Additionally, we implement SEAL <sup>1</sup> and GraphMAE <sup>2</sup> from the official GitHub
repository under MIT License. Both BSD license and MIT license can be used or redistributed under
stipulated conditions. Moreover, we conduct experiments on four benchmark datasets from Open
Graph Benchmark (OGB). OGB is released under MIT License. We visualize significant results by
using Matplotlib where the license is based on Python Software Foundation (PSF) license.

# 493 F Broader Impact

LapDiff aims to capture latent factors for graph representation learning. It provides a strong foundation for future models that aim to understand the rich information of relational data in graphs. Our model would not only be capable of discerning the latent structures within graph data but also adept at applying this knowledge across a broad spectrum of applications. With its Laplacian smoothing noise which is structure-aware, LapDiff contributes to the ongoing discourse on data privacy. By generating representations that respect the underlying structure of data without compromising individual privacy, LapDiff aligns with the ethical use of data in AI. Also, our model's versatility suggests broad applicability beyond traditional domains, offering potential breakthroughs in any field that benefits from understanding complex networks, including neuroscience, epidemiology, and environmental studies. However, LapDiff needs to be used carefully for graph representation learning tasks such as link prediction or node classification in social networks where privacy and anonymity are important.

https://github.com/facebookresearch/SEAL\_OGB

<sup>&</sup>lt;sup>2</sup>https://github.com/THUDM/GraphMAE2

# NeurIPS Paper Checklist

#### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction explicitly state the development of the LapDiff with Laplacian-based diffusion, and the claims match the theoretical derivations and experiments reported (see Section 1 and 2)

#### Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the
  contributions made in the paper and important assumptions and limitations. A No or
  NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals
  are not attained by the paper.

#### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Limitations are mentioned regarding scalability and assumptions in the diffusion process. (See Section 5)

#### Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

#### 3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

#### Answer: [Yes]

Justification: The derivations of the ELBO and loss function are presented in detail in the Appendix. Assumptions about Markov properties and Laplacian smoothing are clearly stated.

#### Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if
  they appear in the supplemental material, the authors are encouraged to provide a short
  proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Experimental settings, datasets (OGB), and baselines are described in the main text and Appendix. Hyperparameters and complexity analysis are also provided in Appendix. We provide codes as Supplementary.

#### Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
- (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The paper states that implementations are based on PyTorch and PyTorch Geometric with clear licenses (Appendix), but the code is released only to reviewers at submission time.

#### Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be
  possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not
  including code, unless this is central to the contribution (e.g., for a new open-source
  benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how
  to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new
  proposed method and baselines. If only a subset of experiments are reproducible, they
  should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

#### 6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Training procedures, loss components, and hyperparameter sensitivity analyses are described in Appendix. Optimizer details are also outlined in Algorithm 1.

#### Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
  material.

### 7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Guidelines:

Justification: The results are provided with mean and standard deviation under 10 fixed runs.

# • The answer NA means that the paper does not include experiments.

• The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

#### 8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702

703

704

705

706

707 708

709

710

711

712

713

714

Justification: Computational and memory complexity analyses, GPU are detailed in Appendix.

#### Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

#### 9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: The datasets (OGB) and libraries used are all properly licensed. No ethical violations are apparent in the methodology.

#### Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a
  deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

#### 10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: In Appendix, we discussed privacy implications and broad applicability. Negative societal impacts such as misuse are only briefly mentioned, but at least acknowledged.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
  - If the authors answer NA or No, they should explain why their work has no societal
    impact or why the paper does not address societal impact.
  - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
  - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
  - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
  - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

#### 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The model is not released as a pretrained asset with high misuse risk. Only standard datasets and code bases are used.

#### Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
  necessary safeguards to allow for controlled use of the model, for example by requiring
  that users adhere to usage guidelines or restrictions to access the model or implementing
  safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

#### 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Appendix E provides details about BSD/MIT/PSF licenses of all used assets and repositories.

#### Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the
  package should be provided. For popular datasets, paperswithcode.com/datasets
  has curated licenses for some datasets. Their licensing guide can help determine the
  license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

#### 13. New assets

768

769

770

771

772

773

774

775

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792 793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810

812

813

815

816

817

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new datasets or assets are introduced.

#### Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

#### 14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve human subjects or crowdsourcing.

#### Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

# 15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The research does not involve human subjects.

#### Guidelines:

 The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

#### 16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: LLMs were not used as part of the methodology.

#### Guidelines

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.