

# 000 PRINCIPLED LATENT DIFFUSION FOR GRAPHS VIA LAPLA- 001 CIAN AUTOENCODERS 002 003 004

005 **Anonymous authors**

006 Paper under double-blind review

## 007 ABSTRACT 008 009

011 Graph diffusion models achieve state-of-the-art performance in graph generation but  
012 suffer from quadratic complexity in the number of nodes—and much of their capacity is  
013 wasted modeling the absence of edges in sparse graphs. Inspired by latent diffusion in  
014 other modalities, a natural idea is to compress graphs into a low-dimensional latent  
015 space and perform diffusion there. However, unlike images or text, graph generation  
016 requires nearly lossless reconstruction, as even a single error in decoding an adjacency  
017 matrix can render the entire sample invalid. This challenge has remained largely  
018 unaddressed. We propose **LG-Flow**, a latent graph diffusion framework that directly  
019 overcomes these obstacles. A permutation-equivariant autoencoder maps each node into  
020 a fixed-dimensional embedding from which the full adjacency is provably recoverable,  
021 enabling near-lossless reconstruction for both undirected graphs and DAGs. The size of  
022 this latent representation scales linearly with the number of nodes, eliminating the  
023 quadratic bottleneck and making it feasible to train larger and more expressive models.  
024 In this latent space, we train a Diffusion Transformer with flow matching, enabling  
025 efficient and expressive graph generation. Our approach achieves competitive results  
026 against state-of-the-art graph diffusion models, while achieving up to 1000 $\times$  speed-up.

## 027 1 INTRODUCTION 028

029 Generative modeling of graphs has advanced rapidly in recent years, with diffusion models achieving state-  
030 of-the-art performance across various domains, including molecules (Igashov et al., 2024), combinatorial  
031 optimization (Sun and Yang, 2023), and neural architecture search. However, most existing graph  
032 diffusion models operate directly in graph space, which entails quadratic computational complexity in the  
033 number of nodes (Vignac et al., 2023). These models typically learn both node and edge representations,  
034 since predicting edges solely from node features has proven insufficient across benchmarks (Qin et al.,  
035 2023). As a result, scalability remains limited. Quadratic complexity also leads to disproportionate  
036 memory consumption relative to the number of parameters, which restricts model capacity compared to  
037 image-based diffusion. In addition, attempts to adapt scalable auto-regressive sequence models to graphs  
038 break permutation equivariance when graphs are linearized, e.g., Chen et al. (2025), which limits their  
039 applicability.

040 In contrast, latent diffusion has advanced image generation by shifting the generative process into a  
041 compressed latent space (Esser et al., 2024; Rombach et al., 2022). A VAE maps data into low-dimensional  
042 representations, where Gaussian diffusion or flow models operate more efficiently, avoiding wasted  
043 capacity on imperceptible details. This yields efficient training, faster inference, and supports high-capacity  
044 architectures.

045 Adapting this principle to graphs is highly appealing but presents unique challenges. Unlike in images,  
046 where minor decoding errors are visually negligible, a single mistake in reconstructing an adjacency matrix  
047 can destroy structural validity—for instance, by deleting a chemical bond or misclassifying its type.  
048 Compression must therefore not come at the expense of reconstruction fidelity. The decoder must achieve  
049 near-lossless reconstruction to guarantee valid samples. Unfortunately, graph autoencoders are rarely  
050 evaluated for reconstruction quality (Kipf and Welling, 2016; Lee and Min, 2022; Li et al., 2020), with the  
051 notable exception of Boget et al. (2024).

052 Despite these obstacles, the latent diffusion paradigm is particularly promising for graphs. Most real-world  
053 graphs are sparse (see Table 1), which means diffusion models trained in graph space devote most  
of their capacity to modeling the absence of edges. A latent approach can alleviate this inefficiency,

reduce quadratic complexity, and enable the training of larger and more expressive generative models. This motivates designing latent graph representations whose size—understood as the total number of dimensions—scales linearly with the number of nodes, for example, by assigning each node a fixed-dimensional latent vector.

Latent diffusion offers modularity and reusability: instead of designing graph-specific diffusion models, one can employ generic denoising architectures, reusing methods and optimization techniques from the thriving image generation community such as the *Diffusion Transformer* (DiT) (Peebles and Xie, 2023). Once the encoder is trained, the diffusion stage is uniform across modalities (see, e.g., Joshi et al. (2025)). In the graph domain, this approach unifies undirected graph (Qin et al., 2025; Vignac et al., 2023) and *directed acyclic graph* (DAG) generation (Carballo-Castro et al., 2025; Li et al., 2025a), which are typically treated separately. The latter is especially relevant to highly impactful applications like chip design and circuit synthesis (Mishchenko and Miyasaka, 2023; Li et al., 2025b; Phothilimthana et al., 2023b). This modular view shifts the focus to designing strong, cheaper-to-train graph autoencoders, which rely on linear-complexity (*Message-Passing*) *Graph Neural Networks* (MPNNs) (Gilmer et al., 2017; Scarselli et al., 2009).

From these considerations, we derive four requirements for a principled latent graph diffusion model. **(R1)** The autoencoder must provide provable reconstruction guarantees that ensure the structural validity of decoded graphs. **(R2)** The size of a graphs’ latent representations should scale linearly with the number of nodes, enabling scalability to large graphs. **(R3)** Graph-specific inductive biases should be encapsulated in the autoencoder, allowing the diffusion model itself to remain generic. **(R4)** The overall framework should maintain competitive generative performance while avoiding the quadratic complexity of graph-space diffusion.

However, as argued above, existing latent graph diffusion models fail to satisfy all four requirements. Most lack reconstruction guarantees (Yang et al., 2024b; Fu et al., 2024), while others forfeit efficiency by relying on ad-hoc transformer designs (Nguyen et al., 2024) or a quadratic number of edge tokens (Zhou et al., 2024).

**Present work** Hence, to address the above requirements **R1-R4**, we propose a latent graph diffusion framework. Concretely, we

1. derive the *Laplacian Graph Variational Autoencoder* (LG-VAE), a principled permutation-equivariant graph autoencoder with provable reconstruction guarantees, extending theory from positional encodings in *Graph Transformers* (GTs) to both undirected graphs and DAGs.
2. Our autoencoder maps each node to a fixed-dimensional embedding from which the adjacency matrix is recoverable, yielding compact latent representations that scale linearly with graph size.
3. In this latent space, we apply flow matching with a DiT architecture, achieving competitive performance across benchmarks while maintaining scalability and architectural simplicity. Our approach (**LG-Flow**) delivers competitive results on both synthetic and real-world benchmarks, while achieving speed-ups ranging from  $10\times$  to  $1000\times$ .

*Overall, our **LG-Flow** enables diffusion-based graph generative models to scale to larger graph sizes with significantly improved inference efficiency, paving the way for their application to domains previously out of reach.*

## 2 RELATED WORK

Here, we discuss related work. A more detailed discussion is provided in Appendix A, including an overview of prior graph diffusion models.

**Graph generation** Graph generation methods are typically categorized into two main groups. *Autoregressive models* build graphs by progressively adding nodes and edges (You et al., 2018; Jang et al., 2024). Their main advantage is computational efficiency, since they do not need to materialize the full adjacency matrix. However, they depend on an arbitrary node ordering to transform the graph into a

Table 1: Average density of training graphs across common benchmarks.

Dataset	Density
Planar	0.09
Tree	0.03
Ego	0.04
Moses	0.11
Guacamol	0.09
TPU Tile	0.06

108 sequence, either learned or defined through complex algorithms, which breaks permutation equivariance. In  
 109 contrast, *one-shot models* generate the entire graph at once, thereby preserving permutation equivariance.  
 110 Many approaches have been explored in this class, starting with GVAEs (Kipf and Welling, 2016) and  
 111 GANs (Martinkus et al., 2022). Following their success in other modalities such as images, diffusion  
 112 models (Ho et al., 2020; Song et al., 2021; Lipman et al., 2023) quickly emerged as the dominant approach  
 113 for graph generation.

114  
 115 **Latent diffusion models** Latent diffusion models (LDMs) (Rombach et al., 2022) generate data by  
 116 applying the diffusion process in a compressed latent space rather than pixel space, typically using a VAE  
 117 for encoding and decoding. This reduces computation while preserving detail, and underlies models such  
 118 as Stable Diffusion. More recently, DiTs (Peebles and Xie, 2023) emerged as effective denoisers across  
 119 modalities (Esser et al., 2024; Joshi et al., 2025). For graphs, however, latent diffusion has not matched the  
 120 performance of discrete diffusion models (Fu et al., 2024; Nguyen et al., 2024; Yang et al., 2024a). Only  
 121 Nguyen et al. (2024) evaluate autoencoder reconstruction, and only on small QM9 graphs, with an ad-hoc  
 122 transformer that materializes the full attention matrix, reducing efficiency. Similarly, Zhou et al. (2024)  
 123 achieve competitive results on MOSES but rely on  $n^2$  edge tokens from the full adjacency matrix, also  
 124 limiting efficiency. Other works target only molecule generation (Ketata et al., 2025; Wang et al., 2024;  
 125 Bian et al., 2024) or tasks like link prediction (Fu et al., 2024) and regression (Zhou et al., 2024). In  
 126 contrast, our approach ensures the adjacency matrix is provably recoverable from the latent space.  
 127  
 128

### 3 BACKGROUND

131 In this section, we introduce theoretical tools used throughout the paper, including graphs, adjacency-  
 132 identifying positional encodings, and Laplacian positional encodings. For completeness, the formal  
 133 introduction of flow matching is deferred to Appendix B. Note that since flow matching and Gaussian  
 134 diffusion are essentially equivalent, we use the two terms interchangeably (Gao et al., 2024).

135  
 136 **Notations** Let  $\mathbb{N} := \{1, 2, \dots\}$  and for  $n \in \mathbb{N}$  define  $[n] := \{1, \dots, n\} \subset \mathbb{N}$ . We denote by  $\mathbb{R}^{+*}$  the  
 137 set of strictly positive real numbers. We consider *node- and edge-labeled undirected graphs*  $G$ , i.e.,  
 138 tuples  $(V(G), E(G), \ell_x, \ell_e)$  with up to  $n$  nodes and  $m$  edges, without self-loops or isolated nodes.  
 139 Here,  $V(G)$  is the set of nodes,  $E(G)$  the set of edges,  $\ell_x: V(G) \rightarrow [a]$ , for  $a \in \mathbb{N}$ , assigns each  
 140 node one of  $a$  discrete labels, and  $\ell_e: E(G) \rightarrow [b]$ , for  $b \in \mathbb{N}$ , assigns each edge one of  $b$  discrete  
 141 labels. An undirected edge  $\{u, v\} \in E(G)$  is written either  $(u, v)$  or  $(v, u)$ . We also consider *node- and edge-labeled directed acyclic graphs* (DAGs), defined analogously except that  $(u, v)$  denotes a  
 142 directed edge from  $u$  to  $v$  and the graph contains no directed cycles. Throughout, we assume an arbitrary  
 143 but fixed ordering of nodes so that each node corresponds to a number in  $[n]$ . The adjacency matrix  
 144 of  $G$  is denoted  $\mathbf{A}(G) \in \{0, 1\}^{n \times n}$ , where  $\mathbf{A}(G)_{ij} = 1$  if and only if nodes  $i$  and  $j$  are connected.  
 145 We construct a node feature vector  $\mathbf{X} \in [a]^n$  consistent with  $\ell_x$ , so that  $X_i = X_j$  if and only if  
 146  $\ell_x(i) = \ell_x(j)$ . Similarly, we define an edge feature matrix  $\mathbf{E} \in [b]^m$  consistent with  $\mathbf{A}$  and  $\ell_e$ , so that  
 147  $E_{ij} = E_{kl}$  if and only if  $\ell_e((i, j)) = \ell_e((k, l))$ . For any matrix  $\mathbf{M}$ , we denote by  $\mathbf{M}_i$  its  $i$ -th row. For a  
 148 real-valued matrix  $\mathbf{M} \in \mathbb{R}^{n \times n}$ ,  $\mathbf{M}^\top$  denotes its transpose, and for a complex matrix  $\mathbf{N}$ ,  $\mathbf{N}^*$  denotes its  
 149 Hermitian transpose. For two binary matrices  $\mathbf{P}$  and  $\mathbf{Q}$ , we denote their logical disjunction  $\mathbf{P} \vee \mathbf{Q}$  where  
 150  $(\mathbf{P} \vee \mathbf{Q})_{ij} = 1$  if  $\mathbf{P}_{ij} = 1$  or  $\mathbf{Q}_{ij} = 1$ , and 0 otherwise. Let  $\mathbf{u}$  and  $\mathbf{v}$  be two vectors;  $\mathbf{u} \odot \mathbf{v}$  denotes the  
 151 element-wise product between them. Finally, the *graph Laplacian* is  $\mathbf{L} := \mathbf{D} - \mathbf{A}$ , where  $\mathbf{D} \in \mathbb{N}^{n \times n}$  is  
 152 the degree matrix. Its eigenvalue decomposition is  $\mathbf{L} := \mathbf{U} \Lambda \mathbf{U}^\top$ , where  $\Lambda := (\lambda_1, \dots, \lambda_n)$  is the vector  
 153 of eigenvalues (possibly repeated) and  $\mathbf{U} \in \mathbb{R}^{n \times n}$  the matrix of eigenvectors, with the  $i$ -th column  $\mathbf{U}_i^\top$   
 154 corresponding to eigenvalue  $\lambda_i$ . We denote by  $\mathbf{I}_d$  the  $d$ -dimensional identity matrix.  
 155  
 156

#### 3.1 ADJACENCY-IDENTIFYING POSITIONAL ENCODINGS

157 The design of our autoencoder extends *adjacency-identifying positional encodings*, introduced in the  
 158 literature on positional encodings for GTs (Müller and Morris, 2024). Intuitively, such encodings guarantee  
 159 that the adjacency matrix of an *undirected* graph can be reconstructed by the attention mechanism from  
 160 node embeddings, provided the positional encoding is chosen appropriately.

162 Formally, let  $G$  be a directed graph  $(V(G), E(G))$  with  $n$  nodes, and let  $\mathbf{P} \in \mathbb{R}^{n \times d}$  denote a matrix of  
 163  $d$ -dimensional node embeddings. Define

$$164 \quad 165 \quad \tilde{\mathbf{P}} := \frac{1}{\sqrt{d}} \mathbf{P} \mathbf{W}^Q (\mathbf{P} \mathbf{W}^K)^\top, \quad (1)$$

166 where  $\mathbf{W}^Q, \mathbf{W}^K \in \mathbb{R}^{d \times d}$ . We say that  $\mathbf{P}$  is *adjacency-identifying* if there exist  $\mathbf{W}^Q, \mathbf{W}^K$  such that for  
 167 every row  $i$  and column  $j$ ,  $\tilde{\mathbf{P}}_{ij} = \max_k \tilde{\mathbf{P}}_{ik} \iff A(G)_{ij} = 1$ , where  $A(G)$  denotes the adjacency  
 168 matrix of  $G$ . Note that in connected graphs, each row has between one and  $(n - 1)$  maxima. We further say  
 169 that a matrix  $\mathbf{Q} \in \mathbb{R}^{n \times d}$  is *sufficiently adjacency-identifying* if it approximates an adjacency-identifying  
 170 matrix  $\mathbf{P}$  to arbitrary precision, i.e.,  $\|\mathbf{P} - \mathbf{Q}\|_F < \epsilon$ , for all  $\epsilon > 0$ .  
 171

172 **LPE** The *Laplacian Positional Encoding* (LPE), introduced in Dwivedi et al. (2022); Müller and Morris  
 173 (2024), is an example of a sufficiently adjacency-identifying positional encoding, constructed as follows.  
 174 Let  $\epsilon \in \mathbb{R}^l$  be a learnable vector initialized to zero. Let  $\phi: \mathbb{R}^2 \rightarrow \mathbb{R}^d$  denote a feedforward network,  
 175 **applied row-wise**, and let  $\rho: \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^d$  denote a permutation-equivariant neural network applied  
 176 row-wise. For an eigenvector matrix  $\mathbf{U}$  of the graph Laplacian and its corresponding eigenvalues  $\Lambda$ , we  
 177 define

$$178 \quad \phi(\mathbf{U}, \Lambda) := \left[ \phi(\mathbf{U}_1 \parallel \Lambda + \epsilon), \dots, \phi(\mathbf{U}_n \parallel \Lambda + \epsilon) \right] \in \mathbb{R}^{n \times n \times d}, \quad (2)$$

179 where  $\mathbf{U}_i \parallel \Lambda + \epsilon \in \mathbb{R}^{n \times 2}$  denotes row-wise concatenation. The network  $\phi$  is applied on the last  
 180 dimension, which yields  $\phi(\mathbf{U}_i \parallel \Lambda + \epsilon) \in \mathbb{R}^{n \times d}$ , for each  $i \in [n]$ . Stacking these outputs over all nodes  
 181 produces  $\phi(\mathbf{U}, \Lambda)$ .

182 The LPE is then obtained by aggregating contributions across all eigenvectors and eigenvalues. Specifically,

$$184 \quad 185 \quad \text{LPE}(\mathbf{U}, \Lambda) := \rho \left( \sum_{i=1}^n \phi(\mathbf{U}, \Lambda)_i^\top \right), \quad (3)$$

186 where  $\phi(\mathbf{U}, \Lambda)_i^\top \in \mathbb{R}^{n \times d}$  denotes the  $i$ -th column of  $\phi(\mathbf{U}, \Lambda) \in \mathbb{R}^{n \times n \times d}$ . Intuitively, this means that  
 187 for each node we combine the information from all Laplacian (linearly independent) eigenvectors and  
 188 eigenvalues, and then apply a DeepSet (Zaheer et al., 2017) over the (multi)set  $\{\phi(\mathbf{U}_i \parallel \Lambda + \epsilon)\}_{i=1}^n$ .  
 189

## 190 4 PRINCIPLED LATENT DIFFUSION FOR GRAPHS

192 We now present our latent diffusion framework for graphs. The core idea is to design an autoencoder with  
 193 provable reconstruction guarantees (see Figure 1 for an overview), covering both undirected graphs and  
 194 DAGs, and train a diffusion model in the resulting latent space. An overview of the whole framework is  
 195 provided in Figure 2.  
 196

### 197 4.1 LAPLACIAN GRAPH VARIATIONAL AUTOENCODER

199 In latent diffusion models, the expressive power of the decoder sets an upper bound on the quality of  
 200 generated samples. For graphs, even a single incorrect edge can invalidate the entire structure, e.g., by  
 201 breaking chemical validity in molecules, so strong reconstruction guarantees are essential. However, most  
 202 existing graph autoencoders are not benchmarked for reconstruction accuracy, raising concerns about their  
 203 reliability (Fu et al., 2024; Nguyen et al., 2024).

204 We address this by building on LPEs; see Section 3.1. In fact, we require node embeddings that are  
 205 expressive enough to preserve all structural information, so that the adjacency matrix can be provably  
 206 recovered from them. The LPE is a natural choice: by construction, it yields embeddings that are *sufficiently*  
 207 *adjacency-identifying*, i.e., they preserve all information necessary for reconstructing the adjacency matrix.  
 208 This property also provides a natural recipe for the decoder: adjacency can be recovered through a bilinear  
 209 layer (as in Equation (1)), followed by a row-wise argmax. This design satisfies two of our requirements:  
 210 **(R1)** provable reconstruction guarantees and **(R2)** graph latent representations size that scales linearly with  
 211 the number of nodes.

212 Building on this foundation, we introduce the *Laplacian Graph Variational Autoencoder* (LG-VAE).  
 213 Following the VAE framework, the encoder  $\mathcal{E}$  maps a graph  $G$  with node features  $\mathbf{X} \in \mathbb{R}^n$ , edge features  
 214  $\mathbf{E} \in \mathbb{R}^m$ , and Laplacian  $\mathbf{L} = \mathbf{U} \Lambda \mathbf{U}^\top$  to the parameters of a Gaussian posterior with mean  $\mu_Z$  and  
 215 log-variance  $\log \sigma_Z^2$ . Latents  $\mathbf{Z} \in \mathbb{R}^{n \times d}$  are then sampled using the reparameterization trick and passed  
 to the decoder  $\mathcal{D}$ , which reconstructs  $\hat{G} = \mathcal{D}(\mathbf{Z})$ ; see Figure 1 for an overview.

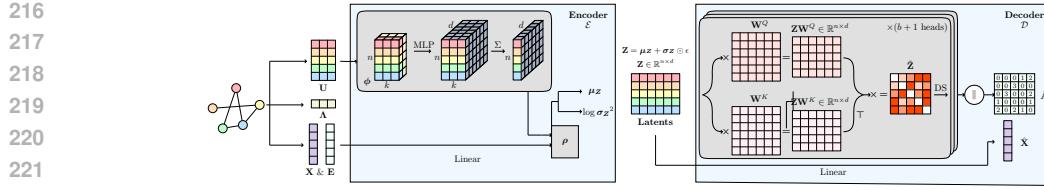


Figure 1: Overview of the LG-VAE. The encoder  $\mathcal{E}$  (left) encodes structure via  $\phi$  and node/edge labels, then aggregates them with  $\rho$ . Latents  $\mathbf{Z}$  are sampled using the reparameterization trick and passed to the decoder  $\mathcal{D}$  (right). Node labels are decoded directly, while adjacency is decoded by (a) computing scores via bilinear dot products for  $b + 1$  heads, (b) processing scores with a row-wise DeepSet, and (c) concatenating outputs across heads. The final adjacency  $\hat{\mathbf{A}}$  is obtained with an argmax.

**Encoder** The encoder extends the LPE to incorporate node and edge labels. We compute

$$\mathbf{H}_\phi := \sum_{i=1}^n \phi(\mathbf{U}, \mathbf{\Lambda})_i^\top \in \mathbb{R}^{n \times d}, \quad \mathbf{H}_X := \mathbf{X} \mathbf{W}^X + \mathbf{b}^X \in \mathbb{R}^{n \times d}, \quad \mathbf{H}_E := \mathbf{E} \mathbf{W}^E + \mathbf{b}^E \in \mathbb{R}^{m \times d},$$

where  $\mathbf{W}^X, \mathbf{W}^E$  are projection matrices, and  $\mathbf{b}^X, \mathbf{b}^E$  the associated bias, and feed them into  $\rho$ , yielding

$$\mu_{\mathbf{Z}}, \log \sigma_{\mathbf{Z}}^2 := \mathcal{E}(G, \mathbf{X}, \mathbf{E}, \mathbf{U}, \mathbf{\Lambda}) = \rho(\mathbf{H}_X + \mathbf{H}_\phi, \mathbf{H}_E). \quad (4)$$

Here  $\mathbf{H}_\phi$  is treated as an additional node feature, and  $\rho$  can be any permutation-equivariant network. In our experiments, we use GIN (Xu et al., 2019) and GINE (Hu et al., 2020). A latent sample is then obtained as

$$\mathbf{Z} = \mu_{\mathbf{Z}} + \sigma_{\mathbf{Z}} \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}_d).$$

Note that, as detailed in Appendix D.1, and depicted in Figure 1, we only use the eigenvectors associated to the  $k$  lowest eigenvalues, as using the full matrix eigenvectors might be impractical for large graphs.

**Decoder** The decoder reconstructs both node and edge labels as well as the adjacency matrix. Node labels are predicted as

$$\hat{\mathbf{X}} := \mathcal{D}_x(\mathbf{Z}) = \text{softmax}(\mathbf{Z} \mathbf{W}^{D_x} + \mathbf{b}^{D_x}), \quad \hat{\mathbf{X}} \in \mathbb{R}^{n \times a},$$

where  $\mathbf{W}^{D_x} \in \mathbb{R}^{d \times a}$  is a projection matrix and  $\mathbf{b}^{D_x}$  is a bias. To reconstruct the adjacency matrix, we first process  $\mathbf{Z}$  to produce bilinear scores,

$$\tilde{\mathbf{Z}} := \frac{1}{\sqrt{d}} \mathbf{Z} \mathbf{W}^Q (\mathbf{Z} \mathbf{W}^K)^\top \in \mathbb{R}^{n \times n}.$$

In principle, because the LPE is sufficiently adjacency-identifying, we should be able to detect edges by looking at multiple maxima over the rows of  $\tilde{\mathbf{Z}}$ , which is impractical. Instead, we treat the detection of maxima—and thus of edges—as a binary classification problem. Since rows may have variable lengths, as graphs have varying sizes, and since we need to preserve permutation equivariance, we instantiate this classifier as a DeepSet. It takes the rows of  $\tilde{\mathbf{Z}}$  as inputs and outputs individual logits for each  $\hat{A}_{ij}$ . Formally, our decoder can be written as

$$\hat{\mathbf{A}} = \mathcal{D}_e(\mathbf{Z}) = \sigma \left( \text{DeepSet} \left( \tilde{\mathbf{Z}} \right) \right), \quad \hat{\mathbf{A}} \in \mathbb{R}^{n \times n},$$

where  $\text{DeepSet}: \mathbb{R}^n \rightarrow \mathbb{R}^n$  is applied row-wise and  $\sigma(\cdot)$  is a point-wise applied sigmoidal function. To decode edge labels along with the graph structure, we extend the bilinear layer to a multi-headed version; see Appendix B.1 for details.

**Loss function** Following standard practice in VAEs, the training objective combines node and edge reconstruction losses, with KL regularization:

$$\mathcal{L}(\hat{G}) = \mathcal{L}_{\text{node}}(\mathbf{X}, \hat{\mathbf{X}}) + \mathcal{L}_{\text{edge}}(\mathbf{A}, \hat{\mathbf{A}}) + \beta \mathcal{L}_{\text{KL}}(\mu_{\mathbf{Z}}, \sigma_{\mathbf{Z}}), \quad \beta > 0,$$

where  $\mathcal{L}_{\text{node}}(\mathbf{X}, \hat{\mathbf{X}}) = \text{CrossEntropy}(\mathbf{X}, \hat{\mathbf{X}})$ ,  $\mathcal{L}_{\text{edge}}(\mathbf{A}, \hat{\mathbf{A}}) = \text{CrossEntropy}(\mathbf{A}, \hat{\mathbf{A}})$ , and  $\mathcal{L}_{\text{KL}}(\mu_{\mathbf{Z}}, \sigma_{\mathbf{Z}}) = D_{\text{KL}}(\mathcal{N}(\mathbf{Z}; \mu_{\mathbf{Z}}, \sigma_{\mathbf{Z}}) \parallel \mathcal{N}(0, \mathbf{I}_d))$ . Note that the decoder can be readily adapted to continuous features by outputting a scalar instead of a categorical vector and replacing the cross-entropy loss with a standard mean-squared error.

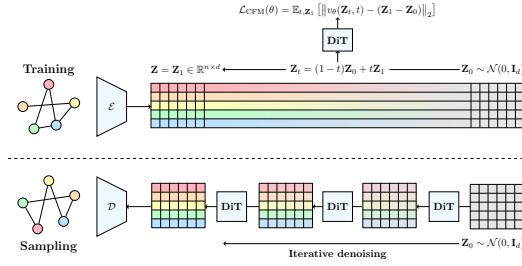


Figure 2: Overview of **LG-Flow**. During training (top), the frozen encoder maps graphs into latent representations  $\mathbf{Z}$ . Noisy latents  $\mathbf{Z}_t$  are sampled along a linear interpolation path and passed through the DiT, which predicts  $v_\theta$  and is optimized with the conditional flow matching loss. During sampling (bottom), noise is generated and iteratively denoised using the trained DiT. The final latents are then decoded into synthetic graphs using the frozen decoder.

## 4.2 LATENT FLOW MATCHING

Once the LG-VAE is trained, we freeze its parameters and train a generative model in the latent space of the LG-VAE. We adopt FM (Lipman et al., 2023), implemented with a DiT (Peebles and Xie, 2023) and name our latent flow matching model **LG-Flow**. Given latents  $\mathbf{Z} \in \mathbb{R}^{n \times d}$ , we sample  $t \in [0, 1]$  and from  $\mathbf{Z}_t = (1-t)\mathbf{Z} + t\mathbf{Z}_0$  with  $\mathbf{Z}_0 \sim \mathcal{N}(0, \mathbf{I}_d)$ . We follow the logit-normal time distribution (Esser et al., 2024), which emphasizes intermediate timesteps, and optimize the FM objective

$$\mathcal{L}_{\text{CFM}}(\theta) = \mathbb{E}_{t, \mathbf{Z}} [\|v_\theta(\mathbf{Z}_t, t) - (\mathbf{Z} - \mathbf{Z}_0)\|_2].$$

This latent formulation yields several benefits. First, training is efficient since DiT operates on low-dimensional node embeddings rather than directly predicting the full adjacency matrix. While the Transformer architecture has quadratic complexity in theory, efficient implementations yield near-linear scaling. Consequently, the only step with quadratic complexity is decoding, handled by our lightweight decoder. Secondly, using a modality-agnostic FM backbone highlights modularity, i.e., the encoder handles all graph-specific complexity, while the diffusion stage remains generic.

## 4.3 EXTENDING TO DAGS

The LG-VAE described so far applies to undirected graphs. The central role of DAGs in applications such as chip design and logical circuit synthesis motivates extending the LG-VAE and **LG-Flow** framework beyond undirected graphs to the directed setting. Extending it to DAGs is non-trivial, since the LPE relies on the eigenvalue decomposition of the symmetric Laplacian, which discards edge directionality when applied to directed graphs. To preserve this information, we adopt the *magnetic Laplacian* (Forman, 1993; Shubin, 1994; Colin de Verdière, 2013; Furutani et al., 2019; Geisler et al., 2023), a Hermitian matrix that generalizes the Laplacian to directed graphs. It is defined as  $\mathbf{L}_q := \mathbf{D}_s - \mathbf{A}_s \odot e^{i\Theta^{(q)}}$ , where  $\mathbf{A}_s := \mathbf{A} \vee \mathbf{A}^\top$  is the symmetrized adjacency,  $\mathbf{D}_s$  its degree matrix,  $\Theta^{(q)} = 2\pi q(A_{ij} - A_{ji})$ , and  $q \in \mathbb{R}^{++}$  is a fixed real parameter. This matrix admits an eigendecomposition  $\mathbf{L}_q = \mathbf{U} \Lambda \mathbf{U}^*$  with eigenvectors in  $\mathbb{C}^n$ , which we use to introduce the *magnetic LPE* (mLPE).

The mLPE mirrors the construction of the LPE but separates the real and imaginary parts of the eigenvectors, allowing the use of standard real-valued neural networks. Concretely, let  $\mathbf{U}^R = \Re(\mathbf{U})$  and  $\mathbf{U}^I = \Im(\mathbf{U})$  be the real and imaginary part of  $\mathbf{U}$ , respectively. We then define

$$\phi(\mathbf{U}^R, \mathbf{U}^I, \Lambda) = [\phi(\mathbf{U}_1^R \| \mathbf{U}_1^I \| \Lambda + \epsilon), \dots, \phi(\mathbf{U}_n^R \| \mathbf{U}_n^I \| \Lambda + \epsilon)] \in \mathbb{R}^{n \times n \times d},$$

where  $\phi: \mathbb{R}^3 \rightarrow \mathbb{R}^d$  is applied row-wise, on the last dimension of  $\mathbf{U}_1^R \| \mathbf{U}_1^I \| \Lambda + \epsilon \in \mathbb{R}^{n \times n \times 3}$  and  $\epsilon \in \mathbb{R}^l$  is a learnable, zero-initialized vector. The mLPE is then obtained by aggregating across eigenvectors as

$$\text{mLPE}(\mathbf{U}^R, \mathbf{U}^I, \Lambda) = \rho \left( \sum_{i=1}^n \phi(\mathbf{U}^R, \mathbf{U}^I, \Lambda)_i^\top \right),$$

with  $\rho: \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^d$  is an equivariant feed-forward neural networks.

To characterize adjacency recovery in the directed setting, we introduce the notion of *out-adjacency-identifiability*. Given node embeddings  $\mathbf{P} \in \mathbb{R}^{n \times d}$ , we compute

$$\tilde{\mathbf{P}}^R = \frac{1}{\sqrt{d}}(\mathbf{P}\mathbf{W}^{K_R})(\mathbf{P}\mathbf{W}^{Q_R})^\top, \quad \tilde{\mathbf{P}}^I = \frac{1}{\sqrt{d}}(\mathbf{P}\mathbf{W}^{K_I})(\mathbf{P}\mathbf{W}^{Q_I})^\top,$$

with learnable matrices  $\mathbf{W}^{K_R}, \mathbf{W}^{Q_R}, \mathbf{W}^{K_I}, \mathbf{W}^{Q_I} \in \mathbb{R}^{d \times d}$ , and combine them as

$$\tilde{\mathbf{P}} := \tilde{\mathbf{P}}^R + \frac{2-c}{s} \tilde{\mathbf{P}}^I, \quad c = \cos(2\pi q), \quad s = \sin(2\pi q).$$

We say that  $\mathbf{P}$  is out-adjacency-identifying if for each node  $i$ ,

$$\tilde{\mathbf{P}}_{ij} = \max_k \tilde{\mathbf{P}}_{ik} \iff A_{ij} = 1,$$

We further say that a matrix  $\mathbf{Q} \in \mathbb{R}^{n \times d}$  is *sufficiently out-adjacency-identifying* if it approximates an adjacency-identifying matrix  $\mathbf{P}$  to arbitrary precision, i.e.,  $\|\mathbf{P} - \mathbf{Q}\|_F < \epsilon$ , for all  $\epsilon > 0$ .

Now, the result below shows that the mLPE produces node embeddings from which the full directed adjacency of a DAG can be recovered with arbitrarily high accuracy. In other words, no structural information is lost in the encoding step.

**Theorem 1.** *The magnetic Laplacian positional encoding (mLPE) is sufficiently out-of-adjacency-identifying.*

Proofs and the full DAG autoencoder specification are provided in Appendix C.2. This result ensures that, by replacing the LPE with the mLPE, the LG-VAE naturally extends to DAGs while preserving node-level latent representations that are compatible with latent diffusion.

## 5 EXPERIMENTAL STUDY

We now investigate the performance of our approach in relation to the requirements we derived in the introduction. Specifically, we answer the following questions.

**Q1** How faithfully does LG-VAE reconstruct the graph’s structure?

**Q2** How does **LG-Flow** compare to prior graph diffusion models in generative performance?

**Q3** How efficient is **LG-Flow** in inference time compared to state-of-the-art graph diffusion models?

**Datasets** We evaluate the generative performance of our approach on six datasets: three synthetic datasets—EXTENDED PLANAR, EXTENDED TREE, and EGO—, two molecular datasets—MOSES and GUACAMOL—, and TPU TILE, a DAG generation datasets. An extensive description of those datasets is available in Appendix E.1.

### 5.1 Q1: AUTOENCODER RECONSTRUCTION ABILITY

We begin by empirically evaluating our autoencoder’s ability to reconstruct adjacency matrices faithfully. We focus on two simple, unlabeled datasets with strong structural constraints—EXTENDED PLANAR and EXTENDED TREE. For these distributions, near-lossless reconstruction is crucial to ensure that decoded graphs maintain their structural properties, specifically planarity and tree structure.

**Baselines** We evaluate LG-VAE on EXTENDED PLANAR and TREE, and compare it against two prior graph autoencoders: DGVAE (Boget et al., 2024) and the autoencoder used in GLAD (Nguyen et al., 2024). We also conduct ablations of LG-VAE’s components, (i) removing the positional encoder  $\phi$  to obtain a raw GNN ENCODER, and (ii) replacing our decoder with the original GVAE DECODER, yielding a GVAE-like variant that retains  $\phi$  as positional encoder.

**Metrics** Our primary evaluation metric is SAMPLE ACCURACY, i.e., the fraction of test graphs that the

Table 2: **Autoencoder reconstruction performance.** Best results are in **bold**. Only LG-VAE achieves near-lossless reconstruction, as shown by the SAMPLE ACCURACY metric.

Method	EXTENDED PLANAR		EXTENDED TREE	
	Edge Acc.	Sample Acc.	Edge Acc.	Sample Acc.
GLAD	0.9129	0.	0.9866	0.
DGVAE	0.9547	0.	0.9777	0.
Ours	0.9000	0.	0.6351	0.
GVAE DECODER	0.9977	0.7734	0.2539	0.
GNN ENCODER	<b>1.0000</b>	<b>0.9961</b>	<b>1.0000</b>	<b>0.9844</b>
LG-VAE	<b>1.0000</b>	<b>0.9961</b>	<b>1.0000</b>	<b>0.9844</b>

378 decoder reconstructs exactly, including node/edge labels when present. We also report EDGE ACCURACY,  
 379 i.e., the proportion of correctly reconstructed edges in the test set.  
 380

381 **Results** We report results in Table 2. Except for LG-VAE, no method achieves near-lossless reconstruction  
 382 in terms of SAMPLE ACCURACY. However, all methods attain EDGE ACCURACY above 0.9, only  
 383 LG-VAE and its raw GNN ENCODER variant obtain non-zero SAMPLE ACCURACY.

384 Consequently, these findings highlight the limitations of current VAEs, whose latent representations are  
 385 less expressive than those of LG-VAE. They struggle to reconstruct samples from latents with high  
 386 accuracy, casting doubt on their ability to recover the required structural properties. For instance, even if a  
 387 latent diffusion model perfectly matches the latent distribution, the decoder is not guaranteed to yield  
 388 planar graphs or trees.

## 389 390 5.2 Q2 & Q3: ASSESSING GENERATIVE PERFORMANCE AND INFERENCE EFFICIENCY

391 In this section, we evaluate our latent diffusion model on various benchmarks, including synthetic graphs,  
 392 molecules, and DAGs.  
 393

394 **Baselines** We evaluate **LG-Flow** against several prior graph diffusion models, two discrete-time methods,  
 395 DiGRESS (Vignac et al., 2023) and its sparse counterpart, SPARSEDIFF (Qin et al., 2023), two continuous-  
 396 time methods, COMETH (Siraudeau et al., 2024) and DEFoG (Qin et al., 2025). Note that DISCO (Xu et al.,  
 397 2024) and COMETH are highly similar methods, and either could have been evaluated. We chose to focus  
 398 on COMETH due to the inherent high computational cost of discrete graph diffusion models. On DAGs, we  
 399 assess our approach against diffusion models tailored to DAGs, namely LAYERDAG, an autoregressive  
 400 diffusion model that generates DAGs layer by layer, and DIRECTO, a recent adaptation of DEFoG to DAG  
 401 generation.

402 **Metrics** For synthetic graphs and DAGs, we evaluate generation quality using the standard *Maximum*  
 403 *Mean Discrepancy* (MMD) metrics between test and generated sets, computed on various statistics, e.g.,  
 404 degree or clustering coefficient. Validity for trees, planar graphs, and DAGs is defined by structural  
 405 constraints (planarity, tree property, or acyclicity). For molecular datasets, we assess validity based on  
 406 fundamental chemical rules, uniqueness, and novelty, complemented by benchmark-specific scores. A  
 407 detailed description of all metrics is deferred to Appendix E.1.

408 Across all datasets, we also measure sampling time (**Time**) to assess inference efficiency. **All models were**  
 409 **sampled using equal batch sizes whenever possible, or otherwise with batch sizes chosen to maximize GPU**  
 410 **utilization.**

411 **Results on synthetic graphs** Our results are shown in Table 3a. Our approach is competitive across all  
 412 distribution metrics and achieves strong performance on V.U.N. In terms of sampling time, it delivers a  
 413  $50\times$  speed-up on EXTENDED PLANAR and a  $10\times$  speed-up on EXTENDED TREE. On EGO, our method  
 414 matches the performance of DEFoG while reducing sampling time by up to  $1000\times$ .

415 To further demonstrate the efficiency of **LG-Flow**, we plot memory requirements during sampling as a  
 416 function of batch size in Table 3. Due to the DiT’s efficient implementation, the memory usage of our  
 417 method scales linearly with the batch size, allowing the entire test batch of 151 samples to fit into memory.  
 418 In contrast, DEFoG runs out of memory once the batch size exceeds 32.

419 **Results on molecular generation** Our results are summarized  
 420 in Table 4. On GUACAMOL, our method achieves strong  
 421 performance on validity, uniqueness, and novelty, consistently  
 422 outperforming DiGRESS and DISCO. On **KL div**, we reach  
 423 state-of-the-art results, with performance close to DEFoG.  
 424 Most notably, we surpass all baselines on **FCD** by a large  
 425 margin, highlighting the method’s ability to capture the  
 426 underlying chemical distribution.

427 On MOSES, **LG-Flow** remains competitive across all metrics,  
 428 outperforming DiGRESS and DISCO on validity again.  
 429 Crucially, while maintaining competitive accuracy, our method  
 430 delivers substantial inference efficiency gains, achieving  
 431 a **36.4 $\times$**  speedup on GUACAMOL and a  $58\times$  speedup on  
 MOSES compared to the state-of-the-art DEFoG. To put

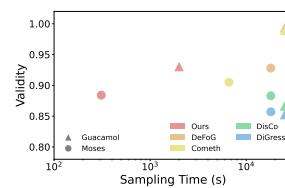


Figure 3: We compare molecular validity and sampling time. The ideal lies in the top-left corner, with high validity and low cost. Our method offers a better trade-off than prior graph-space diffusion models.

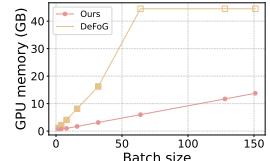
Table 3: **Generation results on synthetic graphs.** In the table, best results are in **bold**, second best are underlined. In Table 3, empty squares indicate out-of-memory errors. While our method scales linearly with batch size and can accommodate the entire test batch of 151 samples on a single GPU, DEFOG runs out of memory at batch size 64. **MMD** metrics are reported with a  $10^3$  factor.

(a) Generation results on the Extended Planar and Extended Tree datasets.

Dataset											
Extended Planar											
Method	Deg. (↓)	Orbit (↓)	Cluster. (↓)	Spec. (↓)	V.U.N. (↑)	Time (↓)	Method	Deg. (↓)	Orbit (↓)	Cluster. (↓)	Spec. (↓)
SparseDiff	0.8±0.1	0.9±0.1	20.9±1.3	1.9±0.2	74.1±2.1	$12.1 \times 10^2$	0.±0.	0.±0.	0.±0.	<b>1.2</b> ±0.2	$95.3 \pm 1.5$
Cometh	<u>0.7</u> ±0.1	<u>2.2</u> ±0.3	20.0±1.1	<u>1.3</u> ±0.2	96.7±0.8	$9.7 \times 10^2$	0.±0.	0.±0.	0.±0.	<u>1.4</u> ±0.4	$93.8 \pm 1.6$
DeFog	<b>0.3</b> ±0.1	<b>0.1</b> ±0.1	<u>9.7</u> ±0.3	<u>1.3</u> ±0.1	<b>99.6</b> ±0.4	$8.9 \times 10^2$	<u>0.1</u> ±0.	0.±0.	0.±0.	<b>1.2</b> ±0.2	<b>98.1</b> ±0.7
<b>LG-Flow</b>	<b>0.3</b> ±0.1	<b>0.5</b> ±0.2	<u>10.1</u> ±3.1	<b>1.0</b> ±0.2	99.0±0.2	<b>4.6</b> ±0.	0.1±0.	0.±0.	0.±0.	1.6±0.3	$93.3 \pm 0.8$

(b) Generation results on large graphs: Ego.

Dataset						
Ego						
Method	Deg. (↓)	Orbit (↓)	Cluster. (↓)	Spec. (↓)	Ratio (↓)	Time (s) (↓)
Training Set	0.2	6.8	7.1	1.0	1.0	—
DiGress	8.9±1.6	30±3	54±4	<u>19</u> ±3.2	19±3.1	—
SparseDiff	3.7±0.4	20±4	32±1	<b>5.6</b> ±0.8	7.9±0.9	$5.3 \times 10^2$
Cometh	8.1±0.1	<u>19.6</u> ±2.0	37.4±1.5	12.9±1.6	15.4±1.7	$1.3 \times 10^4$
DeFog	<b>1.8</b> ±0.5	21.4±2.6	<u>30.7</u> ±3.3	<u>6.2</u> ±1.2	<u>5.2</u> ±0.8	$1.3 \times 10^4$
<b>LG-Flow</b>	<u>1.9</u> ±0.5	<u>15.5</u> ±3.8	<u>22.4</u> ±2.8	<u>6.2</u> ±1.1	<u>5.3</u> ±1.0	<b>11.6</b> ±0.6



(c) Comparison of memory use for DEFOG and our method on the Ego dataset.

Table 4: **Generation on molecular datasets : MOSES & GuacaMol.** Best results are denoted in **bold** and second best are underlined.

Dataset											
GuacaMol											
Model	V. (↑)	V.U. (↑)	V.U.N. (↑)	KL div (↑)	FCD (↑)	Time (s) (↓)	Val. (↑)	Unique. (↑)	Novel. (↑)	Filters (↑)	FCD (↓)
Training Set	100.0	100.0	0.0	99.9	92.8	—	100.0	100.0	0.0	100.0	0.01
DiGress	85.2	85.2	85.1	92.9	68.0	—	85.7	<b>100.0</b>	95.0	97.1	<b>1.19</b>
DisCo	86.6	86.6	86.5	92.6	59.7	—	88.3	<b>100.0</b>	<b>97.7</b>	95.6	1.44
Cometh	98.9	98.9	97.6	96.7	72.7	$4.0 \times 10^4$	90.5	99.9	92.6	<b>99.1</b>	<u>1.27</u>
DeFog	<b>99.0</b>	<b>99.0</b>	<b>97.9</b>	<b>97.7</b>	<u>73.8</u>	$4.0 \times 10^4$	<b>92.8</b>	99.9	92.1	98.9	1.95
<b>LG-Flow</b>	93.0	93.0	91.5	<u>97.2</u>	<b>81.8</b>	$1.1 \times 10^3$	88.4	99.9	90.5	98.9	1.43

these results into perspective, we plot validity against sampling efficiency in Figure 3. Our method exhibits a significantly better trade-off than prior diffusion models that operate directly in graph space.

**Results on DAG generation** Our results are reported in Table 5. As already noted in Carballo-Castro et al. (2025), LAYERDAG’s strong performance on V.U.N is misleading: it collapses on distributional metrics, clearly failing to capture the underlying distributional characteristics.

In contrast, we outperform DIRECTO on **Cluster.**, **Spec.**, **Wave.**, and, most importantly, on **Valid**. **LG-Flow** shows slightly higher memorization than DIRECTO, likely due to the limited diversity of the training data. Most importantly, our approach is dramatically more efficient at inference, achieving up to a  $680 \times$  speed-up over DIRECTO.

Table 5: **Generation on DAGs: TPU Tile dataset.** We highlight best results in **bold**.

Method	Out Deg.	In Deg.	Cluster.	Spec.	Wave.	Valid	Unique	Novel	V.U.N	Time
Training Set	0.3	0.3	0.7	0.6	0.2	1.0	1.0	0.	0.	—
LayerDAG	$193.3 \pm 90.5$	$222.5 \pm 39.5$	$151.2 \pm 52.2$	$50.1 \pm 20.6$	$76.5 \pm 25.1$	<b>100.0</b> ±0.0	<b>99.5</b> ±1.0	<b>98.5</b> ±3.0	<b>98.5</b> ±3.0	—
Directo	<b>3.9</b> ±1.7	<b>37.6</b> ±5.1	$21.1 \pm 11.7$	$12.6 \pm 2.2$	$6.2 \pm 0.9$	$90.5 \pm 3.3$	$90.5 \pm 4.6$	$97.5 \pm 3.2$	$80.5 \pm 4.6$	$3.4 \times 10^2$
<b>LG-Flow</b>	$12.8 \pm 3.6$	$81.8 \pm 6.3$	<b>5.9</b> ±2.5	<b>8.1</b> ±1.3	<b>4.0</b> ±1.3	$98.5 \pm 2.6$	$94.0 \pm 3.3$	$77.5 \pm 3.9$	$74.5 \pm 3.8$	<b>0.5</b> ±0.3

## 6 CONCLUSION

We proposed a latent diffusion framework for graphs, built on a permutation-equivariant Laplacian autoencoder with provable reconstruction guarantees for undirected graphs and DAGs. By mapping nodes to fixed-dimensional embeddings, our method removes the quadratic cost of graph-space diffusion and supports high-capacity backbones such as Diffusion Transformers. Experiments show near-lossless reconstruction, competitive or superior generative performance, and sampling up to three orders of magnitude faster. This framework shifts the focus from graph-specific diffusion design to cheaper, more expressive autoencoders, unifies graphs and DAGs generation, enables transferring conditional methods from image generation, and paves the way for scaling generative models to much larger graphs.

486 REFERENCES  
487

488 J. Austin, D. D. Johnson, J. Ho, D. Tarlow, and R. van den Berg. Structured denoising diffusion models in  
489 discrete state-spaces. In M. Ranzato, A. Beygelzimer, Y. N. Dauphin, P. Liang, and J. W. Vaughan,  
490 editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural  
Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, 2021.

492 M. Bechler-Speicher, M. Eliasof, C.-B. Schönlieb, R. Gilad-Bachrach, and A. Globerson. Towards  
493 invariance to node identifiers in graph neural networks. *CoRR*, 2025.

495 T. Bian, Y. Niu, H. Chang, D. Yan, J. Huang, Y. Rong, T. Xu, J. Li, and H. Cheng. Hierarchical graph latent  
496 diffusion model for conditional molecule generation. In *Proceedings of the 33rd ACM International  
497 Conference on Information and Knowledge Management, CIKM '24*, 2024. ISBN 9798400704369.

498 Y. Boget, M. Gregorova, and A. Kalousis. Discrete graph auto-encoder. *Transactions on Machine Learning  
499 Research*, 2024. ISSN 2835-8856.

501 N. Brown, M. Fiscato, M. H. Segler, and A. C. Vaucher. Guacamol: benchmarking models for de novo  
502 molecular design. *Journal of Chemical Information and Modeling*, (3), 2019.

503 A. Campbell, J. Benton, V. D. Bortoli, T. Rainforth, G. Deligiannidis, and A. Doucet. A continuous time  
504 framework for discrete denoising models. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho,  
505 and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on  
506 Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 -  
507 December 9, 2022*, 2022.

509 A. Carballo-Castro, M. Madeira, Y. Qin, D. Thanou, and P. Frossard. Generating directed graphs with dual  
510 attention and asymmetric encoding. *ArXiv preprint*, 2025.

512 T. Chen, R. Zhang, and G. E. Hinton. Analog bits: Generating discrete data using diffusion models with  
513 self-conditioning. In *Proc. of ICLR*, 2023a.

514 X. Chen, J. He, X. Han, and L. Liu. Efficient and degree-guided graph generation via discrete diffusion  
515 modeling. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, editors,  
516 *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*,  
517 Proceedings of Machine Learning Research, 2023b.

518 X. Chen, Y. Wang, J. He, Y. Du, S. Hassoun, X. Xu, and L. Liu. Graph generative pre-trained transformer.  
519 In *Forty-second International Conference on Machine Learning*, 2025.

521 Y. Colin de Verdière. Magnetic interpretation of the nodal defect on graphs. *Analysis & PDE*, (5), 2013.

523 V. P. Dwivedi, A. T. Luu, T. Laurent, Y. Bengio, and X. Bresson. Graph neural networks with learnable  
524 structural and positional representations. In *Proc. of ICLR*, 2022.

525 P. Esser, S. Kulal, A. Blattmann, R. Entezari, J. Müller, H. Saini, Y. Levi, D. Lorenz, A. Sauer, F. Boesel,  
526 D. Podell, T. Dockhorn, Z. English, and R. Rombach. Scaling rectified flow transformers for  
527 high-resolution image synthesis. In *Forty-first International Conference on Machine Learning, ICML  
528 2024, Vienna, Austria, July 21-27, 2024*, 2024.

530 R. Forman. Determinants of laplacians on graphs. *Topology*, (1), 1993.

532 X. Fu, Y. Gao, Y. Wei, Q. Sun, H. Peng, J. Li, and X. Li. Hyperbolic geometric latent diffusion model for  
533 graph generation. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna,  
534 Austria, July 21-27, 2024*, 2024.

535 S. Furutani, T. Shibahara, M. Akiyama, K. Hato, and M. Aida. Graph signal processing for directed graphs  
536 based on the hermitian laplacian. In *Joint European conference on machine learning and knowledge  
537 discovery in databases*. Springer, 2019.

538 R. Gao, E. Hoogeboom, J. Heek, V. D. Bortoli, K. P. Murphy, and T. Salimans. Diffusion meets flow  
539 matching: Two sides of the same coin. 2024.

540 I. Gat, T. Remez, N. Shaul, F. Kreuk, R. T. Q. Chen, G. Synnaeve, Y. Adi, and Y. Lipman. Discrete  
 541 flow matching. In A. Globersons, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. M. Tomczak, and  
 542 C. Zhang, editors, *Advances in Neural Information Processing Systems 38: Annual Conference on*  
 543 *Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 -*  
 544 *15, 2024*, 2024.

545 S. Geisler, Y. Li, D. J. Mankowitz, A. T. Cemgil, S. Günnemann, and C. Paduraru. Transformers meet  
 546 directed graphs. In A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, editors,  
 547 *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA,*  
 548 *Proceedings of Machine Learning Research*, 2023.

550 J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum  
 551 chemistry. In D. Precup and Y. W. Teh, editors, *Proc. of ICML*, Proceedings of Machine Learning  
 552 Research, 2017.

553 J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato,  
 554 R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33:*  
 555 *Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12,*  
 556 *2020, virtual*, 2020.

558 W. Hu, B. Liu, J. Gomes, M. Zitnik, P. Liang, V. S. Pande, and J. Leskovec. Strategies for pre-training  
 559 graph neural networks. In *Proc. of ICLR*, 2020.

560 I. Igashov, A. Schneuing, M. H. S. Segler, M. M. Bronstein, and B. E. Correia. Retrobridge: Modeling  
 561 retrosynthesis with markov bridges. In *Proc. of ICLR*, 2024.

563 Y. Jang, S. Lee, and S. Ahn. A simple and scalable representation for graph generation. In *Proc. of ICLR*,  
 564 2024.

566 J. Jo, S. Lee, and S. J. Hwang. Score-based generative modeling of graphs via the system of stochastic  
 567 differential equations. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvári, G. Niu, and S. Sabato,  
 568 editors, *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore,*  
 569 *Maryland, USA, Proceedings of Machine Learning Research*, 2022.

570 C. K. Joshi, X. Fu, Y.-L. Liao, V. Gharakhanyan, B. K. Miller, A. Sriram, and Z. W. Ulissi. All-atom  
 571 diffusion transformers: Unified generative modelling of molecules and materials. In *Forty-second*  
 572 *International Conference on Machine Learning*, 2025.

574 M. A. Ketata, N. Gao, J. Sommer, T. Wollschläger, and S. Günnemann. Lift your molecules: Molecular  
 575 graph generation in latent euclidean space. In *The Thirteenth International Conference on Learning*  
 576 *Representations*, 2025.

577 T. N. Kipf and M. Welling. Variational graph auto-encoders. *ArXiv preprint*, 2016.

579 M. Krimmel, J. Wiens, K. M. Borgwardt, and D. Chen. Towards fast graph generation via autoregressive  
 580 noisy filtration modeling. *ArXiv preprint*, 2025.

581 N. Laabid, S. Rissanen, M. Heinonen, A. Solin, and V. Garg. Equivariant denoisers cannot copy  
 582 graphs: Align your graph diffusion models. In *The Thirteenth International Conference on Learning*  
 583 *Representations*, 2025.

585 M. Lee and K. Min. Mgcvae: multi-objective inverse design via molecular graph conditional variational  
 586 autoencoder. *Journal of chemical information and modeling*, (12), 2022.

587 J. Li, J. Yu, J. Li, H. Zhang, K. Zhao, Y. Rong, H. Cheng, and J. Huang. Dirichlet graph variational  
 588 autoencoder. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in*  
 589 *Neural Information Processing Systems 33: Annual Conference on Neural Information Processing*  
 590 *Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.

592 M. Li, V. Shitole, E. Chien, C. Man, Z. Wang, S. Sridharan, Y. Zhang, T. Krishna, and P. Li. LayerDAG: A  
 593 layerwise autoregressive diffusion model for directed acyclic graph generation. In *International*  
 594 *Conference on Learning Representations*, 2025a.

594 X. Li, X. Li, L. Chen, X. Zhang, M. Yuan, and J. Wang. Circuit transformer: A transformer that preserves  
 595 logical equivalence. In *The Thirteenth International Conference on Learning Representations*, 2025b.  
 596

597 D. Lim, J. D. Robinson, L. Zhao, T. E. Smidt, S. Sra, H. Maron, and S. Jegelka. Sign and basis invariant  
 598 networks for spectral graph representation learning. In *Proc. of ICLR*, 2023.

599 Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le. Flow matching for generative modeling.  
 600 In *Proc. of ICLR*, 2023.

601 Y. Lipman, M. Havasi, P. Holderrieth, N. Shaul, M. Le, B. Karrer, R. T. Chen, D. Lopez-Paz, H. Ben-Hamu,  
 602 and I. Gat. Flow matching guide and code. *CoRR*, 2024.

603

604 A. Lou, C. Meng, and S. Ermon. Discrete diffusion modeling by estimating the ratios of the data  
 605 distribution. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria,  
 606 July 21-27, 2024*, 2024.

607 K. Martinkus, A. Loukas, N. Perraudeau, and R. Wattenhofer. SPECTRE: spectral conditioning helps to  
 608 overcome the expressivity limits of one-shot graph generators. In K. Chaudhuri, S. Jegelka, L. Song,  
 609 C. Szepesvári, G. Niu, and S. Sabato, editors, *International Conference on Machine Learning, ICML  
 610 2022, 17-23 July 2022, Baltimore, Maryland, USA*, Proceedings of Machine Learning Research, 2022.

611

612 D. Mendez, A. Gaulton, A. P. Bento, J. Chambers, M. De Veij, E. Félix, M. P. Magariños, J. F. Mosquera,  
 613 P. Mutowo, M. Nowotka, et al. Chembl: towards direct deposition of bioassay data. *Nucleic acids  
 614 research*, (D1), 2019.

615 G. Minello, A. Bicciato, L. Rossi, A. Torsello, and L. Cosmo. Generating graphs via spectral diffusion. In  
 616 *The Thirteenth International Conference on Learning Representations*, 2025.

617

618 A. Mishchenko and Y. Miyasaka. Problems and results of iwls 2023 programming contest, 2023.

619 M. Morris, B. C. Grau, and I. Horrocks. Orbit-equivariant graph neural networks. In *Proc. of ICLR*, 2024.

620

621 L. Müller and C. Morris. Aligning transformers with weisfeiler-leman. In *Forty-first International  
 622 Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, 2024.

623 V. K. Nguyen, Y. Boget, F. Lavda, and A. Kalousis. Glad: improving latent graph generative modeling with  
 624 simple quantization. In *Proceedings of the ICML 2024 Workshop SPIGM*. OpenReview. net, 2024.

625

626 C. Niu, Y. Song, J. Song, S. Zhao, A. Grover, and S. Ermon. Permutation invariant graph generation via  
 627 score-based generative modeling. In S. Chiappa and R. Calandra, editors, *The 23rd International  
 628 Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo,  
 629 Sicily, Italy]*, Proceedings of Machine Learning Research, 2020.

630

631 W. Peebles and S. Xie. Scalable diffusion models with transformers. In *IEEE/CVF International  
 Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 2023.

632

633 M. Phothilimthana, S. Abu-El-Haija, K. Cao, B. Fatemi, M. Burrows, C. Mendis, and B. Perozzi.  
 634 Tpugraphs: A performance prediction dataset on large tensor computational graphs. *Advances in Neural  
 635 Information Processing Systems*, 36:70355–70375, 2023a.

636

637 M. M. Phothilimthana, S. Abu-El-Haija, K. Cao, B. Fatemi, M. Burrows, C. Mendis, and B. Perozzi.  
 638 Tpugraphs: A performance prediction dataset on large tensor computational graphs. In A. Oh,  
 639 T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information  
 640 Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS  
 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023b.

641

642 D. Polykovskiy, A. Zhebrak, B. Sanchez-Lengeling, S. Golovanov, O. Tatanov, S. Belyaev, R. Kurbanov,  
 643 A. Artamonov, V. Aladinskiy, M. Veselov, et al. Molecular sets (moses): a benchmarking platform for  
 644 molecular generation models. *Frontiers in pharmacology*, 2020.

645

646 Y. Qin, C. Vignac, and P. Frossard. Sparse training of discrete diffusion models for graph generation. *arXiv  
 647 preprint*, 2023.

648

649 Y. Qin, M. Madeira, D. Thanou, and P. Frossard. Defog: Discrete flow matching for graph generation. In  
 650 *International Conference on Machine Learning (ICML)*, 2025.

648 R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with  
 649 latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
 650 *recognition*, 2022.

651 E. Rossi, B. Charpentier, F. D. Giovanni, F. Frasca, S. Günnemann, and M. M. Bronstein. Edge  
 652 directionality improves learning on heterophilic graphs. In *Learning on Graphs Conference*, 2023.

653 F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model.  
 654 *IEEE Transactions on Neural Networks*, (1), 2009.

655 M. A. Shubin. Discrete magnetic laplacian. *Communications in mathematical physics*, (2), 1994.

656 A. Siraudin, F. D. Malliaros, and C. Morris. Cometh: A continuous-time discrete-state graph diffusion  
 657 model. *Transactions on Machine Learning Research*, 2024.

658 Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative  
 659 modeling through stochastic differential equations. In *Proc. of ICLR*, 2021.

660 T. Sterling and J. J. Irwin. Zinc 15 – ligand discovery for everyone. *Journal of Chemical Information and*  
 661 *Modeling*, (11), 2015. PMID: 26479676.

662 Z. Sun and Y. Yang. DIFUSCO: graph-based diffusion solvers for combinatorial optimization. In  
 663 A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural*  
 664 *Information Processing Systems 36: Annual Conference on Neural Information Processing Systems*  
 665 *2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023.

666 C. Vignac, I. Krawczuk, A. Siraudin, B. Wang, V. Cevher, and P. Frossard. Digress: Discrete denoising  
 667 diffusion for graph generation. In *Proc. of ICLR*, 2023.

668 C. Wang, H. H. Ong, S. Chiba, and J. C. Rajapakse. Gldm: hit molecule generation with constrained graph  
 669 latent diffusion model. *Briefings in Bioinformatics*, (3), 2024. ISSN 1477-4054.

670 K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? In *Proc. of ICLR*,  
 671 2019.

672 Z. Xu, R. Qiu, Y. Chen, H. Chen, X. Fan, M. Pan, Z. Zeng, M. Das, and H. Tong. Discrete-state  
 673 continuous-time diffusion for graph generation. In A. Globersons, L. Mackey, D. Belgrave, A. Fan,  
 674 U. Paquet, J. M. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*  
 675 *38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC,*  
 676 *Canada, December 10 - 15, 2024*, 2024.

677 L. Yang, Z. Huang, Z. Zhang, Z. Liu, S. Hong, W. Zhang, W. Yang, B. Cui, and L. Zhang. Graphusion:  
 678 Latent diffusion for graph generation. *TKDE*, (11), 2024a.

679 L. Yang, Z. Huang, Z. Zhang, Z. Liu, S. Hong, W. Zhang, W. Yang, B. Cui, and L. Zhang. Graphusion:  
 680 Latent diffusion for graph generation. *TKDE*, (11), 2024b.

681 J. You, R. Ying, X. Ren, W. L. Hamilton, and J. Leskovec. Graphrnn: Generating realistic graphs with deep  
 682 auto-regressive models. In J. G. Dy and A. Krause, editors, *Proc. of ICML*, Proceedings of Machine  
 683 Learning Research, 2018.

684 J. You, R. Ying, and J. Leskovec. Position-aware graph neural networks. In K. Chaudhuri and  
 685 R. Salakhutdinov, editors, *Proc. of ICML*, Proceedings of Machine Learning Research, 2019.

686 M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. Salakhutdinov, and A. J. Smola. Deep sets. In  
 687 I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett,  
 688 editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural*  
 689 *Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 2017.

690 C. Zhou, X. Wang, and M. Zhang. Unifying generation and prediction on graphs with latent graph diffusion.  
 691 In A. Globersons, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. M. Tomczak, and C. Zhang, editors,  
 692 *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information*  
 693 *Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024.

694 J. Zhu, L. He, C. Gao, D. Hou, Z. Su, P. S. Yu, J. Kurths, and F. Hellmann. Sdmg: Smoothing your  
 695 diffusion models for powerful graph representation learning. In *Forty-second International Conference*  
 696 *on Machine Learning*, 2025.

# 702 Appendices

## 703 704 705 A EXTENDED RELATED WORK

706 707 708 Here, we discuss more related work.

709 710 **Graph generation** Graph generation methods are typically categorized into two main groups. 711 *Autoregressive models* build graphs by progressively adding nodes and edges (You et al., 2018; Jang et al., 712 2024). Their main advantage is computational efficiency, since they do not need to materialize the full 713 adjacency matrix. However, they depend on an arbitrary node ordering to transform the graph into a 714 sequence, either learned or defined through complex algorithms, which breaks permutation equivariance. In 715 contrast, *one-shot models* generate the entire graph at once, thereby preserving permutation equivariance. 716 Many approaches have been explored in this class, starting with GVAEs (Kipf and Welling, 2016) and 717 GANs (Martinkus et al., 2022). Following their success in other modalities such as images, diffusion 718 models (Ho et al., 2020; Lipman et al., 2023; Song et al., 2021) quickly emerged as the dominant approach 719 for graph generation.

720 **Graph diffusion models** Diffusion models can be grouped according to two criteria: whether they 721 operate in continuous or discrete state spaces, and whether they train on a continuous or discrete time 722 axis. Discrete-time diffusion models (Austin et al., 2021; Ho et al., 2020) were successfully adapted 723 to graph generation, with discrete-state models (Chen et al., 2023b; Vignac et al., 2023) showing an 724 advantage over their continuous-state counterparts (Jo et al., 2022; Niu et al., 2020). However, these 725 models relied on a fixed time discretization during training, which limited their flexibility. They were later 726 extended to continuous-time, both in continuous (Lipman et al., 2023; Song et al., 2021) and discrete state 727 spaces (Campbell et al., 2022; Gat et al., 2024; Lou et al., 2024), and also adapted to graph generation, 728 once again confirming the superiority of discrete-state approaches (Sirauidin et al., 2024; Qin et al., 2025; 729 Xu et al., 2024). Previous graph diffusion models relied on Laplacian decompositions but pursue objectives 730 that differ from ours. Regarding Minello et al. (2025), the approach is fundamentally different from ours: it 731 generates approximate eigenvalues and eigenvectors and then reconstructs the adjacency matrix through a 732 separate prediction module, whereas we embed structural information in the latent space by feeding 733 eigenvectors and eigenvalues to the encoder. Concerning Zhu et al. (2025), the scope of the paper is 734 different from ours. The authors focus on representation learning using graph diffusion models, while our 735 method is evaluated solely for generation. In contrast to traditional graph diffusion models that target 736 the whole adjacency matrix, they train their model to reconstruct only the low-frequency Laplacian 737 eigenvectors. Overall, these methods aim to reconstruct spectral components of the Laplacian, whereas we 738 use them to encode structural information into the latent space.

739 **Latent diffusion models** *Latent diffusion models* (LDMs) are a popular framework, first introduced for 740 image generation (Rombach et al., 2022). They extend traditional diffusion models by operating in a 741 compressed latent space rather than directly in pixel space. An autoencoder—typically a VAE—is used to 742 encode high-resolution data into a lower-dimensional latent representation. The diffusion process then 743 unfolds within this latent space, substantially reducing computational cost, while the decoder reconstructs 744 fine-grained details. This paradigm underlies widely used models such as Stable Diffusion. More recently, 745 DiTs (Peebles and Xie, 2023) have been introduced as the main denoising architecture and have achieved 746 strong results across modalities (Esser et al., 2024; Joshi et al., 2025). Although some works have attempted 747 to adapt latent diffusion to graph generation, none have demonstrated competitive performance compared 748 to prior state-of-the-art discrete diffusion models (Fu et al., 2024; Nguyen et al., 2024; Yang et al., 2024a). 749 To the best of our knowledge, only Nguyen et al. (2024) evaluate the reconstruction performance of their 750 autoencoder, and only on the QM9 dataset, which contains very small graphs. Additionally, they employ an 751 ad-hoc transformer architecture that materializes the full attention score matrix, which limits computational 752 efficiency. While Zhou et al. (2024) appears to report competitive results on MOSES, it encodes the full 753 adjacency matrix through  $n^2$  edge tokens, with  $n$  the number of nodes, which prevents efficiency gains. 754 Other works are specific to molecule generation (Bian et al., 2024; Ketata et al., 2025; Wang et al., 2024). 755 Finally, some focus on tasks other than graph generation, such as link prediction (Fu et al., 2024) or 756 regression (Zhou et al., 2024). The key difference between these works and ours is that we design our 757 autoencoder to ensure that the adjacency matrix can be provably recovered from the latent space.

756 **B ADDITIONAL BACKGROUND AND NOTATIONS**  
 757

758 Here, we provide additional notation and background.  
 759

760 **B.1 NOTATIONS**  
 761

762 Let  $\mathbb{N} := \{1, 2, \dots\}$  and for  $n \in \mathbb{N}$  define  $[n] := \{1, \dots, n\} \subset \mathbb{N}$ . We denote by  $\mathbb{R}^{+*}$  the set of  
 763 strictly positive real numbers. We consider *node- and edge-labeled undirected graphs*  $G$ , i.e., tuples  
 764  $(V(G), E(G), \ell_x, \ell_e)$  with up to  $n$  nodes and  $m$  edges, without self-loops or isolated nodes. Here,  $V(G)$   
 765 is the set of nodes,  $E(G)$  the set of edges,  $\ell_x: V(G) \rightarrow [a]$ , for  $a \in \mathbb{N}$ , assigns each node one of  $a$   
 766 discrete labels, and  $\ell_e: E(G) \rightarrow [b]$ , for  $b \in \mathbb{N}$ , assigns each edge one of  $b$  discrete labels. An undirected  
 767 edge  $\{u, v\} \in E(G)$  is written either  $(u, v)$  or  $(v, u)$ . We also consider *node- and edge-labeled directed*  
 768 *acyclic graphs* (DAGs), defined analogously except that  $(u, v)$  denotes a directed edge from  $u$  to  $v$   
 769 and the graph contains no directed cycles. Throughout, we assume an arbitrary but fixed ordering  
 770 of nodes so that each node corresponds to a number in  $[n]$ . The adjacency matrix of  $G$  is denoted  
 771  $\mathbf{A}(G) \in \{0, 1\}^{n \times n}$ , where  $\mathbf{A}(G)_{ij} = 1$  if and only if nodes  $i$  and  $j$  are connected. We construct a node  
 772 feature vector  $\mathbf{X} \in [a]^n$  consistent with  $\ell_x$ , so that  $X_i = X_j$  if and only if  $\ell_x(i) = \ell_x(j)$ . Similarly,  
 773 we define an edge feature matrix  $\mathbf{E} \in [b]^m$  consistent with  $\mathbf{A}$  and  $\ell_e$ , so that  $E_{ij} = E_{kl}$  if and only  
 774 if  $\ell_e((i, j)) = \ell_e((k, l))$ . For any matrix  $\mathbf{M}$ , we denote by  $\mathbf{M}_i$  its  $i$ -th row. For a real-valued matrix  
 775  $\mathbf{M} \in \mathbb{R}^{n \times n}$ ,  $\mathbf{M}^\top$  denotes its transpose, and for a complex matrix  $\mathbf{N}$ ,  $\mathbf{N}^*$  denotes its Hermitian transpose.  
 776 For two binary matrices  $\mathbf{P}$  and  $\mathbf{Q}$ , we denote their logical disjunction  $\mathbf{P} \vee \mathbf{Q}$  where  $(\mathbf{P} \vee \mathbf{Q})_{ij} = 1$  if  
 777  $\mathbf{P}_{ij} = 1$  or  $\mathbf{Q}_{ij} = 1$ , and 0 otherwise. Let  $\mathbf{u}$  and  $\mathbf{v}$  be two vectors;  $\mathbf{u} \odot \mathbf{v}$  denotes the element-wise  
 778 product between them. Finally, the *graph Laplacian* is  $\mathbf{L} := \mathbf{D} - \mathbf{A}$ , where  $\mathbf{D} \in \mathbb{N}^{n \times n}$  is the degree  
 779 matrix. Its eigenvalue decomposition is  $\mathbf{L} := \mathbf{U} \Lambda \mathbf{U}^\top$ , where  $\Lambda := (\lambda_1, \dots, \lambda_n)$  is the vector of  
 780 eigenvalues (possibly repeated) and  $\mathbf{U} \in \mathbb{R}^{n \times n}$  the matrix of eigenvectors, with the  $i$ -th column  $\mathbf{U}_i^\top$   
 781 corresponding to eigenvalue  $\lambda_i$ . We denote by  $\mathbf{I}_d$  the  $d$ -dimensional identity matrix.  
 782

783 **Flow Matching** *Flow Matching* (FM) (Lipman et al., 2023) is a generative modeling framework that  
 784 transports samples from a base distribution  $p_0$  to a target distribution  $p_1$  by gradually denoising them.  
 785 The method relies on a time-dependent velocity field  $v_t: [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ , which defines a *flow map*  
 786  $\psi_t: [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$  through the ordinary differential equation  
 787

$$\frac{d}{dt} \psi_t(\mathbf{x}) = v_t(\psi_t(\mathbf{x})), \quad \psi_0(\mathbf{x}) := \mathbf{x}.$$

788 The flow induces a family of intermediate distributions  $(p_t)_{t \in [0, 1]}$ , where  $p_t$  is the *pushforward* of  $p_0$  by  
 789  $\psi_t$ . In other words, if  $\mathbf{x}_0 \sim p_0$ , then  $\psi_t(\mathbf{x}_0) \sim p_t$ , and in particular  $\psi_1(\mathbf{x}_0) \sim p_1$ . The *probability path*  
 790  $(p_t)$  is arbitrary as long as it is differentiable in time and satisfies the boundary conditions at  $t = 0$  and  
 791  $t = 1$ .  
 792

793 In practice, the true velocity field  $v_t$  is intractable. FM therefore relies on conditional probability paths  
 794  $p_t(\mathbf{x} | \mathbf{x}_1)$  for which the conditional velocity  $u_t(\mathbf{x} | \mathbf{x}_1)$  can be computed in closed form. A standard  
 795 choice is the linear interpolation  
 796

$$\mathbf{x}(t) := (1 - t)\mathbf{x}_0 + t\mathbf{x}_1, \quad \mathbf{x}_0 \sim p_0, \quad \mathbf{x}_1 \sim p_1,$$

797 with corresponding velocity  
 798

$$u_t(\mathbf{x} | \mathbf{x}_1) := \frac{\mathbf{x}_1 - \mathbf{x}}{1 - t}.$$

800 To train a generative model, the intractable  $v_t$  is replaced by a neural network  $v_\theta(t, \mathbf{x})$ , with parameters  
 801  $\theta$  that learns to match the conditional velocities  $u_t$ . This is achieved by minimizing the *conditional*  
 802 *flow-matching loss*  
 803

$$\mathcal{L}(\theta) := \mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_1} \left[ \|v_\theta(t, (1 - t)\mathbf{x}_0 + t\mathbf{x}_1) - (\mathbf{x}_1 - \mathbf{x}_0)\|^2 \right].$$

804 This formulation avoids integrating the ODE during training, which makes FM both efficient and  
 805 straightforward to implement. After training, new samples are obtained by combining the learned velocity  
 806 field  $v_\theta$  from  $t = 0$  to  $t = 1$ , yielding  $\mathbf{x}_1 \sim p_1$  deterministically. We refer to Lipman et al. (2024) for a  
 807 comprehensive description of the framework. Since flow matching and Gaussian diffusion are essentially  
 808 equivalent, we use the two terms interchangeably (Gao et al., 2024).  
 809

**Variational Autoencoders** A *Variational Autoencoder* (VAE), also referred to as a KL-regularized autoencoder, is a latent-variable generative model that extends the autoencoding framework with a probabilistic formulation. Given data  $\mathbf{x} \in \mathbb{R}^{n \times d}$ , the encoder parameterizes an approximate posterior  $q_\phi(\mathbf{z} | \mathbf{x})$ , where  $\mathbf{z}$  denotes the latent,  $\mathbf{z} \in \mathbb{R}^{n \times p}$ ,  $p \ll d$ . The decoder defines the conditional likelihood  $p_\theta(\mathbf{x} | \mathbf{z})$ . Training maximizes the *Evidence Lower Bound* (ELBO),

$$\mathcal{L}(\theta, \phi; \mathbf{x}) := \mathbb{E}_{q_\phi(\mathbf{z} | \mathbf{x})} [\log p_\theta(\mathbf{x} | \mathbf{z})] - D_{\text{KL}}(q_\phi(\mathbf{z} | \mathbf{x}) \| p(\mathbf{z})) ,$$

where the first term enforces reconstruction fidelity and the second regularizes the posterior toward a prior  $p(\mathbf{z})$ , typically  $\mathcal{N}(0, \mathbf{I}_p)$ . To enable backpropagation through the latents, the reparameterization trick is employed, i.e, the encoder predicts the mean and log-variance of the posterior distribution,  $\boldsymbol{\mu}_\phi$  and  $\log \boldsymbol{\sigma}_\phi$ , and  $\mathbf{z}$  is sampled according to

$$\mathbf{z} = \boldsymbol{\mu}_\phi(\mathbf{x}) + \boldsymbol{\sigma}_\phi(\mathbf{x}) \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}_p).$$

This formulation yields a continuous and structured latent space, supporting interpolation and generative sampling.

**GIN** A *Graph Isomorphism Network* (GIN) is a message-passing graph neural network designed to maximize expressive power under the neighborhood aggregation framework (Xu et al., 2019). Let each node  $i$  have features  $\mathbf{h}_i^{(l)}$  at layer  $l > 0$ , and let  $N(i)$  denote its neighbors. GIN updates node representations via

$$\mathbf{h}_i^{(l+1)} := \text{MLP}^{(l+1)} \left( (1 + \varepsilon) \mathbf{h}_i^{(l)} + \sum_{j \in N(i)} \mathbf{h}_j^{(l)} \right),$$

and  $\mathbf{h}_i^{(0)}$  is the initial node feature after being fed through an MLP. A natural extension, termed GINE, incorporates edge features into this aggregation by modifying each neighbor’s contribution, formally,

$$\mathbf{h}_i^{(l+1)} := \text{MLP}^{(l+1)} \left( (1 + \varepsilon) \mathbf{h}_i^{(l)} + \sum_{j \in N(i)} (\mathbf{h}_j^{(l)} + \mathbf{e}_{j,i}^{(l)}) \right),$$

where  $\mathbf{e}_{j,i}^{(l)}$  denotes the edge feature between nodes  $j$  and  $i$ ; this model retains the aggregation mechanism of GIN while leveraging edge-level information.

**Diffusion Transformers (DiT)** A *Diffusion Transformer* (DiT) is a diffusion-based generative model in which the denoising backbone is a Transformer model. Importantly, conditioning signals—such as the timestep  $t$  or additional context  $c$ —are incorporated via adaptive layer normalization (AdaLN) within each block, following a formulation of the form,

$$\text{AdaLN}(\mathbf{h}, c) := \gamma(c) \text{LayerNorm}(\mathbf{h}) + \beta(c),$$

where  $\gamma(c)$  and  $\beta(c)$  are learned affine modulations derived from  $c$ . This injection method preserves the scalability and expressive capacity of standard Transformers while enabling effective conditional generation across data modalities.

## C METHOD DETAILS AND PROOFS

Here, we describe additional details and outline missing proof from the main paper.

### C.1 DECODING EDGE LABELS

As a practical solution to decode edge labels along with the graph structure, we use a multi-headed version of the bilinear layer, where  $(\mathbf{W}_h^Q, \mathbf{W}_h^K) \in \mathbb{R}^{d \times d}$ ,  $h \in [b + 1]$  define  $b$  parallel projection heads, corresponding to the  $b$  possible edge labels, and an additional head for the absence of an edge. Formally,

$$\tilde{\mathbf{A}}_h := \frac{1}{\sqrt{d}} \mathbf{Z} \mathbf{W}_h^Q (\mathbf{Z} \mathbf{W}_h^K)^\top \in \mathbb{R}^{n \times n}, \quad h \in [b + 1],$$

$$\tilde{\mathbf{A}} := \left[ \text{DeepSet} \left( \tilde{\mathbf{A}}_0 \right), \dots, \text{DeepSet} \left( \tilde{\mathbf{A}}_b \right) \right] \in \mathbb{R}^{n \times n \times (b+1)},$$

$$\hat{\mathbf{A}} := \mathcal{D}_e(\mathbf{Z}) = \text{softmax}(\tilde{\mathbf{A}}) \in \mathbb{R}^{n \times n \times (b+1)},$$

where the softmax is applied over the last dimension of  $\hat{\mathbf{A}}$ . At inference, the node (respectively edge) labels are given by the argmax over the scores across the  $a$  (resp.  $b + 1$ ) classes.

864 C.2 FORMAL DEFINITION OF DAG AUTO-ENCODER  
865866 In this section, we formally define the architecture of our auto-encoder for DAGs.  
867868 **Encoder** We instantiate  $\rho$ , as GIN architecture, wrapped into a PYTORCH GEOMETRIC’s  
869 DirGNNConv (Rossi et al., 2023), which use one layer for in-neighbors and one for out-neighbors. We  
870 implement  $\phi$  as a row-wise MLP.  
871872 The encoder is defined in a similar way to that for undirected graphs, except it takes both the real and  
873 imaginary parts of the eigenvectors as input. Given a DAG  $G$  with node features  $\mathbf{X} \in \mathbb{R}^n$ , edge features  
874  $\mathbf{E} \in \mathbb{R}^{n \times n}$ , and magnetic Laplacian  $\mathbf{L}_q = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^*$ , the encoder produces the mean  $\boldsymbol{\mu}_Z \in \mathbb{R}^d$  and  
875 log-variance  $\log \sigma_Z^2 \in \mathbb{R}^d$  of a Gaussian variational posterior:  
876

877 
$$\mathbf{H}_\phi := \sum_{i=1}^n \phi(\mathbf{U}^R, \mathbf{U}^I, \mathbf{\Lambda})_i^\top \in \mathbb{R}^{n \times d},$$
  
878 
$$\mathbf{H}_X := \mathbf{X} \mathbf{W}^X + \mathbf{b}^X \in \mathbb{R}^{n \times d},$$
  
879 
$$\mathbf{H}_E := \mathbf{E} \mathbf{W}^E + \mathbf{b}^E \in \mathbb{R}^{m \times d},$$
  
880 
$$\boldsymbol{\mu}_Z, \log \sigma_Z^2 = \mathcal{E}(G, \mathbf{X}, \mathbf{E}, \mathbf{U}^R, \mathbf{U}^I, \mathbf{\Lambda}) = \rho(\mathbf{H}_X, \mathbf{H}_\phi, \mathbf{H}_E).$$
  
881

882 A latent sample  $\mathbf{Z} \in \mathbb{R}^{n \times d}$  is then obtained using the reparameterization trick,  
883

884 
$$\mathbf{Z} = \boldsymbol{\mu}_Z + \sigma_Z \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}_d),$$
  
885

886 where  $\sigma_Z = \exp(\frac{1}{2} \log \sigma_Z^2)$ .  
887888 **Decoder** Similarly to the undirected case, the node decoder  $\mathcal{D}_x$  reconstructs node labels directly from the  
889 latents  $\mathbf{Z}$ . Then an edge decoder  $\mathcal{D}_e$  is defined consistently with our definition of out-adjacency-identifying  
890 encodings, as follows,  
891

892 
$$\hat{\mathbf{X}} := \mathcal{D}_x(\mathbf{Z}) = \mathbf{Z} \mathbf{W}^{D_x} + \mathbf{b}^{D_x}, \quad \hat{\mathbf{X}} \in \mathbb{R}^{n \times a},$$
  
893 
$$\tilde{\mathbf{Z}}_R := \frac{1}{\sqrt{d}} \mathbf{Z} \mathbf{W}_R^Q (\mathbf{Z} \mathbf{W}_R^K)^\top$$
  
894 
$$\tilde{\mathbf{Z}}_I := \frac{1}{\sqrt{d}} \mathbf{Z} \mathbf{W}_I^Q (\mathbf{Z} \mathbf{W}_I^K)^\top$$
  
895 
$$\tilde{\mathbf{Z}} := \tilde{\mathbf{Z}}_R + \frac{2 - c}{s} \tilde{\mathbf{Z}}_I$$
  
896 
$$\hat{\mathbf{A}} := \mathcal{D}_e(\mathbf{Z}) = \sigma \left( \text{DeepSet} \left( \tilde{\mathbf{Z}} \right) \right), \quad \hat{\mathbf{A}} \in \mathbb{R}^{n \times n}$$
  
897

898 where  $c := \cos(2\pi q)$  and  $s := \sin(2\pi q)$ ,  $\mathbf{W}^X \in \mathbb{R}^{d \times a}$  and  $\mathbf{W}_R^Q, \mathbf{W}_R^K, \mathbf{W}_I^Q, \mathbf{W}_I^K \in \mathbb{R}^{d \times d}$  are a  
899 learnable weight matrices and the MLP is applied row-wise.  
900901 C.3 PROOFS  
902903 We first prove the following lemma, which provides a sufficient condition for node embeddings to be  
904 out-adjacency-identifying.  
905906 **Lemma 2.** *Let  $G$  be a directed graph with magnetic Laplacian  $\mathbf{L}_q$ . Let  $\mathbf{P} \in \mathbb{R}^{n \times d}$  a matrix of node  
907 embeddings. If there exists  $\mathbf{W}^{K_R}, \mathbf{W}^{Q_R}, \mathbf{W}^{K_I}, \mathbf{W}^{Q_I} \in \mathbb{R}^{d \times d}$  such that*  
908

909 
$$\tilde{\mathbf{P}}^R = \frac{1}{\sqrt{d}} (\mathbf{P} \mathbf{W}^{K_R}) (\mathbf{P} \mathbf{W}^{Q_R})^\top = \Re(\mathbf{L}_q),$$
  
910 
$$\tilde{\mathbf{P}}^I = \frac{1}{\sqrt{d}} (\mathbf{P} \mathbf{W}^{K_I}) (\mathbf{P} \mathbf{W}^{Q_I})^\top = \Im(\mathbf{L}_q),$$
  
911

912 then  $\mathbf{P}$  is out-adjacency-identifying.  
913

918 *Proof.* Note that if there exist  $\mathbf{W}^{K_R}, \mathbf{W}^{Q_R}, \mathbf{W}^{K_I}, \mathbf{W}^{Q_I} \in \mathbb{R}^{d \times d}$  such that

$$\begin{aligned} 919 \quad \tilde{\mathbf{P}}^R &= \Re(\mathbf{L}_q), \\ 920 \quad \tilde{\mathbf{P}}^I &= \Im(\mathbf{L}_q), \end{aligned}$$

923 there also exists  $\mathbf{W}_*^{K_R}, \mathbf{W}_*^{Q_R}, \mathbf{W}_*^{K_I}, \mathbf{W}_*^{Q_I} \in \mathbb{R}^{d \times d}$  such that

$$\begin{aligned} 924 \quad \tilde{\mathbf{P}}^R &= -\Re(\mathbf{L}_q), \\ 925 \quad \tilde{\mathbf{P}}^I &= -\Im(\mathbf{L}_q). \end{aligned}$$

927 The negative magnetic Laplacian is  $-\mathbf{L}_q = \mathbf{A}_s \odot e^{i\Theta^{(q)}} - \mathbf{D}_s$ . Note that since  $\mathbf{D}_s$  is a diagonal matrix,  
928 it does not affect the off-diagonal coefficients of  $-\mathbf{L}_q$ . Denote  $c := \cos(2\pi q)$  and  $s := \sin(2\pi q)$ , and  
929 consider the different cases.

- 931 1.  $(i, j) \in E(G)$ :  $-\mathbf{L}_{q,ij} = \mathbf{A}_{s,ij} e^{i2\pi q} = e^{i2\pi q}$ . Therefore,  $-\Re(\mathbf{L}_q)_{ij} = c$  and  $-\Im(\mathbf{L}_q)_{ij} = s$ .  
932 In that case,  $\tilde{\mathbf{P}}_{ij} = c + \frac{2-c}{s}s = 2$ .
- 934 2.  $(j, i) \in E(G)$ :  $-\mathbf{L}_{q,ij} = \mathbf{A}_{s,ij} e^{-i2\pi q} = e^{-i2\pi q}$ . Therefore,  $-\Re(\mathbf{L}_q)_{ij} = c$  and  
935  $-\Im(\mathbf{L}_q)_{ij} = -s$ . In that case,  $\tilde{\mathbf{P}}_{ij} = c - \frac{2-c}{s}s = 2(c-1) < 0$ .
- 937 3.  $(i, j), (j, i) \notin E(G)$ :  $-\mathbf{L}_{q,ij} = 0 = -\Re(\mathbf{L}_q)_{ij} = -\Im(\mathbf{L}_q)_{ij}$ . In that case,  $\tilde{\mathbf{P}}_{ij} = 0$ .
- 939 4. Since we consider DAGs, the case where  $(i, j) \in E(G)$  and  $(j, i) \in E(G)$ , i.e.,  $G$  contains a  
940 2-node cycle, never occurs.

941 The maximum over a row of  $\tilde{\mathbf{P}}$  is therefore 2, and we have that

$$943 \quad \tilde{\mathbf{P}}_{ij} = \max_k \tilde{\mathbf{P}}_{ik} \iff \mathbf{A}_{ij} = 1,$$

945 which concludes the proof. □

947 The following uses Lemma 2 as a sufficient condition for identifying out-adjacency.

948 **Lemma 3.** *Let  $G$  be a directed graph with magnetic Laplacian  $\mathbf{L}_q$ , with eigendecomposition  $\mathbf{L}_q = \mathbf{U} \mathbf{A} \mathbf{U}^*$   
949 and let  $\mathbf{X} + i\mathbf{Y} = \mathbf{U} \mathbf{A}^{1/2}$ . Then for any permutation matrices  $\mathbf{M}_X, \mathbf{M}_Y \in \mathbb{R}^{n \times n}$ , the matrix  
950  $\mathbf{P} = [\mathbf{X} \mathbf{M}_X, \mathbf{Y} \mathbf{M}_Y] \in \mathbb{R}^{n \times 2n}$  is out-adjacency-identifying.*

952 *Proof.* The idea is to find a pair of suitable bilinear projections to invoke Lemma 2. Note that since  $\mathbf{M}_X$   
953 and  $\mathbf{M}_Y$  are permutation matrices,  $\mathbf{M}_X \mathbf{M}_X^\top = \mathbf{M}_Y \mathbf{M}_Y^\top = \mathbf{I}_n$ . For  $(\mathbf{W}^{K_R}, \mathbf{W}^{Q_R})$ , we choose

$$\begin{aligned} 955 \quad \mathbf{W}^{K_R} &= \sqrt{d} \mathbf{I}_{2n}, \\ 956 \quad \mathbf{W}^{Q_R} &= \mathbf{I}_{2n}. \end{aligned}$$

957 Therefore,  $\frac{1}{\sqrt{d}} (\mathbf{P} \mathbf{W}^{K_R}) (\mathbf{P} \mathbf{W}^{Q_R})^\top = \mathbf{X} \mathbf{M}_X \mathbf{M}_X^\top \mathbf{X}^\top + \mathbf{Y} \mathbf{M}_Y \mathbf{M}_Y^\top \mathbf{Y}^\top = \mathbf{X} \mathbf{X}^\top + \mathbf{Y} \mathbf{Y}^\top =$   
958  $\Re(\mathbf{L}_q)$ . For  $(\mathbf{W}^{K_I}, \mathbf{W}^{Q_I})$ , we choose

$$\begin{aligned} 961 \quad \mathbf{W}^{K_I} &= \sqrt{d} \begin{bmatrix} \mathbf{0} & \mathbf{M}_X^\top \\ \mathbf{M}_Y^\top & \mathbf{0}, \end{bmatrix}, \\ 963 \quad \mathbf{W}^{Q_I} &= \begin{bmatrix} \mathbf{M}_X^\top & \mathbf{0} \\ \mathbf{0} & -\mathbf{M}_Y^\top \end{bmatrix}. \end{aligned}$$

965 In that case,  $\frac{1}{\sqrt{d}} (\mathbf{P} \mathbf{W}^{K_I}) (\mathbf{P} \mathbf{W}^{Q_I})^\top = [\mathbf{Y} \mathbf{M}_Y \mathbf{M}_Y^\top, \mathbf{X} \mathbf{M}_X \mathbf{M}_X^\top] [\mathbf{X} \mathbf{M}_X \mathbf{M}_X^\top, -\mathbf{Y} \mathbf{M}_Y \mathbf{M}_Y^\top]^\top =$   
966  $[\mathbf{Y}, \mathbf{X}] [\mathbf{X}, -\mathbf{Y}]^\top = \mathbf{Y} \mathbf{X}^\top - \mathbf{X} \mathbf{Y}^\top = \Im(\mathbf{L}_q)$ .

968 Applying Lemma 2 concludes the proof. □

970 We are now ready to prove the main theorem.

971 **Theorem 4.** *Theorem 1, restated) The mLPE is sufficiently out-adjacency-identifying.*

972 *Proof.* Let us start by recalling some notations. Let  $G$  be a  $n$ -order directed acyclic graph, with magnetic  
 973 Laplacian  $\mathbf{L}_q = \mathbf{U} \Lambda \mathbf{U}^*$ . Further write  $\mathbf{U} \Lambda^{1/2} = \mathbf{X} + i\mathbf{Y}$ , and denote  $\mathbf{U}^R = \Re(\mathbf{U}) \in \mathbb{R}^n$  the real  
 974 part of  $\mathbf{U}$  and  $\mathbf{U}^I = \Im(\mathbf{U}) \in \mathbb{R}^n$  the imaginary part of  $\mathbf{U}$ . In what follows,  $j$  refers to a column of  $\mathbf{U}^R$   
 975 or  $\mathbf{U}^I$ , and  $i$  refers to the index of a node. We denote  $u_{ij}^R$  the  $i$ -th component of  $j$ -th eigenvector of  $\mathbf{U}^R$ ,  
 976 and similarly for  $u_{ij}^I$ .  
 977

978 We divide up the  $d$ -dimensional space of the mLPE into two  $\frac{d}{2}$ -dimensional sub-spaces,  $\text{adj}_{Re}$  and  $\text{adj}_{Im}$ ,  
 979 i.e. we write, for a node  $v_i$ ,

$$980 \text{mLPE}(v) = [\text{adj}_{Re}(v_i), \text{adj}_{Im}(v_i)] \in \mathbb{R}^d, \\ 981$$

982 where

$$983 \text{adj}_{Re}(v_i) = \rho_{Re} \left( \sum_{j=1}^n \phi_{Re}(u_{ij}^R, u_{ij}^I, \lambda_j + \epsilon_j) \right), \\ 984 \\ 985 \text{adj}_{Im}(v_i) = \rho_{Im} \left( \sum_{j=1}^n \phi_{Im}(u_{ij}^R, u_{ij}^I, \lambda_j + \epsilon_j) \right), \\ 986 \\ 987 \\ 988$$

989 and where we have chosen  $\rho$  to be a sum over the first dimension of its input, followed by a feed-forward  
 990 neural network  $\rho_{Re}$  and  $\rho_{Im}$ . For convenience, we also fix an arbitrary ordering over the nodes in  $V(G)$   
 991 and define  
 992

$$993 \mathbf{Q}^R := \begin{bmatrix} \text{adj}_{Re}(v_1) \\ \vdots \\ \text{adj}_{Re}(v_n) \end{bmatrix} \\ 994 \\ 995 \mathbf{Q}^I := \begin{bmatrix} \text{adj}_{Im}(v_1) \\ \vdots \\ \text{adj}_{Im}(v_n) \end{bmatrix}, \\ 996 \\ 997 \\ 998 \\ 999$$

1000 where  $v_i$  is the  $i$ -th node in our ordering. Note that both  $\text{adj}_{Re}(v_i)$  and  $\text{adj}_{Im}(v_i)$  are DeepSets over the  
 1001 multiset

$$1002 M_i := \{(u_{ij}^R, u_{ij}^I, \lambda_j + \epsilon_j)\}.$$

1003 To prove that  $[\mathbf{Q}^R, \mathbf{Q}^I]$  is sufficiently out-adjacency-identifying, we show that  $[\mathbf{Q}^R, \mathbf{Q}^I]$  approximates  
 1004  $[\mathbf{X} \mathbf{M}_X, \mathbf{Y} \mathbf{M}_Y]$  arbitrarily close, for  $\mathbf{M}_X, \mathbf{M}_Y$  some permutation matrices. To that extent, we  
 1005 demonstrate how  $\text{adj}_{Re}(v_i)$  (resp.  $\text{adj}_{Im}(v_i)$ ) approximates  $\mathbf{X}_i$  (respectively  $\mathbf{Y}_i$ ) arbitrarily close for all  $i$ .  
 1006 Since  $\text{adj}_{Re}$  and  $\text{adj}_{Im}$  are DeepSets, they can approximate any continuous permutation-invariant function  
 1007 over a compact set over the real numbers. Therefore, it remains to show that there exists a permutation  
 1008 invariant function  $f^R$  (respectively  $f^I$ ) such that  $f^R(M_i) = \mathbf{X}_i \mathbf{M}_X$  (respectively  $f^I(M_i) = \mathbf{Y}_i \mathbf{M}_Y$ )  
 1009 for all  $i$  and some permutation matrix  $\mathbf{M}_X$  (respectively  $\mathbf{M}_Y$ ). To this end, note that the magnetic  
 1010 Laplacian graph  $G$  of order  $n$  has at most  $n$  distinct eigenvalues. Hence, we can choose  $\epsilon_j$  such that  
 1011  $\lambda_j + \epsilon_j$  is unique, for all  $j$ . In particular, let

$$1012 \epsilon_j = j \cdot \delta, \\ 1013$$

1014 where we choose  $\delta < \min_{l,o} |\lambda_l - \lambda_o| > 0$ , i.e., the smallest non-zero difference between two eigenvalues.  
 1015 Now let

$$1016 f^R(M_i) := [u_{is_1}^R \cdot \sqrt{\lambda_1 + s_1 \cdot \delta}, \dots, u_{is_n}^R \cdot \sqrt{\lambda_n + t_n \cdot \delta}], \\ 1017 \\ 1018 f^I(M_i) := [u_{it_1}^I \cdot \sqrt{\lambda_1 + t_1 \cdot \delta}, \dots, u_{it_n}^I \cdot \sqrt{\lambda_n + t_n \cdot \delta}], \\ 1019$$

1020 where the  $\{s_k\}_{k=1}^n$  (respectively  $\{t_k\}_{k=1}^n$ ) are indices in  $[n]$  such that the  $\lambda_1 + s_k \cdot \delta$  (respectively  
 1021  $\lambda_1 + t_k \cdot \delta$ ) are sorted in ascending order. This ordering is reflected in some permutation matrix  $\mathbf{M}_X$   
 1022 (respectively  $\mathbf{M}_Y$ ).

1023 Then, note that, since  $\Lambda^{1/2}$  is a real-valued, diagonal matrix,  $\mathbf{X} = \Re(\mathbf{U} \Lambda^{1/2}) = \Re(\mathbf{U}) \Lambda^{1/2} = \mathbf{U}^R \Lambda^{1/2}$ ,  
 1024 similarly  $\mathbf{Y} = \Im(\mathbf{U} \Lambda^{1/2}) = \Im(\mathbf{U}) \Lambda^{1/2} = \mathbf{U}^I \Lambda^{1/2}$ . The  $i$ -th row of  $\mathbf{U}^R \Lambda^{1/2}$  is equal to  
 1025

$$1026 \mathbf{X}_i = (\mathbf{U}^R \Lambda^{1/2})_i = [u_{i1}^R \cdot \sqrt{\lambda_1}, \dots, u_{in}^R \cdot \sqrt{\lambda_n}],$$

1026 and, similarly,

$$1027 \quad \mathbf{Y}_i = (\mathbf{U}^T \Lambda^{1/2})_i = [u_{i1}^I \cdot \sqrt{\lambda_1}, \dots, u_{in}^I \cdot \sqrt{\lambda_n}].$$

1029 Since we can choose  $\delta$  arbitrarily small, we can choose it such that:

$$1030 \quad \|f^R(M_i) - X_i\| < \gamma_R,$$

$$1032 \quad \|f^I(M_i) - Y_i\| < \gamma_I,$$

1033 for any  $\gamma_R, \gamma_I > 0$ . Since  $f^R$  and  $f^I$  are continuous permutation-invariant functions over a compact set, a  
1034 Deepset can approximate them arbitrarily close, and as a result, we have  
1035

$$1036 \quad \|\text{adj}_{Re}(v_i) - \mathbf{X}_i \mathbf{M}_X\| < \gamma_R,$$

$$1037 \quad \|\text{adj}_{Im}(v_i) - \mathbf{Y}_i \mathbf{M}_Y\| < \gamma_I,$$

1038 for all  $i$  and arbitrarily small  $\gamma^R, \gamma^I$ . In matrix form, it gives  
1039

$$1040 \quad \|\mathbf{Q}^R - \mathbf{X} \mathbf{M}_X\| < \gamma_R,$$

$$1041 \quad \|\mathbf{Q}^I - \mathbf{Y} \mathbf{M}_Y\| < \gamma_I,$$

1042 and, further  
1043

$$1044 \quad \|[\mathbf{Q}^R, \mathbf{Q}^I] - [\mathbf{X} \mathbf{M}_X, \mathbf{Y} \mathbf{M}_Y]\| < \gamma,$$

1045 where  $\gamma = \max(\gamma_R, \gamma_I)$ . We know by Lemma 3 that  $[\mathbf{X} \mathbf{M}_X, \mathbf{Y} \mathbf{M}_Y]$  is out-adjacency-identifying.  
1046 Since the mLPE can approximate it arbitrarily close, we can, in turn, conclude that mLPE is sufficiently  
1047 out-adjacency-identifying.  $\square$

## 1048 D IMPLEMENTATION DETAILS

1050 Here, we outline implementation details.

### 1053 AUTOENCODER

1055 **Number of eigenvectors** Our definition of the LPE and mLPE assumes that we have access to all  $n$   
1056 eigenvectors. However, this might be impractical for large graphs. In our implementation, we only take the  
1057  $k$  smallest eigenvalues and their corresponding eigenvectors. In practice,  $\phi$  is therefore instantiated as

$$1058 \quad \phi(\mathbf{U}_{:,k}, \mathbf{A}_{:k}) := \left[ \phi(\mathbf{U}_{1,:k} \parallel \mathbf{A}_{:k} + \epsilon), \dots, \phi(\mathbf{U}_{n,:k} \parallel \mathbf{A}_{:k} + \epsilon) \right] \in \mathbb{R}^{n \times k \times d},$$

1060 where  $\mathbf{U}_{:,k}$  denotes the  $k$  first columns of  $\mathbf{U}$ ,  $\mathbf{A}_{:k}$  the  $k$  first eigenvalues,  $\mathbf{U}_{i,:k}$  the  $k$  first columns of  $i$ -th  
1061 row of  $\mathbf{U}$ .

1062 **Sign invariance** The eigenvectors of the graph Laplacian are not unique, e.g., if  $\mathbf{v}$  is an eigenvector, then  
1063 so is  $-\mathbf{v}$  (Lim et al., 2023). Hence, GNNs should ideally be *sign invariant*, i.e., represent  $\mathbf{v}$  and  $-\mathbf{v}$   
1064 identically. To account for this symmetry, we adopt a SignNet-like mechanism (Lim et al., 2023), which  
1065 randomly flips the signs of eigenvectors during training.

1067 **Decoder DeepSet** Recall the definition of our edge decoder,

$$1069 \quad \tilde{\mathbf{Z}} := \frac{1}{\sqrt{d}} \mathbf{Z} \mathbf{W}^Q (\mathbf{Z} \mathbf{W}^K)^\top$$

$$1071 \quad \hat{\mathbf{A}} := \sigma(\text{DeepSet}(\tilde{\mathbf{Z}})), \quad \hat{\mathbf{A}} \in \mathbb{R}^{n \times n}.$$

1072 Here, we describe the architecture of the DeepSet. Let  
1073

$$1075 \quad \mathbf{C} := \frac{1}{n} \sum_{i=1}^n (\tilde{\mathbf{Z}} \mathbf{W}^C + \mathbf{b}^C)_{:,i} \in \mathbb{R}^{n \times d_{DS}}$$

$$1078 \quad \text{DeepSet}(\tilde{\mathbf{Z}}) := \mathbf{W}^{\text{out}} \left( \text{GeLU} \left( \tilde{\mathbf{Z}} \mathbf{W}^{\text{in}} + \mathbf{C} \mathbf{W}^{\text{ctx}} + \mathbf{b}^{\text{in}} \right) \right) + \mathbf{b}^{\text{out}},$$

1079 where

- $\mathbf{W}^C \in \mathbb{R}^{1 \times d_{DS}}$  and  $\mathbf{b}^C \in \mathbb{R}^{d_{DS}}$  are learnable parameters; the summation is taken over the projected input rows.
- $\mathbf{W}^{\text{in}} \in \mathbb{R}^{1 \times d_{DS}}$ ,  $\mathbf{W}^{\text{ctx}} \in \mathbb{R}^{d_{DS} \times d_{DS}}$ , and  $\mathbf{b}^{\text{in}} \in \mathbb{R}^{d_{DS}}$  are learnable parameters. We view  $\tilde{\mathbf{Z}}$  as a tensor in  $\mathbb{R}^{d_{DS} \times d_{DS} \times 1}$ , yielding a projection in  $\mathbb{R}^{n \times n \times d_{DS}}$ , with  $\mathbf{C}\mathbf{W}^{\text{ctx}}$  broadcast across columns.
- $\mathbf{W}^{\text{out}} \in \mathbb{R}^{d_{DS} \times 1}$  and  $\mathbf{b}^{\text{out}} \in \mathbb{R}$  are the final learnable parameters.

**On the problem of structurally equivalent nodes** Our encoder assigns latent representations to nodes. To avoid handling permutations in the decoder and the loss function, we preserve the nodes’ ordering when feeding graphs to the autoencoder. Consequently, a reconstructed edge is considered correct by the loss function only if it exactly matches an edge in the original graph. However, in some situations, this assumption becomes problematic.

Consider the case illustrated in Figure 4. Structurally equivalent nodes are colored identically; they belong to the same *orbit*. Structurally equivalent nodes also correspond to identical rows in the eigenvector matrix  $\mathbf{U}$ , which in turn means that  $\phi$  will assign them identical representations. Since GNNs aggregate information based on neighborhoods, such nodes inevitably receive identical embeddings from  $\rho$  as well. For example, the encoder assigns the same representation to all four yellow nodes. The same applies to the two red nodes, which also form an orbit. Consequently, when decoding edges between yellow and red nodes, the decoder should treat a yellow node as equally likely to connect to either red node. Although both predictions are structurally correct, the preserved node ordering forces the loss function to mark one of them as incorrect. Similar observations were made in (Laabid et al., 2025; Morris et al., 2024).

We experimentally found that this would lead to poor reconstruction performance, particularly in trees, where cases similar to Figure 4 are common. Therefore, we designed a principled way to perturb the input of nodes belonging to non-trivial orbits (i.e., non-singleton orbits) to make them distinguishable in the latent space and ease the decoder’s task.

During preprocessing, we identify non-trivial orbits using the final node colorings from the 1-WL algorithm. Although 1-WL does not always distinguish non-equivalent nodes, prior work (Morris et al., 2024) suggests that identical 1-WL colorings serve as a good practical solution to detect structural equivalence. Before feeding eigenvectors to  $\phi$ , we slightly perturb the rows of  $\mathbf{U}$  corresponding to nodes in non-trivial orbits by applying a low-magnitude Gaussian modulation. Formally,

$$\mathbf{U}_{\text{perturb}} = (1 + \eta \mathbf{M} \boldsymbol{\epsilon}) \mathbf{U},$$

where  $\eta$  is a small modulation coefficient (typically  $\eta = 0.005$ ),  $\mathbf{M} \in \mathbb{R}^{n \times n}$  is a mask with rows equal to 1 if a node belongs to a non-trivial orbit and 0 otherwise, and  $\boldsymbol{\epsilon}$  is a matrix whose rows are independent Gaussian random vectors.

In simple terms, this modulation injects Gaussian noise into the rows of  $\mathbf{U}$  corresponding to nodes in non-trivial orbits. Since nodes in the same orbit have different inputs to  $\phi$ , the corresponding latents will be slightly different. This provides a simple, parsimonious way to make structurally equivalent nodes distinguishable in the latent space. In particular, we found it to yield much better reconstruction results than previously proposed symmetry-breaking methods, such as assigning node identifiers (Bechler-Speicher et al., 2025; You et al., 2019).

## D.2 DiT

**Self-conditioning** We implement self-conditioning (Chen et al., 2023a) in all our experiments. Self-conditioning provides the denoiser at step  $t$  with its own prediction from the previous step as an additional input. During training, this input is dropped with probability 0.5. When conditioning is active, the model first produces a prediction without conditioning, which is then reused as input for a second pass. Gradients are not backpropagated through the first prediction. We find that this approach consistently improves the quality of our results.

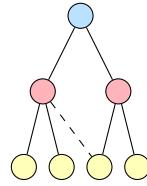


Figure 4: Example of failure case for the decoder. Solid lines denote original edges, while the dashed line denotes a structurally correct prediction that will be considered as incorrect by the loss function.

1134   **Logit-normal time step sampling** Following Esser et al. (2024), instead of sampling diffusion times  $t$   
 1135   uniformly from  $[0, 1]$ , we draw them from a logit-normal distribution. Formally,

$$1137 \quad \text{Logit}(t) = \log\left(\frac{t}{1-t}\right) \sim \mathcal{N}(\mu, \sigma^2), \quad t \in (0, 1).$$

1139   In practice, this is implemented by sampling

$$1140 \quad z \sim \mathcal{N}(\mu, \sigma^2), \quad t = \frac{1}{1 + e^{-z}}.$$

1142   The resulting density is

$$1144 \quad p(t) \propto \frac{1}{t(1-t)} \exp\left(-\frac{(\text{Logit}(t) - \mu)^2}{2\sigma^2}\right),$$

1146   which biases training towards intermediate time steps, the most informative region for flow matching  
 1147   models.

## 1149   E EXPERIMENTAL DETAILS

1152   Here, we outline details on the experiments.

### 1154   E.1 DATASET DETAILS

#### 1156   E.1.1 SYNTHETIC DATASETS

1157   **Description** We experiment on three synthetic datasets: EXTENDED TREE, EXTENDED PLANAR, and  
 1158   EGO. The first two extend the benchmarks introduced in Martinkus et al. (2022), namely TREE and  
 1159   PLANAR, which originally consisted of random trees and planar graphs with a fixed size of 64 nodes.  
 1160   These original datasets contained only 200 samples, split into 128/32/40 for train/validation/test, which is  
 1161   too small and leads to high variance during evaluation. Following Krimmel et al. (2025), we extend these  
 1162   benchmarks to 8,192 training samples and 256 samples each for validation and test. The EGO dataset,  
 1163   proposed in You et al. (2018), consists of 757 3-hop ego networks extracted from the Citeseer network,  
 1164   with up to 400 nodes.

1165   **Metrics** On synthetic graphs, we evaluate our method using the *Maximum Mean Discrepancy* (MMD)  
 1166   metric, as introduced in Martinkus et al. (2022). For each setting, we compute graph statistics on a set of  
 1167   generated samples and the test set, embed them via a kernel, and compute the MMD between the resulting  
 1168   distributions. We report MMD scores for the degree distribution (**Deg.**), counts of non-isomorphic  
 1169   four-node orbits (**Orbit**), clustering coefficient (**Cluster.**), and Laplacian spectrum (**Spec.**). For the  
 1170   EXTENDED TREE and EXTENDED PLANAR datasets, we also report the proportion of valid, unique, and  
 1171   novel graphs (**V.U.N.**), where validity is defined by the dataset constraints (e.g., planarity or tree structure).

#### 1173   E.1.2 MOLECULAR DATASETS

1174   **Description** We next evaluate our method on the Moses benchmark (Polykovskiy et al., 2020), derived  
 1175   from the ZINC Clean Leads collection (Sterling and Irwin, 2015), which comprises molecules with 8 to 27  
 1176   heavy atoms, filtered according to specific criteria. We also consider the Guacamol benchmark (Brown  
 1177   et al., 2019), based on the ChEMBL 24 database (Mendez et al., 2019). This dataset comprises synthesized  
 1178   molecules tested against biological targets, ranging in size from 2 to 88 heavy atoms. Consistent with prior  
 1179   work (Vignac et al., 2023; Siraudin et al., 2024), the GUACAMOL dataset undergoes a filtering procedure:  
 1180   SMILES strings are converted into graphs and then mapped back to SMILES, discarding molecules for  
 1181   which this conversion fails. We utilize the implementation presented in Vignac et al. (2023) for this process.  
 1182   The standard dataset splits are used.

1184   **Metrics** For both datasets, we evaluate validity (**Valid**), uniqueness (**Unique**), and novelty (**Novel**).  
 1185   Following Qin et al. (2025), for GUACAMOL we instead report the percentages of valid (**V.**), valid and  
 1186   unique (**V.U.**), and valid, unique, and novel (**V.U.N.**) samples. Each benchmark also includes specific  
 1187   metrics: MOSES compares generated molecules to a scaffold test set, reporting the **Fréchet ChemNet**  
 1188   **Distance (FCD)**, the **Similarity to the Nearest Neighbor (SNN)**—the average Tanimoto similarity

1188  
1189  
1190 Table 6: Hyperparameters of LG-VAE on the different datasets.  
1191  
1192  
1193  
1194  
1195  
1196  
1197  
1198  
1199  
1200  
1201  
1202  
1203  
1204  
1205  
1206  
1207  
1208  
1209  
1210  
1211  
1212  
1213  
1214  
1215  
1216  
1217  
1218  
1219  
1220  
1221  
1222  
1223  
1224  
1225  
1226  
1227  
1228  
1229  
1230  
1231  
1232  
1233  
1234  
1235  
1236  
1237  
1238  
1239  
1240  
1241

Dataset	Planar	Tree	Ego	Moses	GuacaMol	TPU Tile
Number of eigenvectors	16	64	64	16	32	64
Latent dim. $d$	16	24	24	24	24	24
Modulation magnitude	0	0.005	0.005	0.005	0.05	0.05
Number of layers $\phi$	2	2	2	2	2	2
Embedding dim $\phi$	256	256	256	256	384	256
Number of layers $\rho$	16	16	16	16	16	12
Embedding dim $\rho$	256	256	256	256	384	256
Decoder DeepSet embedding dim	16	16	16	24	24	24

between fingerprints of generated molecules and those of a reference set—and **Scaffold similarity (Scaf)**, which measures the frequency match of Bemis–Murcko scaffolds between generated and reference sets. Finally, **Filters** indicates the percentage of generated molecules passing the dataset’s construction filters. GUACAMOL reports two metrics: the **Fréchet ChemNet Distance (FCD)** and the **KL divergence (KL)** between distributions of physicochemical descriptors in the training and generated sets.

### E.1.3 DAG DATASET

**Description** We evaluate our DAG generation variant using the TPU Tiles dataset (Phothilimhana et al., 2023a). It consists of computational graphs extracted from Tensor Processing Unit workloads. We split the dataset into 5 040 training, 630 validation, and 631 test graphs.

**Metrics** We compute MMD over statistics analogous to those used for undirected graphs, replacing the degree distribution with incoming (**In Deg.**) and outgoing (**Out Deg.**) degree distributions, and adding features from the graph wavelet transform (**Wave.**). Samples are considered valid if they satisfy the DAG property.

### E.2 RESOURCES

All our models, both autoencoders and DiTs, are trained on two Nvidia L40 GPUs with 40GB VRAM. All models were sampled using a single Nvidia L40 GPU, and we optimized the batch size to fully utilise its memory.

### E.3 TRAINING DETAILS

All our autoencoders are trained using the AdamW optimizer, a weight decay of  $1e^{-4}$ , a KL coefficient of  $1e^{-6}$ , and a cosine learning rate schedule. We train our DiTs using a constant learning rate of  $2e^{-4}$ , and maintain Exponential Moving Average (EMA) weights for evaluation.

### E.4 HYPERPARAMETERS

**Autoencoder** We use a different number of eigenvectors depending on the dataset, as well as a different value for  $d$ . Our hyperparameters are listed in Table 6.

**DiT** We experiment using three variants of the DiT: two variants coming from the original DiT paper, DiT Small (DiT-S) and Base (DiT-B), and one smaller variant, coined DiT Tiny (DiT-T), since some of our datasets are smaller than those used in image generation. We detail their hyperparameters in Table 7.

### E.5 ADDITIONAL EXPERIMENTAL RESULTS

**Reconstruction performance** See Table 8 for results on reconstruction performance. Note that reconstruction performance on EGO is noticeably lower than on the other datasets. This can be attributed to two factors: the small dataset size (only 606 training samples) and the inclusion of the validation set within the training data, which limits reliable assessment of generalization during training. Nevertheless, EGO

1242

Table 7: Hyperparameters of DiT variants used in our experiments.

1243

1244

1245

1246

1247

1248

1249

1250

1251

does not impose strong structural constraints, and the **Edge. Acc.** remains high (0.9957). As a result, the latent diffusion model can still be trained and generate samples that statistically resemble the training distribution.

1254

1255

1256

1257

Table 8: Reconstruction accuracy of our autoencoder across datasets. **Edge Sample Acc.** and **Node Sample Acc.** denote the proportion of graphs in which, all edges and all nodes, respectively, are correctly reconstructed. For datasets without node labels, **Edge Sample Acc.** coincides with **Sample Acc.**.

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

Datasets	Edge Sample Acc.	Node Sample Acc.	Sample Acc.
Extended Planar	0.9961	—	0.9961
Extended Tree	0.9844	—	0.9844
Ego	0.0861	—	0.0861
Moses	0.9967	1.0	0.9967
GuacaMol	0.9973	1.0	0.9973
TPU Tile	0.9493	1.0	0.9493

**Ablation on the number of eigenvectors  $k$**  We conducted an ablation on the number of eigenvectors used in LG-VAE to better assess the influence of this hyperparameter on reconstruction and generative performance. To this end, we ran experiments on Extended Tree and Extended Planar with  $k \in [4, 8, 16, 32, 64]$ . Our results are presented in Tables 9 and 10.

On EXTENDED PLANAR, the number of eigenvectors  $k$  has only a minor influence on reconstruction accuracy, although MMD metrics worsen for small  $k$ , consistent with the loss of high-frequency structural information. In contrast, on EXTENDED TREE,  $k$  has a much stronger impact on reconstruction quality; interestingly, the best results are obtained with only 25–50% of the eigenvectors, showing that including high-frequencies do not necessarily yield better performance.

1278

1279

Table 9: **Ablation study** : number of eigenvectors on the EXTENDED PLANAR dataset.

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

Number of eigenvectors	Sample Acc.	Deg. (↓)	Orbit (↓)	Cluster. (↓)	Spec. (↓)	V.U.N. (↑)
4	0.9922	0.3±0.0	0.7±0.2	13.4±2.2	1.2±0.3	<b>99.4</b> ±0.7
8	0.9844	0.3±0.2	1.0±0.3	12.9±3.2	<b>1.0</b> ±0.1	99.9±0.1
16	<b>0.9961</b>	0.3±0.1	0.5±0.2	10.1±3.1	<b>1.0</b> ±0.2	99.0±0.2
32	0.9844	<b>0.2</b> ±0.0	<b>0.1</b> ±0.1	<b>7.6</b> ±1.2	1.1±0.1	99.1±0.3
64	0.9961	0.3±0.1	0.5±0.4	12.8±2.3	1.1±0.1	99.0±0.4

Table 10: **Ablation study** : number of eigenvectors on the EXTENDED TREE dataset.

1290

1291

1292

1293

1294

1295

Number of eigenvectors	Sample Acc.	Deg. (↓)	Orbit (↓)	Cluster. (↓)	Spec. (↓)	V.U.N. (↑)
4	0.4648	1.3±0.1	0.1±0.0	<b>0.0</b> ±0.0	6.1±0.4	42.3±3.2
8	0.4688	0.7±0.0	<b>0.0</b> ±0.0	<b>0.0</b> ±0.0	5.1±0.6	44.1±1.3
16	0.9727	<b>0.1</b> ±0.0	<b>0.0</b> ±0.0	<b>0.0</b> ±0.0	<b>1.1</b> ±0.1	92.3±1.2
32	<b>0.9883</b>	<b>0.1</b> ±0.0	<b>0.0</b> ±0.0	<b>0.0</b> ±0.0	1.7±0.2	<b>92.9</b> ±1.3
64	0.6680	<b>0.1</b> ±0.1	0.1±0.0	<b>0.0</b> ±0.0	1.6±0.2	49.9±2.8

1296 **Experiments on larger planar graphs** See results in Table 11. To evaluate scalability, we trained  
 1297 LG-VAE and LG-Flow on planar graphs of increasing sizes (128, 256). We denote by PLANAR 128 the  
 1298 dataset of planar graphs with 128 nodes, and similarly for 256 and 512. Recall that EXTENDED PLANAR  
 1299 originally contained graphs with 64 nodes.

1300 LG-Flow preserves high performance as graph size grows, achieving strong V.U.N scores and low MMDs,  
 1301 while its inference time increases sub-quadratically. In contrast, DeFoG’s performance degrades sharply  
 1302 with graph size, especially on V.U.N. Its quadratic complexity restricts batch sizes during training (6 for  
 1303 PLANAR 256, 2 for PLANAR 512), leading to slow convergence and long training times. LG-Flow, by  
 1304 contrast, supports substantially larger batches (256 for PLANAR 128 and PLANAR 256, 128 for PLANAR  
 1305 512). This efficiency yields speedups of  $340\times$ ,  $262\times$ , and  $470\times$  on the corresponding datasets.

1307 Table 11: **Reconstruction and generation on larger planar graphs.** We highlight best results in bold.

$ V $	Model	Sample Acc.	Deg. ( $\downarrow$ )	Orbit ( $\downarrow$ )	Cluster. ( $\downarrow$ )	Spec. ( $\downarrow$ )	V.U.N. ( $\uparrow$ )	Time (s) ( $\downarrow$ )
64	DeFog	–	<b>0.3</b> $\pm$ 0.1	<b>0.1</b> $\pm$ 0.1	<b>9.7</b> $\pm$ 0.3	1.3 $\pm$ 0.1	<b>99.6</b> $\pm$ 0.4	$8.9 \times 10^2$
	LG-Flow	0.9661	<b>0.3</b> $\pm$ 0.1	0.5 $\pm$ 0.2	10.1 $\pm$ 3.1	<b>1.0</b> $\pm$ 0.2	99.0 $\pm$ 0.2	<b>4.6</b> $\pm$ 0.
128	DeFog	–	<b>0.0</b> $\pm$ 0.0	0.7 $\pm$ 0.2	<b>3.1</b> $\pm$ 0.8	1.4 $\pm$ 0.2	91.9 $\pm$ 1.2	$3.8 \times 10^3$
	LG-Flow	1.0000	0.1 $\pm$ 0.0	<b>0.3</b> $\pm$ 0.2	4.4 $\pm$ 1.1	<b>0.4</b> $\pm$ 0.1	<b>99.1</b> $\pm$ 0.5	<b>11.2</b> $\pm$ 0.
256	DeFog	–	0.8 $\pm$ 0.1	28.6 $\pm$ 1.8	49.6 $\pm$ 3.1	2.9 $\pm$ 0.1	16.8 $\pm$ 2.7	$1.5 \times 10^4$
	LG-Flow	0.9727	<b>0.2</b> $\pm$ 0.1	<b>2.0</b> $\pm$ 0.8	<b>17.1</b> $\pm$ 8.0	<b>0.2</b> $\pm$ 0.0	<b>97.9</b> $\pm$ 0.5	<b>57.2</b> $\pm$ 0.0
512	DeFog	–	1.2 $\pm$ 0.1	56.5 $\pm$ 2.9	114.8 $\pm$ 7.7	1.5 $\pm$ 0.0	1.4 $\pm$ 0.8	$6.1 \times 10^4$
	LG-Flow	0.8945	<b>0.2</b> $\pm$ 0.1	<b>1.1</b> $\pm$ 0.4	<b>20.3</b> $\pm$ 6.3	<b>0.1</b> $\pm$ 0.0	<b>84.4</b> $\pm$ 3.2	<b>1.3</b> $\times 10^2$

1318 **Additional DAG dataset** In addition to TPU TILE, we evaluated LG-Flow on a synthetic dataset of  
 1319 DAGs generated using Price’s model, the DAG analogue of the Barabási–Albert model. We compare our  
 1320 generation results against Directo in Table 12. Both models were sampled with batch size 256 for a fair  
 1321 evaluation of inference efficiency. LG-Flow performs on par with Directo on this dataset while yielding a  
 1322  $244\times$  inference speed-up.

1324 Table 12: **Generation on DAGs: Price’s dataset.**

Method	Out Deg.	In Deg.	Cluster.	Spec.	Wave.	V.U.N	Time
Directo	2.5 $\pm$ 0.2	0.2 $\pm$ 0.0	9.1 $\pm$ 0.6	2.8 $\pm$ 0.4	0.5 $\pm$ 0.1	98.9 $\pm$ 0.3	$1.1 \times 10^3$
LG-Flow	2.7 $\pm$ 0.4	0.3 $\pm$ 0.1	12.6 $\pm$ 5.5	5.5 $\pm$ 4.2	0.6 $\pm$ 0.5	98.8 $\pm$ 0.7	4.5 $\pm$ 0.0

## 1330 E.6 EXPERIMENTAL DETAILS ON BASELINES

1332 For fair inference-time evaluation, we used a fixed batch size for all baselines; when this was not feasible,  
 1333 we tuned the batch size to optimize GPU usage. On EXTENDED PLANAR and EXTENDED TREE, all  
 1334 models were sampled with a batch size of 256. On Ego, discrete diffusion models could not fit the full  
 1335 151-sample test batch on GPU and were sampled with a batch size of 32. On GuacaMol and MOSES, all  
 1336 models were sampled with a batch size of 1000. On TPU Tile, we sampled both Directo and LG-Flow  
 1337 using a batch size of 40, i.e. the full test batch.

## 1339 F LIMITATIONS

1341 The primary limitation of our approach is the requirement for sufficient data to train the autoencoder to  
 1342 achieve nearly lossless reconstruction. Yet, most existing graph generation datasets are relatively small, as  
 1343 they were primarily designed for diffusion models operating directly in graph space—a computationally  
 1344 costly setting. Since our model scales much more efficiently, it opens the possibility of leveraging  
 1345 substantially larger datasets, which in turn mitigates this limitation. Training on parametric distributions  
 1346 such as Stochastic Block Models also proves difficult, since nodes are stochastically indistinguishable and  
 1347 thus structurally equivalent. Nevertheless, we argue that such datasets offer limited practical value, as these  
 1348 simplistic random models fail to capture the complexity of real-world graph distributions.