



PDF Download
3746252.3761205.pdf
01 January 2026
Total Citations: 0
Total Downloads: 59

Latest updates: <https://dl.acm.org/doi/10.1145/3746252.3761205>

RESEARCH-ARTICLE

GraphRCG: Self-Conditioned Graph Generation

SONG WANG, University of Virginia, Charlottesville, VA, United States

ZHEN TAN, Arizona State University, Tempe, AZ, United States

XINYU ZHAO, The University of North Carolina at Chapel Hill, Chapel Hill, NC, United States

TIANLONG CHEN, The University of North Carolina at Chapel Hill, Chapel Hill, NC, United States

HUAN LIU, Arizona State University, Tempe, AZ, United States

JUNDONG LI, University of Virginia, Charlottesville, VA, United States

Open Access Support provided by:

The University of North Carolina at Chapel Hill

Arizona State University

University of Virginia

Published: 10 November 2025

[Citation in BibTeX format](#)

CIKM '25: The 34th ACM International Conference on Information and Knowledge Management
November 10 - 14, 2025
Seoul, Republic of Korea

Conference Sponsors:
SIGWEB
SIGIR

GraphRCG: Self-conditioned Graph Generation

Song Wang
University of Virginia
Charlottesville, VA, USA
sw3wv@virginia.edu

Zhen Tan
Arizona State University
Tempe, AZ, USA
ztan36@asu.edu

Xinyu Zhao
UNC at Chapel Hill
Chapel Hill, NC, USA
xinyuzh@cs.unc.edu

Tianlong Chen
UNC at Chapel Hill
Chapel Hill, NC, USA
tianlong@cs.unc.edu

Huan Liu
Arizona State University
Tempe, AZ, USA
huanliu@asu.edu

Jundong Li
University of Virginia
Charlottesville, VA, USA
jundong@virginia.edu

Abstract

Graph generation aims to create new graphs that closely align with a target graph distribution. Existing works often implicitly capture this distribution by aligning the output of a generator with each training sample. As such, the overview of the entire distribution is not explicitly captured and used for graph generation. In contrast, in this work, we propose a novel self-conditioned graph generation framework designed to explicitly model graph distributions and employ these distributions to guide the generation process. We first perform self-conditioned modeling to capture the graph distributions by transforming each graph sample into a low-dimensional representation and optimizing a representation generator to create new representations reflective of the learned distribution. Subsequently, we leverage these bootstrapped representations as self-conditioned guidance for the generation process, thereby facilitating the generation of graphs that more accurately reflect the learned distributions. We conduct extensive experiments on generic and molecular graph datasets. Our framework, GraphRCG, demonstrates superior performance over existing state-of-the-art graph generation methods in terms of graph quality and fidelity to training data.

CCS Concepts

• Information systems → Data mining.

Keywords

Graph Neural Networks; Graph Generation; Diffusion Models

ACM Reference Format:

Song Wang, Zhen Tan, Xinyu Zhao, Tianlong Chen, Huan Liu, and Jundong Li. 2025. GraphRCG: Self-conditioned Graph Generation. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM '25)*, November 10–14, 2025, Seoul, Republic of Korea. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3746252.3761205>

1 Introduction

The task of generating graphs that align with a specific distribution plays a crucial role in various fields such as drug discovery [43], public health [12], and traffic modeling [60, 62]. In recent times,

deep generative models have been prevalently studied to address the problem of graph generation [21, 25, 31]. Unlike conventional methods that rely on random graph models, recent methods generally learn graph distributions through advanced deep generative models, e.g., variational autoencoders (VAEs) [13, 57], generative adversarial networks (GANs) [6, 10], normalizing flows [31, 65], and diffusion models [25, 34, 53]. These models excel at capturing complex structural patterns in graphs, enabling the creation of new graphs with desirable characteristics.

Despite these advances, the precise modeling and utilization of graph distributions, although crucial for high-fidelity generation, remains underexplored. In fact, it is essential to accurately capture and utilize important patterns in the training data for generation [22, 23], particularly in complex scenarios like molecular graph generation [8]. For example, precise modeling of molecular properties is key to optimizing molecular structures while maintaining similarity to known molecules. However, the prevalent strategy is to use reconstruction loss to implicitly embed graph distribution within the generator, which may compromise effectiveness. Moreover, the utilization of the captured graph distributions is also less investigated. Ideally, generators should be designed to explicitly guide the generative process, ensuring that the output graphs closely follow the defined graph distributions. Nonetheless, existing research tends to rely on simple features to control generation, such as molecular characteristics [53] or degree information [4]. Such a strategy requires domain knowledge to design the specific features, while also lacking more comprehensive modeling of the entire distribution. Therefore, the study of graph generation is confronted with two crucial research questions (RQs): **(RQ1) Capturing Distributions.** How to precisely capture the graph distribution with rich information helpful for the generation process? **(RQ2) Utilizing Distributions.** How to adeptly harness these distributions as direct guidance for the generation of graphs? Addressing these challenges is essential for high-fidelity graph generation.

In practice, however, the above research questions present significant challenges due to the intricate nature of graph data. (1) **Complex Dataset Patterns.** Real-world graphs, such as social networks [48, 55] and molecular structures [25, 54, 56], exhibit highly complex patterns. These include varying degrees of sparsity, inconsistent clustering coefficients, and specific distributions of node and edge attributes [16]. Capturing these complex patterns accurately through generative models can be particularly challenging. (2) **Progressive Alignment to Training data.** Unlike images, where generation is often treated as a pixel-wise or patch-wise



This work is licensed under a Creative Commons Attribution 4.0 International License. *CIKM '25, Seoul, Republic of Korea*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2040-6/2025/11
<https://doi.org/10.1145/3746252.3761205>

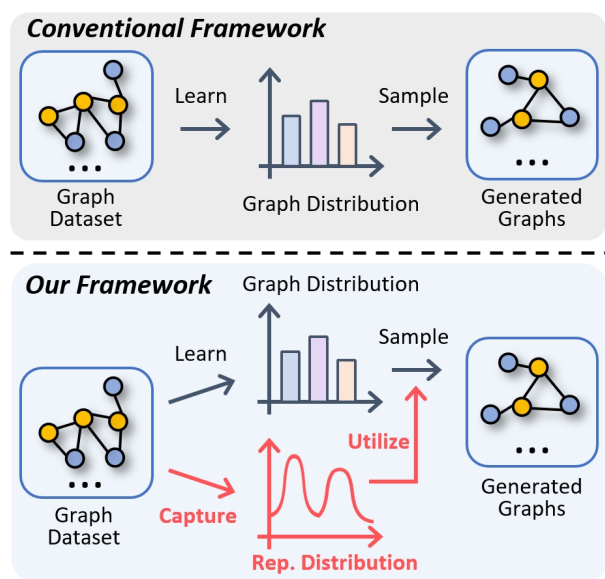


Figure 1: The comparison between the conventional framework and GraphRCG. Instead of directly learning graph distributions, we encode graphs into representations, and learn their distributions for further utilization during generation.

process [7, 15, 38], the generation of graphs is inherently sequential [34, 57, 64]. That being said, graphs are generated through a sequence of steps, each with significant implications, such as modifying chemical properties via the addition or removal of atoms and bonds [22, 23, 43]. As a result, it is suboptimal to directly guide generation toward true distributions, particularly in the initial steps of the process with graphs largely deviating from the learned distribution.

To deal with these two challenges, in this work, we introduce a novel graph generation framework named GraphRCG, which targets at **Graph Representation-Conditioned Generation**. As presented in Fig. 1, our framework is designed to first encode graphs into representations, and then capture and utilize such representation distributions for graph generation. By operating on representations instead of directly on graphs, we manage to effectively distill complex, dataset-specific knowledge into these representations and also enable their further utilization for graph generation. GraphRCG is built upon two integral components: (1) **Self-Conditioned Modeling**. To capture graph distributions while addressing the issue of complex dataset patterns in RQ1, we propose to model the essence of graph distribution through a representation generator, which could produce bootstrapped representations that authentically reflect the learned distribution. This strategy is able to enhance the quality of captured graph distributions by encapsulating the complex patterns in a parametrized manner. Moreover, the design also enables the subsequent utilization of captured distributions for generation through sampling diverse representations. (2) **Self-Conditioned Guidance**. We utilize the acquired distributions to guide graph generation to ensure the fidelity of generated graphs regarding the learned distributions. Regarding RQ2, to overcome the challenge of discrete sequential generation, we introduce a novel strategy

of step-wise guidance. This strategy employs bootstrapped representations with varying degrees of noise throughout different steps of graph generation, guiding each step closer to the learned distributions in a progressive manner while obviating the need for additional human intervention. In summary, our contributions are as follows:

- **Exploration.** In this work, we explore the potential and importance of explicitly capturing and utilizing training data distributions for graph generation to enhance performance by generating graphs that are more closely aligned with the training distributions.
- **Algorithm.** We innovatively propose a self-conditioned graph generation framework, namely GraphRCG, to capture and utilize training data distributions via bootstrapped representations with our devised self-conditioned modeling and self-conditioned guidance, respectively.
- **Experiments.** We perform a systematic study to evaluate the performance of our framework in a variety of real-world and synthetic datasets. The results demonstrate the effectiveness of our framework in comparison to other state-of-the-art baselines.

2 Related Works

2.1 Denoising Diffusion Models for Generation

Recently, denoising diffusion models [14, 45, 46] have been widely adopted in various fields of generation tasks. Their notable capabilities span various domains, including image generation [7, 15, 38], text generation [2, 3, 27], and temporal data modeling [30, 49]. In general, denoising diffusion models are probabilistic models designed to learn a data distribution by denoising variables with normal distributions [24, 47]. In particular, they first create noisy data by progressively adding and intensifying the noise in the clean data, in a Markovian manner. Subsequently, these models learn a denoising network to backtrack each step of the perturbation. During training, the denoising networks are required to predict the clean data or the added noise, given the noisy data after perturbation. After optimization, the denoising networks could be used to generate new data via iterative denoising of noise sampled from a prior distribution [40, 50].

2.2 Graph Generation

Based on the strategies used, graph generation methods could be classified into two categories: (1) *One-shot Generation*. In this category, the models generate all edges among a defined node set in one single step. One-shot generation models are typically built upon the Variational Autoencoder (VAE) or the Generative Adversarial Network (GAN) structure, aiming to generate edges independently based on the learned latent embeddings. Normalizing flow models [31, 65] propose to estimate the graph density, by establishing an invertible and deterministic function to map latent embeddings to the graphs. More recently, diffusion models have also been adopted for graph generation [21, 25]. To deal with the discrete nature of graph data, DiGress [53] leverages discrete diffusion by considering node and edge types as states in the Markovian transition matrix. (2) *Sequential Generation*. This strategy entails generating graphs through a series of sequential steps, typically by incrementally

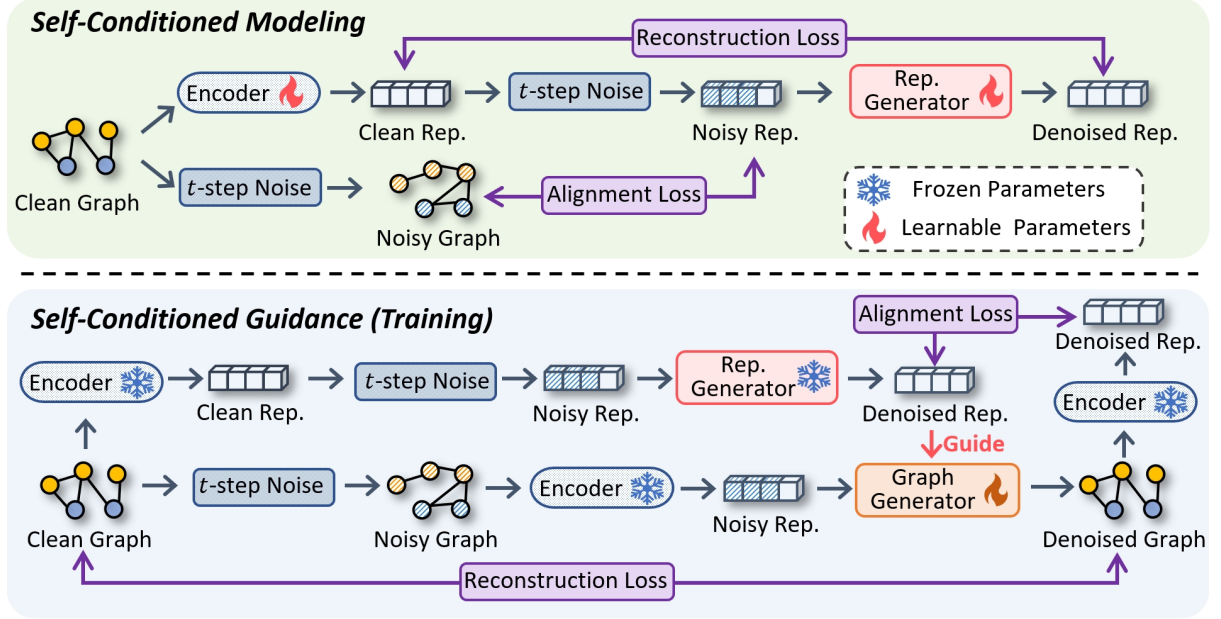


Figure 2: The overall process of GraphRCG during training. Specifically, we learn from the training data a representation generator that outputs a representation based on noise sampled from a standard Gaussian distribution. After that, a graph generator is trained to generate new graphs from noisy graphs under the guidance of the representations.

adding nodes and connecting them with edges. Models in this category often utilize recurrent networks [28, 64] or Reinforcement Learning (RL) [63] to guide the generation process [1, 43]. Sequential generation is particularly suitable for generating graphs with specific desired properties [67].

2.3 Conditional Generation

Recent works have also explored various strategies to condition generation on specific classes or features. For example, ARM [7] proposes to utilize gradients from classifiers to guide the generation process within each step. LDM [38] enables the incorporation of external information, such as text [37] and semantic maps [20], with a specific encoder to learn the representations. The representations are then incorporated into the underlying UNet backbone [39]. In RCG [26], the authors employ a self-conditioned strategy to condition generation on representations learned from a pre-trained encoder. Despite the advancements in image generation, it still presents significant difficulty when applying these methods in graph generation, due to the complex dataset patterns [4, 16, 25] and discrete sequential generation [34, 43, 57]. In contrast, our framework deals with these challenges with self-conditioned modeling and guidance to enhance graph generation efficacy.

3 Self-Conditioned Graph Generation

Our self-conditioned graph generation framework consists of two modules: self-conditioned modeling and self-conditioned guidance for RQ1 and RQ2, respectively. As illustrated in Fig. 2, in self-conditioned modeling, we employ a representation generator to capture graph distributions by learning to denoise representations

with noise. The alignment loss between noisy representations and noisy graphs acts as a self-supervised loss to train the encoder. With the optimized representation generator, we train the graph generator by denoising graphs with added noise. At each generation step, we perform self-conditioned guidance via bootstrapped representations with noise at the same timestep. We further enhance guidance by aligning the denoised graph with the clean representations.

In this work, we represent a graph as $G = (X, E)$, where $X \in \mathbb{R}^{n \times a}$ and $E \in \mathbb{R}^{n \times n \times b}$ contain all the one-hot encodings of nodes and edges, respectively. Here a and b are the numbers of node types and edge types, respectively. We consider the state of “no edge” as an edge type. n is the number of nodes in G . Notably, although we focus on discrete categorical features (i.e., types) in this paper, our work can be easily extended to scenarios with continuous features.

3.1 Self-Conditioned Modeling

GraphRCG aims to capture graph distributions (RQ1) by learning a low-dimensional representation distribution, such that the representations could be subsequently used for guidance during the generation process. Therefore, this requires the representation generator to comprehensively capture the complex patterns in graph datasets, in terms of both node features and structures.

To transform graphs into representations, we first employ a graph encoder h_η (parametrized by η) to transform an input graph $G = (X, E)$, into a low-dimensional representation: $\mathbf{h} = h_\eta(X, E)$.

Representation Generator. Given the representations of all training data provided by the graph encoder, the representation generator is required to learn to generate representations with the same distribution. To enhance generation performance, following [26],

we utilize the Representation Diffusion Model (RDM) architecture, which generates representations from Gaussian noise, based on the process of Denoising Diffusion Implicit Models (DDIM) [46]. In particular, RDM utilizes a backbone comprising a fully connected network with multiple residual blocks. Each block is composed of an input layer, a timestep embedding projection layer, and an output layer. The number of residual blocks and the hidden dimension size both act as hyper-parameters. During training, given a representation of a graph sample, learned by the graph encoder, we first perturb it by adding random noise as follows:

$$\mathbf{h}_t = \sqrt{\alpha_t} \mathbf{h}_0 + \sqrt{1 - \alpha_t} \epsilon, \quad \text{where } \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (1)$$

where \mathbf{h}_0 is the clean representation learned from G by the graph encoder, i.e., $\mathbf{h}_0 = h_\eta(\mathbf{X}, \mathbf{E})$. \mathbf{h}_t is the noisy version of \mathbf{h}_0 at timestep t . $\alpha_{1:T} \in (0, 1]^T$ is a decreasing sequence, where T is the total number of timesteps. Then the representation generator f , parameterized by γ , is trained to denoise the perturbed representation to obtain a clean one. In this way, the corresponding objective for the representation generator could be formulated as follows:

$$\mathcal{L}_{RG} = \mathbb{E}_{\mathbf{h}_0, \epsilon \sim \mathcal{N}(0, \mathbf{I}), t} \left[\|\mathbf{h}_0 - f_\gamma(\mathbf{h}_t, t)\|_2^2 \right], \quad (2)$$

where \mathbf{h}_0 is sampled from representations of graphs in training data, and t is uniformly sampled from $\{1, 2, \dots, T\}$. The target of the representation generator f_γ is to capture the representation distribution and learn to generate representations from random noise. After optimization, the representation generator could perform sampling from random noise, following the DDIM strategy [46], to obtain new representations.

On the other hand, the noisy representation \mathbf{h}_t , i.e., the input to the representation generator, is not related to any noisy graph. That being said, the noisy representation might not faithfully represent the noise added to an actual graph. Therefore, to enhance such consistency, we propose an alignment loss to train the encoder, which is formulated as follows:

$$\mathcal{L}_{AR} = \mathbb{E}_{G, t} \left[\|\mathbf{h}_t - h_\eta(\mathbf{X}_t, \mathbf{E}_t)\|_2^2 \right], \quad (3)$$

where $G_t = (\mathbf{X}_t, \mathbf{E}_t)$ denotes the noisy graph with the same timestep t . The loss \mathcal{L}_{AR} is designed to align the noise processes of representations and graphs, such that the noisy representation \mathbf{h}_t preserves the same information as the noisy graph at the same timestep G_t . In the following, we detail the process of adding noise to any graph. **Adding Noise to Graphs.** In discrete diffusion, adding noise equates to transitioning between states, that is, choosing a state based on a categorical distribution. For each timestep t , the probability of moving from one state to another is defined by a Markov transition matrix \mathbf{Q}_t , where $\mathbf{Q}_t[i, j]$ represents the likelihood of transitioning from state i to state j . For graph generation, these states generally represent specific node types or edge types. Particularly, an edge-type state represents the absence of an edge. The process of adding noise is performed in a disentangled manner, which operates independently across nodes and edges with separate noise perturbations. A step of the noise process could be expressed as

$$q(G_t|G_{t-1}) = (\mathbf{X}_{t-1} \mathbf{Q}_t^X, \mathbf{E}_{t-1} \mathbf{Q}_t^E), \quad (4)$$

where t is the timestep. Moreover, \mathbf{Q}_t^X and \mathbf{Q}_t^E represent the transition matrices for the nodes and edges at timestep t , respectively.

Given the Markovian nature of the noise model, the noise addition is not cumulative, as the probability $q(G_t|G_0)$ could be directly calculated from all the respective Markov transition matrices:

$$q(G_t|G_0) = (\mathbf{X}_0 \prod_{i=1}^t \mathbf{Q}_i^X, \mathbf{E}_0 \prod_{i=1}^t \mathbf{Q}_i^E), \quad (5)$$

where $\mathbf{X}_0 = \mathbf{X}$ and $\mathbf{E}_0 = \mathbf{E}$ denote the input node and edge types of G , respectively. This formulation captures the essence of the discrete diffusion process, i.e., applying independent, state-specific transitions at each timestep for both nodes and edges in graph generation. In this manner, we could obtain $\mathbf{X}_t = \mathbf{X}_0 \prod_{i=1}^t \mathbf{Q}_i^X$ and $\mathbf{E}_t = \mathbf{E}_0 \prod_{i=1}^t \mathbf{Q}_i^E$.

To specify the Markov transition matrices, previous works have explored several feasible choices. The most prevalent choices in the literature have been uniform transitions [2, 59] and absorbing transitions [4, 23]. However, these do not contain the distribution information and thus could not benefit the capturing of graph distribution in our framework. Therefore, in our approach, we leverage the marginal transitions [18, 53], in which the probability of transitioning to any given state is directly related to its marginal probability observed in the dataset. In this manner, the transition matrices are modeled in a way that mirrors the natural distribution of states in graph data. As the edges are generally sparse in the graph data, the probability of jumping to the state of “no edge” is significantly higher than that of other states. To present the noise process, we first define $\mathbf{p}^X \in \mathbb{R}^a$ and $\mathbf{p}^E \in \mathbb{R}^b$ as the marginal distributions for the node and edge types, respectively. The marginal transition matrices for nodes and edges are formulated as

$$\mathbf{Q}_t^X = \alpha^t \mathbf{I} + \beta^t \mathbf{1}_a (\mathbf{p}^X)^\top, \quad \text{and} \quad \mathbf{Q}_t^E = \alpha^t \mathbf{I} + \beta^t \mathbf{1}_b (\mathbf{p}^E)^\top. \quad (6)$$

The above formulation ensures that $\lim_{t \rightarrow \infty} \prod_{i=1}^t \mathbf{Q}_i^X = \mathbf{P}^X$, where each column in \mathbf{P}^X is \mathbf{p}^X . The equation also holds true for the edges. More details of the noise model are provided in Appendix A. The overall process of our self-conditioned modeling module is presented in Algorithm 1.


3.2 Self-Conditioned Guidance (Training)

In self-conditioned guidance for RQ2, we optimize a graph generator to create new graphs conditioned on bootstrapped representations from the representation generator. Existing works have explored various methods to perform generation conditioned on specific properties, such as high activity in molecular graph generation [16]. However, these properties are generally several scalar values, which could capture only a small fraction of the information in the dataset. Moreover, these approaches often adopt a classifier or regressor to guide the generation process, which could not exploit the useful information in representations. In contrast, using representations as guidance for generation could largely benefit from the learned distribution in our representation generator.

As our representation generator is implemented by a diffusion model, we aim to utilize the information from not only the generated representation but also its generation process. With this in mind, we implement the generator as a denoising diffusion model, so that each step of diffusion could benefit from the representation of the same denoising step in a bootstrapped manner. Our graph generator g_θ , implemented as a denoising network and parametrized by θ , is

Table 1: Comparison of generation results on SBM, Planar, and Ego. The best results are shown in bold.

Model	SBM				Planar				Ego			
	Deg. ↓	Clus. ↓	Orbit ↓	Spec. ↓	Deg. ↓	Clus. ↓	Orbit ↓	Spec. ↓	Deg. ↓	Clus. ↓	Orbit ↓	Spec. ↓
Training	0.0008	0.0332	0.0255	0.0063	0.0002	0.0310	0.0005	0.0052	0.0002	0.0100	0.0120	0.0014
GraphRNN	0.0055	0.0584	0.0785	0.0065	0.0049	0.2779	1.2543	0.0459	0.0768	1.1456	0.1087	-
GRAN	0.0113	0.0553	0.0540	0.0054	0.0007	0.0426	0.0009	0.0075	0.5778	0.3360	0.0406	-
SPECTRE	0.0015	0.0521	0.0412	0.0056	0.0005	0.0785	0.0012	0.0112	-	-	-	-
DiGress	0.0013	0.0498	0.0433	-	0.00027	0.0563	0.0098	0.0062	0.0708	0.0092	0.1205	-
HiGen	0.0019	0.0498	0.0352	0.0046	-	-	-	-	0.0472	0.0031	0.0387	0.0062
SparseDiff	0.0016	0.0497	0.0346	0.0043	0.0007	0.0447	0.0017	0.0068	0.0019	0.0537	0.0209	0.0050
EDP-GNN	-	-	-	-	-	-	-	-	0.0520	0.0930	0.0070	-
HOG-Diff	-	-	-	-	-	-	-	-	0.0150	0.0270	0.0040	-
GraphRCG	0.0011	0.0475	0.0378	0.0038	0.00025	0.0341	0.0010	0.0059	0.0015	0.0448	0.0183	0.0042

**Figure 4: The generated graphs from the Planar dataset with different interpolation ratios between two representations.**

4 Experiments

We evaluate GraphRCG on realistic molecular and synthetic non-molecular datasets, in comparison to baselines, including autoregressive models: GRAN [29] and GraphRNN [64], a GAN-based model: SPECTRE [32], and diffusion models: EDP-GNN [34], DiGress [53], GDSS [21], GraphARM [23], HiGen [22], SparseDiff [36], EB-GFN [66], GPrinFlowNet [33], and HOG-Diff [17].

For generic graph generation, we evaluate the quality of the generated graphs with structure-based evaluation metrics. We follow previous work [64] and calculate the MMD (Maximum Mean Discrepancy) between the graphs in the test set and the generated graphs, regarding (1) degree distributions, (2) clustering coefficients distributions, (3) the number of orbits with four nodes, and (4) the spectra of the graphs obtained from the eigenvalues of the normalized graph Laplacian [4, 64]. For molecular graph generation, following previous works [21, 23], we evaluate the generated molecular graphs with several key metrics: (1) Fréchet ChemNet Distance (FCD) [35], which quantifies the discrepancy between the distributions of training and generated graphs, based on the activation values obtained from the penultimate layer in ChemNet. (2) Neighborhood Subgraph Pairwise Distance Kernel (NSPDK) MMD [5], which measures the Maximum Mean Discrepancy (MMD) between the generated and the test molecules, considering both node and edge features. (3) Validity, which refers to the proportion of generated molecules that are structurally valid without requiring valency adjustments. (4) Uniqueness, which measures the proportion of unique molecules that are distinct from each other.

For the representation generator, we implement it as an RDM (Representation Diffusion Model) used in RCG [26]. The representation dimension size is set as 256 for all datasets. The learning rate of the RDM is set as 10^4 , and the weight decay rate is set as 0.01. Our code is provided at <https://github.com/SongW-SW/GraphRCG>.

4.1 Comparative Results

Generic Graph Generation. In this subsection, we further evaluate our framework on generic graph datasets with relatively larger sizes than molecular graphs. In particular, we consider two synthetic datasets: SBM, drawn from stochastic block models [32], with a maximum size of 200 nodes, and Planar, containing planar graphs with a fixed size of 64 [53]. In addition, we consider a realistic citation dataset Ego [42], originated from Citeseer [11]. Further details of these datasets are provided in Appendix ?? . From the results presented in Table 1, we could obtain the following observations: (1) GraphRCG outperforms other baselines on all three datasets across various metrics for graph generation, demonstrating the effectiveness of our framework in precisely capturing graph distributions and utilizing them for generation guidance. (2) The performance improvement over other methods is more substantial on the Planar dataset. Nevertheless, our framework learns graph distributions with a representation generator, which enables the modeling of complex underlying patterns. (3) GraphRCG is particularly competitive in the MMD score regarding degree distributions and the number of orbits. This observation demonstrates that GraphRCG could authentically capture the complex graph distribution of the training samples with the help of self-conditioned modeling.

Molecular Graph Generation. To evaluate our framework on molecular graph generation, we select popular datasets QM9 [58] and ZINC250k [19], with details and provided in Appendix ?? . We present the molecular graph generation results on QM9 in Table 2. Specifically, GraphRCG demonstrates competitive performance on QM9 across various metrics, compared to other state-of-the-art baselines. The best FCD values on QM9 and ZINC250k indicate that GraphRCG effectively captures the chemical property distributions in the dataset. Furthermore, the outstanding validity score on QM9

Table 2: Results of various methods on the QM9 and ZINC250k Datasets. The best results are shown in bold.

Model	QM9 Dataset				ZINC250k Dataset			
	Validity \uparrow	NSPDK \downarrow	FCD \downarrow	Unique \uparrow	Validity \uparrow	NSPDK \downarrow	FCD \downarrow	Unique \uparrow
EDP-GNN	47.52	0.005	2.68	99.25	82.97	0.049	16.74	99.79
SPECTRE	87.33	0.163	47.96	35.7	90.20	0.109	18.44	67.05
GDSS	95.72	0.003	2.9	98.46	97.01	0.019	14.66	99.64
DiGress	99.01	0.0005	0.36	96.66	91.02	0.082	23.06	81.23
GraphARM	90.25	0.002	1.22	95.62	88.23	0.055	16.26	99.46
GraphRCG	99.12	0.0008	0.28	98.39	92.38	0.041	13.48	96.15

also signifies that our framework GraphRCG is capable of generating valid molecules that are more closely aligned with the training data distribution.

4.2 Representation Interpolation

As our self-conditioned guidance leverages representations, we could manually perform linear interpolation for two representations to generate graphs that represent the properties of both representations. With our step-wise incorporation strategy, we extract two series of representations generated by our representation generator for all timesteps, i.e., $t = 1, 2, \dots, T$, and perform interpolation for each timestep. We denote α as the interpolation ratio, where $\alpha = 0$ and $\alpha = 1$ indicate that we entirely utilize one of the representations. We provide the visualization results in Fig. 4, from which we observe that the generated graphs guided by interpolated representations remain meaningful with different interpolation ratios. This result demonstrates that our representation generator is capable of capturing smooth graph distributions with rich information. Furthermore, our design also enables further applications with specific representations as guidance.

4.3 Ablation Study

Effect of Self-Conditioned Modeling. Previous works [53] have also explored various strategies to capture and model distributions for the guidance of graph generation. For example, DiGress leverages manually designed structural features (e.g., the number of cycles) for graph generation. However, such features contain less information and require domain knowledge. In this subsection, we compare several variants of our framework with these structural features to evaluate the efficacy of self-conditioned modeling in encapsulating the graph distribution. We consider using the following features (representations) to replace our self-conditioned modeling module: (1) structural features proposed in DiGress, (2) representations of true training samples, which limit the uniqueness of generation, (3) a mixture of true representations, and (4) distributions learned by a Gaussian Mixture Model (GMM). We further considered the variant of directly using a dot-product decoder to replace self-conditioned generation. As the structural features involve molecular information, we conduct experiments on the QM9 dataset with explicit hydrogens, which is a more complex setting with larger graphs.

From the results presented in Table 3, we first observe that the inclusion of structural features significantly enhances the performance of DiGress, especially in the metric of molecule stability, which benefits from the integration of molecular characteristics.

Table 3: The ablation study results regarding self-conditioned modeling on dataset QM9 with explicit hydrogen.

Model	Valid \uparrow	Unique \uparrow	Atom S. \uparrow	Mol S. \uparrow
Dataset	97.8	100	98.5	87.0
DiGress w/o A	92.3	97.9	97.3	66.8
DiGress w/ A	95.4	97.6	98.1	79.8
GraphRCG w/ A	92.9	95.4	93.1	76.6
GraphRCG w/ T	91.2	91.5	90.5	72.3
GraphRCG w/ T+M	94.6	97.0	90.2	74.9
GraphRCG w/ GMM	96.4	94.1	91.5	77.1
GraphRCG w/o Self	90.5	96.6	90.1	75.5
GraphRCG	96.9	98.1	97.2	81.9

Table 4: The ablation study results on different variants of our framework GraphRCG regarding self-conditioned guidance on the dataset Ego.

Dataset	Ego			
	Deg. \downarrow	Clus. \downarrow	Orbit \downarrow	Spec. \downarrow
GraphRCG-Gradient	0.0134	0.0955	0.0464	0.0151
GraphRCG w/o \mathcal{L}_{AG}	0.0088	0.0647	0.0252	0.0092
GraphRCG-fixed	0.0053	0.0653	0.0310	0.0125
GraphRCG	0.0015	0.0448	0.0183	0.0042

The substitution of self-conditioned modeling with structural features, however, results in a notable decline in performance, highlighting the importance of distributions with comprehensive information for effective guidance. Furthermore, among the variants that directly operate on training sample representations, it is evident that the mere utilization of pre-existing representations yields suboptimal outcomes, especially concerning the metric of uniqueness. This suggests that relying solely on existing representations captures only a small portion of the authentic distribution, thereby adversely impacting the overall performance. Furthermore, the variant of directly using a decoder shows worse performance, which demonstrates the necessity of utilizing a diffusion-based generator in self-conditioned guidance.

Effect of Self-Conditioned Guidance. In this subsection, we investigate the effect of self-conditioned guidance in our framework. We first replace the entire guidance strategy with a gradient-based method, which leverages computed gradients as guidance. Therefore, the graph generator does not involve any representation during training. For the second variant, we remove the alignment loss \mathcal{L}_{AG} described in Eq. (11). Without this loss, the graph generator is not well-aligned with the representations produced by the representation generator, thereby affecting the guidance performance. For

Table 5: The performance of various methods on three generic datasets.

Model	Community			Cora			Enzymes		
	Deg. ↓	Clus. ↓	Orbit ↓	Deg. ↓	Clus. ↓	Orbit ↓	Deg. ↓	Clus. ↓	Orbit ↓
SPECTRE	0.048	0.049	0.016	0.021	0.080	0.007	0.136	0.195	0.125
GDSS	0.045	0.086	0.007	0.160	0.376	0.187	0.026	0.061	0.009
DiGress	0.047	0.041	0.026	0.044	0.042	0.223	0.004	0.083	0.002
GraphARM	0.034	0.082	0.004	0.273	0.138	0.105	0.029	0.054	0.015
EB-GFN	-	-	-	-	-	-	0.079	0.213	0.227
GPrinFlowNet	-	-	-	-	-	-	0.027	0.062	0.046
HOG-Diff	-	-	-	-	-	-	0.011	0.061	0.007
GraphRCG	0.040	0.053	0.029	0.038	0.036	0.173	0.004	0.079	0.002

Table 6: Comparison of different denoising diffusion models for graph generation. $G(n, p)$ denotes the Erdős-Rényi graph model [9], where p is the probability of an edge existing between two nodes, and n is the graph size.

Model	Diffusion Type	Convergence	Conditional Generation	Featured Generation
EDP-GNN	Continuous	$\mathcal{N}(0, 1)$	-	-
GDSS	Continuous	$\mathcal{N}(0, 1)$	-	✓
DiscDDPM	Discrete	$G(n, 0.5)$	-	-
DiGress	Discrete	Empirical	Gradients from a classifier	✓
SparseDiff	Discrete	Empirical	-	✓
EDGE	Discrete	$G(n, 0)$	Degree sequence	✓
Ours	Continuous & Discrete	$\mathcal{N}(0, 1)$ & Empirical	Representation	✓

the third variant, we directly use the fixed representation, i.e., the clean representation, to guide generation. In this case, the step-wise guidance strategy is removed, resulting in the infeasibility of progressive guidance. From the results presented in Table 4, we could first observe that our framework outperforms all other variants in most evaluation metrics, demonstrating the effectiveness of our self-conditioned guidance strategy. Moreover, GraphRCG-Gradient and GraphRCG w/o \mathcal{L}_{AG} exhibit significantly poorer performance, as evidenced by higher scores across all metrics. This deterioration, particularly in Clus. and Orbit values, underscores the difficulty in preserving graph distributions without representation guidance or alignment loss. GraphRCG-fixed shows improved performance relative to the other variants, demonstrating the benefits of representation guidance even with a fixed one. Nevertheless, the results also indicate that the step-wise self-conditioned guidance is more beneficial for progressively guiding the generation process.

4.4 Results on Additional Generic Datasets

In this subsection, we conduct additional experiments on the three generic datasets : Community [64], Cora [42], and Enzymes [41], to further improve the integrity of our evaluation. We provide the results in Table 5. The results of other baselines are obtained from GraphARM [23]. From the results, we could observe that our framework also achieves competitive results, compared to other methods on various datasets. The performance is particularly better on dataset Cora and Enzymes, indicating that our framework is generalizable to various datasets.

4.5 Comparison with Other Methods

In this subsection, we provide a detailed comparison regarding the specific technical details of our method GraphRCG and other

baselines in Table 6. Typically, existing diffusing methods for graph generation rely on either continuous or discrete diffusion. As a classic example, GDSS [21] adopts Gaussian transition kernels to perform continuous diffusion with a score-matching strategy. DiGress, on the other hand, performs discrete diffusion while considering categorical features of nodes and edges. Nevertheless, although continuous representations could capture complex structural patterns in graph distributions, they are less effective in generating discrete graph data. In contrast, our framework is capable of leveraging continuous diffusion to guide discrete diffusion, thereby combining the strengths of both and allowing for the generation of a wider range of graph structures. Regarding the convergence type, which refers to the pure noise state, we also combine both continuous and discrete noise to facilitate the optimization of both our representation generator and graph generator. Comparing the conditional generation type, our framework is conditioned on representations, which preserve richer information of complex distributions of each dataset. As a result, the generation process could significantly benefit from the guidance of representations.

5 Conclusion

In this work, we investigate the importance of capturing and utilizing distributions for graph generation. We propose a novel self-conditioned generation, that encompasses self-conditioned modeling and self-conditioned guidance. Instead of directly learning from graph distributions, we encode all graphs into representations to capture graph distributions, which could preserve richer information. Our self-conditioned guidance module further guides the generation process in each timestep with representations with different degrees of noise. We conduct extensive experiments to evaluate our framework, and the results demonstrate the efficacy

of our framework in graph generation. Future work along this line can explore scaling our framework to handle large-scale graph generation efficiently remains an open challenge, where techniques such as memory-efficient architectures could be considered. Additionally, extending GraphRCG to multi-modal graph generation, where node and edge attributes are explicitly modeled alongside structural information, could further improve the applicability of our framework in real-world scenarios.

Acknowledgements

This work is supported in part by the National Science Foundation under grants (IIS-2006844, IIS-2144209, IIS-2223769, CNS-2154962, BCS2228534, and CMMI2411248), the Commonwealth Cyber Initiative Awards under grant (VV-1Q25-004), and the research gift funding from Netflix and Snap.

A Noise Model

In this section, we provide additional information on the noise model used during the training of our graph generator. Particularly, the graphs at timestep t could be represented as $\mathbf{X}_t = \mathbf{X}_0 \prod_{i=1}^t \mathbf{Q}_i^X$ and $\mathbf{E}_t = \mathbf{E}_0 \prod_{i=1}^t \mathbf{Q}_i^E$. Here, \mathbf{Q}_i^X and \mathbf{Q}_i^E are defined in Eq. (13), i.e.,

$$\mathbf{Q}_t^X = \alpha^t \mathbf{I} + \beta^t \mathbf{1}_a (\mathbf{p}^X)^\top, \quad \text{and} \quad \mathbf{Q}_t^E = \alpha^t \mathbf{I} + \beta^t \mathbf{1}_b (\mathbf{p}^E)^\top. \quad (13)$$

In practice, the noise process is not cumulative. By defining $\bar{\alpha}^t = \prod_{i=1}^t \alpha^i$ and $\bar{\beta}^t = 1 - \bar{\alpha}^t$, we could achieve

$$\prod_{i=1}^t \mathbf{Q}_i^X = \bar{\alpha}^t \mathbf{I} + \bar{\beta}^t \mathbf{1}_a (\mathbf{p}^X)^\top, \quad \text{and} \quad \prod_{i=1}^t \mathbf{Q}_i^E = \bar{\alpha}^t \mathbf{I} + \bar{\beta}^t \mathbf{1}_b (\mathbf{p}^E)^\top. \quad (14)$$

Following [53], we set $\bar{\alpha}^t = \cos(0.5\pi(t/T+s)/(1+s))^2$ with s being a small value. In this manner, when computing the noisy graphs $G_t = (\mathbf{X}_t, \mathbf{E}_t)$, we do not need to recursively multiply \mathbf{Q}_i^X . Instead, we utilize the value of $\bar{\alpha}^t$ to directly obtain G_t .

The noise model for the representation generator is similar and in a simpler form. We first sample $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ and compute the noisy representations based on $\mathbf{h}_t = \sqrt{\bar{\alpha}_t} \mathbf{h}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$.

Algorithm 1 The training process of self-conditioned modeling.

Input: A training graph distribution \mathcal{G} , maximum number of denoising steps T , number of training epochs T_{tr} .

Output: Optimized graph encoder and representation generator.

```

1: for  $i = 1, 2, \dots, T_{tr}$  do
2:   Sample  $G = (\mathbf{X}, \mathbf{E})$  from  $\mathcal{G}$ ;
3:   Sample  $t \sim \mathcal{U}(1, 2, \dots, T)$  and  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ ;
   // Training the representation generator
4:    $\mathbf{h}_0 \leftarrow h_\eta(\mathbf{X}, \mathbf{E})$ ;
5:    $\mathbf{h}_t \leftarrow \sqrt{\bar{\alpha}_t} \mathbf{h}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ ;
6:    $\tilde{\mathbf{h}}_t = f_Y(\mathbf{h}_t, t)$ ;
7:   Optimize  $f_Y$  with  $\mathcal{L}_{RG}$  in Eq. (2); // Training the graph encoder
8:   Sample  $G_t = (\mathbf{X}_t, \mathbf{E}_t)$  from  $\mathbf{X} \prod_{i=1}^t \mathbf{Q}_i^X \times \mathbf{E} \prod_{i=1}^t \mathbf{Q}_i^E$ ;
9:    $\bar{\mathbf{h}}_t = h_\eta(\mathbf{X}_t, \mathbf{E}_t)$ ;
10:  Optimize  $h_\eta$  with  $\mathcal{L}_{AR}$  in Eq. (3);
11: end for

```

Algorithm 2 Detailed training and sampling process of our self-conditioned guidance module.

Input: A training graph distribution \mathcal{G} , maximum number of denoising steps T , number of training epochs T_{tr} .

Output: A generated graph that aligns with the distribution \mathcal{G} .

```

// Training phase
1: for  $i = 1, 2, \dots, T_{tr}$  do
2:   Sample  $G = (\mathbf{X}, \mathbf{E})$  from  $\mathcal{G}$ ;
3:   Sample  $t \sim \mathcal{U}(1, 2, \dots, T)$  and  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ ;
4:   Sample  $G_t = (\mathbf{X}_t, \mathbf{E}_t)$  from  $\mathbf{X} \prod_{i=1}^t \mathbf{Q}_i^X \times \mathbf{E} \prod_{i=1}^t \mathbf{Q}_i^E$ ;
5:    $\mathbf{h}_0 \leftarrow h_\eta(\mathbf{X}, \mathbf{E})$  and  $\mathbf{h}_t \leftarrow \sqrt{\bar{\alpha}_t} \mathbf{h}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ ;
6:    $(\tilde{\mathbf{p}}^X, \tilde{\mathbf{p}}^E) \leftarrow g_\theta(\mathbf{X}_t, \mathbf{E}_t, \tilde{\mathbf{h}}_t, t)$ ;
7:    $\tilde{\mathbf{h}}_t = f_Y(\mathbf{h}_t, t)$ ;
8:   Optimize  $g_\theta$  with  $\mathcal{L}_{GG}$  in Eq. (10) and  $\mathcal{L}_{AG}$  in Eq. (11);
9: end for
// Sampling phase
10: Sample  $n$  from the training data distribution of graph sizes;
11: Sample  $G_T = (\mathbf{X}_T, \mathbf{E}_T) \sim \mathbf{p}^X \times \mathbf{p}^E$ ;
12: for  $t = T, T-1, \dots, 1$  do
13:    $\mathbf{h}_t = h_\eta(\mathbf{X}_t, \mathbf{E}_t)$ ;
14:    $\mathbf{h}_{t-1} = f_Y(\mathbf{h}_t, t)$ ;
15:    $(\tilde{\mathbf{p}}^X, \tilde{\mathbf{p}}^E) \leftarrow g_\theta(\mathbf{X}_t, \mathbf{E}_t, \mathbf{h}_{t-1}, t)$ ;
16:    $G_{t-1} \sim p_\theta(G_{t-1} | G_t, \mathbf{h}_{t-1}) = \prod_{i=1}^n \tilde{\mathbf{p}}_i^X \prod_{i=1}^n \prod_{j=1}^n \tilde{\mathbf{p}}_{i,j}^E$ ;
17: end for
18: return  $G_0$ .

```

B Complexity Analysis

We analyze the computational complexity of our self-conditioned graph generation framework, which involves both the representation generator and the graph generator.

Representation Generator. Given an input graph $G = (\mathbf{X}, \mathbf{E})$ with n nodes and m edges, the graph encoder h_η processes the graph into a low-dimensional vector representation $\mathbf{h} \in \mathbb{R}^d$. Assuming a message-passing encoder, the time complexity is $O(L_h \cdot (m + n))$, where L_h is the number of layers. The representation generator f_Y is implemented as a residual MLP with B blocks and hidden dimension d_h , leading to $O(B \cdot d_h^2)$ cost per denoising step. Since diffusion requires T steps, the total complexity for representation modeling is: $O(L_h(m + n) + T \cdot B d_h^2)$.

Graph Generator. The graph generator g_θ is based on a Transformer with L layers applied to node and edge embeddings. For each layer, the self-attention mechanism incurs $O(n^2 d_x)$ cost for nodes and $O(m d_e)$ for edges, where d_x and d_e are embedding dimensions. Including feedforward projections, the per-layer cost is $O(n^2 d_x + m d_e)$. Over L layers and T diffusion steps, the total complexity becomes $O(T \cdot L(n^2 d_x + m d_e))$.

Overall. The overall training and sampling complexity is therefore $O(L_h(m + n) + T \cdot B d_h^2 + T \cdot L(n^2 d_x + m d_e))$. In practice, n and m are small in most molecular and benchmark datasets (e.g., $n < 200$), making the quadratic term $n^2 d_x$ tractable. Moreover, both the MLP blocks in f_Y and the attention layers in g_θ are highly parallelizable on GPUs, so the framework is computationally feasible despite requiring two diffusion processes.

GenAI Usage Disclosure

Generative AI (such as Large Language Models) is used to help refine the writing of the paper and is not used in any other parts of the work. The usage of GenAI in this paper is to make the writing style more formal and compliant with scientific research requirements.

References

- [1] Sungsoo Ahn, Binghong Chen, Tianzhe Wang, and Le Song. 2021. Spanning tree-based graph generation for molecules. In *International Conference on Learning Representations*.
- [2] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. 2021. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems* 34 (2021), 17981–17993.
- [3] Ting Chen, Ruixiang ZHANG, and Geoffrey Hinton. 2022. Analog Bits: Generating Discrete Data using Diffusion Models with Self-Conditioning. In *The Eleventh International Conference on Learning Representations*.
- [4] Xiaohui Chen, Jiaying He, Xu Han, and Li-Ping Liu. 2023. Efficient and Degree-Guided Graph Generation via Discrete Diffusion Modeling. *arXiv preprint arXiv:2305.04111* (2023).
- [5] Fabrizio Costa and Kurt De Grave. 2010. Fast neighborhood subgraph pairwise distance kernel. In *Proceedings of the 26th International Conference on Machine Learning*. Omnipress; Madison, WI, USA, 255–262.
- [6] Nicola De Cao and Thomas Kipf. 2018. MolGAN: An implicit generative model for small molecular graphs. *arXiv:1805.11973* (2018).
- [7] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* 34 (2021), 8780–8794.
- [8] Yuanqi Du, Xian Liu, Nilay Mahesh Shah, Shengchao Liu, Jieyu Zhang, and Bolei Zhou. 2022. ChemSpace: Interpretable and Interactive Chemical Space Exploration. *Transactions on Machine Learning Research* (2022).
- [9] Paul Erdős, Alfréd Rényi, et al. 1960. On the evolution of random graphs. *Publ. math. inst. hung. acad. sci* 5, 1 (1960), 17–60.
- [10] Anuththari Gamage, Eli Chien, Jianhao Peng, and Olgica Milenkovic. 2020. Multi-motifgan (mmsgan): Motif-targeted graph generation and prediction. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 4182–4186.
- [11] C Lee Giles, Kurt D Bollacker, and Steve Lawrence. 1998. CiteSeer: An automatic citation indexing system. In *Proceedings of the third ACM conference on Digital libraries*. 89–98.
- [12] Xiaojie Guo, Yuanqi Du, Sivani Tadepalli, Liang Zhao, and Amarda Shehu. 2021. Generating tertiary protein structures via interpretable graph variational autoencoders. *Bioinformatics Advances* 1, 1 (2021), vbab036.
- [13] Xiaojie Guo, Yuanqi Du, and Liang Zhao. 2020. Property controllable variational autoencoder via invertible mutual dependence. In *International Conference on Learning Representations*.
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *NeurIPS* (2020).
- [15] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. 2022. Cascaded diffusion models for high fidelity image generation. *The Journal of Machine Learning Research* 23, 1 (2022), 2249–2281.
- [16] Han Huang, Leilei Sun, Bowen Du, and Weifeng Lv. 2022. Conditional Diffusion Based on Discrete Graph Structures for Molecular Graph Generation. In *NeurIPS 2022 Workshop on Score-Based Methods*.
- [17] Yiming Huang and Tolga Birdal. 2025. HOG-Diff: Higher-Order Guided Diffusion for Graph Generation. *arXiv preprint arXiv:2502.04308* (2025).
- [18] John B Ingraham, Max Baranov, Zak Costello, Karl W Barber, Wujie Wang, Ahmed Ismail, Vincent Frappier, Dana M Lord, Christopher Ng-Thow-Hing, Erik R Van Vlack, et al. 2023. Illuminating protein space with a programmable generative model. *Nature* 623, 7989 (2023), 1070–1078.
- [19] John J Irwin, Teague Sterling, Michael M Mysinger, Erin S Bolstad, and Ryan G Coleman. 2012. ZINC: a free tool to discover chemistry for biology. *Journal of chemical information and modeling* 52, 7 (2012), 1757–1768.
- [20] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
- [21] Jaehyeong Jo, Seul Lee, and Sung Ju Hwang. 2022. Score-based generative modeling of graphs via the system of stochastic differential equations. In *International Conference on Machine Learning*. PMLR, 10362–10383.
- [22] Mahdi Karami. 2023. HiGen: Hierarchical Graph Generative Networks. *arXiv preprint arXiv:2305.19337* (2023).
- [23] Linghai Kong, Jiaming Cui, Haotian Sun, Yuchen Zhuang, B Aditya Prakash, and Chao Zhang. 2023. Autoregressive diffusion model for graph generation. In *International Conference on Machine Learning*. PMLR, 17391–17408.
- [24] Zhifeng Kong and Wei Ping. 2021. On Fast Sampling of Diffusion Probabilistic Models. In *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*.
- [25] Seul Lee, Jaehyeong Jo, and Sung Ju Hwang. 2023. Exploring chemical space with score-based out-of-distribution generation. In *International Conference on Machine Learning*. PMLR, 18872–18892.
- [26] Tianhong Li, Dina Katabi, and Kaiming He. 2023. Self-conditioned Image Generation via Generating Representations. *arXiv preprint arXiv:2312.03701* (2023).
- [27] Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. 2022. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems* 35 (2022), 4328–4343.
- [28] Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. 2018. Learning deep generative models of graphs. *arXiv preprint arXiv:1803.03324* (2018).
- [29] Renjie Liao, Yujia Li, Yang Song, Shenlong Wang, Will Hamilton, David K Duvenaud, Raquel Urtasun, and Richard Zemel. 2019. Efficient graph generation with graph recurrent attention networks. *Advances in neural information processing systems* 32 (2019).
- [30] Juan Miguel Lopez Alcaraz and Nils Strodthoff. 2023. Diffusion-based time series imputation and forecasting with structured atate space models. *Transactions on machine learning research* (2023), 1–36.
- [31] Youzhi Luo, Keqiang Yan, and Shuiwang Ji. 2021. Graphdf: A discrete flow model for molecular graph generation. In *International Conference on Machine Learning*. PMLR, 7192–7203.
- [32] Karolis Martinkus, Andreas Loukas, Nathanaël Perraudin, and Roger Wattenhofer. 2022. Spectre: Spectral conditioning helps to overcome the expressivity limits of one-shot graph generators. In *International Conference on Machine Learning*. PMLR, 15159–15179.
- [33] Zhanfeng Mo, Tianze Luo, and Sinno Jialin Pan. 2024. Graph principal flow network for conditional graph generation. In *Proceedings of the ACM Web Conference 2024*.
- [34] Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. 2020. Permutation invariant graph generation via score-based generative modeling. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 4474–4484.
- [35] Kristina Preuer, Philipp Renz, Thomas Unterthiner, Sepp Hochreiter, and Gunter Klambauer. 2018. Fréchet ChemNet distance: a metric for generative models for molecules in drug discovery. *Journal of Chemical Information and Modeling* (2018).
- [36] Yiming Qin, Clement Vignac, and Pascal Frossard. 2023. Sparse Training of Discrete Diffusion Models for Graph Generation. *arXiv preprint arXiv:2311.02142* (2023).
- [37] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. 2016. Generative adversarial text to image synthesis. In *International conference on machine learning*. PMLR, 1060–1069.
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10684–10695.
- [39] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. Springer, 234–241.
- [40] Robin San-Roman, Eliya Nachmani, and Lior Wolf. 2021. Noise estimation for generative diffusion models. *arXiv preprint arXiv:2104.02600* (2021).
- [41] Ida Schomburg, Antje Chang, Christian Ebeling, Marion Gremse, Christian Heldt, Gregor Huhn, and Dietmar Schomburg. 2004. BRENDA, the enzyme database: updates and major new developments. *Nucleic acids research* 32, suppl_1 (2004), D431–D433.
- [42] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* 29, 3 (2008), 93–93.
- [43] Chence Shi, Minkai Xu, Zhaocheng Zhu, Weinan Zhang, Ming Zhang, and Jian Tang. 2019. GraphAF: a Flow-based Autoregressive Model for Molecular Graph Generation. In *International Conference on Learning Representations*.
- [44] Yunsheng Shi, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjin Wang, and Yu Sun. 2020. Masked label prediction: Unified message passing model for semi-supervised classification. *arXiv preprint arXiv:2009.03509* (2020).
- [45] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*.
- [46] Jiaming Song, Chenlin Meng, and Stefano Ermon. 2021. Denoising Diffusion Implicit Models. In *ICLR*.
- [47] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2020. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.

- [48] Zhen Tan, Song Wang, Kaize Ding, Jundong Li, and Huan Liu. 2022. Transductive Linear Probing: A Novel Framework for Few-Shot Node Classification. In *LoG*.
- [49] Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. 2021. Csd: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in Neural Information Processing Systems* 34 (2021), 24804–24816.
- [50] Arash Vahdat, Karsten Kreis, and Jan Kautz. 2021. Score-based generative modeling in latent space. *Advances in Neural Information Processing Systems* 34 (2021), 11287–11302.
- [51] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*.
- [52] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph attention networks. In *ICLR*.
- [53] Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. 2022. DiGress: Discrete Denoising diffusion for graph generation. In *The Eleventh International Conference on Learning Representations*.
- [54] Song Wang, Chen Chen, and Jundong Li. 2022. Graph Few-shot Learning with Task-specific Structures. In *NeurIPS*.
- [55] Song Wang, Kaize Ding, Chuxu Zhang, Chen Chen, and Jundong Li. 2022. Task-Adaptive Few-shot Node Classification. In *SIGKDD*.
- [56] Song Wang, Yushun Dong, Xiao Huang, Chen Chen, and Jundong Li. 2022. FAITH: Few-Shot Graph Classification with Hierarchical Task Graphs. In *IJCAI*.
- [57] Shiyu Wang, Xiaojie Guo, and Liang Zhao. 2022. Deep generative model for periodic graphs. *Advances in Neural Information Processing Systems* 35 (2022).
- [58] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. 2018. MoleculeNet: a benchmark for molecular machine learning. *Chemical science* 9, 2 (2018), 513–530.
- [59] Dongchao Yang, Jianwei Yu, Helin Wang, Wen Wang, Chao Weng, Yuexian Zou, and Dong Yu. 2023. Diffsound: Discrete diffusion model for text-to-sound generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* (2023).
- [60] Jinhui Yi, Huan Yan, Haotian Wang, Jian Yuan, and Yong Li. 2023. Deepsta: A spatial-temporal attention network for logistics delivery timely rate prediction in anomaly conditions. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 4916–4922.
- [61] Jinhui Yi, Huan Yan, Haotian Wang, Jian Yuan, and Yong Li. 2024. Learning to Estimate Package Delivery Time in Mixed Imbalanced Delivery and Pickup Logistics Services. In *Proceedings of the 32nd ACM International Conference on Advances in Geographic Information Systems*. 432–443.
- [62] Jinhui Yi, Huan Yan, Haotian Wang, Jian Yuan, and Yong Li. 2024. RCCNet: A Spatial-Temporal Neural Network Model for Logistics Delivery Timely Rate Prediction. *ACM Transactions on Intelligent Systems and Technology* 15, 6 (2024), 1–21.
- [63] Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. 2018. Graph convolutional policy network for goal-directed molecular graph generation. *Advances in neural information processing systems* 31 (2018).
- [64] Jiaxuan You, Rex Ying, Xiang Ren, William Hamilton, and Jure Leskovec. 2018. Graphrnn: Generating realistic graphs with deep auto-regressive models. In *ICML*.
- [65] Chengxi Zang and Fei Wang. 2020. Moflow: an invertible flow model for generating molecular graphs. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 617–626.
- [66] Dinghuai Zhang, Nikolay Malkin, Zhen Liu, Alexandra Volokhova, Aaron Courville, and Yoshua Bengio. 2022. Generative flow networks for discrete probabilistic modeling. In *International Conference on Machine Learning*.
- [67] Yanqiao Zhu, Yuanqi Du, Yinkai Wang, Yichen Xu, Jieyu Zhang, Qiang Liu, and Shu Wu. 2022. A survey on deep graph generation: Methods and applications. In *Learning on Graphs Conference*. PMLR, 47–1.