

Neural Graph Generator: Feature-Conditioned Graph Generation using Latent Diffusion Models

Iakovos Evdaimon¹, Giannis Nikolentzos², Michail Chatzianastasis¹, Hadi Abdine¹, and Michalis Vazirgiannis^{1,3}

¹École Polytechnique, IP Paris, France

²University of Peloponnese, Greece

³Mohamed bin Zayed University of Artificial Intelligence, United Arab Emirates

Abstract

Graph generation has emerged as a crucial task in machine learning, with significant challenges in generating graphs that accurately reflect specific properties. Existing methods often fall short in efficiently addressing this need as they struggle with the high-dimensional complexity and varied nature of graph properties. In this paper, we introduce the **Neural Graph Generator** (NGG), a novel approach which utilizes conditioned latent diffusion models for graph generation. NGG demonstrates a remarkable capacity to model complex graph patterns, offering control over the graph generation process. NGG employs a variational graph autoencoder for graph compression and a diffusion process in the latent vector space, guided by vectors summarizing graph statistics. We demonstrate NGG’s versatility across various graph generation tasks, showing its capability to capture desired graph properties and generalize to unseen graphs. This work signifies a significant shift in graph generation methodologies, offering a more practical and efficient solution for generating diverse types of graphs with specific characteristics.

1 Introduction

In recent years, the field of machine learning on graphs has witnessed an extensive growth, mainly due to the availability of large amounts of data represented as graphs. Indeed, graphs arise naturally in several application domains such as in social networks, in chemo-informatics and in bio-informatics. One of the most challenging tasks of machine learning on graphs is that of graph generation [Zhu et al., 2022]. Graph generation has attracted a lot of attention recently and its main objective is to create novel and realistic graphs. For instance, in chemo-informatics, graph generative models are employed to generate novel, realistic molecular graphs which also exhibit desired properties (e.g., high drug-likeness) [Jin et al., 2018, Zang and Wang, 2020].

Recently, there is a surge of interest in developing new graph generative models, and most of the proposed models typically fall into one of the following five families of models: (1) Auto-Regressive models; (2) Variational Autoencoders; (3) Generative Adversarial Networks; (4) Normalizing Flows; and (5) Diffusion models. These models can capture the complex structural and semantic information of graphs, but focus mainly on specific types of graphs such as molecules [Hoogetboom et al., 2022], proteins [Ingraham et al., 2019], computer programs [Brockschmidt et al., 2019] or patient trajectories [Nikolentzos et al., 2023]. Traditionally, in different application domains, there is a need for generating graphs that exhibit specific properties (e.g., degree distribution, node triangle participation, community structure, etc.).

The different models that are commonly employed to generate these graphs, such as the Erdős–Rényi model and the Barabási-Albert model, capture one or more network properties but neglect others. For example, the Barabási-Albert model can generate graphs with a specific degree distribution, but ignores the rest of the graph’s properties.

Recent graph generative models have improved in accurately representing various properties of real-world graphs. Nonetheless, to produce graphs adhering to specific properties, it is necessary first to compile a dataset of graphs that display these characteristics before proceeding to train the model using this dataset. Thus, to generate different types of graphs, different models need to be trained, which is impractical. This work aims to fill this gap. Specifically, we develop a neural network model, so-called Neural Graph Generator (NGG), to generate graphs that exhibit specific properties. We capitalize on recent advances in latent diffusion models to perform vector-conditioned graph generation. The model applies the diffusion process not to the graph data but instead to an encoded latent representation of the graph. We produce vectors that contain a summary of the statistics of each graph (e. g., number of nodes, number of edges, clustering coefficient, etc.). Those vectors guide the denoising part of the diffusion process. We first pre-train a variational graph autoencoder which we use to map graphs into vectors and map vectors into graphs. The diffusion is performed in the space of vectors. To generate new graphs that exhibit specific properties, we sample a vector from the standard normal distribution, and we denoise this vector, while denoising is guided by the vector of graph properties. Once the noise has been removed from the vector, we feed the emerging vector into the decoder of the pre-trained variational graph autoencoder to produce a graph. The proposed model is evaluated in the task of graph generation conditioned on graph properties. Our results indicate that the NGG model can indeed generate graphs whose properties are similar to the ones that are given as input to the model. To summarize, our work makes the following contributions:

- (i) We introduce **Neural Graph Generator**, a novel graph generative model which leverages latent diffusion for conditional graph generation. This model represents a significant shift from traditional graph generation methods, focusing on prompting with a vector that includes a set of diverse properties of the graph.
- (ii) We introduce a large-scale dataset of synthetic graphs that covers several different types of graphs on which our model was trained. This dataset can be used for pre-training any graph generative model in the future.
- (iii) We extensively evaluate our model across various graph generation tasks, demonstrating its effectiveness in capturing specific graph properties, generalizing to larger graphs, and generating graphs from subsets of properties.
- (iv) We release the pre-trained autoencoder, the pre-trained latent diffusion model, and the synthetic dataset of 1M graphs which are likely to be useful for both practitioners and the scientific community ¹.

2 Related Work

Graph Generative Models. For the problem of graph generation, as discussed above, most models belong to one of the following five families: (1) Auto-regressive models; (2) Variational autoencoders; (3) Generative adversarial networks; (4) Normalizing flows; and (5) Diffusion

¹<https://github.com/iakovosevdaimon/Graph-Generator/>

models. Auto-regressive models assume a specific node ordering and generate graphs in a sequential fashion. GraphRNN [You et al., 2018], GraphGen [Goyal et al., 2020] and GraphGen-Redux [Bacciu and Podda, 2021] are three examples of auto-regressive models. GraphRNN generates nodes and its associated edges sequentially, while the other two models generate edges and their endpoint nodes sequentially. Variational autoencoders consist of two modules. First, the encoder which maps the input data to a space that corresponds to the parameters of a Gaussian distribution. Typically, this distribution is encouraged to be similar to the standard normal distribution. Second, the decoder which maps from the latent space to the input space. In the graph domain, GraphVAE [Simonovsky and Komodakis, 2018] is a model whose encoder is an instance of a message passing graph neural network, while its decoder is a simple multi-layer perceptron. There are models that embed graphs into distributions different from Gaussian such as DGVAE [Li et al., 2020] which utilizes the Dirichlet distributions as priors on the latent variables and the latent variables represent graph cluster memberships. Variational autoencoders have also achieved success in the field of chemoinformatics [Jin et al., 2018, Samanta et al., 2020]. Similar to variational autoencoders, generative adversarial networks also consist of two main components, a generator whose objective is to generate realistic samples and a discriminator whose objective is to distinguish real samples from synthetic ones. Most models from this family typically consist of a generator that takes vectors sampled from a standard normal distribution and processes them with a multi-layer perceptron to generate a graph, and of a discriminator which is an instance of a message passing graph neural network [De Cao and Kipf, 2018, Pölsterl and Wachinger, 2021]. Architectures based on normalizing flows explicitly model a probability distribution by leveraging a method that uses the change-of-variable law of probabilities to transform a simple distribution into a complex one by a sequence of invertible and differentiable mappings. Normalizing flows can be applied to graphs by designing message passing mechanisms that are exactly reversible [Liu et al., 2019]. Such models have been applied to molecular data. For example, MoFlow [Zang and Wang, 2020] generates chemical bonds through a Glow-based model, and atoms given bonds using a graph conditional flow. Diffusion models gradually add noise to the input data in the forward diffusion process, and then learn to remove the noise in the reverse diffusion process. There exist models that add Gaussian noise to the graph’s adjacency matrix and binarize the continuous values of the output of the reverse diffusion process to produce valid graphs [Niu et al., 2020, Jo et al., 2022]. More recent models apply a discrete diffusion process that progressively adds noise by adding or removing edges and changing the nodes’ and/or edges’ types [Vignac et al., 2023].

Conditional Graph Generation. One of the main objectives of graph generative models is conditional generation, which is the process of generating a graph that satisfies a specific label or property. Typically, a conditional code is introduced that explicitly controls the property of generated graphs. Such models are popular in the field of chemo-informatics where novel molecules that possess desired chemical properties (e.g., high binding affinity against a target protein) need to be generated. For variational autoencoders, a simple approach is to feed the concatenation of the conditional code and the latent graph representation to the decoder [Simonovsky and Komodakis, 2018]. The conditional code can also be concatenated to node representations sampled from a single distribution [Yang et al., 2019]. In models that generate new nodes and edges sequentially, once a new node is created, its feature vector must be initialized, and the conditional code could be concatenated to that feature vector [Li et al., 2018a]. There exist models whose message passing scheme is modified to include the conditional code [Li et al., 2018b]. MOOD [Lee et al., 2023] utilizes the gradient of a property prediction network to guide the sampling process to domains that are likely to satisfy

specific properties. There are also models that employ regularization schemes to learn disentangled representations where each dimension focuses on a specific property [Du et al., 2022].

Latent Diffusion Models. Diffusion models are generative models that have recently gained significant attention [Ho et al., 2020, Sohl-Dickstein et al., 2015, Vincent, 2011]. Diffusion models have achieved state-of-the-art performance in several tasks such as in conditional image generation [Dhariwal and Nichol, 2021, Song et al., 2020], image colorization [Saharia et al., 2022a] and image super resolution [Saharia et al., 2022b], just to name a few. Latent diffusion models were introduced to reduce the computational complexity of diffusion models. These models use some autoencoding component to project input data (e.g., images) into some latent space and perform the diffusion process in that space [Rombach et al., 2022]. Latent diffusion models can thus be trained on limited computational resources while retaining the quality of diffusion models. Latent diffusion models have been applied to different problems such as image synthesis [Rombach et al., 2022], video synthesis [Blattmann et al., 2023] and image reconstruction [Takagi and Nishimoto, 2023]. Recently, latent diffusion models have been introduced in the molecular geometry domain, and have been evaluated in the task of 3D molecule generation [Xu et al., 2023], but also for generating protein backbone structures [Fu et al., 2023]. These two works are the closest to our method. Even though architecture-wise our model is similar to the models presented in these papers, we follow an entirely different research direction and focus on conditional generation of general graphs (similar to traditional graph generators) instead of certain classes of graphs such as molecules and proteins.

3 Neural Graph Generator

In this section, we first introduce some key notation for graphs, and we then present the two main components of the proposed NGG model: (1) the variational graph autoencoder which produces a compressed latent representation for each graph; and (2) the diffusion model which performs diffusion in the latent space and which can be conditioned on various inputs (vector of graph properties in our case). An overview of the proposed model is given in Figure 1.

3.1 Notation

Let $[n] = \{1, \dots, n\} \subset \mathbb{N}$ for $n \geq 1$. Let $G = (V, E)$ be an undirected graph, where V is the set of nodes and E is the edge set. We will denote by n the number of vertices and by m the number of edges, i.e., $n = |V|$ and $m = |E|$. The neighborhood $\mathcal{N}(v)$ of a node v is the set of all vertices adjacent to v . Hence, $\mathcal{N}(v) = \{u \mid (v, u) \in E\}$ where (v, u) is an edge between the vertices v and u of V . The adjacency matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ of a graph G is a symmetric (typically sparse) matrix that is used to encode edge information in the graph. The element of the i^{th} row and j^{th} column is equal to the weight of the edge between the vertices v_i and v_j if such an edge exists, and 0 otherwise. The degree $d(v)$ of a node v is equal to the number of edges that are adjacent to the node, i.e., $d(v) = |\mathcal{N}(v)|$. In different applications, nodes of graphs are annotated with feature vectors. We use $\mathbf{X} \in \mathbb{R}^{n \times d}$ to denote the node features where d is the feature dimension size. The feature of a given node v_i corresponds to the i^{th} row of \mathbf{X} .

3.2 Graph Compression

To map graphs into low-dimensional representations, we capitalize on previous work, and we use an autoencoder [Simonovsky and Komodakis, 2018]. More precisely, given a graph G , the en-

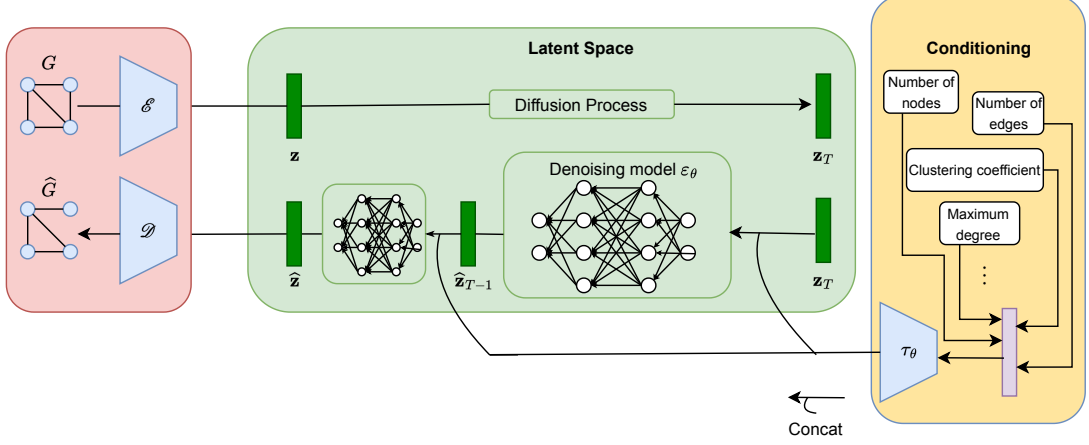


Figure 1: Overview of the proposed architecture. The variational graph autoencoder is responsible for generating a compressed latent representation for each graph. Those representations are fed to the diffusion model which operates in that latent space. The denoising process is conditioned on a vector that contains the graph’s properties. The output of the diffusion model is passed on to the decoder which generates a graph.

coder \mathcal{E} encodes the graph into a latent representation $\mathbf{z} = \mathcal{E}(G)$, and the decoder \mathcal{D} reconstructs the graph from the latent representation, giving $\hat{G} = \mathcal{D}(\mathbf{z}) = \mathcal{D}(\mathcal{E}(G))$, where $\mathbf{z} \in \mathbb{R}^d$.

The encoder \mathcal{E} corresponds to a message passing neural network which consists of GIN layers [Xu et al., 2019]. Node representations are updated as follows:

$$\mathbf{h}_v^{(k)} = \text{MLP}^{(k)} \left(\left(1 + \epsilon^{(k)}\right) \mathbf{h}_v^{(k-1)} + \sum_{u \in \mathcal{N}(v)} \mathbf{h}_u^{(k-1)} \right)$$

where k denotes the layer and ϵ is a trainable parameter. Note that the nodes are initially annotated with features extracted from the eigenvectors of the normalized Laplacian matrix. These features are stored in matrix \mathbf{X} , while $\mathbf{h}_v^{(0)}$ is equal to the row of matrix \mathbf{X} that corresponds to node v . To produce a representation for the entire graph, we use the sum operator as follows:

$$\mathbf{h}_G = \sum_{v \in V} \mathbf{h}_v^{(K)}$$

where K is the number of layers. To avoid arbitrarily high-variance latent spaces, we actually use a variational autoencoder. We thus use fully-connected layers $\text{MLP}_\mu, \text{MLP}_\sigma$, to embed the graphs into the parameters of a Gaussian distribution. We also add the standard regularization term into the loss function which imposes a slight penalty towards a standard normal distribution on the learned latent representations. Then, the vector \mathbf{z} is sampled from the learned Gaussian distribution:

$$\boldsymbol{\mu} = \text{MLP}_\mu(\mathbf{h}_G), \quad \boldsymbol{\sigma} = \text{MLP}_\sigma(\mathbf{h}_G), \quad \mathbf{z} = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \epsilon$$

where \odot denotes element-wise multiplication, and ϵ is a random noise vector sampled from $\mathcal{N}(0, \mathbf{I})$.

The decoder \mathcal{D} takes the vector \mathbf{z} as input and reconstructs the input graph. We experiment with a simple MLP decoder that consists of a series of fully connected layers. Specifically, the decoder is implemented as follows:

$$\hat{\mathbf{A}} = \text{MLP}_{\mathcal{D}}(\mathbf{z})$$

Note that $\hat{\mathbf{A}}$ is an $(n_{\max} \times n_{\max})$ -dimensional matrix where n_{\max} is a pre-defined maximum graph size. In fact, the MLP outputs the upper triangular part of this matrix, and from this $\hat{\mathbf{A}}$ is constructed. Similar to prior work [De Cao and Kipf, 2018], we found that pre-defining the maximum graph size makes the model significantly faster and also easier to optimize. Note that $\hat{\mathbf{A}}$ contains continuous values and has a probabilistic interpretation since each element represents the probability that two nodes are connected by an edge. To produce a discrete object, we use the straight-through Gumbel-Softmax [Jang et al., 2017], i.e., we use a sample from a categorical distribution in the forward pass, while the relaxed values are utilized in the backward pass. The matrix $\hat{\mathbf{A}}$ is then compared with the adjacency matrix of the input graph \mathbf{A} to compute the reconstruction loss.

To compute the reconstruction loss, following previous work, we could employ some graph matching algorithm to compare each input graph against the corresponding reconstructed graph [Simonovsky and Komodakis, 2018]. However, such algorithms are expensive to compute, while we also found empirically that we can achieve similar (or even higher) levels of performance if we just impose a specific ordering on the nodes of each graph and compare the reconstructed adjacency against the adjacency that follows this ordering. In general, a node ordering that is effective for one type of graphs might not be effective for other types of graphs. For instance, in the task of molecule generation, certain canonical orderings are more effective than others [Li et al., 2018a]. Since our dataset consists of graphs that exhibit very different properties from each other, in order to create a general graph generative model, we experimented with the following node orderings: (1) nodes are sorted by degree (from higher to lower); (2) ordering corresponds to BFS/DFS tree’s default ordering, which is descending by node degree and rooted at the node with the highest degree; and (3) nodes are sorted by Pagerank scores (from higher to lower).

3.3 Latent Diffusion Model

Once the autoencoder is trained, we can use the encoder \mathcal{E} to embed graphs into a low-dimensional latent space. Those embeddings are expected to capture both local and global properties of input graphs. The main advantage of latent diffusion models over standard diffusion models is their efficiency. The smaller latent space representation makes executing the diffusion process much faster. This allows the models to be trained on a single or a few GPUs instead of hundreds of GPUs. Indeed, recent diffusion models for graphs [Vignac et al., 2023] operate on adjacency tensors (i.e., tensors of dimension $n \times n \times d$), and thus are expensive to train, while it is also expensive to generate new graphs.

Latent diffusion models consist of two main components: (1) a noise model; and (2) a denoising neural network. The noise model q progressively corrupts the latent representation of the input graph \mathbf{z} to create a sequence of increasingly noisy vectors $(\mathbf{z}_1, \dots, \mathbf{z}_T)$. Specifically, the forward process q samples some noise from a Gaussian distribution at each time step t , which is added to the representation of the previous time step as follows:

$$q(\mathbf{z}_t | \mathbf{z}_{t-1}) = \mathcal{N}(\mathbf{z}_t; \sqrt{1 - \beta_t} \mathbf{z}_{t-1}, \beta_t \mathbf{I})$$

where β_t is a known variance schedule and $0 < \beta_1 < \dots < \beta_T < 1$. Note that the forward process q is fixed, and thus, \mathbf{z}_t can be obtained directly from \mathbf{z} during training. Specifically, we have that:

$$q(\mathbf{z}_t | \mathbf{z}) = \mathcal{N}(\mathbf{z}_t; \sqrt{\bar{a}_t} \mathbf{z}, (1 - \bar{a}_t) \mathbf{I})$$

where $\bar{a}_t = \prod_{i=1}^t a_i$ and $a_t = 1 - \beta_t$. Note that if the schedule is set appropriately, \mathbf{z}_T is pure Gaussian noise. The second component of the latent diffusion model (i.e., the denoising neural

network) is responsible for predicting the added noise for a given time step t . To train the denoising neural network ϵ_θ , the following function is minimized:

$$\begin{aligned} L_{\text{LDM}} &= \mathbb{E}_{\mathcal{E}(\mathcal{G}), \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_\theta(\mathbf{z}_t, t)\|_2^2 \right] \\ &= \mathbb{E}_{\mathcal{E}(\mathcal{G}), \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_\theta(\sqrt{\bar{a}_t}\mathbf{z} + \sqrt{(1-\bar{a}_t)}\epsilon, t)\|_2^2 \right] \end{aligned}$$

We implement the denoising neural network ϵ_θ as an MLP. First, we use sinusoidal positional embeddings to produce a unique representation for each time step t . Then, we feed vector \mathbf{z}_t to the MLP and in each hidden layer of the MLP, we add the positional embedding to the latent representation.

Once the denoising neural network ϵ_θ is trained, to generate new data, we sample a vector of pure noise \mathbf{z}_T from a Gaussian distribution, and then use the neural network to gradually denoise it (using the conditional probability it has learned). The denoised vector can then be transformed into a graph with a single pass through the decoder \mathcal{D} .

3.4 Conditioning

The ability to condition graph generation on local and global properties of graphs is crucial for the success of our work. Latent diffusion models are capable of modeling conditional distributions. Specifically, text, images, or other inputs can be encoded into the latent space and used to condition the model to generate outputs with desired properties. Given some condition code \mathbf{c} , diffusion models are capable of modeling the conditional distribution $p(\mathbf{z}|\mathbf{c})$ by using a conditional neural network $\epsilon_\theta(\mathbf{z}_t, t, \tau_\theta(\mathbf{c}))$ as follows:

$$L_{\text{LDM}} = \mathbb{E}_{\mathcal{E}(\mathcal{G}), \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_\theta(\mathbf{z}_t, t, \tau_\theta(\mathbf{c}))\|_2^2 \right]$$

where ϵ_θ and τ_θ are neural network models that are jointly optimized. The architecture of τ_θ depends on the conditioning modality. As already discussed, the condition code is a vector of properties of the graph. Thus, we use an MLP as the condition encoder τ_θ to compute \mathbf{y} as follows:

$$\mathbf{y} = \text{MLP}_c(\mathbf{c})$$

Then, we concatenate vectors \mathbf{z} and \mathbf{y} and feed them into the denoising neural network. We experimented with other operations (e.g., addition, element-wise product), but in preliminary experiments, concatenation led to better results than the other operations.

The condition code consists of 15 local and global properties of the input graph. Those properties are listed in Table 4 in Appendix A. Those specific 15 properties were chosen mainly because they cover a broad range of graph properties and they can give insights about the structure of the graph. Furthermore, the properties can be computed efficiently in polynomial time.

4 Experimental Evaluation

We train and evaluate the NGG model on a dataset that contains 1M synthetic graphs. We consider that the evaluation encompasses three distinct scenarios. Initially, we assess the model’s performance on our original dataset. Subsequently, we evaluate its capabilities when trained on graphs containing up to 50 nodes, test its adaptability to larger graphs with more than 50 nodes, and examine its generalization performance. Lastly, we randomly conceal between 1 to 8 properties of some of the input graphs (i.e., number of nodes, edges, density, etc.), treating

them as missing values not provided by the user. Subsequently, we retrain the model and assess its performance on the test set, repeating the process of concealing a random number of graph properties.

Dataset. We train and evaluate the NGG model on a dataset that contains 1M synthetic graphs, where each graph contains at most 100 nodes. To create the dataset, we use different types of graph generators. This allows us to construct graphs from different families and thus, cover a wide range of values of the considered properties. Each synthetic graph belongs to one of the following 17 families of graphs: (1) paths; (2) cycles; (3) wheels; (4) stars; (5) ladders; (6) lollipops; (7) Erdős-Rényi random graphs; (8) Newman-Watts-Strogatz small-world graphs; (9) Watts-Strogatz small-world graphs; (10) random d -degree regular graphs; (11) Barabási-Albert graphs; (12) dual Barabási-Albert graphs; (13) extended Barabási-Albert graphs; (14) graphs generated using the Holme and Kim algorithm; (15) random lobsters; (16) stochastic block model graphs; and (17) random partition graphs. More details about the constructed dataset are given in Appendix A. The generated graphs in our dataset are devoid of self-loops, isolated nodes, and multigraphs are also excluded.

Experimental setup. With regard to the variational autoencoder, the encoder \mathcal{E} consists of 2 layers, the hidden dimension size is set equal to 64, while the input graphs are embedded into 32-dimensional vectors. We imposed an ordering on the nodes based on the BFS tree’s default ordering. The decoder \mathcal{D} consists of 3 layers and the hidden dimension size is set to 256. The output of \mathcal{D} is a 100^2 -dimensional vector and is transformed into an adjacency matrix of dimension (100×100) . Thus, the model can handle graphs of up to 100 nodes. We also initially annotate the nodes of each graph with 10-dimensional spectral features (from the eigenvectors associated with the 10 smallest eigenvalues of the normalized Laplacian matrix). For the latent diffusion model, we set the number of timesteps T equal to 500, while we use a cosine schedule for adding noise at each time step. The denoising neural network ϵ_θ consists of 3 layers, the hidden dimension size is equal to 512, while the size of the output is set to 128. To train both models, we use the Adam optimizer and we set the learning rate to 0.001. The batch size is set to 256. We set the number of epochs of the variational autoencoder to 200, while the number of epochs of the latent diffusion model to 100. Overall, the variational autoencoder consists of 2,640,492 parameters, while the diffusion model consists of 973,088 parameters. Thus, in total, the NGG model consists of 3,613,580 parameters.

Evaluation and baseline. Given a collection of graphs $\{G_1, G_2, \dots, G_N\}$ that belong to the test set, we use the trained generative model to produce another collection of graphs $\{\hat{G}_1, \hat{G}_2, \dots, \hat{G}_N\}$. To generate each graph of the new collection, we sample pure noise from a Gaussian distribution and use the neural network to gradually denoise it, while the whole denoising process is conditioned on the vector of properties of the corresponding graph of the test set G_i . Once the noise has been removed, the emerging vector is transformed into a graph \hat{G}_i with a single pass through \mathcal{D} . Note, however, that we are not interested in comparing the generated graphs against the corresponding graphs of the test set, but we are interested in finding whether the properties of the generated graphs match the properties of the corresponding graphs of the test set. Therefore, for each generated graph \hat{G}_i , we compute a vector where each component is the value of each one of the 15 considered properties. Then, this vector is compared against the vector of properties of G_i . We use two different metrics: (1) the mean absolute error (MAE); and (2) the symmetric mean absolute percentage error (SMAPE), which we use instead of the mean absolute percentage error (MAPE) because we have values equal to zero or close to zero. SMAPE values range from 0% to 100%, while lower values indicate better

Table 1: Within distribution performance of the proposed NGG model and the baseline model in terms of the considered properties.

Property	VGAE		NGG	
	MAE	SMAPE	MAE	SMAPE
# nodes	24.22	25.18	2.63	3.09
# edges	701.99	66.48	62.33	8.44
Density	0.32	52.13	0.04	7.23
Min. degree	13.95	64.52	11.61	49.46
Max. degree	14.59	22.32	1.59	3.55
Avg. degree	18.99	58.60	1.64	6.68
Assortativity coefficient	0.30	61.32	0.11	39.10
# triangles	10,356.20	99.91	1,026.44	24.38
Avg. # triangles formed by an edge	8.85	66.52	9.44	68.32
Max. # triangles formed by an edge	539.03	83.97	49.26	16.66
Avg. local clustering coefficient	0.29	35.04	0.08	15.42
Global clustering coefficient	0.36	58.12	0.05	14.07
Max. k -core	15.07	54.51	1.66	8.61
# communities	1.74	21.86	0.96	12.34
Diameter	3.73	31.31	2.40	15.96
All	0.80	55.96	0.23	21.05

performance, as they represent a smaller percentage difference between the predicted and actual values. We calculate the aforementioned metrics by taking into account all 15 properties and applying z-score normalization to them.

We also use a variant of the variational autoencoder (VGAE) as a baseline where the process is conditioned on vector of properties \mathbf{c} for controlled sampling by concatenating the vector with the graph’s latent representation.

4.1 Within Distribution Performance

In this experiment, we split the constructed dataset into a training, a validation and a test set. Then, we train the proposed NGG model and the baseline on the training set. We choose the models that achieve the lowest loss on the validation set and we evaluate their performance on the test set. More specifically, the 1M samples are split into training, validation, and test sets with a 80 : 10 : 10 split ratio, respectively. Thus, the training set contains 800,000 samples, while each of the validation and test sets contain 100,000 samples. Each one of the three sets preserves approximately the proportions of each graph type. Table 1 illustrates the MAE and SMAPE achieved by NGG and VGAE for each property individually and for all properties together (last row). First of all, we observe that the proposed NGG model outperforms VGAE in terms of almost all considered properties. Secondly, we notice that NGG accurately captures most of the properties, while it struggles with the minimum degree, the number of triangles, and the average and the maximum number of triangles formed by an edge.

4.2 Out of Distribution Generalization

In many scenarios, machine learning models fail to generalize to unseen data. Thus, in this second experiment, we study whether the proposed NGG model and the baseline model can generate graphs whose properties’ values are different from the ones the models were trained on (i.e. out of distribution). Specifically, we train the two models on graphs of up to 40 nodes and


Table 2: Performance Comparison of NGG and Baseline Model under Different Conditions: Out-of-Distribution Performance (trained on graphs with up to 50 nodes and evaluated on larger graphs) and Within-Distribution Performance with Masking Applied to some Condition Vector Elements.

Property	Out of Distribution				Masked			
	VGAE		NGG		VGAE		NGG	
	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE	MAE	SMAPE
# nodes	15.48	18.76	6.68	6.66	25.70	26.05	27.30	27.54
# edges	355.79	60.12	99.09	11.31	785.20	49.83	832.26	52.90
Density	0.28	48.10	0.07	10.74	0.27	34.04	0.29	36.65
Min. degree	12.59	67.12	7.95	41.49	13.96	57.71	13.36	57.50
Max degree	4.46	7.95	2.93	4.64	23.97	35.17	26.04	38.06
Avg. degree	14.34	51.23	3.11	9.27	18.83	40.84	20.17	43.77
Assortativity coefficient	0.29	65.66	0.72	47.57	0.46	60.98	0.43	65.18
# triangles	3,913.93	83.79	1,482.23	27.21	15,101.24	74.92	15,956.96	79.93
Avg . # triangles formed by an edge	7.78	54.99	17.93	74.95	9.11	49.11	17.73	77.18
Max . # triangles formed by an edge	278.36	64.51	83.15	17.71	709.65	66.55	762.29	71.23
Avg. local clustering coefficient	0.26	26.01	0.09	15.91	0.33	40.78	0.36	44.69
Global clustering coefficient	0.27	46.16	0.07	14.44	0.32	41.89	0.34	44.34
Max k-core	11.39	52.05	2.28	9.99	16.93	42.96	17.99	45.76
# communities	2.64	27.47	1.02	12.32	1.86	22.98	1.96	23.40
Diameter	3.52	29.98	2.55	16.61	3.40	23.82	3.31	24.19
All	0.96	56.13	0.54	28.89	0.77	42.14	0.78	42.88


also on graphs that consist of more than 60 nodes. Then, the models are evaluated on the rest of the graphs (i. e., graphs of size from 41 nodes to 59 nodes). The results of this experiment are illustrated in Table 2 (Left). We observe that the proposed NGG model still outperforms the VGAE model in this experiment. However, the VGAE model can generate graphs that better capture the average number of triangles formed by an edge than NGG. Furthermore, we can also see that both models achieve a bit lower levels of performance compared to their corresponding performance reported in Table 1. This is not surprising since the model is trained on graphs that are larger or smaller than the ones on which it is evaluated.

4.3 Conditioning on Subset of Properties

Since a practitioner might be interested only in a handful of properties, and not in all 15 of them, we also investigate whether the proposed model can generate graphs providing only a subset of the properties. Within each batch, we randomly choose if the properties of each sample would be masked, treating some of the elements of their associated condition vector as missing values as follows. For each one of the chosen samples, we randomly choose i elements (without replacement) where $i \in \{1, 2, \dots, 8\}$ and replace the values of those elements with a value equal to -100 . Note that we also corrupt similarly the samples of the test set at inference time. We then compare the observed properties of the condition vector against the corresponding properties of the generated samples to compute the three evaluation metrics. The obtained results are shown in Table 2 (Right). We observe that the baseline model outperforms the proposed NGG model, only by a small margin though. Furthermore, the model achieves significantly worse levels of performance compared to the ones reported in Table 1. While the proposed model generally proves effective in the task of graph generation, its performance diminishes when only a subset of the considered properties are available, failing to achieve the



(a) Generated graph that emerged from condition vector \mathbf{c}_1



(b) Generated graph that emerged from condition vector \mathbf{c}_2

Figure 2: Example of two graphs generated by the proposed NGG model given condition vectors \mathbf{c}_1 and \mathbf{c}_2 .

same levels of performance as when all properties are present.

4.4 Examples of Generated Graphs

We will next give two examples of graphs generated by the proposed NGG model. To generate the two graphs, we used the model that was trained on 80% of the entire dataset. For ease of presentation, we decided to generate two graphs that both consist of 15 nodes. However, the second graph is much more dense than the first graph. Specifically, we utilized the following two condition vectors:

$$\mathbf{c}_1 = (15 \ 34 \ 0.32 \ 2 \ 8 \ 4.5 \ -0.046 \ 17 \ 1.5 \ 8 \ 0.4 \ 0.35 \ 4 \ 4 \ 4)^\top$$

$$\mathbf{c}_2 = (15 \ 94 \ 0.89 \ 10 \ 14 \ 12.5 \ -0.149 \ 329 \ 10.5 \ 80 \ 0.9 \ 0.9 \ 10 \ 2 \ 2)^\top$$

where each dimension corresponds to one of the 15 properties respectively. The above two vectors gave rise to the two graphs that are shown in Figure 2. The corresponding communities for each generated graph are also illustrated (node color denotes the community to which the node belongs). We observe that one generated graph indeed consists of 15 nodes, while the other consists of 16 nodes. Furthermore, the graph shown in Figure 2(a) is indeed sparse, while the one shown in Figure 2(b) is dense. We also computed the values of the considered properties for both graphs and these are given below:

$$\hat{\mathbf{c}}_1 = (15 \ 25 \ 0.23 \ 1 \ 7 \ 3.3 \ -0.380 \ 8 \ 0.96 \ 3 \ 0.43 \ 0.32 \ 2 \ 3 \ 4)^\top$$

$$\hat{\mathbf{c}}_2 = (16 \ 108 \ 0.9 \ 3 \ 15 \ 13.5 \ -0.148 \ 458 \ 12.7 \ 93 \ 0.97 \ 0.97 \ 14 \ 2 \ 2)^\top$$

We can see that for most properties, their actual values are close to the ones that were utilized for generating the two graphs. To summarize, those results provide a qualitative validation of the proposed NGG model. We also visually inspected several other graphs and it turns out that the proposed model can generate graphs that approximately exhibit the desired properties.

4.5 Uniqueness of Generated Graphs

We also investigated whether given a vector of properties, the proposed NGG model can generate non-isomorphic graphs. Specifically, we randomly chose 50 condition codes, and for each code,

we used the model to produce 100 graphs. We then tested every pair of these graphs for isomorphism (9900/2 pairs in total for each condition code). We found that all generated graphs are unique, i.e., there was no pair of isomorphic graphs. Therefore, the NGG model seems to also produce diverse sets of graphs.

5 Conclusion

We introduced the Neural Graph Generator, a conditional latent diffusion model for efficient and accurate graph generation. NGG represents a significant new paradigm, adeptly generating graphs with specific, user-defined properties, a task that has long challenged traditional models. We demonstrated the ability of NGG to generalize in graphs than those in the training set, and the ability to handle missing values in the condition vector. Nevertheless, GNN models, while powerful, struggle to capture specific properties like triangles or cycles, thus affecting the performance of our model. Future work will focus on exploring the model’s applicability in various real-world scenarios.

References

- [Bacciu and Podda, 2021] Bacciu, D. and Podda, M. (2021). GraphGen-Redux: A Fast and Lightweight Recurrent Model for labeled Graph Generation. In *Proceedings of the 2021 International Joint Conference on Neural Networks*, pages 1–8.
- [Blattmann et al., 2023] Blattmann, A., Rombach, R., Ling, H., Dockhorn, T., Kim, S. W., Fidler, S., and Kreis, K. (2023). Align your Latents: High-Resolution Video Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22563–22575.
- [Brockschmidt et al., 2019] Brockschmidt, M., Allamanis, M., Gaunt, A. L., and Polozov, O. (2019). Generative Code Modeling with Graphs. In *7th International Conference on Learning Representations*.
- [De Cao and Kipf, 2018] De Cao, N. and Kipf, T. (2018). MolGAN: An implicit generative model for small molecular graphs. *arXiv preprint arXiv:1805.11973*.
- [Dhariwal and Nichol, 2021] Dhariwal, P. and Nichol, A. (2021). Diffusion Models Beat GANs on Image Synthesis. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, pages 8780–8794.
- [Du et al., 2022] Du, Y., Guo, X., Shehu, A., and Zhao, L. (2022). Interpretable Molecular Graph Generation via Monotonic Constraint. In *Proceedings of the 2022 SIAM International Conference on Data Mining*, pages 73–81.
- [Fu et al., 2023] Fu, C., Yan, K., Wang, L., Au, W. Y., McThrow, M., Komikado, T., Maruhashi, K., Uchino, K., Qian, X., and Ji, S. (2023). A Latent Diffusion Model for Protein Structure Generation. In *Proceedings of the 2nd Learning on Graphs Conference*.
- [Goyal et al., 2020] Goyal, N., Jain, H. V., and Ranu, S. (2020). GraphGen: A Scalable Approach to Domain-agnostic Labeled Graph Generation. In *Proceedings of The Web Conference 2020*, pages 1253–1263.

- [Ho et al., 2020] Ho, J., Jain, A., and Abbeel, P. (2020). Denoising Diffusion Probabilistic Models. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 6840–6851.
- [Hoogeboom et al., 2022] Hoogeboom, E., Satorras, V. G., Vignac, C., and Welling, M. (2022). Equivariant Diffusion for Molecule Generation in 3D. In *Proceedings of the 39th International Conference on Machine Learning*, pages 8867–8887.
- [Ingraham et al., 2019] Ingraham, J., Garg, V. K., Barzilay, R., and Jaakkola, T. (2019). Generative models for graph-based protein design. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 15820–15831.
- [Jang et al., 2017] Jang, E., Gu, S., and Poole, B. (2017). Categorical Reparameterization with Gumbel-Softmax. In *5th International Conference on Learning Representations*.
- [Jin et al., 2018] Jin, W., Barzilay, R., and Jaakkola, T. (2018). Junction Tree Variational Autoencoder for Molecular Graph Generation. In *Proceedings of the 35th International Conference on Machine Learning*, pages 2323–2332.
- [Jo et al., 2022] Jo, J., Lee, S., and Hwang, S. J. (2022). Score-based Generative Modeling of Graphs via the System of Stochastic Differential Equations. In *Proceedings of the 39th International Conference on Machine Learning*, pages 10362–10383.
- [Lee et al., 2023] Lee, S., Jo, J., and Hwang, S. J. (2023). Exploring Chemical Space with Score-based Out-of-distribution Generation. In *Proceedings of the 40th International Conference on Machine Learning*, pages 18872–18892.
- [Li et al., 2020] Li, J., Yu, J., Li, J., Zhang, H., Zhao, K., Rong, Y., Cheng, H., and Huang, J. (2020). Dirichlet Graph Variational Autoencoder. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 5274–5283.
- [Li et al., 2018a] Li, Y., Vinyals, O., Dyer, C., Pascanu, R., and Battaglia, P. (2018a). Learning Deep Generative Models of Graphs. *arXiv preprint arXiv:1803.03324*.
- [Li et al., 2018b] Li, Y., Zhang, L., and Liu, Z. (2018b). Multi-objective de novo drug design with conditional graph generative model. *Journal of Cheminformatics*, 10:1–24.
- [Liu et al., 2019] Liu, J., Kumar, A., Ba, J., Kiros, J., and Swersky, K. (2019). Graph Normalizing Flows. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 13578–13588.
- [Nikolentzos et al., 2023] Nikolentzos, G., Vazirgiannis, M., Xypolopoulos, C., Lingman, M., and Brandt, E. G. (2023). Synthetic electronic health records generated with variational graph autoencoders. *NPJ Digital Medicine*, 6(1):83.
- [Niu et al., 2020] Niu, C., Song, Y., Song, J., Zhao, S., Grover, A., and Ermon, S. (2020). Permutation Invariant Graph Generation via Score-Based Generative Modeling. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, pages 4474–4484.
- [Pölsterl and Wachinger, 2021] Pölsterl, S. and Wachinger, C. (2021). Adversarial learned molecular graph inference and generation. In *Proceedings of the 2020 European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 173–189.

- [Rombach et al., 2022] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695.
- [Saharia et al., 2022a] Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., and Norouzi, M. (2022a). Palette: Image-to-Image Diffusion Models. In *ACM SIGGRAPH 2022 Conference Proceedings*, pages 1–10.
- [Saharia et al., 2022b] Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D. J., and Norouzi, M. (2022b). Image Super-Resolution via Iterative Refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4713–4726.
- [Samanta et al., 2020] Samanta, B., De, A., Jana, G., Gómez, V., Chattaraj, P., Ganguly, N., and Gomez-Rodriguez, M. (2020). NeVAE: A Deep Generative Model for Molecular Graphs. *Journal of Machine Learning Research*, 21(114):1–33.
- [Simonovsky and Komodakis, 2018] Simonovsky, M. and Komodakis, N. (2018). GraphVAE: Towards Generation of Small Graphs Using Variational Autoencoders. In *Proceedings of the 27th International Conference on Artificial Neural Networks*, pages 412–422.
- [Sohl-Dickstein et al., 2015] Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. (2015). Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 2256–2265.
- [Song et al., 2020] Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. (2020). Score-Based Generative Modeling through Stochastic Differential Equations. In *8th International Conference on Learning Representations*.
- [Takagi and Nishimoto, 2023] Takagi, Y. and Nishimoto, S. (2023). High-resolution image reconstruction with latent diffusion models from human brain activity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14453–14463.
- [Vignac et al., 2023] Vignac, C., Krawczuk, I., Siraudin, A., Wang, B., Cevher, V., and Frossard, P. (2023). Digress: Discrete denoising diffusion for graph generation. In *11th International Conference on Learning Representations*.
- [Vincent, 2011] Vincent, P. (2011). A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674.
- [Xu et al., 2019] Xu, K., Hu, W., Leskovec, J., and Jegelka, S. (2019). How Powerful are Graph Neural Networks? In *7th International Conference on Learning Representations*.
- [Xu et al., 2023] Xu, M., Powers, A. S., Dror, R. O., Ermon, S., and Leskovec, J. (2023). Geometric Latent Diffusion Models for 3D Molecule Generation. In *Proceedings of the 40th International Conference on Machine Learning*, pages 38592–38610.
- [Yang et al., 2019] Yang, C., Zhuang, P., Shi, W., Luu, A., and Li, P. (2019). Conditional Structure Generation through Graph Variational Generative Adversarial Nets. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 1340–1351.
- [You et al., 2018] You, J., Ying, R., Ren, X., Hamilton, W., and Leskovec, J. (2018). GraphRNN: Generating Realistic Graphs with Deep Auto-regressive Models. In *Proceedings of the 35th International Conference on Machine Learning*, pages 5708–5717.

- [Zang and Wang, 2020] Zang, C. and Wang, F. (2020). MoFlow: An Invertible Flow Model for Generating Molecular Graphs. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 617–626.
- [Zhu et al., 2022] Zhu, Y., Du, Y., Wang, Y., Xu, Y., Zhang, J., Liu, Q., and Wu, S. (2022). A Survey on Deep Graph Generation: Methods and Applications. In *Proceedings of the 1st Learning on Graphs Conference*, pages 47:1–47:21.

A Dataset Details

The distribution of the samples per family of graphs is illustrated in Table 3.

Table 3: Distribution of samples per graph type

Type of graph	Number of samples	Proportion
Barabási–Albert random graph	250,136	25.01%
Watts–Strogatz small-world graph	204,280	20.43%
stochastic block model graph	125,468	12.55%
Erdős–Rényi graph	122,568	12.26%
dual Barabási–Albert random graph	122,568	12.26%
extended Barabási–Albert model graph	122,568	12.26%
Newman–Watts–Strogatz small-world graph	122,567	12.26%
Holme and Kim algorithm for growing graphs with power-law degree distribution and approximate average clustering	122,567	12.26%
Random Lobster graph	81,713	8.17%
random partition graph	81,713	8.17%
Random d -degree regular graph	7,000	0.7%
Lollipop graph	4,145	0.41%
Path graph	91	0.01%
Cycle graph	91	0.01%
Star graph	91	0.01%
Wheel graph	91	0.01%
Ladder graph	46	0.0046%

Table 4: Description of the 15 considered graph properties.

Property	Range	Description
# nodes	[2,100]	Count of individual vertices or entities in a graph.
# edges	[1,4950]	Count of connections or relationships between nodes in the graph.
Density	[0.01,1]	A measure of how interconnected a graph is, calculated as the number of actual edges over the total possible edges.
Min. degree	[1,99]	Minimum degree is the smallest degree among all nodes.
Max. degree	[1,99]	Maximum degree is the largest degree among all nodes.
Avg. degree	[1,99]	Mean degree across all nodes.
Assortativity coefficient	[-1,1]	A measure of the tendency of nodes to connect with others of similar degree. Positive values suggest a preference for connections between nodes of similar degrees.
# triangles	[0, 161700]	Number of sets of three nodes that form a closed loop (triangle) in the network.
Avg. # triangles formed by an edge	[0, 98]	Average number of triangles each edge participates in.
Max. # triangles formed by an edge	[0, 4851]	Highest number of triangles a single edge is part of.
Avg. local clustering coefficient	[0, 1]	A measure of the extent to which nodes in a neighborhood tend to form clusters or cliques.
Global clustering coefficient	[0, 1]	A measure of the overall tendency of nodes to form clusters in the entire graph, reflecting the global structure of the network.
Max. k -core	[0, 99]	Largest subgraph where each node has at least k neighbors within the subgraph.
# communities	[0, 50]	Number of clusters of nodes (communities) in the graph.
Diameter	[0, 99]	Maximum distance between any pair of nodes in the graph.

The 15 considered graph properties and their description are listed in Table 4.