
HiGen: Hierarchical Graph Generative Networks

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Most real-world graphs exhibit a hierarchical structure, which is often overlooked
2 by existing graph generation methods. To address this limitation, we propose a
3 novel graph generative network that captures the hierarchical nature of graphs
4 and successively generates the graph sub-structures in a coarse-to-fine fashion. At
5 each level of hierarchy, this model generates communities in parallel, followed by
6 the prediction of cross-edges between communities using a separate model. This
7 modular approach results in a highly scalable graph generative network. More-
8 over, we model the output distribution of edges in the hierarchical graph with
9 a multinomial distribution and derive a recursive factorization for this distribu-
10 tion, enabling us to generate sub-graphs with integer-valued edge weights in an
11 autoregressive approach. Empirical studies demonstrate that the proposed gen-
12 erative model can effectively capture both local and global properties of graphs
13 and achieves state-of-the-art performance in terms of graph quality on various
14 benchmarks.

15 1 Introduction

16 Graphs play a fundamental role in representing relationships and are widely applicable in various
17 domains. The task of generating graphs from data holds immense value for diverse applications but
18 also poses significant challenges (Dai et al., 2020). Some of the applications include: the exploration
19 of novel molecular and chemical structures (Jin et al., 2020), document generation (Blei et al., 2003),
20 circuit design (Mirhoseini et al., 2021), the analysis and synthesis of realistic data networks, as well
21 as the synthesis of scene graphs in computer (Manolis Savva et al., 2019) (Ramakrishnan et al., 2021).

22 In all the aforementioned domains, a common observation is the presence of locally heterogeneous
23 edge distributions in the graph representing the system, leading to the formation of clusters or
24 communities and hierarchical structures. These clusters represent groups of nodes characterized by
25 a high density of edges within the group and a comparatively lower density of edges connecting
26 the group with the rest of the graph. In a hierarchical structure that arise from graph clustering, the
27 communities in the lower levels capture the local structures and relationships within the graph. These
28 communities provide insights into the fine-grained interactions among nodes. On the other hand, the
29 higher levels of the hierarchy reflect the broader interactions between communities and characterize
30 global properties of the graph. Therefore, in order to generate realistic graphs, it is essential for graph
31 generation models to learn this multi-scale structure, and be able to capture the cross-level relations.
32 While hierarchical multi-resolution generative models were developed for specific data types such as
33 voice (Oord et al., 2016), image (Reed et al., 2017; Karami et al., 2019) and molecular motifs (Jin
34 et al., 2020), these methods rely on domain-specific priors that are not applicable to general graphs
35 with unordered nature. To the best of our knowledge, there exists no data-driven generative models
36 specifically designed for generic graphs that can effectively incorporate hierarchical structure.

37 Graph generative models have been extensively studied in the literature. Classical methods based on
38 random graph theory, such as those proposed in (Erdos & Rényi (1960) and (Barabási & Albert (1999)),

39 can only capture a limited set of hand-engineered graph statistics. Leskovec et al. (2010) leveraged
 40 the Kronecker product of matrices but the resulting generative model is very limited in modeling
 41 the underlying graph distributions. With recent advances in graph neural networks, a variety of
 42 deep neural network models have been introduced that are based on variational autoencoders (VAE)
 43 (Kingma & Welling, 2013) or generative adversarial networks (GAN) (Goodfellow et al., 2020).
 44 Some examples of such models include (De Cao & Kipf, 2018; Simonovsky & Komodakis, 2018;
 45 Kipf & Welling, 2016; Ma et al., 2018; Liu et al., 2019; Bojchevski et al., 2018; Yang et al., 2019).
 46 The major challenge in VAE based models is that they rely on heuristics to solve a graph matching
 47 problem for aligning the VAE’s input and sampled output, limiting them to small graphs. On the other
 48 hand, GAN-based methods circumvent the need for graph matching by using a permutation invariant
 49 discriminator. However, they can still suffer from convergence issues and have difficulty capturing
 50 complex dependencies in graph structures for moderate to large graphs (Li et al., 2018); (Martinkus
 51 et al., 2022). To address these limitations, (Martinkus et al., 2022) recently proposed using spectral
 52 conditioning to enhance the expressivity of GAN models in capturing global graph properties.

53 On the other hand, autoregressive models approach graph generation as a sequential decision-making
 54 process. Following this paradigm, (Li et al., 2018) proposed generative model based on GNN but it
 55 has high complexity of $\mathcal{O}(mn^2)$. In a distinct approach, GraphRNN (You et al., 2018) modeled graph
 56 generation with a two-stage RNN architecture for generating new nodes and their links, respectively.
 57 However, traversing all elements of the adjacency matrix in a predefined order results in $\mathcal{O}(n^2)$
 58 time complexity making it non-scalable to large graphs. In contrast, GRAN (Liao et al., 2019)
 59 employs a graph attention network and generates the adjacency matrix row by row, resulting in a
 60 $\mathcal{O}(n)$ complexity sequential generation process. To improve the scalability of generative models,
 61 (Dai et al., 2020) proposed an algorithm for sparse graphs that decreases the training complexity
 62 to $\mathcal{O}(\log n)$, but at the expense of increasing the generation time complexity to $\mathcal{O}((n + m) \log n)$.
 63 Despite their improvement in capturing complex statistics of the graphs, autoregressive models highly
 64 rely on an appropriate node ordering and do not take into account the community structures of the
 65 graphs. Additionally, due to their recursive nature, they are not fully parallelizable.

66 A new family of diffusion model for graphs has emerged recently. Continuous denoising diffusion
 67 was developed by (Jo et al., 2022), which adds Gaussian noise to the graph adjacency matrix and
 68 node features during the diffusion process. However, since continuous noise destroys the sparsity
 69 and structural properties of the graph, discrete denoising diffusion models have been developed as
 70 a solution in (Haefeli et al., 2022; Vignac et al., 2022). These models progressively edit graphs by
 71 adding or removing edges in the diffusion process, and then denoising graph neural networks are
 72 trained to reverse the diffusion process. While the denoising diffusion models can offer promising
 73 results, their main drawback is the requirement of a long chain of reverse diffusion, which can result
 74 in relatively slow sampling.

75 In his work, we introduce HiGen, a **H**ierarchical **G**raph **G**enerative Network to address the limitations
 76 of existing generative models by incorporating community structures and cross-level interactions.
 77 This approach involves generating graphs in a coarse-to-fine manner, where graph generation at each
 78 level is conditioned on a higher level (lower resolution) graph. The generation of communities at lower
 79 levels is performed in parallel, followed by the prediction of cross-edges between communities using a
 80 separate model. This parallelized approach enables high scalability. To capture hierarchical relations,
 81 our model allows each node at a given level to depend not only on its neighbouring nodes but also on
 82 its corresponding super-node at the higher level. Furthermore, we address the generation of integer-
 83 valued edge weights of the hierarchical structure by modeling the output distribution of edges using
 84 a multinomial distribution. We show that multinomial distribution can be factorized successively,
 85 enabling the autoregressive generation of each community. This property makes the proposed
 86 architecture well-suited for generating graphs with integer-valued edge weights. Furthermore, by
 87 breaking down the graph generation process into the generation of multiple small partitions that are
 88 conditionally independent of each other, HiGen reduces its sensitivity to a predefined initial ordering
 89 of nodes.

90 2 Background

91 A graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a collection of nodes (vertices) \mathcal{V} and edges \mathcal{E} with corresponding sizes
 92 $n = |\mathcal{V}|$ and $m = |\mathcal{E}|$ and an adjacency matrix \mathbf{A}^π for the node ordering π . The node set can
 93 be partitioned into c communities (a.k.a. cluster or modules) using a graph partitioning function

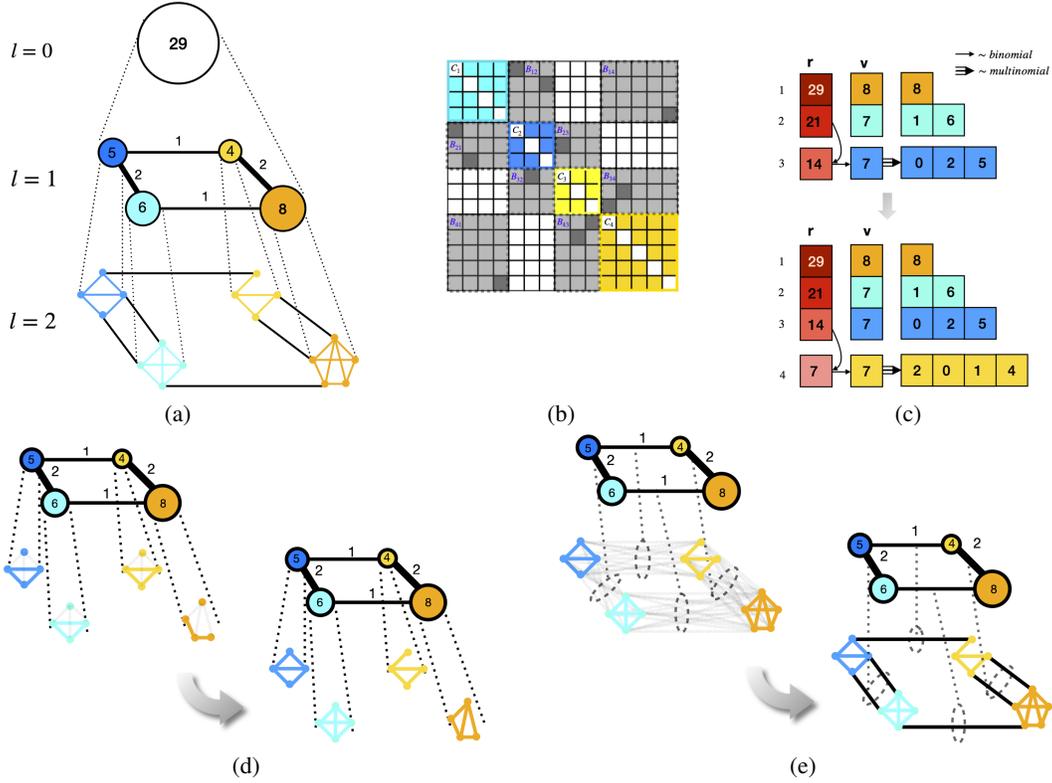


Figure 1: (a) A sample hierarchical graph with 2 levels is shown. Communities are shown in different colors and the weight of a node and the weight of an edge in a higher level, represent the sum of the edges in the corresponding community and bipartite, respectively. Node size and edge width indicate their weights. (b) The matrix shows corresponding adjacency of the graph \mathcal{G}^2 matrix where each of its sub-graphs corresponds to a block in the adjacency matrix, communities are shown in different colors and bipartites are colored in gray. (c) Decomposition of multinomial distribution as a recursive *stick-breaking* process where at each iteration, first a fraction of the remaining weights w_m is allocated to the m -th node in the sub-graph and then this fraction v_m is distributed among that row of lower triangular adjacency matrix, \hat{A} . (d) Parallel generation of communities. (e) Parallel prediction of bipartites. Shaded lines are the *augmented edges* representing candidate edges at each step.

94 $\mathcal{F} : \mathcal{V} \rightarrow \{1, \dots, c\}$, where each cluster of nodes forms a sub-graph denoted by $\mathcal{C}_i = (\mathcal{V}(\mathcal{C}_i), \mathcal{E}(\mathcal{C}_i))$
 95 with adjacency matrix \mathbf{A}_i . The cross-links between neighboring communities form a *bipartite*
 96 *graph*, denoted by $\mathcal{B}_{ij} = (\mathcal{V}(\mathcal{C}_i), \mathcal{V}(\mathcal{C}_j), \mathcal{E}(\mathcal{B}_{ij}))$ with adjacency matrix \mathbf{A}_{ij} . Each community
 97 is aggregated to a super-node and each bipartite corresponds to a super-edge linking neighboring
 98 communities, which induces a coarser graph at the higher (a.k.a. parent) level. Herein, the levels are
 99 indexed by superscripts. Formally, each community at level l , \mathcal{C}_i^l , is mapped to a node at the higher
 100 level graph, also called its parent node, $v_i^{l-1} := Pa(\mathcal{C}_i^l)$ and each bipartite at level l is represented by
 101 an edge in the higher level, also called its parent edge, $e_{ij}^{l-1} = Pa(\mathcal{B}_{ij}^l) = (v_i^{l-1}, v_j^{l-1})$. The weights
 102 of the self edges and the weights of the cross-edges in the parent level are determined by the sum of
 103 the weights of the edges within their corresponding community and bipartite, respectively. Therefore, the
 104 edges in the induced graphs at the higher levels have integer-valued weights: $w_{ii}^{l-1} = \sum_{e \in \mathcal{E}(\mathcal{C}_i^l)} w_e$
 105 and $w_{ij}^{l-1} = \sum_{e \in \mathcal{E}(\mathcal{B}_{ij}^l)} w_e$, moreover sum of all edge weights remains constant in all levels so
 106 $w_0 := \sum_{e \in \mathcal{E}(\mathcal{G}^l)} w_e = |\mathcal{E}|, \forall l \in [0, \dots, L]$.

107 This clustering process continues recursively in a bottom-up approach until a single node graph \mathcal{G}^0 is
 108 obtained, producing a *hierarchical graph*, defined by the set of graphs in all levels of abstractions,
 109 $\mathcal{HG} := \{\mathcal{G}^0, \dots, \mathcal{G}^{L-1}, \mathcal{G}^L\}$. This forms a dendrogram tree with \mathcal{G}^0 being the root and \mathcal{G}^L being the
 110 final graph that is generated at the leaf level. An \mathcal{HG} is visualized in figure [1a](#). The hierarchical tree

111 structure enables modeling of both local and long-range interactions among nodes, as well as control
 112 over the flow of information between them, across multiple levels of abstraction. This is a key aspect
 113 of our proposed generative model.

114 3 Hierarchical Graph Generation

115 In graph generative networks, the objective is to learn a generative model, $p(\mathcal{G})$ given a set of training
 116 graphs. This work aims to establish a hierarchical multi-resolution framework for generating graphs in
 117 a coarse-to-fine fashion. In this framework, we assume that the graphs do not have node attributes, so
 118 the generative model only needs to characterize the graph topology. Given a particular node ordering
 119 π , and a hierarchical graph $\mathcal{HG} := \{\mathcal{G}^0, \dots, \mathcal{G}^{L-1}, \mathcal{G}^L\}$, produced by recursively applying a graph
 120 partitioning function, \mathcal{F} , we can factorize the generative model using the chain rule of probability as:

$$\begin{aligned} p(\mathcal{G} = \mathcal{G}^L, \pi) &= p(\{\mathcal{G}^L, \mathcal{G}^{L-1}, \dots, \mathcal{G}^0\}, \pi) = p(\mathcal{G}^L, \pi \mid \{\mathcal{G}^{L-1}, \dots, \mathcal{G}^0\}) \dots p(\mathcal{G}^1, \pi \mid \mathcal{G}^0) p(\mathcal{G}^0) \\ &= \prod_{l=0}^L p(\mathcal{G}^l, \pi \mid \mathcal{G}^{l-1}) \times p(\mathcal{G}^0) \end{aligned} \quad (1)$$

121 In other words, the generative process involves specifying the probability of the graph at each level
 122 conditioned on its parent level graph in the hierarchy. This process is iterated recursively until the
 123 lowest level, or leaf level, is reached. Here, the distribution of the root $p(\mathcal{G}^0) = p(\mathbf{w}^0 = w_0)$ can be
 124 simply estimated using the empirical distribution of the number of edges $|\mathcal{E}|$ of graphs in the training
 125 set.

126 Based on the partitioned structure within each level of \mathcal{HG} , the conditional generative probability
 127 $p(\mathcal{G}^l \mid \mathcal{G}^{l-1})$ can be decomposed into the probability of its communities and bipartite graphs as:

$$\begin{aligned} p(\mathcal{G}^l \mid \mathcal{G}^{l-1}) &= p(\{C_i^l \mid \forall i \in \mathcal{V}(\mathcal{G}^{l-1})\} \cup \{\mathcal{B}_{ij}^l \mid \forall (i, j) \in \mathcal{E}(\mathcal{G}^{l-1})\} \mid \mathcal{G}^{l-1}) \\ &\cong \prod_{i \in \mathcal{V}(\mathcal{G}^{l-1})} p(C_i^l \mid \mathcal{G}^{l-1}) \times \prod_{(i, j) \in \mathcal{E}(\mathcal{G}^{l-1})} p(\mathcal{B}_{ij}^l \mid \mathcal{G}^{l-1}) \end{aligned} \quad (2)$$

128 The approximation in this decomposition becomes an equivalence when each community C_i^l or
 129 bipartite graph \mathcal{B}_{ij}^l is assumed to be independent of all other components in its level conditioned on
 130 the parent graph \mathcal{G}^{l-1} . \square Since the integer-valued weights of the edges in each level can be modeled
 131 by a multinomial distribution, we can leverage the properties of multinomial distribution to prove the
 132 conditional independence of the components.

133 **Theorem 3.1.** *Let the random vector $\mathbf{w} := [w_e]_{e \in \mathcal{E}(\mathcal{G}^l)}$ denote the set of weights of all edges of \mathcal{G}^l
 134 such that their sum is $w_0 = \mathbf{1}^T \mathbf{w}$. The joint probability of \mathbf{w} can be described by a multinomial
 135 distribution: $\mathbf{w} \sim \text{Mu}(\mathbf{w} \mid w_0, \boldsymbol{\theta}^l)$. By observing that the sum of edge weights within each community
 136 C_i^l and bipartite graph \mathcal{B}_{ij}^l are determined by the weights of their parent edges in the higher level,
 137 w_{ii}^{l-1} and w_{ij}^{l-1} respectively, we can establish that these components are conditionally independent
 138 and each of them follow a multinomial distribution:*

$$p(\mathcal{G}^l \mid \mathcal{G}^{l-1}) \sim \prod_{i \in \mathcal{V}(\mathcal{G}^{l-1})} \text{Mu}([w_e]_{e \in C_i^l} \mid w_{ii}^{l-1}, \boldsymbol{\theta}_{ii}^l) \times \prod_{(i, j) \in \mathcal{E}(\mathcal{G}^{l-1})} \text{Mu}([w_e]_{e \in \mathcal{B}_{ij}^l} \mid w_{ij}^{l-1}, \boldsymbol{\theta}_{ij}^l) \quad (3)$$

139 where $\{\boldsymbol{\theta}_{ij}^l[e] \in [0, 1], \text{ s.t. } \mathbf{1}^T \boldsymbol{\theta}_{ij}^l = 1 \mid \forall (i, j) \in \mathcal{E}(\mathcal{G}^{l-1})\}$ are the multinomial model parameters.

140 *Proof.* The detailed proof can be found in Appendix [A.1](#) \square

141 Therefore, given the parent graph at a higher level, the generation of graph at its subsequent level
 142 can be reduced to generation of its partition and bipartite sub-graphs. As illustrated in figure, this
 143 decomposition enables parallel generation of the communities in each level which can be followed by
 144 predicting all bipartite sub-graphs in that level at one pass. Each of these sub-graphs corresponds
 145 to a block in the adjacency matrix, as visualized in figure [1b](#) so the proposed hierarchical model
 146 generates adjacency matrix in a blocks-wise fashion and constructs the final graph topology.

¹Indeed, this assumption implies that the cross dependency between communities are primarily encoded by their parent abstract graph which is reasonable where the nodes' dependencies are mostly local and are within community rather than being global.

147 **3.1 Community Generation**

148 Based on the equation (3), the edge weights within each community can be jointly modeled using
 149 a multinomial distribution. Our objective is to model the generative probability of communities in
 150 each level as an autoregressive process. To accomplish this, we need to factorize the multinomial
 151 distribution accordingly. Toward this goal, we present two different approaches in the following.

152 **Lemma 3.2.** *A random counting vector $\mathbf{w} \in \mathbb{Z}_+^E$ with a multinomial distribution can be recursively
 153 decomposed into a sequence of binomial distributions as follows:*

$$Mu(\mathbf{w}_1, \dots, \mathbf{w}_E \mid w, [\theta_1, \dots, \theta_E]) = \prod_{e=1}^E Bi(\mathbf{w}_e \mid w - \sum_{i<e} \mathbf{w}_i, \hat{\theta}_e), \quad (4)$$

$$\text{where: } \hat{\theta}_e = \frac{\theta_e}{1 - \sum_{i<e} \theta_i}$$

154 *This decomposition is known as a stick-breaking process, where $\hat{\theta}_e$ is the fraction of the remaining
 155 probabilities we take away every time and allocate to the e -th component (Linderman et al. 2015).*

156 This lemma enable us to model the generation of a community as an edge-by-edge autoregressive
 157 process, similar to existing algorithms such as GraphRNN (You et al. 2018) or DeepGMG (Li et al.
 158 2018) with $\mathcal{O}(|\mathcal{V}_C|^2)$ generation steps. However, inspired by GRAN (Liao et al. 2019), a community
 159 can be generated more efficiently by generating one node at a time. This requires decomposing the
 160 generative probability of edges in a group-wise form, where the candidate edges between the t -th
 161 node and the already generated graph are grouped together. In other words, this model completes the
 162 lower triangle adjacency matrix one row at a time conditioned on the already generated sub-graph and
 163 the parent-level graph. The following theorem formally derives this decomposition for multinomial
 164 distributions.

165 **Theorem 3.3.** *For a random counting vector $\mathbf{w} \in \mathbb{Z}_+^E$ with a multinomial distribution $Mu(\mathbf{w} \mid w, \theta)$,
 166 let's split it into M disjoint groups $\mathbf{w} = [\mathbf{u}_1, \dots, \mathbf{u}_M]$ where $\mathbf{u}_m \in \mathbb{Z}_+^{E_m}$, $\sum_{m=1}^M E_m = E$, and
 167 also split the probability vector accordingly as $\theta = [\theta_1, \dots, \theta_M]$. Additionally, let's define sum of all
 168 variables in the m -th group by a random count variable $v_m := \sum_{e=1}^{E_m} \mathbf{u}_{m,e}$. Then, the multinomial
 169 distribution can be factorized as a chain of binomial and multinomial distributions:*

$$Mu(\mathbf{w} = [\mathbf{u}_1, \dots, \mathbf{u}_M] \mid w, \theta = [\theta_1, \dots, \theta_M]) = \prod_{m=1}^M Bi(v_m \mid w - \sum_{i<m} v_i, \eta_{v_m}) Mu(\mathbf{u}_m \mid v_m, \lambda_m),$$

$$\text{where: } \eta_{v_m} = \frac{\mathbf{1}^T \theta_m}{1 - \sum_{i<m} \mathbf{1}^T \theta_i}, \quad \lambda_m = \frac{\theta_m}{\mathbf{1}^T \theta_m}. \quad (5)$$

170 *Here, the probability of binomial, η_{v_m} , is the fraction of the remaining probability mass that is
 171 allocated to v_m , i.e. the sum of all weights in the m -th group. The vector parameter λ_m is the
 172 normalized multinomial probabilities of all count variables in the m -th group. Intuitively, this
 173 decomposition of multinomial distribution can be viewed as a recursive stick-breaking process where
 174 at each step, first a binomial distribution is used to determine how much probability mass to allocate
 175 to the current group, and a multinomial distribution is used to distribute that probability mass
 176 among the variables in the group. The resulting distribution is equivalent to the original multinomial
 177 distribution.*

178 *Proof.* Refer to appendix A.2 for the proof. □

179 Let $\hat{\mathcal{C}}_{i,t}^l$ denote an already generated sub-graph, at the t -th step, augmented with the set of candidate
 180 edges, from the new node, $v_t(\mathcal{C}_i^l)$, to its preceding node denoted by $\hat{\mathcal{E}}_t(\hat{\mathcal{C}}_{i,t}^l) := \{(t, j) \mid j < t\}$.
 181 We collect the weights of these edges in the random vector $\mathbf{u}_t := [w_e]_{e \in \hat{\mathcal{E}}_t(\hat{\mathcal{C}}_{i,t}^l)}$ (that is the t -th
 182 row of the lower triangle of adjacency matrix $\hat{\mathbf{A}}_i^l$), where the sum of the candidate edge weights is
 183 v_t . Based on theorem 3.3 the probability of \mathbf{u}_t can be characterized by the product of a binomial
 184 and a multinomial distribution. This process is illustrated in figure 1c. We further increase the

185 expressiveness of the generative network by extending this probability to a mixture model with K
 186 mixtures:

$$p(\mathbf{u}_t) = \sum_{k=1}^K \beta_k^l \text{Bi}(v_t | w_{ii}^{l-1} - \sum_{i < t} v_i, \eta_{t,k}^l) \text{Mu}(\mathbf{u}_t | v_t, \boldsymbol{\lambda}_{t,k}^l) \quad (6)$$

$$\boldsymbol{\lambda}_{t,k}^l = \text{softmax} \left(\text{MLP}_{\boldsymbol{\theta}}^l \left(\left[\Delta \mathbf{h}_{\hat{\mathcal{E}}_t(\hat{\mathcal{C}}_{i,t}^l)} \parallel h_{Pa(\mathcal{C}_i^l)} \right] \right) \right) [k, :] \quad (7)$$

$$\eta_{t,k}^l = \text{sigmoid} \left(\text{MLP}_{\eta}^l \left(\left[\text{pool}(\mathbf{h}_{\hat{\mathcal{C}}_{i,t}^l}) \parallel h_{Pa(\mathcal{C}_i^l)} \right] \right) \right) [k]$$

$$\beta^l = \text{softmax} \left(\text{MLP}_{\beta}^l \left(\left[\text{pool}(\mathbf{h}_{\hat{\mathcal{C}}_{i,t}^l}) \parallel h_{Pa(\mathcal{C}_i^l)} \right] \right) \right)$$

187 Where $\Delta \mathbf{h}_{\hat{\mathcal{E}}_t(\hat{\mathcal{C}}_{i,t}^l)}$ is a $|\hat{\mathcal{E}}_t(\hat{\mathcal{C}}_{i,t}^l)| \times d_h$ dimensional matrix, consisting of the set of edge features
 188 $\{\Delta h_{(t,s)} := h_t - h_s \mid \forall (t,s) \in \hat{\mathcal{E}}_t(\hat{\mathcal{C}}_{i,t}^l)\}$, $\mathbf{h}_{\hat{\mathcal{C}}_{i,t}^l}$ is a $t \times d_h$ matrix of node features in the
 189 augmented community graph. The mixture weights are denoted by β^l . Here, the node features are
 190 learned by GNN models and the graph level representation is obtained by the `addpool()` aggregation
 191 function. In order to produce $K \times |\mathcal{E}_t(\mathcal{C}_i^l)|$ dimensional matrix of multinomial probabilities, the
 192 $\text{MLP}_{\boldsymbol{\theta}}^l()$ network acts at the edge level, while $\text{MLP}_{\eta}^l()$ and $\text{MLP}_{\beta}^l()$ act at the graph level to produce
 193 the binomial probabilities and K dimensional arrays for K mixture models, respectively. All of these
 194 MLP networks are built by two hidden layers with `ReLU()` activation functions.

195 During the generation process of each community \mathcal{C}_i^l , the node features of its parent node $h_{Pa(\mathcal{C}_i^l)}$
 196 are used as the context. This context is concatenated to the node and edge feature matrices using the
 197 operation $[x \parallel y]$, which concatenates vector y to each row of matrix x . The purpose of this context
 198 is to enrich the node and edge features by capturing long-range interactions and encoding the global
 199 structure of the graph, which is important for generating local components.

200 3.2 Bipartite Generation

201 Once all the communities in level l are generated, the edges of all bipartite graphs at that level can
 202 be predicted simultaneously. An augmented graph $\hat{\mathcal{G}}^l$ composed of all the communities, $\{\mathcal{C}_i^l \mid \forall i \in$
 203 $\mathcal{V}(\mathcal{G}^{l-1})\}$, and the candidate edges of all bipartites, $\{\mathcal{B}_{ij}^l \mid \forall (i,j) \in \mathcal{E}(\mathcal{G}^{l-1})\}$, is used as the input
 204 of a GNN to obtain node and edge features. We similarly extend the multinomial distribution of a
 205 bipartite, (12), using a mixture model to express its generative probability:

$$p(\mathbf{w} := \hat{\mathcal{E}}(\mathcal{B}_{ij}^l)) = \sum_{k=1}^K \beta_k^l \text{Mu}(\mathbf{w} \mid w_{ij}^{l-1}, \boldsymbol{\theta}_{ij,k}^l)$$

$$\boldsymbol{\theta}_{ij,k}^l = \text{softmax} \left(\text{MLP}_{\boldsymbol{\theta}}^l \left(\left[\Delta \mathbf{h}_{\hat{\mathcal{E}}(\mathcal{B}_{ij}^l)} \parallel \Delta h_{Pa(\mathcal{B}_{ij}^l)} \right] \right) \right) [k, :] \quad (8)$$

$$\beta^l = \text{softmax} \left(\text{MLP}_{\beta}^l \left(\left[\text{pool}(\Delta \mathbf{h}_{\hat{\mathcal{E}}(\mathcal{B}_{ij}^l)}) \parallel \Delta h_{Pa(\mathcal{B}_{ij}^l)} \right] \right) \right)$$

206 where the random vector $\mathbf{w} := [w_e]_{e \in \hat{\mathcal{E}}(\mathcal{B}_{ij}^l)}$ is the set of weights of all candidate edges in bipartite
 207 \mathcal{B}_{ij}^l and $\Delta \mathbf{h}_{Pa(\mathcal{B}_{ij}^l)}$ are the parent edge features of the bipartite graph.

208 **Node Feature Encoding:** To encode node features, we extend GraphGPS proposed by Rampásek
 209 et al. (2022). GraphGPS combines local message-passing with global attention mechanism and uses
 210 positional and structural encoding for nodes and edges to construct a more expressive and a scalable
 211 graph transformer (GT) (Dwivedi & Bresson, 2020). To apply GraphGPS on augmented graphs,
 212 we use distinct initial edge features to distinguish augmented (candidate) edges from real edges.
 213 Furthermore, for bipartite generation, the attention scores in the Transformers of the augmented graph
 214 $\hat{\mathcal{G}}^l$ are masked to restrict attention only to connected communities. The details of model architecture
 215 are provided in appendix B

216 4 Related Work

217 In order to deal with hierarchical structures in molecular graphs, a generative process was proposed
218 by Jin et al. (2020) which recursively selects motifs, the basic building blocks, from a set and
219 predicts their attachment to the emerging molecule. However, this method requires prior domain-
220 specific knowledge and relies on molecule-specific graph motifs. Additionally, the graphs are
221 only abstracted into two levels, and component generation cannot be performed in parallel. In
222 (Kuznetsov & Polykovskiy 2021), a hierarchical normalizing flow model for molecular graphs was
223 introduced, where new molecules are generated from a single node by recursively dividing each
224 node into two. However, the merging and splitting of pairs of nodes in this model is based on the
225 node’s neighborhood, and do not consider the diverse community structure of graphs, therefore the
226 hierarchical generation of this model is inherently limited.

227 5 Experiments

228 In our empirical studies, we compare the proposed hierarchical graph generative network against
229 state-of-the-art autoregressive models: GRAN and GraphRNN models, diffusion models: DiGress
230 (Vignac et al. 2022) and GDSS (Jo et al. 2022) and a GAN-based model: SPECTRE (Martinkus
231 et al. 2022), on a range of synthetics and real datasets of various sizes.

232 **Datasets:** We used 4 different benchmark graph datasets: (1) the synthetic *Stochastic Block Model*
233 (*SBM*) dataset consisting of 200 graphs with 2-5 communities each with 20-40 nodes, used in a
234 previous work (Martinkus et al. 2022); (2) the *Protein* including 918 protein graphs, each has 100 to
235 500 nodes representing amino acids that are linked if they are closer than 6 Angstroms (Dobson &
236 Doig, 2003), (3) the *Enzyme* that has 587 protein graphs of 10-125 nodes, representing protein tertiary
237 structures of the enzymes from the BRENDA database (Schomburg et al. 2004) and (4) the *Ego*
238 dataset containing 757 3-hop ego networks with 50-300 nodes extracted from the CiteSeer dataset,
239 where nodes represent documents and edges represent citation relationships (Sen et al., 2008).

240 **Graph Partitioning** Different algorithms approach the problem of graph partitioning (clustering)
241 using various clustering quality functions. Two commonly used families of such metrics are modu-
242 larity and cut-based metrics (Tsitsulin et al., 2020). Although optimizing modularity metric is an
243 NP-hard problem, it is well-studied in the literature and several graph partitioning algorithm based on
244 this metric have been proposed. For example, the Louvain algorithm (Blondel et al., 2008) starts with
245 each node as its community and then repeatedly merges communities based on the highest increase
246 in modularity until no further improvement can be made. This heuristic algorithm is computationally
247 efficient and scalable to large graphs for community detection. Moreover, a spectral relaxation of
248 modularity metrics has been proposed in Newman (2006a b) which results in an analytically solution
249 for graph partitioning. Additionally, an unsupervised GNN-based pooling method inspired by this
250 spectral relaxation was proposed for partitioning graphs with node attributes (Tsitsulin et al., 2020).
251 As the modularity metric is based on the graph structure, it is well-suited for our problem. Therefore,
252 we employed the Louvain algorithm to hierarchically cluster the graph datasets in our experiments
253 and then spliced out the intermediate levels to achieve HGs with uniform depth of $L = 2$.

254 **Model Architecture** In the experiments, the GNN models consist of 8 layers of GraphGPS layers
255 (Rampášek et al., 2022). The input node feature of GNNs is augmented with positional and structural
256 encoding, where the first 8 eigenvectors corresponding to the smallest non-zero eigenvalues of the
257 Laplacian and diagonal of the random-walk matrix up to 8-steps are used. Each level has its own
258 GNN and output models. The details of the model architecture are presented in Appendix B and C.

259 We conducted experiments using the proposed hierarchical graph generative network (HiGen) model
260 with two variants for the output distribution of the leaf edges: 1) **HiGen**: the probability of the
261 community edges’ weights at the leaf level are modeled by mixture of Bernoulli, using sigmoid()
262 activation in equation 7 since the leaf levels in our experiments have binary edges weights, while
263 higher levels use mixture of multinomials. 2) **HiGen-m**: the model uses a mixture of multinomial
264 distributions (6) to describe the output distribution for all levels. In this case, we observed that
265 modeling the probability parameters of edge weights of the leaf level, denoted as $\lambda_{t,k}$ in (7), by a multi-
266 hot activation function, defined as $\sigma(\mathbf{z})_i := \text{sigmoid}(z_i) / \sum_{j=1}^K \text{sigmoid}(z_j)$ where $\sigma : \mathbb{R}^K \rightarrow (K - 1)$ -
267 simplex, provided slightly better performance than the standard softmax() function. However, for

Table 1: Comparison of generation metrics on benchmark datasets. The baseline results for SBM and Protein graphs are obtained from (Martinkus et al. 2022, Vignac et al. 2022), and the results for enzyme graphs (except for GRAN, which we implemented) are obtained from (Jo et al. 2022), while we implemented them for Ego. "-": not applicable due to resource issue or not reported in the reference papers.

Model	Stochastic block model				Protein			
	Deg. ↓	Clus. ↓	Orbit ↓	Spec. ↓	Deg. ↓	Clus. ↓	Orbit ↓	Spec. ↓
Training set	0.0008	0.0332	0.0255	0.0063	0.0003	0.0068	0.0032	0.0009
GraphRNN	0.0055	0.0584	0.0785	0.0065	0.0040	0.1475	0.5851	0.0152
GRAN	0.0113	0.0553	0.0540	0.0054	0.0479	0.1234	0.3458	0.0125
SPECTRE	0.0015	0.0521	0.0412	0.0056	0.0056	0.0843	0.0267	0.0052
DiGress	0.0013	0.0498	0.0433	-	-	-	-	-
HiGen-m	0.0017	0.0503	0.0604	0.0068	0.0041	0.109	0.0472	0.0061
HiGen	0.0019	0.0498	0.0352	0.0046	0.0012	0.0435	0.0234	0.0025

Model	Enzyme			Ego			
	Deg. ↓	Clus. ↓	Orbit ↓	Deg. ↓	Clus. ↓	Orbit ↓	Spec. ↓
Training set	0.0011	0.0025	3.7e-4	2.2e-4	0.010	0.012	1.4e-3
GraphRNN	0.017	0.062	0.046	0.024	0.34	0.14	0.089
GRAN	0.054	0.087	0.033	0.032	0.17	0.026	0.046
GDSS	0.026	0.061	0.009	-	-	-	-
HiGen-m	0.027	0.157	1.2e-3	0.011	0.063	0.021	0.013
HiGen	0.012	0.038	7.2e-4	1.9e-3	0.049	0.029	0.004

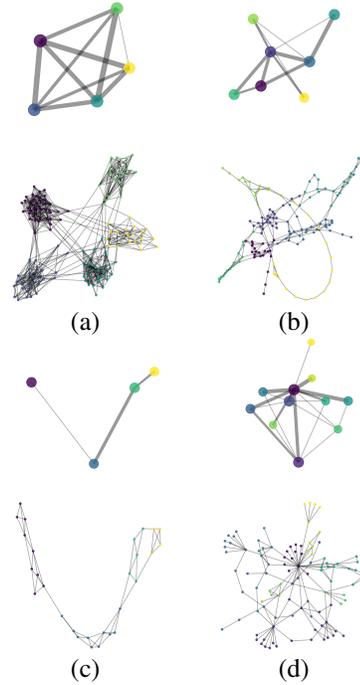


Figure 2: Samples from HiGen. a) SBM, b) Protein, c) Enzyme and d) Ego. Communities are distinguished with different colors and both levels are depicted.

268 both HiGen and HiGen-m, the probabilities of the integer-valued edges at the higher levels are still
 269 modeled by the standard softmax() function²

270 For training, HiGen models used the Adam optimizer (Kingma & Ba (2014)) with a learning rate of
 271 $5e-4$ and its default settings of $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon=1e-8$.

272 **Metrics** To evaluate the graph generative models, we adopt the approach proposed in (Liu et al.
 273 2019, Liao et al. 2019), which compares the distributions of four different graph statistics between
 274 the ground truth and generated graphs: (1) degree distributions, (2) clustering coefficient distributions,
 275 (3) the number of occurrences of all orbits with four nodes, and (4) the spectra of the graphs by
 276 computing the eigenvalues of the normalized graph Laplacian. The first three metrics capture local
 277 graph statistics, while the spectra represents global structure. The maximum mean discrepancy
 278 (MMD) score over these statistics are used as the metrics. While (Liu et al. 2019) computed MMD
 279 scores using the computationally expensive Gaussian earth mover’s distance (EMD) kernel, (Liao
 280 et al. 2019) proposed using the total variation (TV) distance as an alternative measure. TV distance
 281 is much faster and still consistent with the Gaussian EMD kernel. Most recently, (O’Bray et al. 2021)
 282 suggested using other efficient kernels such as an RBF kernel, or a Laplacian kernel, or a linear
 283 kernel. Additionally, (Thompson et al. 2022) proposed new evaluation metrics for comparing graph
 284 sets using a random-GNN approach where GNNs are employed to extract meaningful graph features.
 285 However, in this work, we follow the experimental setup and evaluation metrics of (Liao et al. 2019),
 286 except for the enzyme dataset where we use a Gaussian EMD kernel to be consistent with the results
 287 reported in (Jo et al. 2022). GNN-based performance metrics of HiGen model are also reported in
 288 appendix D.2

²As the leaf levels have binary edge weights while the sum of their weights is determined by their parent edge, a possible extension to this work could be using the cardinality potential model (Hajimirsadeghi et al. 2015), which is derived to model the distribution over the set of binary random variables, to model the edge weight at the leaf level.

289 The performance metrics of the proposed HiGen models are reported in Table 1 with generated graph
 290 samples presented in Figure 2. The results demonstrate that HiGen effectively captures graph statistics
 291 and achieves state-of-the-art on all the benchmarks graphs across various generation metrics. This
 292 improvement in both local and global properties of the generated graphs highlights the effectiveness
 293 of the hierarchical graph generation approach, which models communities and cross-community
 294 interactions separately. The visual comparisons of graph samples generated by the HiGen models,
 295 as well as the experimental evaluation of different node ordering and partitioning functions, are
 296 presented in Appendix D.2

297 6 Discussion

298 The GRAN model can generate graphs one block of nodes at a time in an autoregressive fashion where
 299 the block size is fixed and nodes are assigned to blocks based on an ordering. However, the model’s
 300 performance deteriorates as the block size increases, since adjacent nodes in an ordering may not be
 301 relevant and may belong to different clusters. Additionally, intra-block connections are not modeled
 302 separately. In contrast, our proposed method generates blocks of nodes within each community that
 303 have strong relationships and then predicts the cross-links between communities using a separate
 304 model. As a result, this approach enables the model to capture both local relationships between
 305 nodes within a community and global relationships across communities, resulting in improved
 306 expressiveness of the graph generative model.

307 The proposed hierarchical model allows for highly parallelizable training and generation. Specifically,
 308 let n_c be the size of the largest graph cluster, then, it only requires $\mathcal{O}(n_c \log n)$ sequential steps to
 309 generate a graph of size n .

310 **Node ordering sensitivity** The predefined ordering of dimensions can be crucial for training
 311 autoregressive (AR) models Vinyals et al. (2015), and this sensitivity to node orderings is particularly
 312 pronounced in autoregressive graph generative model Liao et al. (2019); Chen et al. (2021). However,
 313 in the proposed approach, the graph generation process is divided into the generation of multiple
 314 small partitions, performed sequentially across the levels, rather than generating the entire graph
 315 by a single AR model. Therefore, given an ordering for the parent level, the graph generation
 316 depends only on the permutation of the nodes within the graph communities rather than the node
 317 ordering of the entire graph. In other words, the proposed method is invariant to a large portion of
 318 possible node permutations, and therefore the set of distinctive adjacency matrices is much smaller
 319 in HiGen. For example, the node ordering $\pi_1 = [v_1, v_2, v_3, v_4]$ with clusters $\mathcal{V}_{G_1} = \{v_1, v_2\}$ and
 320 $\mathcal{V}_{G_2} = \{v_3, v_4\}$ has a similar hierarchical graph as $\pi_2 = [v_1, v_3, v_2, v_4]$, since the node ordering
 321 within the communities is preserved at all levels. Formally, let $\{\mathcal{C}_i^l \mid \forall i \in \mathcal{V}_{G^{l-1}}\}$ be the set of
 322 communities at level l produced by a deterministic partitioning function, where $n_i^l = |\mathcal{V}(\mathcal{C}_i^l)|$ denotes
 323 the size of each partition. The upper bound on the number of distinct node orderings in an HG
 324 generated by the proposed process is then reduced to $\prod_{l=1}^L \prod_i n_i^l!$.³

325 7 Conclusion

326 The proposed HiGen framework generates graphs in a hierarchical and block-wise manner, leveraging
 327 the inherent hierarchical structure present in real-world graphs. By decomposing the generation
 328 process into separate and parallel generation of communities and bipartite sub-graphs, it combines the
 329 benefits of one-shot and AR graph generative models. Experimental results on benchmark datasets
 330 demonstrate that HiGen achieves state-of-the-art performance across various generation metrics. The
 331 hierarchical and block-wise generation strategy of HiGen enables scaling up graph generative models
 332 to large and complex graphs, opening up opportunities to extend it to newer generative paradigms
 333 such as diffusion models.

³It is worth noting that all node permutations do not result in distinctive adjacency matrices due to the automorphism property of graphs Liao et al. (2019); Chen et al. (2021). Therefore, the number of node permutations provides an upper bound rather than an exact count.

334 References

- 335 Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):
336 509–512, 1999.
- 337 David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning*
338 *research*, 3(Jan):993–1022, 2003.
- 339 Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of
340 communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008,
341 2008.
- 342 Aleksandar Bojchevski, Oleksandr Shchur, Daniel Zügner, and Stephan Günnemann. Netgan: Generating graphs
343 via random walks. *arXiv preprint arXiv:1803.00816*, 2018.
- 344 Xiaohui Chen, Xu Han, Jiajing Hu, Francisco JR Ruiz, and Liping Liu. Order matters: Probabilistic modeling of
345 node sequence for graph generation. *arXiv preprint arXiv:2106.06189*, 2021.
- 346 Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos,
347 Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers.
348 *arXiv preprint arXiv:2009.14794*, 2020.
- 349 Hanjun Dai, Azade Nazi, Yujia Li, Bo Dai, and Dale Schuurmans. Scalable deep generative modeling for sparse
350 graphs. In *International Conference on Machine Learning*, pp. 2302–2312. PMLR, 2020.
- 351 Nicola De Cao and Thomas Kipf. Molgan: An implicit generative model for small molecular graphs. *arXiv*
352 *preprint arXiv:1805.11973*, 2018.
- 353 Paul D Dobson and Andrew J Doig. Distinguishing enzyme structures from non-enzymes without alignments.
354 *Journal of molecular biology*, 330(4):771–783, 2003.
- 355 Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. *arXiv preprint*
356 *arXiv:2012.09699*, 2020.
- 357 Paul Erdos and Alfréd Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60,
358 1960.
- 359 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron
360 Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):
361 139–144, 2020.
- 362 Kilian Konstantin Haefeli, Karolis Martinkus, Nathanaël Perraudin, and Roger Wattenhofer. Diffusion models
363 for graphs benefit from discrete state spaces. *arXiv preprint arXiv:2210.01549*, 2022.
- 364 Hossein Hajimirsadeghi, Wang Yan, Arash Vahdat, and Greg Mori. Visual recognition by counting instances: A
365 multi-instance cardinality potential kernel. In *Proceedings of the IEEE conference on computer vision and*
366 *pattern recognition*, pp. 2596–2605, 2015.
- 367 Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Hierarchical generation of molecular graphs using structural
368 motifs. In *International conference on machine learning*, pp. 4839–4848. PMLR, 2020.
- 369 Jaehyeong Jo, Seul Lee, and Sung Ju Hwang. Score-based generative modeling of graphs via the system of
370 stochastic differential equations. In *International Conference on Machine Learning*, pp. 10362–10383. PMLR,
371 2022.
- 372 Mahdi Karami, Dale Schuurmans, Jascha Sohl-Dickstein, Laurent Dinh, and Daniel Duckworth. Invertible
373 convolutional flow. *Advances in Neural Information Processing Systems*, 32, 2019.
- 374 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*,
375 2014.
- 376 Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- 377 Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv*
378 *preprint arXiv:1609.02907*, 2016.
- 379 Maksim Kuznetsov and Daniil Polykovskiy. Molgrow: A graph normalizing flow for hierarchical molecular
380 generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 8226–8234,
381 2021.

- 382 Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos, and Zoubin Ghahramani. Kronecker
383 graphs: An approach to modeling networks. *Journal of Machine Learning Research*, 11(Feb):985–1042,
384 2010.
- 385 Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning deep generative models of
386 graphs. *arXiv preprint arXiv:1803.03324*, 2018.
- 387 Renjie Liao, Yujia Li, Yang Song, Shenlong Wang, Will Hamilton, David K Duvenaud, Raquel Urtasun, and
388 Richard Zemel. Efficient graph generation with graph recurrent attention networks. *Advances in neural
389 information processing systems*, 32, 2019.
- 390 Scott Linderman, Matthew J Johnson, and Ryan P Adams. Dependent multinomial models made easy: Stick-
391 breaking with the pólya-gamma augmentation. *Advances in Neural Information Processing Systems*, 28,
392 2015.
- 393 Jenny Liu, Aviral Kumar, Jimmy Ba, Jamie Kiros, and Kevin Swersky. Graph normalizing flows, 2019.
- 394 Tengfei Ma, Jie Chen, and Cao Xiao. Constrained generation of semantically valid graphs via regularizing
395 variational autoencoders. *arXiv preprint arXiv:1809.02630*, 2018.
- 396 Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub,
397 Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A Platform for Embodied AI
398 Research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- 399 Karolis Martinkus, Andreas Loukas, Nathanaël Perraudin, and Roger Wattenhofer. Spectre: Spectral conditioning
400 helps to overcome the expressivity limits of one-shot graph generators. In *International Conference on
401 Machine Learning*, pp. 15159–15179. PMLR, 2022.
- 402 Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe Wenjie Jiang, Ebrahim Songhori, Shen Wang, Young-Joon
403 Lee, Eric Johnson, Omkar Pathak, Azade Nazi, et al. A graph placement methodology for fast chip design.
404 *Nature*, 594(7862):207–212, 2021.
- 405 Mark EJ Newman. Finding community structure in networks using the eigenvectors of matrices. *Physical review
406 E*, 74(3):036104, 2006a.
- 407 Mark EJ Newman. Modularity and community structure in networks. *Proceedings of the national academy of
408 sciences*, 103(23):8577–8582, 2006b.
- 409 Leslie O’Bray, Max Horn, Bastian Rieck, and Karsten Borgwardt. Evaluation metrics for graph generative
410 models: Problems, pitfalls, and practical solutions. *arXiv preprint arXiv:2106.01098*, 2021.
- 411 Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalch-
412 brenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint
413 arXiv:1609.03499*, 2016.
- 414 Santhosh Kumar Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alexander Clegg,
415 John M Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, Manolis Savva,
416 Yili Zhao, and Dhruv Batra. Habitat-matterport 3d dataset (HM3d): 1000 large-scale 3d environments for
417 embodied AI. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks
418 Track (Round 2)*, 2021.
- 419 Ladislav Rampásek, Michael Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini.
420 Recipe for a general, powerful, scalable graph transformer. *Advances in Neural Information Processing
421 Systems*, 35:14501–14515, 2022.
- 422 Scott Reed, Aäron Oord, Nal Kalchbrenner, Sergio Gómez Colmenarejo, Ziyu Wang, Yutian Chen, Dan Belov,
423 and Nando Freitas. Parallel multiscale autoregressive density estimation. In *International Conference on
424 Machine Learning*, pp. 2912–2921. PMLR, 2017.
- 425 Ida Schomburg, Antje Chang, Christian Ebeling, Marion Gremse, Christian Heldt, Gregor Huhn, and Dietmar
426 Schomburg. Brenda, the enzyme database: updates and major new developments. *Nucleic acids research*, 32
427 (suppl_1):D431–D433, 2004.
- 428 Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective
429 classification in network data. *AI magazine*, 29(3):93–93, 2008.
- 430 Kyle Siegrist. *Probability, Mathematical Statistics, Stochastic Processes*. LibreTexts, 2017.
431 URL [https://stats.libretexts.org/Bookshelves/Probability_Theory/Probability_](https://stats.libretexts.org/Bookshelves/Probability_Theory/Probability_Mathematical_Statistics_and_Stochastic_Processes_(Siegrist))
432 [Mathematical_Statistics_and_Stochastic_Processes_\(Siegrist\)](https://stats.libretexts.org/Bookshelves/Probability_Theory/Probability_Mathematical_Statistics_and_Stochastic_Processes_(Siegrist))

- 433 Martin Simonovsky and Nikos Komodakis. GraphVAE: Towards generation of small graphs using variational
434 autoencoders. *arXiv preprint arXiv:1802.03480*, 2018.
- 435 Rylee Thompson, Boris Knyazev, Elahe Ghalebi, Jungtaek Kim, and Graham W Taylor. On evaluation metrics
436 for graph generative models. *arXiv preprint arXiv:2201.09871*, 2022.
- 437 Anton Tsitsulin, John Palowitch, Bryan Perozzi, and Emmanuel Müller. Graph clustering with graph neural
438 networks. *arXiv preprint arXiv:2006.16904*, 2020.
- 439 Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress:
440 Discrete denoising diffusion for graph generation. *arXiv preprint arXiv:2209.14734*, 2022.
- 441 Oriol Vinyals, Samy Bengio, and Manjunath Kudlur. Order matters: Sequence to sequence for sets. *arXiv
442 preprint arXiv:1511.06391*, 2015.
- 443 Carl Yang, Peiye Zhuang, Wenhan Shi, Alan Luu, and Pan Li. Conditional structure generation through graph
444 variational generative adversarial nets. *Advances in neural information processing systems*, 32, 2019.
- 445 Jiaxuan You, Rex Ying, Xiang Ren, William Hamilton, and Jure Leskovec. Graphrnn: Generating realistic
446 graphs with deep auto-regressive models. In *ICML*, pp. 5694–5703, 2018.