

---

# GraphDF: A Discrete Flow Model for Molecular Graph Generation

---

Youzhi Luo<sup>1</sup> Keqiang Yan<sup>1</sup> Shuiwang Ji<sup>1</sup>

## Abstract

We consider the problem of molecular graph generation using deep models. While graphs are discrete, most existing methods use continuous latent variables, resulting in inaccurate modeling of discrete graph structures. In this work, we propose GraphDF, a novel discrete latent variable model for molecular graph generation based on normalizing flow methods. GraphDF uses invertible modulo shift transforms to map discrete latent variables to graph nodes and edges. We show that the use of discrete latent variables reduces computational costs and eliminates the negative effect of dequantization. Comprehensive experimental results show that GraphDF outperforms prior methods on random generation, property optimization, and constrained optimization tasks.

## 1. Introduction

A fundamental problem in drug discovery (Stokes et al., 2020; Wang et al., 2020) and chemistry is to design novel molecules with specific properties. This task remains to be challenging because the space of molecules is naturally discrete and very large, with estimated size on the order of  $10^{33}$  (Polishchuk et al., 2013). Recent advances in deep generative models have led to significant progress in this field. Notably, many recent studies represent molecular structures as graphs and propose to generate novel molecular graphs with advanced deep generative models. From a conceptual and computational perspective, these methods first map molecular graphs to vectors in the continuous latent space. When generating, the generative models project a randomly picked continuous vector in the latent space back to the molecular graph space.

To make use of generative models with continuous latent variables, existing methods (Madhawa et al., 2019; Shi\* et al., 2020; Zang & Wang, 2020) convert discrete graph

data to continuous data by adding real-valued noise. However, such dequantization processing prevents models from capturing the original discrete distribution of discrete graph structures, thus increasing the difficulty of model training. This becomes a key limitation that makes it hard to capture the true distribution of graph structures and generate diverse molecules.

In this work, we propose GraphDF, a generative model using discrete latent variables for molecular graph generation. GraphDF generates molecular graphs by sequentially sampling discrete latent variables and mapping them to new nodes and edges via invertible modulo shift transforms. Compared with previous methods, our proposed discrete transform eliminates the cost of computing Jacobian matrix and overcomes the limitations resulted from dequantization. Experimental results show that GraphDF can outperform previous state-of-the-art methods over various molecule generation tasks. The implementation of GraphDF has been integrated into the DIG (Liu et al., 2021) framework.

## 2. Background and Related Work

### 2.1. Molecule Generation

Let  $\mathcal{M} = \{M_i\}_{i=1}^m$  be a set of molecules,  $S(M) \in \mathbb{R}$  be a function computing a specific property score of the molecule  $M$ , and  $\text{SIM}(M, M') \in [0, 1]$  be a function measuring the similarity between two molecules  $M$  and  $M'$ . The three molecule generation tasks can be defined as below.

- Learning a generation model  $p_\theta(\cdot)$  from  $\mathcal{M}$ , where  $p_\theta(M)$  is the probability of sampling the molecule  $M$  from the model.
- Learning a property optimization model  $p_\theta(\cdot)$  with respect to  $S$  by maximizing  $\mathbb{E}_{M \sim p_\theta}[S(M)]$ .
- Learning a constrained optimization model  $p_\theta(\cdot|M)$ , so as to maximize  $\mathbb{E}_{M'|M \sim p_\theta}[S(M')]$  while satisfying  $\text{SIM}(M, M') > \delta$ , where  $M$  is a given molecule,  $M'$  is the optimized molecule, and  $\delta$  is a similarity threshold.

Early attempts for molecule generation (Gómez-Bombarelli et al., 2018; Kusner et al., 2017; Dai et al., 2018) represented molecules as SMILES strings (Weininger, 1988) and developed sequence generation models. These models may not

---

<sup>1</sup>Department of Computer Science & Engineering, Texas A&M University, TX, USA. Correspondence to: Shuiwang Ji <sjj@tamu.edu>.

easily learn the complicated grammatical rules of SMILES and thus could not generate syntactically valid sequences well. More recent studies represent molecules as graphs, and significant progress has been made in molecular graph generation with deep generative models. A variety of molecular graph generation methods have been developed based on variational auto-encoders (VAE) (Kingma & Welling, 2013). Some of these methods generate molecular graphs from a tree structure (Jin et al., 2018; Kajino, 2019); others add a node or an edge sequentially one at a time (Liu et al., 2018), or generate node and adjacency matrices directly in one-shot sampling (Simonovsky & Komodakis, 2018; Ma et al., 2018). For VAE-based models, Bayesian optimization can be applied in the latent space to search molecules with optimized property scores (Jin et al., 2018; Kajino, 2019). Generative adversarial networks (GAN) (Goodfellow et al., 2014) have also been used in some molecular graph generators (De Cao & Kipf, 2018; You et al., 2018a). Similar to SeqGAN (Yu et al., 2017), these models are trained with reinforcement learning.

In addition to methods based on VAE and GAN, normalizing flow models have been used in various generation tasks (Dinh et al., 2014; 2016; Rezende & Mohamed, 2015). These models define invertible transformations between latent variables and data samples, thus allowing for exact likelihood computation. It has been demonstrated that these models are able to capture the density of molecular graphs more accurately (Shi\* et al., 2020; Zang & Wang, 2020). Our proposed discrete flow model is based on the normalizing flow method.

## 2.2. Flow Models

A flow model is a parametric invertible function  $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^d$ , defining an invertible mapping from the latent variable  $z \in \mathbb{R}^d$  to the data point  $x \in \mathbb{R}^d$ , where  $x \sim p_X(x)$  and  $z \sim p_Z(z)$  are random variables. The latent distribution  $p_Z$  is a pre-defined probability distribution, e.g., a Gaussian distribution. The data distribution  $p_X$  is unknown. But given an arbitrary data point  $x \in \mathbb{R}^d$ , we can use the change-of-variable theorem to compute its log-likelihood as

$$\log p_X(x) = \log p_Z(f_\theta^{-1}(x)) + \log \left| \det \frac{\partial f_\theta^{-1}(x)}{\partial x} \right|, \quad (1)$$

where  $\frac{\partial f_\theta^{-1}(x)}{\partial x} \in \mathbb{R}^{d \times d}$  is the Jacobian matrix. To train a flow model  $f_\theta$  on a dataset  $\mathcal{X} = \{x_i\}_{i=1}^m$ , we use Eqn. (1) to compute the exact log-likelihoods of all data points and update the parameter  $\theta$  by maximizing the log-likelihoods via gradient descent. To randomly sample a data point from  $f_\theta$ , a latent variable  $z$  is first sampled from the pre-defined latent distribution  $p_Z$ . Then we can obtain a data point  $x$  by performing the feedforward transformation  $x = f_\theta(z)$ . Therefore, the mapping  $f_\theta$  should be invertible, and the

computation of the Jacobian matrix  $\frac{\partial f_\theta^{-1}(x)}{\partial x}$  needs to be tractable in order to enabling efficient training and sampling. A common operation satisfying these two requirements is the affine coupling layer (Dinh et al., 2016), which has been used frequently in deep flow models.

Recently, flow models have been used in molecule or graph generation tasks. Liu et al. (2019) proposes an invertible message passing neural network, known as GNF, for simple community graph generation. Following GraphVAE (Simonovsky & Komodakis, 2018), several recent studies generate molecular graphs by generating node and adjacency matrices in a one-shot fashion via flow models, including GraphNVP (Madhawa et al., 2019), GRF (Honda et al., 2019), and MoFlow (Zang & Wang, 2020). GraphAF (Shi\* et al., 2020) generates molecular graphs by adding nodes and edges sequentially via an autoregressive flow model (Papamakarios et al., 2017). GraphCNF (Lippe & Gavves, 2021) jointly optimizes a variational inference model and a flow model for graph generation.

## 3. A Discrete Flow Model for Graphs

While all current molecular graph generation methods use continuous latent variables, we argue that such continuous prior distribution is not suitable to model molecular graph data, which are naturally discrete. On the other hand, some recent attempt has been made to develop generative models with discrete latent variables (Tran et al., 2019). It is, however, only able to generate 1D sequences, such as texts, and cannot be used to generate richly structured objects such as graphs.

In this section, we introduce GraphDF, a novel molecular graph generative model with discrete latent variables. To our best knowledge, GraphDF is the first work which demonstrates that discrete latent variable models have strong capacity to model the density of complicated molecular graph data.

### 3.1. An Overview of Sequential Generation

Let  $k$  be the total number of node types and  $c$  be the total number of edge types. We use  $G = (X, A)$  to denote a molecular graph with  $n$  nodes, where  $X \in \{0, 1\}^{n \times k}$  is the node matrix and  $A \in \{0, 1\}^{n \times n \times c}$  is the adjacency tensor.  $X[i, u] = 1$  represents that the  $i$ -th node has type  $u$ , and  $A[i, j, v] = 1$  represents that there is an edge with type  $v$  between the  $i$ -th node and  $j$ -th node. If there is no edge between the  $i$ -th node and  $j$ -th node, then  $A[i, j, v] = 0$  for  $v = 0, \dots, c - 1$ .

Similar to MolecularRNN (Popova et al., 2019), we generate molecular graphs by a sequential decision process. We start from an empty molecular graph  $G_0$ . At step  $i$ , let the current sub-graph be  $G_{i-1}$  which contains  $i - 1$  nodes.

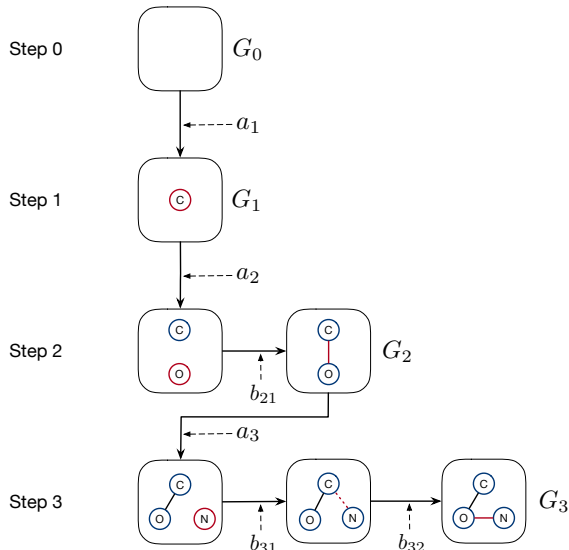


Figure 1. An illustration of sequential generation procedure.

The generative model  $f_\theta$  first generates a new node with type  $a_i$  based on a sampled latent variable  $z_i$ , where  $a_i \in \{0, \dots, k-1\}$ . Afterwards, the edges between the new node and all nodes in  $G_{i-1}$  are generated sequentially by  $f_\theta$ . Specifically, the edge incident to the  $j$ -th node of  $G_{i-1}$  is generated based on the latent variable  $z_{ij}$ , whose edge type  $b_{ij} \in \{0, \dots, c-1, c\}$ , where the extra edge type  $c$  means no edge. If  $b_{ij}$  violates the chemical bond valency rule, it is re-generated by re-sampling a new  $z_{ij}$ . The above process is an autoregressive function of previously generated elements and repeats until no new edge is added at step  $i$ , i.e.,  $b_{ij} = c$  for  $j = 1, \dots, i-1$ . An example of this process is given in Fig. 1. In summary, generating a molecular graph  $G$  is equivalent to generating the sequence representation  $S_G = (a_1, a_2, b_{21}, a_3, b_{31}, b_{32}, a_4, \dots)$  by

$$\begin{aligned} a_i &= f_\theta(a_1, \dots, b_{i-1, i-2}; z_i), \quad i > 0, \\ b_{ij} &= f_\theta(a_1, \dots, b_{i, j-1}; z_{ij}), \quad j = 1, \dots, i-1. \end{aligned} \quad (2)$$

### 3.2. Generation with Discrete Latent Variables

In our method, all latent variables are discrete and sampled from multinomial distributions. Specifically,  $z_i$  is sampled from a multinomial distribution  $p_{Z_a}$  with trainable parameters  $(\alpha_0, \dots, \alpha_{k-1})$  as

$$p_{Z_a}(z_i = s) = \frac{\exp(\alpha_s)}{\sum_{t=0}^{k-1} \exp(\alpha_t)}. \quad (3)$$

Similarly,  $z_{ij}$  is sampled from the multinomial distribution  $p_{Z_b}$  with trainable parameters  $(\beta_0, \dots, \beta_c)$ . We use a discrete flow model to reversibly map discrete latent variables to new nodes and edges. The discrete transform used in the

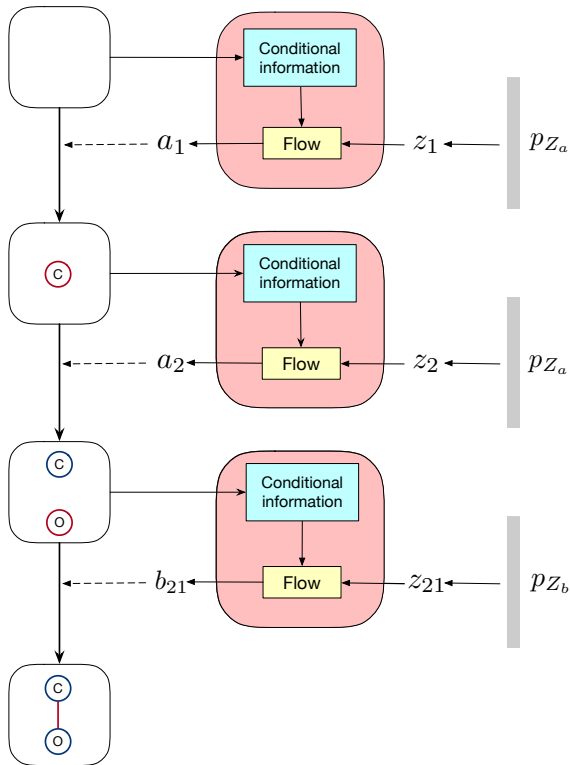


Figure 2. An illustration of generation with GraphDF framework.

discrete flow is a modulo shift transform in the form of

$$q(z) = (z + \mu) \bmod t, \quad (4)$$

where  $t$  is the number of categories, and  $z, \mu \in \{0, \dots, t-1\}$ . Formally, the discrete transform for generating new nodes and edges in Eqn. (2) is the composition of  $D$  modulo shift modules as

$$\begin{aligned} a_i &= q_i^D \circ \dots \circ q_i^1(z_i), \quad i > 0, \\ b_{ij} &= q_{ij}^D \circ \dots \circ q_{ij}^1(z_{ij}), \quad j = 1, \dots, i-1, \end{aligned} \quad (5)$$

where  $f \circ g(\cdot)$  is defined as  $f(g(\cdot))$ , and

$$\begin{aligned} q_i^d(z) &= (z + \mu_i^d) \bmod k, \quad d = 1, \dots, D \\ q_{ij}^d(z) &= (z + \mu_{ij}^d) \bmod (c+1), \quad d = 1, \dots, D. \end{aligned} \quad (6)$$

Here the shift factors  $\mu_i^d$  and  $\mu_{ij}^d$  are functions of  $(a_1, a_2, b_{21}, \dots, b_{i-1, i-2})$  and  $(a_1, a_2, b_{21}, \dots, b_{i, j-1})$ , respectively. They both play the role of capturing conditional information from previously generated elements. The calculation of  $\mu_i^d$  and  $\mu_{ij}^d$  will be described in Sec. 3.4. An illustrative example of generating a simple molecular graph with this framework is given in Fig. 2.

The use of discrete latent variables is the key factor distinguishing our method from other generation methods based

on flow models. Our model has several advantages. The change-of-variable formula of discrete invertible mapping does not involve the Jacobian matrix term in Eqn. (1). Hence the probabilities of sampled nodes and edges can be directly obtained as

$$\begin{aligned} p(a_i|a_1, a_2, b_{21}, \dots, b_{i-1, i-2}) &= p_{Z_a}(z_i), \\ p(b_{ij}|a_1, a_2, b_{21}, \dots, b_{i, j-1}) &= p_{Z_b}(z_{ij}), \end{aligned} \quad (7)$$

leading to dramatic reduction in computational cost.

More importantly, discrete latent variables free the generative model from the defects of dequantization. Previous methods (Madhawa et al., 2019; Shi\* et al., 2020; Zang & Wang, 2020) convert categorical representation of graph nodes and edges to continuous data by adding real-valued noise to their one-hot encodings. Then continuous flow models map such pseudo continuous data to latent variables subject to Gaussian distribution. We argue that such a dequantization operation can be problematic. First, the generative model does not capture the true distribution of original discrete data, but the distorted distribution after dequantization. Second, due to the randomness of noise, the resulting continuous data points can be very different even if they are obtained by performing dequantization on the same graph. Hence during training, a single graph may be mapped to two distant points in the latent space, leading to model oscillation and failure to converge. However, our GraphDF maps the discrete graph data to discrete latent space by discrete invertible flow without introducing any distribution distortion. Hence it can model the underlying distribution of graph structures more accurately, thereby eliminating the impact of two issues caused by dequantization.

Note that our method is fundamentally different from GraphCNF (Lippe & Gavves, 2021), though dequantization is avoided in both methods. GraphCNF maps discrete values to continuous latent space where different categories occupy non-overlapping regions. Such a discrete-to-continuous mapping leads to distribution distortions. In the generation process, discrete values are sampled by applying Bayes, and all graph nodes and edges are generated at once. On the other hand, GraphDF uses a fundamentally different approach to model the discrete distribution. GraphDF completely discards continuous latent space and maps discrete values to discrete latent space by the modulo shift transform. No distribution distortion is introduced here. In addition, compared with GraphCNF, the sequential fashion of GraphDF can capture the graph structure information efficiently and guarantee the chemical validity because of using bond valency check in the generation.

### 3.3. Inference and Training

The inverse process of generation is inference, which involves mapping from graphs to discrete latent variables.

Given a graph  $G = (X, A)$ , we compute its sequence representation  $S_G$  by re-ordering its nodes and edges according to the breadth-first search (BFS) order. Then we can obtain the discrete latent variables corresponding to each element of this sequence by inverting Eqn. (5) as

$$\begin{aligned} z_i &= o_i^1 \circ \dots \circ o_i^D(a_i), \quad i > 0, \\ z_{ij} &= o_{ij}^1 \circ \dots \circ o_{ij}^D(b_{ij}), \quad j = 1, \dots, i - 1, \end{aligned} \quad (8)$$

where

$$\begin{aligned} o_i^d(z) &= (z - \mu_i^d) \bmod k, \quad d = 1, \dots, D \\ o_{ij}^d(z) &= (z - \mu_{ij}^d) \bmod (c + 1), \quad d = 1, \dots, D. \end{aligned} \quad (9)$$

Similar to the generation process, the shift factors  $\mu_i^d$  and  $\mu_{ij}^d$  capture autoregressive information from elements in  $S_G$  before  $a_i$  and  $b_{ij}$ , respectively.

We train a GraphDF model on a molecular graph dataset by maximizing the log-likelihoods of all data samples. Given a molecular graph  $G$ , we first perform inference on  $G$  to compute the latent variables corresponding to its nodes and edges. Then the log-likelihood of  $G$  can be computed as

$$\log p(G) = \sum_{i=1}^n \left( \log p_{Z_a}(z_i) + \sum_{j=1}^{i-1} \log p_{Z_b}(z_{ij}) \right). \quad (10)$$

The parameters of GraphDF models and discrete prior distributions  $p_{Z_a}, p_{Z_b}$  are updated by gradient descent to maximize the log-likelihoods of all data samples.

### 3.4. Invariant and Discrete Conditional Generation

The transformation of discrete flow in Eqn. (5) is autoregressive and requires capturing conditional information into  $\mu_i^d$  and  $\mu_{ij}^d$ . Tran et al. (2019) uses language models, like LSTM (Hochreiter & Schmidhuber, 1997) and Transformer (Vaswani et al., 2017), in discrete flow models to develop sequence generators. However, the same method cannot be used directly to generate  $S_G$  for two reasons as below.

- Language models are only aware of the node and edge type information encoded in  $S_G$ , while rich structure information of graphs is not captured.
- There is no intrinsic order in an intermediately generated sub-graph, but language models assume an order when dealing with  $S_G$ . For instance, for  $(a_1, a_2) = (1, 0)$  and  $(a_1, a_2) = (0, 1)$ , the generative model should give the same estimate  $p(b_{21}|a_1, a_2)$ , but language models cannot provide such a guarantee.

To tackle the two issues, we propose to capture conditional information from intermediate sub-graphs by using graph convolutional networks (GCN) (Kipf & Welling, 2017; Gao

et al., 2018; Gao & Ji, 2019). Specifically, let the sub-graph generated immediately before adding the node  $a_i$  be  $G_{i-1} = (X, A)$ , we use an  $L$ -layer relational GCN (R-GCN) (Schlichtkrull et al., 2018) to obtain node embeddings  $H^L \in \mathbb{R}^{n \times r}$  from  $G_{i-1}$  by performing  $L$  message passing operations. The  $\ell$ -th message passing is performed by:

$$H^\ell = \sum_{v=1}^c \text{ReLU} \left( \tilde{D}_v^{-\frac{1}{2}} \tilde{A}_v \tilde{D}_v^{-\frac{1}{2}} H^{\ell-1} W_v^\ell \right), \quad (11)$$

where  $H^0 = X$ ,  $\tilde{A}_v = A[:, :, v] + I$ ,  $\tilde{D}_v$  is a diagonal matrix and  $\tilde{D}_v[j, j] = \sum_{u=1}^n \tilde{A}_v[j, u]$ ,  $\{W_v^\ell\}_{v=1}^c$  are trainable weight matrices of the  $\ell$ -th layer. Note that the same R-GCN is shared among all modulo shift modules in the discrete flow. Afterwards,  $\mu_i^d$  is calculated by a classification network based on multi-layer perceptrons (MLP) as

$$\begin{aligned} h &= \text{sum}(H^L), \\ \mu_i^d &= \arg \max \text{MLP}_a^d(h), \quad d = 1, \dots, D, \end{aligned} \quad (12)$$

where  $h \in \mathbb{R}^r$  is the graph embedding obtained by performing sum-pooling on node embeddings  $H^L$ ,  $\text{MLP}_a^d$  projects the input to a  $k$ -dimensional logits vector, and  $\arg \max$  outputs the index position of the maximum logits value. Since  $\arg \max$  is not differentiable, we use the softmax-temperature function to approximate the gradient (Jang et al., 2016) during training as

$$\begin{aligned} \gamma^d &= \text{MLP}_a^d(h), \quad d = 1, \dots, D, \\ \frac{d\mu_d}{d\gamma^d} &\approx \frac{d}{d\gamma^d} \text{softmax}\left(\frac{\gamma^d}{\tau}\right), \quad d = 1, \dots, D, \end{aligned} \quad (13)$$

where we fix the temperature  $\tau = 0.1$  in experiments.

Similarly, to obtain  $\mu_{ij}^d$ , we use the R-GCN to extract node embeddings  $H^L$  and graph embedding  $h$  from  $G_{i-1,j}$ , which is the sub-graph generated immediately before adding the edge  $b_{ij}$ . Then we compute  $\mu_{ij}^d$  by MLP as:

$$\mu_{ij}^d = \arg \max \text{MLP}_b^d(\text{Con}(h, H_i^L, H_j^L)), \quad d = 1, \dots, D,$$

where  $H_i^L$  and  $H_j^L$  denote the node embeddings of the  $i$ -th and  $j$ -th node of  $G_{i-1,j}$ , respectively, and  $\text{Con}(\cdot)$  is the concatenation operation. The softmax approximation is also used here to perform back propagation through  $\arg \max$ .

By using our proposed method, the discrete flow model can elegantly extract conditional information from graph structures, which is then incorporated into the mapping from discrete latent space to newly generated graph nodes and edges. The conditional information is guaranteed to be invariant to the order of previously generated elements. We summarize the detailed generation and training algorithms of GraphDF in Appendix A and B, respectively.

### 3.5. Property Optimization with Reinforcement Learning

The above discussions involve molecular graph generation independent of property scores. However, both property optimization and constrained optimization require the generated molecules to have optimal property scores. To this end, we use reinforcement learning to fine-tune a trained GraphDF model to optimize the specified property score.

We can formulate the sequential generation process with GraphDF as a Markov Decision Process (MDP). Formally, at each step, we treat the currently generated sub-graph as the state, and adding a new node or edge as the action. The action probabilities  $p(a_i|G_{i-1})$  and  $p(b_{ij}|G_{i-1,j})$  are thus determined by the GraphDF model in Eqn. (7). The rewards consist of intermediate reward and final reward. The intermediate reward is the penalty for violating chemical bond valency rule. The final reward incorporates the property score and penalty for containing excessive steric strain or violating functional group filters of ZINC (Irwin et al., 2012) in the finally generated molecule.

Following GCPN (You et al., 2018a), we use Proximal Policy Optimization (PPO) (Schulman et al., 2017), a powerful policy optimization algorithm, to fine-tune GraphDF models in the above reinforcement learning environment. The loss function for a generated graph  $G$  is

$$\begin{aligned} L(G, \theta) &= \sum_{i=1}^n \left( L^c(r_i, A_i) + \sum_{j=1}^n L^c(r_{ij}, A_{ij}) \right), \quad (14) \\ L^c(r, A) &= \min\{rA, \text{clip}(r, 1 - \epsilon, 1 + \epsilon)A\}, \end{aligned}$$

where  $r_i = \frac{p(a_i|G_{i-1};\theta)}{p(a_i|G_{i-1};\theta_{\text{old}})}$  and  $r_{ij} = \frac{p(b_{ij}|G_{i-1,j};\theta)}{p(b_{ij}|G_{i-1,j};\theta_{\text{old}})}$  are ratios of action probabilities by new policy and old policy.  $A_i$  and  $A_{ij}$  are estimated advantage functions, for which we use the accumulated rewards of future steps in experiments.

## 4. Experiments

In this section, we evaluate our GraphDF on three tasks of molecule generation, as mentioned in Section 2.1. We show that over most tasks, GraphDF can outperform three strong baselines, JT-VAE (Jin et al., 2018), GCPN (You et al., 2018a) and MRNN (Popova et al., 2019), and other molecular graph generation methods based on flow models including GraphNVP (Madhawa et al., 2019), GRF (Honda et al., 2019), GraphAF (Shi\* et al., 2020), MoFlow (Zang & Wang, 2020) and GraphCNF (Lippe & Gavves, 2021).

### 4.1. Random Generation

**Data.** For random generation of molecular graphs, we train and evaluate GraphDF on three molecule datasets, ZINC250K (Irwin et al., 2012), QM9 (Ramakrishnan

**GraphDF: A Discrete Flow Model for Molecular Graph Generation**

Table 1. Random generation performance on ZINC250K dataset. Results with \* are taken from Shi\* et al. (2020). Results of MoFlow (with  $\diamond$ ) are obtained by running its official source code.

Method	Validity	Validity w/o check	Uniqueness	Novelty	Reconstruction
JT-VAE	<b>100%</b>	n/a	<b>100%*</b>	<b>100%*</b>	76.7%
GCPN	<b>100%</b>	20%*	99.97%*	<b>100%*</b>	n/a
MRNN	<b>100%</b>	65%	99.89%	<b>100%</b>	n/a
GraphNVP	42.6%	n/a	94.8%	<b>100%</b>	<b>100%</b>
GRF	73.4%	n/a	53.7%	<b>100%</b>	<b>100%</b>
GraphAF	<b>100%</b>	68%	99.1%	<b>100%</b>	<b>100%</b>
MoFlow	<b>100%</b> $\diamond$	50.3% $\diamond$	99.99% $\diamond$	<b>100%</b> $\diamond$	<b>100%</b> $\diamond$
GraphCNF	96.35%	n/a	99.98%	99.98%	<b>100%</b>
GraphDF	<b>100%</b>	<b>89.03%</b>	99.16%	<b>100%</b>	<b>100%</b>

Table 2. Random generation performance on QM9 dataset. Results of MoFlow (with  $\diamond$ ) are obtained by running its official source code.

Method	Validity	Validity w/o check	Uniqueness	Novelty	Reconstruction
GraphNVP	83.1%	n/a	<b>99.2%</b>	58.2%	<b>100%</b>
GRF	84.5%	n/a	66%	58.6%	<b>100%</b>
GraphAF	<b>100%</b>	67%	94.15%	88.83%	<b>100%</b>
MoFlow	<b>100%</b> $\diamond$	<b>88.96%</b> $\diamond$	98.53% $\diamond$	96.04% $\diamond$	<b>100%</b> $\diamond$
GraphDF	<b>100%</b>	82.67%	97.62%	<b>98.1%</b>	<b>100%</b>

Table 3. Random generation performance on MOSES dataset. Results with \* are taken from Polykovskiy et al. (2020).

Method	Validity	Validity w/o check	Uniqueness	Novelty	Reconstruction
JT-VAE	<b>100%*</b>	n/a	99.96%*	91.43%*	n/a
GraphCNF	95.66%	n/a	99.98%	<b>100%</b>	n/a
GraphAF	<b>100%</b>	71%	<b>99.99%</b>	<b>100%</b>	<b>100%</b>
GraphDF	<b>100%</b>	<b>87.58%</b>	99.55%	<b>100%</b>	<b>100%</b>

Table 4. Random generation performance on COMMUNITY-SMALL and EGO-SMALL datasets. Following GNF (Liu et al., 2019), we evaluate with MMD metric, the first set of results are obtained using node distribution matching, and the second set of results are obtained from evaluating 1024 generated graphs without node distribution matching.

Method	COMMUNITY-SMALL			EGO-SMALL		
	Degree	Cluster	Orbit	Degree	Cluster	Orbit
GNF	0.20	0.20	0.11	0.03	<b>0.11</b>	<b>0.001</b>
GraphAF	0.18	0.20	<b>0.02</b>	<b>0.03</b>	<b>0.11</b>	<b>0.001</b>
GraphDF	<b>0.06</b>	<b>0.12</b>	0.03	0.04	0.13	0.010
GNF(1024)	0.12	0.15	0.02	<b>0.01</b>	<b>0.03</b>	<b>0.0008</b>
GraphAF(1024)	0.06	<b>0.10</b>	<b>0.015</b>	0.04	0.04	0.008
GraphDF(1024)	<b>0.04</b>	<b>0.10</b>	0.018	0.07	0.10	0.014

et al., 2014), and MOSES (Polykovskiy et al., 2020). All molecules are preprocessed to kekulized form by RDKit (Landrum, 2016). In addition, we conduct experiments to show that GraphDF is very general and can be used to graphs other than molecular graphs. Following the experiments of GNF (Liu et al., 2019), we test GraphDF on two generic graph datasets, COMMUNITY-SMALL and EGO-SMALL. Detailed information about all datasets used in

random generation experiments is provided in Appendix C.

**Setup.** We adopt five widely used metrics to evaluate GraphDF on molecular graph generation. Among all generated molecules, **Validity** is the percentage of molecules which do not violate chemical valency rule. Since we use re-sampling when valency check fails as mentioned in Sec. 3.1, we also report **Validity w/o check**, which is the validity per-

centage when this valency correction is disabled. Among all generated valid molecules, **Uniqueness** and **Novelty** are the percentage of unique molecules and molecules not appearing in the training data, respectively. **Reconstruction** is the percentage of molecules which can be reconstructed from their latent variables. The five metrics are calculated from 10,000 generated molecular graphs. We find that the reported metrics in MoFlow paper (Zang & Wang, 2020) are not calculated from the same 10,000 generated molecules, so we recalculate the metrics of MoFlow by re-running its official source code for fair comparison. As for COMMUNITY-SMALL and EGO-SMALL datasets, we follow GNF (Liu et al., 2019) to calculate the Maximum Mean Discrepancy (MMD) (Gretton et al., 2012) distance of statistics of degrees, clustering coefficients and orbit counts between generated graphs and graphs in the dataset. We calculate the MMD from 1024 generated graphs in the two cases of using and not using node distribution matching.

The flow model in GraphDF consists of 12 modulo shift modules, where the shared R-GCN has 3 message passing layers. The hidden dimension and output dimension of node embeddings are 128. For molecular graphs, the input node embeddings are one-hot embeddings of node types. For graphs in COMMUNITY-SMALL and EGO-SMALL datasets, the input node embeddings are one-hot indicator vectors. After message passing layers, a batch normalization layer and a sum pooling layer are used to get the graph embedding. The two MLPs for calculating  $\mu_i^d$  and  $\mu_{ij}^d$  both have two fully-connected layers, with a hidden dimension of 128 and a non-linear activation function of tanh. A single R-GCN is used among all 12 modulo shift modules in the discrete flow model, while each modulo shift module has an MLP to calculate  $\mu_i^d$  or  $\mu_{ij}^d$ . We use the above model architecture configuration for all experiments in this paper.

Adam (Kingma & Ba, 2015) optimizer is used to train GraphDF models. On each molecule dataset, GraphDF is trained for 10 epochs, where the batch size is 32 and the learning rate is 0.001. On COMMUNITY-SMALL and EGO-SMALL datasets, GraphDF is trained for 1000 epochs, where the batch size is 16 and the learning rate is 0.001. See Appendix D for more training and generation details.

**Results.** Table 1, Table 2 and Table 3 show that GraphDF can outperform the state-of-the-art methods over most metrics on all three molecule datasets. Thanks to the invertible property of flow models and valency correction, GraphDF can always achieve 100% validity and reconstruction rate. In addition, GraphDF achieves reasonably high uniqueness and novelty rate on three datasets, showing that it does not overfit or produce mode collapse. Particularly, GraphDF can outperform GraphAF by a large margin in terms of validity w/o check. Since GraphAF generates molecular graphs in the same sequential fashion but uses continuous flow, we

Table 5. Property optimization performance evaluated by top-3 property scores. The scores of MRNN are taken from Shi\* et al. (2020).

Method	Penalized logP			QED		
	1st	2nd	3rd	1st	2nd	3rd
ZINC (Dataset)	4.52	4.3	4.23	0.948	0.948	0.948
JT-VAE	5.3	4.93	4.49	0.925	0.911	0.910
GCPN	7.98	7.85	7.80	<b>0.948</b>	0.947	0.946
MRNN	8.63	6.08	4.73	0.844	0.796	0.736
GraphAF	12.23	11.29	11.05	<b>0.948</b>	<b>0.948</b>	0.947
MoFlow	n/a	n/a	n/a	<b>0.948</b>	<b>0.948</b>	<b>0.948</b>
GraphDF	<b>13.7</b>	<b>13.18</b>	<b>13.17</b>	<b>0.948</b>	<b>0.948</b>	<b>0.948</b>

believe that the good performance of GraphDF comes from the usage of discrete flow. It strongly demonstrates that discrete latent variables can learn chemical rules and model the underlying distributions more accurately than continuous latent variables. Besides, results in Table 4 show that GraphDF can consistently achieve good performance on generic graph data when compared with GNF and GraphAF.


## 4.2. Property Optimization

**Setup.** In property optimization task, our objective is to generate novel molecules with high property scores. We use penalized logP and QED (Bickerton et al., 2012) as the optimization targets. Penalized logP is the logP score reduced by synthetic accessibility and ring size, and QED is quantitative estimation of drug-likeness. We calculate them using the scripts from the official implementation of GraphAF (Shi\* et al., 2020).


We use reinforcement learning to optimize these two targets. In this process, both intermediate reward and final reward are considered. The intermediate reward is the penalization for violating chemical bond valency rule. At a generation step, if the action in this step is to add a new edge to the sub-graph and the new edge results in the violation of chemical bond valency rule, then a negative reward of -1 is assigned to this step. Besides, when a complete molecular graph  $G$  is generated, the final reward  $R(G)$  is calculated as

$$R(G) = R_p(G) - R_{ss}(G) - R_f(G). \quad (15)$$

$R_{ss}(G)$  is 1 if  $G$  is too sterically strained, i.e., the average angle bend energy of  $G$  exceeds 0.82 kcal/mol, otherwise  $R_{ss}(G)$  is 0.  $R_f(G)$  is 1 when functional group filters of ZINC (Irwin et al., 2012) detect problematic functional groups from  $G$ , otherwise  $R_f(G)$  is 0.  $R_p(G)$  is related to the property score of  $G$ . Denote by  $\log P(\cdot)$  the computation of penalized logP property and  $\text{QED}(\cdot)$  the computation of



(a) Four molecules with top penalized logP scores. Note that they are all long carbon chains, only differing in one atom. These are typical chemical structures with high penalized logP property scores.



(b) Four molecules with top QED scores.

Figure 3. An illustration of generated molecules with best property scores after property optimization.

QED property, we use

$$R_p(G) = \exp(\log P(G)/3.0 - 4.0),$$

and

$$R_p(G) = 2.0 * \text{QED}(G),$$

respectively. The final reward  $R(G)$  is distributed to all intermediate steps with a discount factor  $\gamma = 0.9$ . Specifically, if the total number of steps to generate  $G$  is  $T$ , then the assigned reward is  $\gamma^{T-t}R(G)$  for the  $t$ -th step where  $t = 1, \dots, T$ .




Figure 4. An illustration of two examples in constrained optimization of penalized logP.

We first pretrain GraphDF model on ZINC250K dataset for 1000 epochs and then fine-tune the model with reinforcement learning for 200 iterations. More details about model training are presented in Appendix D. Following previous literatures, we report the top-3 property scores of molecules generated from the model.

**Results.** We summarize the top-3 property scores of ZINC250K dataset and generation methods in Table 5. It shows that GraphDF can outperform all baselines in penalized logP optimization task and achieve comparable performance in QED optimization task. Note that similar reinforcement learning procedure described in Sec. 3.5 is used to optimize the property scores in GCPN (You et al., 2018a) and GraphAF (Shi\* et al., 2020). Hence, we believe that the good performance of GraphDF demonstrates its strong capacity to explore the complicated chemical structure space and generate diverse molecular graphs. Some generated molecules with high penalized logP scores and QED scores are visualized in Figure 3.

### 4.3. Constrained Optimization

**Data.** In this task, we aim to modify the input molecular graph so as to improve its penalized logP score while the similarity between the input and modified molecules is higher than a threshold  $\delta$ . Following JT-VAE (Jin et al., 2018), GCPN (You et al., 2018a), GraphAF (Shi\* et al., 2020), and MoFlow (Zang & Wang, 2020), we use selected 800 molecules from ZINC250K with low penalized logP scores as the input molecules to be optimized. However, we find that the set of 800 molecules used in JT-VAE and GCPN are different from that used in GraphAF and MoFlow, and the official implementation of GCPN does not provide code for re-running constrained optimization on new datasets. Hence, for fair comparison, we train GraphDF and report



Table 6. Constrained optimization performance on 800 molecules used in JT-VAE and GCPN.

$\delta$	JT-VAE			GCPN			GraphDF		
	Improvement	Similarity	Success	Improvement	Similarity	Success	Improvement	Similarity	Success
0.0	1.91±2.04	0.28±0.15	97.5%	4.20±1.28	0.32±0.12	100%	<b>5.93±1.97</b>	0.30±0.12	100%
0.2	1.68±1.85	0.33±0.13	97.1%	4.12±1.19	0.34±0.11	100%	<b>5.62±1.65</b>	0.32±0.10	100%
0.4	0.84±1.45	0.51±0.10	83.6%	2.49±1.30	0.47±0.08	100%	<b>4.13±1.41</b>	0.47±0.07	100%
0.6	0.21±0.71	0.69±0.06	46.4%	0.79±0.63	0.68±0.08	100%	<b>1.72±1.15</b>	0.67±0.06	93%

Table 7. Constrained optimization performance on 800 molecules used in GraphAF and MoFlow.

$\delta$	GraphAF			MoFlow			GraphDF		
	Improvement	Similarity	Success	Improvement	Similarity	Success	Improvement	Similarity	Success
0.0	13.13±6.89	0.29±0.15	100%	8.61±5.44	0.30±0.20	98.88%	<b>14.15±6.86</b>	0.29±0.13	100%
0.2	11.90±6.86	0.33±0.12	100%	7.06±5.04	0.43±0.20	96.75%	<b>12.77±6.59</b>	0.32±0.11	100%
0.4	8.21±6.51	0.49±0.09	99.88%	4.71±4.55	0.61±0.18	85.75%	<b>9.19±6.43</b>	0.48±0.08	99.63%
0.6	<b>4.98±6.49</b>	0.66±0.05	96.88%	2.10±2.86	0.79±0.14	58.25%	4.51±5.80	0.65±0.05	92.13%

results on two datasets, one is used in JT-VAE and GCPN while the other is used in GraphAF and MoFlow.

**Setup.** Similar to the property optimization task, GraphDF is pretrained on ZINC250K dataset for 1000 epochs and then fine-tuned with reinforcement learning. More details about model training can be found in Appendix D.

We use the reward defined in Eqn. (15), where  $R_p(G)$  is defined as

$$R_p(G) = \log P(G) - \log P(G_{in}),$$

where  $G_{in}$  is the input molecular graph. The initial state of a generation process is set to be a sub-graph of the input molecular graph. Specifically, given an input molecular graph  $G_{in}$ , we first re-order its nodes by a breadth-first search, then randomly remove the last  $m$  nodes and edges incident to them, where  $m$  is randomly selected from  $\{0, 1, 2, 3, 4, 5\}$ . The remaining sub-graph is set as the initial state.

We use Tanimoto similarity with Morgan fingerprint (Rogers & Hahn, 2010) to compute the similarity between the input and the optimized molecules. We optimize the input 800 molecules under 4 different similarity constraints of 0.0, 0.2, 0.4, 0.6, respectively. The success rate under each similarity constraint is reported, which is the percentage of molecules that can be successfully optimized over 800 molecules. Over all successfully optimized molecules, we calculate the mean and standard deviation of the largest property improvement, and similarities between them and their corresponding input molecules.

**Results.** Results of constrained optimization on two datasets are summarized in Table 6 and Table 7, respectively. Results show that GraphDF can achieve much higher average property improvement than JT-VAE, GCPN, and

MoFlow under 4 similarity constraints, and outperform GraphAF under 3 of 4 similarity constraints. In addition, the average similarities achieved by GraphDF are comparable with all baselines, and unlike JT-VAE or MoFlow, the success rate of GraphDF does not decrease dramatically as the similarity threshold increases. The good performance further proves that GraphDF has a strong ability to search molecules with high property scores in chemical space. We illustrate several optimization examples in Figure 4.

## 5. Conclusion

We propose GraphDF, the first molecular graph generation framework using discrete latent variables. GraphDF sequentially map discrete latent variables to graph nodes and edges via invertible modulo shift transforms. The use of discrete latent variables results in more accurate modeling of graph structures and stronger capacity to explore the chemical space. More importantly, GraphDF can avoid expensive computation of Jacobian matrix and eliminate the negative effects of dequantization, which exist in many prior methods. We experimentally demonstrate the advantages of GraphDF and achieve new state-of-the-art results in three molecule generation tasks. The major limitations of GraphDF are relying on BFS node ordering during training and the generation speed is slower than one-shot methods. In the future, we would like to apply GraphDF to other types of graph data and extend it to graph translation problems.

## Acknowledgements

We thank Chence Shi for his help on providing technical details on the implementation of reinforcement learning method. This work was supported in part by National Science Foundation grant DBI-2028361.

## References

- Bickerton, G. R., Paolini, G. V., Besnard, J., Muresan, S., and Hopkins, A. L. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2):90–98, 2012.
- Dai, H., Tian, Y., Dai, B., Skiena, S., and Song, L. Syntax-directed variational autoencoder for structured data. In *7th International Conference on Learning Representations*, 2018.
- De Cao, N. and Kipf, T. MolGAN: An implicit generative model for small molecular graphs. *ICML 2018 workshop on Theoretical Foundations and Applications of Deep Generative Models*, 2018.
- Dinh, L., Krueger, D., and Bengio, Y. NICE: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real NVP. In *4th International Conference on Learning Representations*, 2016.
- Gao, H. and Ji, S. Graph U-Nets. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2083–2092. PMLR, 09–15 Jun 2019.
- Gao, H., Wang, Z., and Ji, S. Large-scale learnable graph convolutional networks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '18, pp. 1416–1424, New York, NY, USA, 2018. Association for Computing Machinery.
- Gómez-Bombarelli, R., Wei, J. N., Duvenaud, D., Hernández-Lobato, J. M., Sánchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T. D., Adams, R. P., and Aspuru-Guzik, A. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems*, volume 27, pp. 2672–2680. Curran Associates, Inc., 2014.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Honda, S., Akita, H., Ishiguro, K., Nakanishi, T., and Oono, K. Graph residual flow for molecular graph generation. *arXiv preprint arXiv:1909.13521*, 2019.
- Irwin, J. J., Sterling, T., Mysinger, M. M., Bolstad, E. S., and Coleman, R. G. ZINC: a free tool to discover chemistry for biology. *Journal of chemical information and modeling*, 52(7):1757–1768, 2012.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with Gumbel-Softmax. In *4th International Conference on Learning Representations*, 2016.
- Jin, W., Barzilay, R., and Jaakkola, T. Junction tree variational autoencoder for molecular graph generation. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2323–2332, 2018.
- Kajino, H. Molecular hypergraph grammar with its application to molecular optimization. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, Long Beach, California, USA, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *Proceedings of the 3rd international conference on learning representations*, 2015.
- Kingma, D. P. and Welling, M. Auto-Encoding variational bayes. In *2nd International Conference on Learning Representations*, 2013.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations*, 2017.
- Kusner, M. J., Paige, B., and Hernández-Lobato, J. M. Grammar variational autoencoder. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1945–1954, 2017.
- Landrum, G. RDKit: Open-source cheminformatics. <http://www.rdkit.org>, 2016. Accessed:2021-01-16.
- Lippe, P. and Gavves, E. Categorical normalizing flows via continuous transformations. In *9th International Conference on Learning Representations*, 2021.
- Liu, J., Kumar, A., Ba, J., Kiros, J., and Swersky, K. Graph normalizing flows. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32, pp. 13578–13588. Curran Associates, Inc., 2019.

- Liu, M., Luo, Y., Wang, L., Xie, Y., Yuan, H., Gui, S., Yu, H., Xu, Z., Zhang, J., Liu, Y., et al. DIG: A turnkey library for diving into graph deep learning research. *arXiv preprint arXiv:2103.12608*, 2021.
- Liu, Q., Allamanis, M., Brockschmidt, M., and Gaunt, A. Constrained graph variational autoencoders for molecule design. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31, pp. 7795–7804. Curran Associates, Inc., 2018.
- Ma, T., Chen, J., and Xiao, C. Constrained generation of semantically valid graphs via regularizing variational autoencoders. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31, pp. 7113–7124. Curran Associates, Inc., 2018.
- Madhawa, K., Ishiguro, K., Nakago, K., and Abe, M. Graph-NVP: An invertible flow model for generating molecular graphs. *arXiv preprint arXiv:1905.11600*, 2019.
- Papamakarios, G., Pavlakou, T., and Murray, I. Masked autoregressive flow for density estimation. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30, pp. 2338–2347. Curran Associates, Inc., 2017.
- Polishchuk, P. G., Madzhidov, T. I., and Varnek, A. Estimation of the size of drug-like chemical space based on GDB-17 data. *Journal of computer-aided molecular design*, 27(8):675–679, 2013.
- Polykovskiy, D., Zhebrak, A., Sanchez-Lengeling, B., Golovanov, S., Tatanov, O., Belyaev, S., Kurbanov, R., Artamonov, A., Aladinskiy, V., Veselov, M., Kadurin, A., Johansson, S., Chen, H., Nikolenko, S., Aspuru-Guzik, A., and Zhavoronkov, A. Molecular sets (MOSES): A benchmarking platform for molecular generation models. *Frontiers in Pharmacology*, 11:1931, 2020. ISSN 1663-9812. doi: 10.3389/fphar.2020.565644. URL <https://www.frontiersin.org/article/10.3389/fphar.2020.565644>.
- Popova, M., Shvets, M., Oliva, J., and Isayev, O. MolecularRNN: Generating realistic molecular graphs with optimized properties. *arXiv preprint arXiv:1905.13372*, 2019.
- Ramakrishnan, R., Dral, P. O., Rupp, M., and Von Lilienfeld, O. A. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific data*, 1(1):1–7, 2014.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1530–1538, Lille, France, 2015.
- Rogers, D. and Hahn, M. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.
- Schlichtkrull, M., Kipf, T. N., Bloem, P., van den Berg, R., Titov, I., and Welling, M. Modeling relational data with graph convolutional networks. In Gangemi, A., Navigli, R., Vidal, M.-E., Hitzler, P., Troncy, R., Hollink, L., Tordai, A., and Alam, M. (eds.), *The Semantic Web*, pp. 593–607, Cham, 2018. Springer International Publishing.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Sen, P., Namata, G., Bilgic, M., Getoor, L., Galligher, B., and Eliassi-Rad, T. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.
- Shi\*, C., Xu\*, M., Zhu, Z., Zhang, W., Zhang, M., and Tang, J. GraphAF: a flow-based autoregressive model for molecular graph generation. In *8th International Conference on Learning Representations*, 2020.
- Simonovsky, M. and Komodakis, N. GraphVAE: Towards generation of small graphs using variational autoencoders. In Kůrková, V., Manolopoulos, Y., Hammer, B., Iliadis, L., and Maglogiannis, I. (eds.), *Artificial Neural Networks and Machine Learning – ICANN 2018*, pp. 412–422, Cham, 2018. Springer International Publishing.
- Stokes, J. M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N. M., MacNair, C. R., French, S., Carfrae, L. A., Bloom-Ackerman, Z., et al. A deep learning approach to antibiotic discovery. *Cell*, 180(4):688–702, 2020.
- Tran, D., Vafa, K., Agrawal, K., Dinh, L., and Poole, B. Discrete flows: Invertible generative models of discrete data. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32, pp. 14719–14728. Curran Associates, Inc., 2019.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., and Polosukhin, I. Attention is all you need. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 30, pp. 5998–6008. Curran Associates, Inc., 2017.

- Wang, Z., Liu, M., Luo, Y., Xu, Z., Xie, Y., Wang, L., Cai, L., and Ji, S. Advanced graph and sequence neural networks for molecular property prediction and drug discovery. *arXiv preprint arXiv:2012.01981*, 2020.
- Weininger, D. SMILES, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.
- You, J., Liu, B., Ying, Z., Pande, V., and Leskovec, J. Graph convolutional policy network for goal-directed molecular graph generation. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31, pp. 6410–6421. Curran Associates, Inc., 2018a.
- You, J., Ying, R., Ren, X., Hamilton, W., and Leskovec, J. GraphRNN: Generating realistic graphs with deep auto-regressive models. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 5708–5717, Stockholmsmässan, Stockholm Sweden, 2018b. PMLR.
- Yu, L., Zhang, W., Wang, J., and Yu, Y. SeqGAN: Sequence generative adversarial nets with policy gradient. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.
- Zang, C. and Wang, F. MoFlow: an invertible flow model for generating molecular graphs. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 617–626, 2020.