

GRAPHCG: UNSUPERVISED DISCOVERY OF STEERABLE FACTORS IN GRAPHS

Anonymous authors

Paper under double-blind review

ABSTRACT

Deep generative models have been widely developed for graph data such as molecular graphs and point clouds. Yet, much less investigation has been carried out on understanding the learned latent space of deep graph generative models. Such understandings can open up a unified perspective and provide guidelines for essential tasks like controllable generation. To this end, this work develops a method called GraphCG for unsupervised discovery of steerable factors in latent space of deep graph generative models. We first examine the representation space of the recent deep generative models trained for graph data, and observe that the learned representation space is not perfectly disentangled. Thus, our method is designed for discovering steerable factors of graph data in a model-agnostic and task-agnostic manner. Specifically, GraphCG learns the semantic-rich directions via maximizing the corresponding mutual information, where the edited graph along the same direction will possess certain steerable factors. We conduct experiments on two types of graph data, molecular graphs and point clouds. Both the quantitative and qualitative results show the effectiveness of GraphCG for discovering steerable factors.

1 INTRODUCTION

Graph is a general format for many real-world data. For instance, molecules can be treated as graphs [10, 14] where the chemical atoms and bonds correspond to the topological nodes and edges respectively. Processing point clouds as graphs is also a popular strategy [53, 58], where points are viewed as nodes and edges are built among the nearest neighbors. Many existing works on deep generative models (DGMs) focus on modeling the graph data and improving the synthesis quality. However, understanding DGMs on graph and their learned representations has been much less explored, which may hinder the development of important applications like the controllable generation (also referred to as the data editing) and the discovery of interpretable data structure.

The graph controllable generation task refers to modifying the steerable factors of graph so as to obtain graphs with desired properties [9, 43]. This is an important task in many applications, but traditional methods (*e.g.*, manual editing) possess certain inherent limitations under certain circumstances. A classic example is molecule editing, which aims at modifying the substructures of molecules [38] and can relate to some key tactics in drug discovery like functional group change [13] and scaffold hopping [2, 23]. This is a routine task in pharmaceutical companies, yet, relying on domain experts for manual editing can be subjective or biased [9, 15]. Different from previous works, this paper aims to explore the unsupervised graph editing with DGMs. It can act as a good complementary module to conventional methods and bring many crucial benefits: (1) It enables the efficient graph editing in the large-scale setting. (2) It alleviates the requirements for extensive domain knowledge for factor change labeling. (3) It provides another perspective for editing preference, which reduces biases from the domain experts.

One core property relevant to the general unsupervised data editing using DGMs is the disentanglement. While there does not exist a widely-accepted definition of disentanglement, the key intuition [36] is that a disentangled representation should separate the distinct, informative, and steerable factors of variations in the data. Thus, the controllable generation task would become trivial with the disentangled DGMs as the backbone. Such a disentanglement assumption has been widely used in generative modeling on the image data, *e.g.*, β -VAE [19] learns disentangled representation by forcing the representation to be close to an isotropic unit Gaussian. However, it may introduce extra constraints on the formulations and expressiveness of DGMs [11, 19, 47, 60].

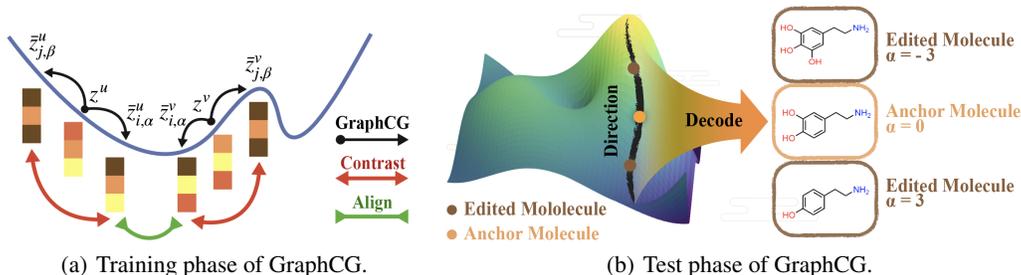


Figure 1: (a) The training phase. Given two latent codes z^u and z^v , we edit the four latent representations along i -th and j -th direction with step size α and β respectively. The goal of GraphCG is to align the positive pair $(\bar{z}_{i,\alpha}^u$ and $\bar{z}_{i,\alpha}^v)$, and contrast them with $\bar{z}_{j,\beta}^u$ and $\bar{z}_{j,\beta}^v$ respectively. (b) The test phase. We will first sample an anchor molecule, and adopt the learned directions in the training phase for editing. With step size $\alpha \in [-3, 3]$, we can generate a sequence of molecules. Specifically, after decoding, there is a functional group change shown up: the number of hydroxyl groups decreases along the sequence in the decoded molecules.

For graph data, one crucial question remains: *is the latent representation space learned from DGMs on graph data disentangled?* In image generation, prior work [36] shows that without inductive bias, the latent space learned by VAEs is not guaranteed to be disentangled. However, the disentanglement property of graph DGMs is much less explored. In Sec. 3, we first study the latent space of DGMs on two typical graph data (molecular graphs and point clouds), and empirically illustrate that the learned space is not perfectly disentangled. This observation then raises the second question: *Given a pretrained DGM with not perfectly disentangled latent space, is there a flexible framework enabling the graph controllable generation in an unsupervised manner?* To tackle this, we propose a model-agnostic and task-agnostic framework called GraphCG for unsupervised graph controllable generation. GraphCG has two main phases, as illustrated in Fig. 1. During the training phase (Fig. 1(a)), GraphCG starts with the assumption that the steerable directions can be learned by maximizing the mutual information (MI) among the semantic directions. We formulate GraphCG with an energy-based model (EBM), which provides a large family of solutions. Then during the test phase, with the learned semantic directions, we can carry out the editing task by moving along the direction with certain step sizes. As the example illustrated in Fig. 1(b), the molecular structure (hydroxyl group) changes consistently along the editing sequence. For evaluation, we visually verify the learned semantic directions on both types of graph data. Further for the molecular graphs, we propose a novel evaluation metric called sequence monotonic ratio (SMR) to measure the output sequences.

We summarize our contributions as follows: (1) We conduct an empirical study on the disentanglement property of three pretrained deep generative models (DGMs) on two types of graph data, molecular graphs and point clouds. We find that the latent space of these pretrained graph DGMs is not perfectly disentangled. (2) We propose a model-agnostic and task-agnostic method called GraphCG for the unsupervised graph controllable generation. GraphCG aims at learning the semantic directions by maximizing their corresponding mutual information, and its outputs are sequences of graphs. (3) We qualitatively evaluate the proposed methods on two types of graph data, molecular graphs and point clouds. Besides, the quantitative results further show the clear improvement over the baselines.

Related work. Recent works leverage the DGMs for various controllable generation tasks [5, 61], where the inherent assumption is that the learned latent representations encode rich semantics, and thus traversal in the latent space can help steer factors of data [17, 26, 52]. Among them, one research direction [40, 52] is using supervised signals to learn the semantic-rich directions, and most works on editing the graph data focus on the supervised setting [27, 57, 64]. However, these approaches can not be applied to many realistic scenarios where extracting the supervised labels is difficult. Another research line [17, 45, 51] considers discovering the latent semantics in an unsupervised manner, but most unsupervised methods are designed to be either model-specific or task-specific, making them not directly applicable to the graph data. More comprehensive discussion is in Appendix B.

2 BACKGROUND AND PROBLEM FORMULATION

Graph and deep generative models (DGMs). Each graph data (including nodes and edges) is denoted as $x \in \mathcal{X}$, where \mathcal{X} is the data space, and DGMs learn the data distribution, *i.e.*, $p(x)$. Our

proposed graph editing method (GraphCG) is model-agnostic, so we briefly introduce the mainstream DGMs for graph data as below. Variational auto-encoder (VAE) [19, 31] measures a variational lower bound of $p(\mathbf{x})$ by introducing a proposal distribution; flow-based model [8, 46] constructs reversible encoding functions such that the data distribution can be deterministically mapped to a prior distribution. Note that these mainstream DGMs, either explicitly or implicitly, contain an encoder ($f(\cdot)$) and a decoder ($g(\cdot)$) parameterized by neural networks as:

$$\mathbf{z} = f(\mathbf{x}), \quad \mathbf{x}' = g(\mathbf{z}), \quad (1)$$

where $\mathbf{z} \in \mathcal{Z}$ is the latent representation, \mathcal{Z} is the latent space, and \mathbf{x}' is the reconstructed output graph. Since in the literature [51, 52], people also call latent representations as latent codes or latent vectors, in what follows, we will use these terms interchangeably.

Semantic direction and step size. In the latent space \mathcal{Z} , we assume there exist D semantically meaningful direction vectors, *i.e.*, \mathbf{d}_i with $i \in \{0, 1, \dots, D-1\}$. There is also a scalar variable, step size α , which controls the degree to edit the sampled data with desired steerable factors (as will be introduced below), and we follow the prior work [51] on taking $\alpha \in [-3, 3]$. Each direction corresponds to one or multiple factors, such that by editing the latent vector \mathbf{z} with direction \mathbf{d}_i and step size α , the reconstructed graph will be augmented with the desired factors, leading to certain structural changes.

Steerable factors. The steerable factors are attributes of DGMs, which usually refer to the semantic information of data that we can explicitly discover from the pretrained DGMs. In this work, we focus on the steerable factors of graph data, which are data- and task-specific. Yet, there is one category of factors that is commonly shared among all the graph data: the **structure** information. Concretely, these steerable factors can be the functional groups or fragments in molecular graphs and shapes or sizes in point clouds. In Appendix C, we provide a detailed description of these steerable factors.

Problem formulation: graph controllable generation. Given a pretrained DGM (*i.e.*, the encoder and decoder are fixed), our goal is to learn the most semantically rich directions (\mathbf{d}_i) in the latent space \mathcal{Z} . Then for each latent code \mathbf{z} , with the i -th semantic direction and a step size α , we can get an edited latent vector $\bar{\mathbf{z}}_{i,\alpha}$ and edited data $\bar{\mathbf{x}}'$, as:

$$\mathbf{z} = f(\mathbf{x}), \quad \bar{\mathbf{z}}_{i,\alpha} = h(\mathbf{z}, \mathbf{d}_i, \alpha), \quad \bar{\mathbf{x}}' = g(\bar{\mathbf{z}}_{i,\alpha}), \quad (2)$$

where \mathbf{d}_i and $h(\cdot)$ are the edit direction and edit function that we want to learn. We expect that $\bar{\mathbf{z}}_{i,\alpha}$ can inherently possess certain steerable factors, which can then be reflected in the graph structure of $\bar{\mathbf{x}}'$.

Energy-based model (EBM). EBM is a flexible framework for distribution modeling:

$$p(\mathbf{x}) = \frac{\exp(-E(\mathbf{x}))}{A} = \frac{\exp(-E(\mathbf{x}))}{\int_{\mathcal{X}} \exp(-E(\mathbf{x}))d\mathbf{x}}, \quad (3)$$

where $E(\cdot)$ is the energy function and A is the partition function. In EBM, the bottleneck is the estimation of partition function A : it is commonly intractable due to the high cardinality of \mathcal{X} . Various methods have been proposed to handle this issue, including but not limited to contrastive divergence [20], noise-contrastive estimation [4, 16], and score matching [24, 54, 56].

3 DISENTANGLEMENT OF LATENT REPRESENTATION

In this section, we quantify the degree of disentanglement of the existing DGMs for graph data. In specific, we adopt six disentanglement measures, and the observed low disentanglement scores show that compared to DGMs for image data, such as StyleGANs [28, 29], the attributes in latent space of graph DGMs are not well disentangled.

The key intuition [36] behind disentanglement is that a disentangled representation space should separate the distinct, informative, and steerable factors of variations in the data. In other words, each latent dimension of the disentangled representation corresponds to one or multiple factors. Therefore, the change of the disentangled dimension can lead to the consistent change in the corresponding factors of the data. This good property has become a foundational assumption in many existing controllable generation methods [17, 51, 52].

Is the latent space of graph DGMs disentangled? In image generation, [36] shows that without inductive bias, the representation learned by VAEs is not perfectly disentangled. To verify if this claim is also valid for the mainstream DGMs on graphs, we conduct the following experiment.

Table 1: The six disentanglement metrics on three pretrained DGMs and two graph types. All measures range from 0 to 1, and higher scores mean more disentangled representation.

Graph Type	Model	Dataset	BetaVAE \uparrow	FactorVAE \uparrow	MIG \uparrow	DCI \uparrow	Modularity \uparrow	SAP \uparrow
Molecular Graph	MoFlow	ZINC250K	0.260	0.175	0.031	0.953	0.620	0.009
	HierVAE	ChEMBL	0.178	0.165	0.022	0.114	0.606	0.026
Point Cloud	PointFlow	Airplane	0.022	0.025	0.029	0.160	0.745	0.022

Steerable factors and experiments on disentanglement measure. There have been a series of works exploring the disentanglement of the latent space in DGMs, and here we take six widely-used ones: BetaVAE [19], FactorVAE [30], MIG [6], DCI [11], Modularity [47], and SAP [32]. Each measure has its own bias, and we put a detailed comparison in Appendix C. Meanwhile, they all share the same high-level idea: given the latent representation from a pretrained DGM, they are proposed to measure how predictive it is to certain steerable factors.

To adapt them to our setting, first we need to extract the steerable factors in graph, which requires the domain knowledge. For instance, in molecular graphs, we can extract some special substructures named fragments or functional groups. These substructures can be treated as steerable factors since they are the key components of the molecules and are closely related to certain molecular properties [50]. We use RDKit [33] to extract 9 most distinguishable fragments as steerable factors for disentanglement measurement. For point clouds, we use PCL tool [48] to extract 75 VFH descriptors [49] as steerable factors, which depicts the geometries and viewpoints accordingly.

Then for measuring the disentanglement, we consider six metrics on two data types with three backbone models. All the metric values range from 0 to 1, and the higher the value, the more disentangled the DGM is. According to Table 1, We can observe that the most of disentanglement scores are quite low, except the DCI [11] on MoFlow. Thus, we can draw the conclusion that generally these graph DGMs are not perfectly disentangled. More details of this experiment (the steerable factors on two data types and six disentanglement metrics) can be found in Appendix C.

4 OUR METHOD

The analysis in Sec. 3 naturally raises the next research question: *given a not well-disentangled representation space, is there a flexible way to do the graph data editing?* The answer is positive. We propose GraphCG, a flexible model-agnostic and task-agnostic framework to learn the semantic directions in an unsupervised manner. It starts with the assumption that the latent representations edited with the same semantic direction and step size should possess similar information (corresponding to the factors) to certain degree, thus by maximizing the mutual information them, we can learn the most semantic-rich directions. Then we formulate this editing task as a density estimation problem with the energy-based model (EBM). As introduced in Sec. 2, EBM covers a broad range of solutions, and we further propose GraphCG-NCE by adopting the noise-contrastive estimation (NCE).

4.1 GRAPH CONTROLLABLE GENERATION WITH MUTUAL INFORMATION

The mutual information (MI) measures the non-linear dependency between variables. To adapt it to our setting, we set the editing condition as containing both the semantic directions and step sizes, and we assume that maximizing the MI between different conditions can maximize the shared information within each condition. The pipeline is as follows.

We first sample two codes in the latent space, z^u and z^v . Then we pick up the i -th semantic direction and one step size α to obtain the edited latent points in \mathcal{Z} . The corresponding points are as

$$\bar{z}_{i,\alpha}^u = h(z^u, d_i, \alpha), \quad \bar{z}_{i,\alpha}^v = h(z^v, d_i, \alpha). \quad (4)$$

Under our assumption, we expect that these two edited points share certain information with respect to the steerable factors. Thus, we want to maximize the MI between $\bar{z}_{i,\alpha}^u$ and $\bar{z}_{i,\alpha}^v$. Since the MI is intractable to compute, we adopt the EBM lower bound [35] as:

$$\mathcal{L}_{\text{MI}}(\bar{z}_{i,\alpha}^u, \bar{z}_{i,\alpha}^v) = \frac{1}{2} \mathbb{E}_{p(\bar{z}_{i,\alpha}^u, \bar{z}_{i,\alpha}^v)} [\log p(\bar{z}_{i,\alpha}^u | \bar{z}_{i,\alpha}^v) + \log p(\bar{z}_{i,\alpha}^v | \bar{z}_{i,\alpha}^u)]. \quad (5)$$

The detailed derivation is in Appendix D. Till this step, we have transformed the graph data editing task into the summation of two conditional log-likelihoods estimation problem.

4.2 GRAPHCG WITH ENERGY-BASED MODEL

Following Eq. (5), maximizing the MI between $I(\bar{z}_{i,\alpha}^u; \bar{z}_{i,\alpha}^v)$ is equivalent to estimating the summation of two conditional log-likelihoods. We then model them using two conditional EBMs. Because these two views are in the mirroring direction, we may as well take one for illustration. For example, for the first conditional log-likelihood, we can model it with EBM as:

$$p(\bar{z}_{i,\alpha}^u | \bar{z}_{i,\alpha}^v) = \frac{\exp(-E(\bar{z}_{i,\alpha}^u, \bar{z}_{i,\alpha}^v))}{\int \exp(-E(\bar{z}_{i,\alpha}^{u'}, \bar{z}_{i,\alpha}^v)) d\bar{z}_{i,\alpha}^{u'}} = \frac{\exp(f(\bar{z}_{i,\alpha}^u, \bar{z}_{i,\alpha}^v))}{A_{ij}}, \quad (6)$$

where $E(\cdot)$ is the energy function, A_{ij} is the intractable partition function, and $f(\cdot)$ is the negative energy. The energy function can be quite flexible, and for simplicity, we use the dot-product:

$$f(\bar{z}_{i,\alpha}^u, \bar{z}_{i,\alpha}^v) = \langle h(\mathbf{z}^u, \mathbf{d}_i, \alpha), h(\mathbf{z}^v, \mathbf{d}_i, \alpha) \rangle, \quad (7)$$

where $h(\cdot)$ is the editing function introduced in Eq. (2). Similarly for the other conditional log-likelihood term, and the objective becomes:

$$\mathcal{L}_{\text{GraphCG}} = \mathbb{E} \left[\log \frac{\exp(f(\bar{z}_{i,\alpha}^u, \bar{z}_{i,\alpha}^v))}{A_{ij}} + \log \frac{\exp(f(\bar{z}_{i,\alpha}^v, \bar{z}_{i,\alpha}^u))}{A_{ji}} \right]. \quad (8)$$

With Eq. (8), we are able to learn the semantically meaningful direction vectors. We name this unsupervised graph controllable generation framework as GraphCG. In specific, GraphCG utilizes EBM for estimation, which yields a wide family of solutions. Next we will introduce an intuitive solution.

4.3 GRAPHCG WITH NOISE CONTRASTIVE ESTIMATION

We solve Eq. (8) using the noise contrastive estimation (NCE) [16]. The high-level idea of NCE is to transform the density estimation problem into a binary classification problem that distinguishes if the data comes from the introduced noise distribution or from the true distribution. NCE has been widely explored for solving EBM [55], and we adopt it as GraphCG-NCE by optimizing:

$$\begin{aligned} \mathcal{L}_{\text{GraphCG-NCE}} = & - \left(\mathbb{E}_{p_n(\bar{z}_{j,\beta}^u | \bar{z}_{i,\alpha}^v)} [\log(1 - \sigma(f(\bar{z}_{j,\beta}^u, \bar{z}_{i,\alpha}^v)))] + \mathbb{E}_{p_{\text{data}}(\bar{z}_{i,\alpha}^u | \bar{z}_{i,\alpha}^v)} [\log \sigma(f(\bar{z}_{i,\alpha}^u, \bar{z}_{i,\alpha}^v))] \right) \\ & + \left(\mathbb{E}_{p_n(\bar{z}_{j,\beta}^v | \bar{z}_{i,\alpha}^u)} [\log(1 - \sigma(f(\bar{z}_{j,\beta}^v, \bar{z}_{i,\alpha}^u)))] + \mathbb{E}_{p_{\text{data}}(\bar{z}_{i,\alpha}^v | \bar{z}_{i,\alpha}^u)} [\log \sigma(f(\bar{z}_{i,\alpha}^v, \bar{z}_{i,\alpha}^u))] \right), \end{aligned} \quad (9)$$

where p_{data} is the data distribution and p_n is the noise distribution (derivations are in Appendix D). Recall that the latent pairs are given, and the noise distribution is on the semantic directions and step sizes. In specific, the step sizes ($\alpha \neq \beta$) are randomly sampled from $[-3, 3]$, and the latent direction indices ($i \neq j$) are randomly sampled from $\{0, 1, \dots, D-1\}$. The objective Eq. (9) is for one latent code pair, and we will take the expectation of it over all the pairs from the dataset. Besides, we would like to consider extra similarity and sparsity constraints as:

$$\mathcal{L}_{\text{sim}} = \mathbb{E}_{i,j} [\text{sim}(d_i, d_j)], \quad \mathcal{L}_{\text{sparsity}} = \mathbb{E}_i [\|d_i\|], \quad (10)$$

where $\text{sim}(\cdot)$ is the similarity function between two latent directions, and we use the dot product. By minimizing these two regularization terms, we can make the learned semantic directions more diverse and sparse. Putting them together, the final objective function is:

$$\mathcal{L} = c_1 \cdot \mathbb{E}_{u,v} [\mathcal{L}_{\text{GraphCG-NCE}}] + c_2 \cdot \mathcal{L}_{\text{sim}} + c_3 \cdot \mathcal{L}_{\text{sparsity}}, \quad (11)$$

where c_1, c_2, c_3 are coefficients, and we treat them as three hyperparameters (check Appendix E). The above pipeline is illustrated in Fig. 1, and for the next we will discuss certain key modules.

Latent pairs, positive and negative views. We consider two options in designing the latent pairs. (1) *Perturbation (GraphCG-P)* is that for each latent variable $z \in \mathcal{Z}$, we apply 2 perturbations (e.g., adding Gaussian noise) on z to get 2 perturbed latent codes as z^u and z^v respectively. (2) *Random sampling (GraphCG-R)* is that we encode two randomly sampled data points from the empirical data distribution as z^u and z^v respectively. Perturbation is one of the widely-used strategies [28] for data augmentation, and random sampling has been widely used in the NCE [55] literature. Then we can define the positive and negative pairs in GraphCG-NCE, where the goal is to align the positives and contrast the negatives. As described in Eq. (9), the positive pairs are latent pairs moving with the same semantic direction and step size, while the negative pairs are the edited latent codes with different semantic directions and/or step sizes.

Semantic direction modeling. We first randomly draw a *basis vector* e_i , and then model the semantic direction d_i as $d_i = \text{MLP}(e_i)$, where $\text{MLP}(\cdot)$ is the multi-layer perceptron network.

Design of editing function. Given the semantic direction and two views, the next task is to design the editing function $h(\cdot)$ in Eq. (2). Since our proposed GraphCG is flexible, and the editing function determines the energy function Eq. (7), we consider both the linear and non-linear editing functions as:

$$\bar{z}_i = z + \alpha \cdot d_i, \quad \bar{z}_i = z + \alpha \cdot d_i + \text{MLP}(z \oplus d_i \oplus [\alpha]), \quad (12)$$

where \oplus is the concatenation of two vectors. Noticing that for the non-linear case, we are adding an extra term by mapping from the latent code, semantic direction, and step-size simultaneously. We expect that this could bring in more modeling expressiveness in the editing function. For more details, *e.g.*, the ablation study to check the effect on the design of the views and editing functions, please refer to [Appendices F and G](#), while more potential explorations are left for future work.

4.4 IMPLEMENTATIONS

During training, the goal of GraphCG is to learn semantically meaningful direction vectors together with an editing function in the latent space, as in Algorithm 1. Then we need to manually annotate the semantic directions with respect to the corresponding factors, using certain post-training evaluation metrics. Finally for the test phase, provided with the pretrained DGM on graph and a selected semantic direction (together with a step size), we can sample a molecule and use GraphCG for editing, as described in Eq. (2). The detailed algorithm is illustrated in Algorithm 2. Next, we highlight several key concepts in GraphCG and briefly discuss the differences to other related methods.

NCE and contrastive representation learning. GraphCG-NCE is applying EBM-NCE, which is essentially a contrastive learning method, and another dominant contrastive loss is the InfoNCE [42]. We summarize their relations as below. (1) Both contrastive methods are doing the same thing: align the positive pairs and contrast the negative pairs. (2) EBM-NCE [18, 35] has been found to outperform InfoNCE on the certain graph

applications like representation learning. (3) What we want to propose here is a flexible framework. Specifically, EBM provides a more general framework by designing the energy functions, and EBM-NCE is just one effective solution. Other promising directions include the denoising score matching or denoising diffusion model [56], while InfoNCE lacks such nice extensibility attribute.

GraphCG and self-supervised learning (SSL). GraphCG shares certain similarities with the self-supervised learning (SSL) method, however there are some inherent differences, as summarized below. (1) SSL aims at learning the data representation [by operating data augmentation on the data space, such as node addition and edge deletion](#). GraphCG aims at learning the semantically meaningful directions [by editing on the latent space \(the representation function is pretrained and fixed\)](#). (2) Based on the first point, SSL aims at using different data points as the negative samples. GraphCG, on the other hand, is using different directions and step-sizes as negatives. Namely, SSL is learning data representation in the inter-data level, and GraphCG is learning the semantic directions in the inter-direction level.

Algorithm 1 Training Phase of GraphCG

- 1: **Input:** Given a pretrained generative model, $f(\cdot)$ and $g(\cdot)$.
 - 2: **Output:** Learned direction vector d_i and function $h(\cdot)$.
 - 3: Select two data points, $z^u, z^v \in \mathcal{Z}$.
 - 4: **for** each step size α and each direction i **do**
 - 5: Set $\bar{z}_{i,\alpha}^u = h(z^u, d_i, \alpha)$.
 - 6: Set $\bar{z}_{i,\alpha}^v = h(z^v, d_i, \alpha)$.
 - 7: Assign positive to pair $(\bar{z}_{i,\alpha}^u, \bar{z}_{i,\alpha}^v)$.
 - 8: **for** step size $\beta \neq \alpha$ and direction $j \neq i$ **do**
 - 9: Set $\bar{z}_{j,\beta}^u = h(z^u, d_j, \beta)$.
 - 10: Set $\bar{z}_{j,\beta}^v = h(z^v, d_j, \beta)$.
 - 11: Assign negative to pair $(\bar{z}_{i,\alpha}^u, \bar{z}_{j,\beta}^v)$.
 - 12: Assign negative to pair $(\bar{z}_{j,\beta}^u, \bar{z}_{i,\alpha}^v)$.
 - 13: **end for**
 - 14: Do SGD w.r.t. GraphCG in Eq. (11).
 - 15: **end for**
-

Algorithm 2 Test Phase of GraphCG

- 1: **Input:** Given a pre-trained generative model ($f(\cdot)$ and $g(\cdot)$), a learned direction vector d .
 - 2: **Output:** A sequence of edited graphs.
 - 3: Sample a molecule with DGM or from a large molecule pool.
 - 4: Encode the molecule to get a latent code z .
 - 5: **for** step size $\alpha \in [-3, 3]$ **do**
 - 6: Do graph edit in the latent space to get $\bar{z}_{i,\alpha} = h(z, d, \alpha)$.
 - 7: Decode to the graph space with $\bar{x}' = g(\bar{z}_{i,\alpha})$.
 - 8: **end for**
 - 9: Output is thus a sequence of edited graphs, $\{\bar{x}'\}$.
-

Output sequence in the discrete space. Recall that during inference time (Algorithm 2), GraphCG takes a DGM and the learned semantic direction to output a sequence of edited graphs. Comparing to the vision domain, where certain models [51, 52] have proven their effectiveness in many tasks, the backbone models in the graph domain have limited discussions. This is challenging because the graph data is in a discrete and structured space, and the evaluation on such space is non-trivial. Meanwhile, GraphCG essentially provides another way to testify the quality of graph generation models. We would like to leave this for future exploration.

5 EXPERIMENTS

In this section, we show both the qualitative and quantitative results of GraphCG on two types of graph data: molecular graphs and point clouds. Due to the page limit, We put the experiment and implementation details in Appendix E. For reproducibility, the code will be public in the near future.

5.1 GRAPH DATA: MOLECULAR GRAPHS

Backbone DGMs. We consider two state-of-the-art DGMs for molecular graph generation. MoFlow [65] is a flow-based generative model on molecules which adopts an invertible mapping between the input molecular graphs and a latent prior. HierVAE [27] is a hierarchical VAE model which encodes and decodes molecule atoms and motifs in a hierarchical manner. Besides, the pretrained checkpoints are also provided, on ZINC250K [25] and ChEMBL [37] dataset respectively.

Editing sequences and anchor molecule. As discussed in Sec. 4, the output of the inference in GraphCG is a sequence of edited molecules with the i -th semantic direction, $\{\bar{x}'\}_i$. We first randomly generate a molecule using the backbone DGMs (without the editing operation), and we name such molecule as the anchor molecule, \bar{x}^* . Then we take 21 step sizes from -3 to 3, with interval 0.3, to obtain a sequence of 21 molecules following Eq. (2). Note that the edited molecule with step size 0 under the linear editing function is the same as the anchor molecule, *i.e.*, \bar{x}^* .

Change of structure factors and evaluation metrics. We are interested in the change of the graph structure (the steerable factors) along the output sequence edited with the i -th semantic direction. To evaluate the structure change, we apply the Tanimoto similarity between each output molecule and the anchor molecule. Besides, for the ease of evaluating the monotonicity, we apply a Tanimoto similarity transformation on the output molecules with positive step sizes by taking the deduction from 2. We call this calibrated Tanimoto similarity (CTS) sequence, marked as $\{s(\bar{x}')\}_i$. An illustration is shown in Fig. 2. Further, we propose a metric called Sequence Monotonic Ratio (SMR), $\phi_{\text{SMR}}(\gamma, \tau)_i$, which measures the monotonic ratio of M generated sequences edited with the i -th direction. It has two arguments: the diversity threshold γ constrains the minimum number of distinct molecules, and the tolerance threshold τ controls the non-monotonic tolerance ratio along each sequence.

Evaluating the diversity of semantic directions. SMR can evaluate the monotonic ratio of output sequences generated by one direction. To better illustrate that GraphCG is able to learn multiple directions with diverse semantic information, we also consider taking the average of top- K SMR to reveal that all the best K directions are semantically meaningful, as in Eq. (15).

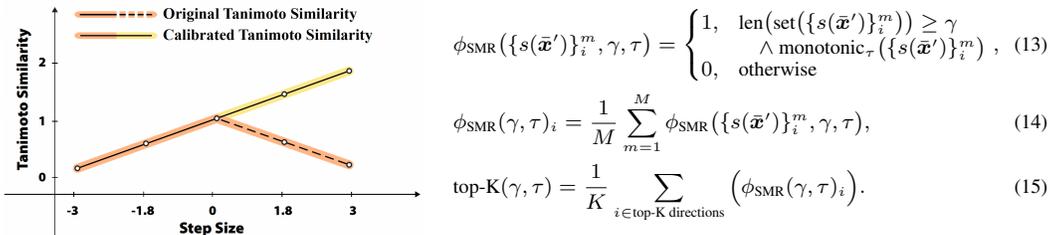


Figure 2: This shows the sequence monotonic ratio (SMR) on calibrated Tanimoto similarity (CTS). Eqs. (13) and (14) are the SMR on each sequence and each direction respectively, where M is the number of generated sequences for the i -th direction and $\{s(\bar{x}')\}_i^m$ is the CTS of the m -th generated sequence with the i -th direction. Eq. (15) is the average of top- K SMR on D directions. More details are in Appendix F.

Table 2: This table lists the sequence monotonic ratio (SMR, %) on calibrated Tanimoto similarity (CTS) with respect to the top-1 and top-3 directions. The best performances are marked in **bold**.

Model	Dataset	diversity γ	Tanimoto top-1				Tanimoto top-3			
			3		4		3		4	
			0	0.2	0	0.2	0	0.2	0	0.2
MoFlow	ZINC250k	Random	23.0	25.0	12.0	15.0	22.0	24.0	11.0	13.7
		Variance	24.0	28.0	12.0	16.0	20.0	25.0	10.0	15.0
		SeFa [51]	4.0	4.0	0.0	0.0	3.3	3.3	0.0	0.0
		DisCo [45]	7.0	14.0	2.0	8.0	5.3	11.7	2.0	7.7
		GraphCG-P	32.0	34.0	16.0	18.0	29.0	31.0	13.7	16.3
		GraphCG-R	25.0	26.0	11.0	14.0	22.0	24.3	10.3	13.3
HierVAE	ChEMBL	Random	14.0	45.0	14.0	43.0	10.0	42.3	9.3	41.7
		Variance	23.0	59.0	19.0	55.0	18.3	52.7	15.7	50.3
		SeFa [51]	4.0	41.0	4.0	41.0	2.3	36.0	2.3	36.0
		GraphCG-P	40.0	73.0	32.0	65.0	36.0	64.3	26.3	57.7
		GraphCG-R	42.0	67.0	30.0	55.0	38.0	62.3	28.7	53.7

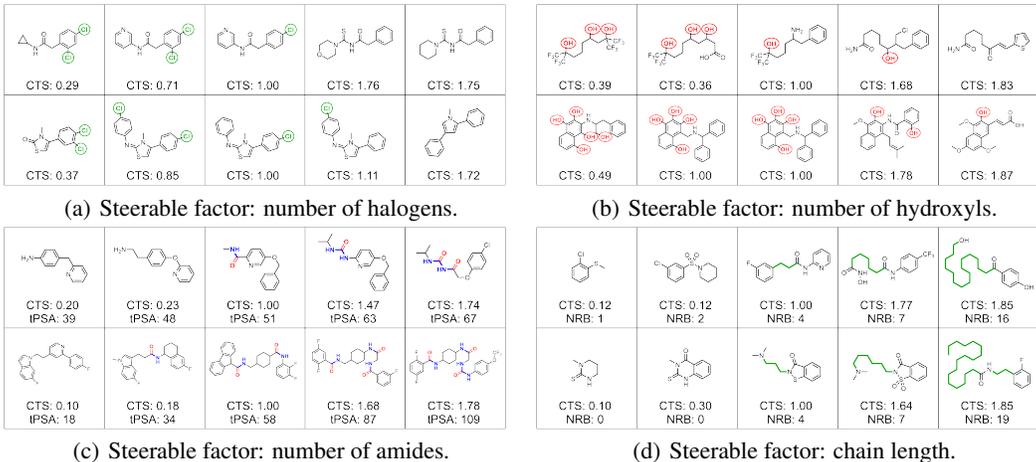


Figure 3: GraphCG for molecular graph editing. We visualize the output molecules and CTS on four directions with two sequences each, where each sequence consists of five steps. The center point is the anchor molecule, and the other four points correspond to step size with -3 , -1.8 , 1.8 , and 3 respectively. Figs. 3(a) to 3(c) show how functional groups in the molecules can be viewed as the steerable factors as they change along the sequence, such as halogen atoms, hydroxyl groups and amides. Fig. 3(d) illustrates the effect on the steerable factor on the length of flexible chains in the molecules. Notably, certain properties change with molecular structures accordingly, like topological polar surface area (tPSA) and number of rotatable bonds (NRB).

Baselines. For baselines, we consider four unsupervised editing methods. (1) The first is Random. It randomly samples a normalized random vector in the latent space as the semantic direction. (2) The second one is Variance. We analyze the variance on each dimension of the latent space, and select the highest one with one-hot encoding as the semantic direction. (3) The third one is SeFa [51]. It first decomposes the latent space into lower dimensions using PCA, and then takes the most principle components (eigenvectors) as the semantic-rich direction vectors. (4) The last one is DisCo [45]. It maps each latent code back to the data space, followed with an encoder for contrastive learning. Namely, it requires the backbone DGMs to be end-to-end and is infeasible for HierVAE.

Quantitative results. We take $D = 10$ to train GraphCG, and the optimal results on 100 sampled sequences are reported in Table 2. We can observe that GraphCG can show consistently better structure change with both top-1 and top-3 directions. This can empirically prove the effectiveness of our proposed GraphCG. More comprehensive results are in Appendix F.

Analysis on steerable factors in molecules: functional group change. For visualization, we sample 8 molecular graph sequences along 4 selected directions in Fig. 3, and the backbone DGM is HierVAE pretrained on ChEMBL. The CTS holds good monotonic trend, and each direction shows certain unique changes in the molecular structures, *i.e.*, the steerable factors in molecules. Some structural changes are reflected in molecular properties. We expand all the details below. In Fig. 3(a) and Fig. 3(b), the number of halogen atoms and hydroxyl groups (in alcohols and phenols) in the

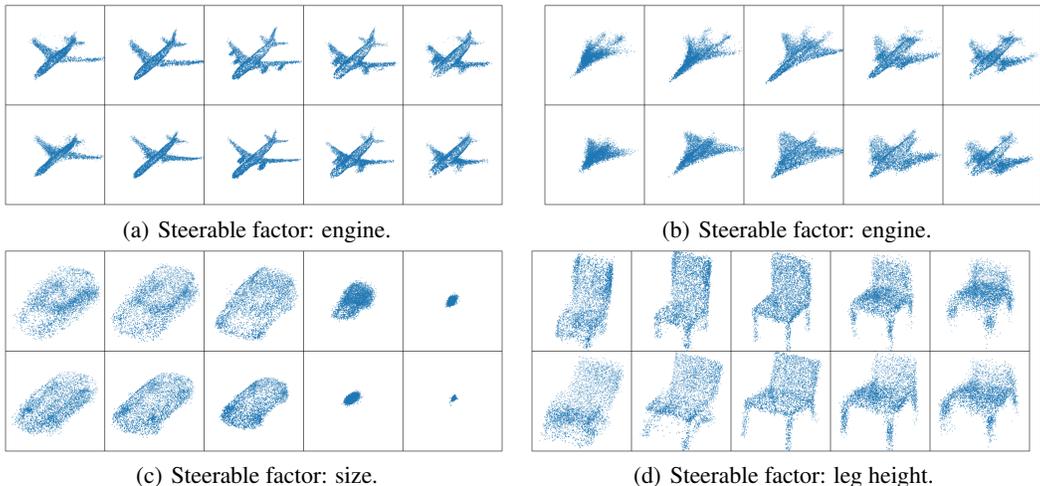


Figure 4: GraphCG for point clouds editing. We show four editing sequences, where each sequence consists of five point clouds, and the center one is the anchor point clouds, *i.e.*, with step size 0. The other four point clouds correspond to step size with -3, -1.8, 1.8, and 3, respectively. Fig. 4(a) and Fig. 4(b) refer the same semantic direction, and they are showing how to steer the factor engine: the number of engines will be decreased and increased with the negative (left) and positive (right) step size respectively. Similarly, Figs. 4(c) and 4(d) illustrate the effect on the steerable factors on the car size and the chair leg height.

molecules decrease from left to right, respectively. In Fig. 3(c), the number of amides in the molecules increases along the path. Because amides are polar functional groups, the topological polar surface area (tPSA) of the molecules also increase accordingly, which is a key molecular property for the prediction of drug transport properties, *e.g.*, permeability [12]. In Fig. 3(d), the flexible chain length, marked by the number of ethylene (CH_2CH_2) units, increases from left to right. Since the number of rotatable bonds (NRB) measures the molecular flexibility, it also increases accordingly [57].

5.2 GRAPH DATA: POINT CLOUDS

Backbone DGMs. We consider one of the latest DGMs on point clouds, PointFlow [62]. It is using the normalizing flow model for estimating the 3D point clouds distribution. Then we consider PointFlow pretrained on three datasets in ShapeNet [3]: Airplane, Car, and Chair. All point clouds are obtained by sampling points uniformly from the mesh surface.

Analysis on steerable factors in point clouds: shape change. To train GraphCG, we take $D = 10$ directions, and we sample 8 point cloud sequences along 3 directions for visualization in Fig. 4. More results are in Appendix G. It is observed that GraphCG can steer the shape of the point clouds, *e.g.*, the size of cars and the height of chair legs. We also find it interesting that GraphCG can steer more finger-trained factors, like modifying the number jet engines of airplanes in Figs. 4(a) and 4(b).

6 CONCLUSION AND DISCUSSION

In this work, we are interested in unsupervised graph editing. It is a well-motivated task for various real-world applications, and we discuss two mainstream data types: molecular graphs and point clouds. We start with exploring the latent space of mainstream deep generative models and propose GraphCG, a model-agnostic and task-agnostic unsupervised method for graph data editing. The key component of GraphCG is EBM, and we take the GraphCG-NCE as the solution for now. For the future work, we may as well extend it to more advanced solutions like denoising diffusion model [21].

One limitation of GraphCG (as well as the solutions to general unsupervised data editing) [17, 45, 51] is that we may need some post-training human selection (as shown in Algorithms 1 and 2) to select the most promising semantic vectors to steer factors. Another issue is the lack of open-sourced evaluation metrics. This requires both a deep understanding of the representation space of deep generative models and domain knowledge of the problem. For instance, activity cliff is a challenging task [22] for editing, while current measures fail to capture it. To set up constructive evaluation metrics can help augment our understandings from both the domain and technique perspectives. This is beyond the scope of our work, yet would be interesting to explore as a future direction.

ETHICS STATEMENT

We authors acknowledge that we have read and commit to adhering to the ICLR Code of Ethics.

REPRODUCIBILITY STATEMENT

To ensure the reproducibility of the empirical results, we provide the implementation details (hyperparameters, dataset specifications, pretrained checkpoints, etc.) in Sec. 5 and appendices C and E to G, and the source code will be released in the future. Besides, the complete derivations of equations and clear explanations are given in Sec. 4 and appendix D.

Specifically, we provide the details for reproducing the results:

- In Table 5, GraphCG-P with Eq. (26) and GraphCG-R with Eq. (24) are reported in Table 2.
- In Table 6, GraphCG-P with Eq. (25) and GraphCG-R with Eq. (25) are reported in Table 2.

For the visualization in Fig. 3, we take the GraphCG-P with Eq. (25), and the backbone generative model is HierVAE pretrained on ChEMBL. Further, we provide an anonymous link [here](#). In these CSV files:

- Direction 0 is the halogen fragment (data 4, 71).
- Direction 5 is the amide fragment (data 95, 61).
- Direction 6 is the chain length (data 57, 14).
- Direction 7 is the alcohol and phenol fragments (data 10, 8).

For the visualization in Fig. 4, we take the GraphCG-R with Eq. (25) on PointFlow, for all three datasets.

REFERENCES

- [1] Michael Arbel, Liang Zhou, and Arthur Gretton. Generalized energy based models. *arXiv:2003.05033*, 2020. 17
- [2] Hans-Joachim Böhm, Alexander Flohr, and Martin Stahl. Scaffold hopping. *Drug discovery today: Technologies*, 1(3):217–224, 2004. 1
- [3] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, et al. Shapenet: An information-rich 3d model repository. *arXiv:1512.03012*, 2015. 9
- [4] Tong Che, Ruixiang Zhang, Jascha Sohl-Dickstein, Hugo Larochelle, Liam Paull, et al. Your gan is secretly an energy-based model and you should use discriminator driven latent sampling. *arXiv:2003.06060*, 2020. 3
- [5] Hongming Chen, Ola Engkvist, Yinhai Wang, Marcus Olivecrona, and Thomas Blaschke. The rise of deep learning in drug discovery. *Drug discovery today*, 23(6):1241–1250, 2018. 2
- [6] Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *NeurIPS*, 2018. 4, 15
- [7] Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal neighbourhood aggregation for graph nets. *arXiv:2004.05718*, 2020. 14
- [8] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv:1410.8516*, 2014. 3
- [9] Jürgen Drews. Drug discovery: a historical perspective. *Science*, 287(5460):1960–1964, 2000. 1, 14
- [10] David Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, et al. Convolutional networks on graphs for learning molecular fingerprints. *arXiv:1509.09292*, 2015. 1, 14
- [11] Cian Eastwood and Christopher KI Williams. A framework for the quantitative evaluation of disentangled representations. In *ICLR*, 2018. 1, 4, 15
- [12] Peter Ertl, Bernhard Rohde, and Paul Selzer. Fast calculation of molecular polar surface area as a sum of fragment-based contributions and its application to the prediction of drug transport properties. *Journal of medicinal chemistry*, 43(20):3714–3717, 2000. 9, 21

- [13] Peter Ertl, Eva Altmann, and Jeffrey M McKenna. The most common functional groups in bioactive molecules and how their popularity has evolved over time. *Journal of medicinal chemistry*, 63(15): 8408–8418, 2020. 1
- [14] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *ICML*, 2017. 1, 14
- [15] Laurent Gomez. Decision making in medicinal chemistry: The power of our intuition. *ACS Medicinal Chemistry Letters*, 9(10):956–958, 2018. 1
- [16] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*, 2010. 3, 5
- [17] Erik Härkönen, Aaron Hertzman, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. In *NeurIPS*, 2020. 2, 3, 9, 14
- [18] Kaveh Hassani and Amir Hosein Khasahmadi. Contrastive multi-view representation learning on graphs. In *ICML*, 2020. 6
- [19] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, et al. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017. 1, 3, 4, 15
- [20] Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002. 3
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 9
- [22] Ye Hu and Bajorath Jurgen. Extending the activity cliff concept: structural categorization of activity cliffs and systematic identification of different types of cliffs in the chembl database. *Journal of chemical information and modeling*, 52(7):1806–1811, 2012. 9
- [23] Ye Hu, Dagmar Stumpfe, and Jurgen Bajorath. Recent advances in scaffold hopping: miniperspective. *Journal of medicinal chemistry*, 60(4):1238–1246, 2017. 1
- [24] Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005. 3
- [25] John J Irwin and Brian K Shoichet. Zinc: a free database of commercially available compounds for virtual screening. *Journal of chemical information and modeling*, 45(1):177–182, 2005. 7
- [26] Ali Jahanian, Lucy Chai, and Phillip Isola. On the "steerability" of generative adversarial networks. In *ICLR*, 2019. 2, 14
- [27] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Hierarchical generation of molecular graphs using structural motifs. In *ICML*, 2020. 2, 7, 14
- [28] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 3, 5, 14
- [29] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, et al. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 3
- [30] Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In *ICML*, 2018. 4, 15
- [31] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv:1312.6114*, 2013. 3
- [32] Abhishek Kumar, Prasanna Sattigeri, and Avinash Balakrishnan. Variational inference of disentangled latent concepts from unlabeled observations. In *ICLR*, 2018. 4, 15
- [33] Greg Landrum et al. RDKit: A software suite for cheminformatics, computational chemistry, and predictive modeling, 2013. 4
- [34] Shengchao Liu, Mehmet Furkan Demirel, and Yingyu Liang. N-gram graph: Simple unsupervised representation for graphs, with applications to molecules. *arXiv:1806.09206*, 2018. 14
- [35] Shengchao Liu, Hanchen Wang, Weiyang Liu, Joan Lasenby, Hongyu Guo, et al. Pre-training molecular graph representation with 3d geometry. In *ICLR*, 2022. 4, 6, 16, 17

- [36] Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, et al. Challenging common assumptions in the unsupervised learning of disentangled representations. In *ICML*, 2019. 1, 2, 3, 15
- [37] David Mendez, Anna Gaulton, A Patrícia Bento, Jon Chambers, Marleen De Veij, et al. ChEMBL: towards direct deposition of bioassay data. *Nucleic acids research*, 47(D1):D930–D940, 2019. 7
- [38] Zlatko Mihalić and Nenad Trinajstić. A graph-theoretical approach to structure-property relationships. *Journal of Chemical Education*, 69(9):701, 1992. 1
- [39] Andriy Mnih and Yee Whye Teh. A fast and simple algorithm for training neural probabilistic language models. *arXiv preprint arXiv:1206.6426*, 2012. 17
- [40] Weili Nie, Arash Vahdat, and Anima Anandkumar. Controllable and compositional generation with latent-space energy-based models. In *NeurIPS*, 2021. 2, 14
- [41] Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics*, 9(1):1–14, 2017. 14
- [42] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018. 6
- [43] Yael Pritch, Eitam Kav-Venaki, and Shmuel Peleg. Shift-map image editing. In *ICCV*, 2009. 1
- [44] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 14
- [45] Xuanchi Ren, Tao Yang, Yuwang Wang, and Wenjun Zeng. Do generative models know disentanglement? contrastive learning is all you need. *arXiv:2102.10543*, 2021. 2, 8, 9, 14
- [46] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *ICML*, 2015. 3
- [47] Karl Ridgeway and Michael C Mozer. Learning deep disentangled embeddings with the f-statistic loss. In *NeurIPS*, 2018. 1, 4, 15
- [48] Radu Bogdan Rusu and Steve Cousins. 3D is here: Point Cloud Library (PCL). In *ICRA*, 2011. 4, 15
- [49] Radu Bogdan Rusu, Gary Bradski, Romain Thibaux, and John Hsu. Fast 3d recognition and pose using the viewpoint feature histogram. In *IROS*, 2010. 4, 15
- [50] Paul G. Seybold, Michael May, and Ujjvala A. Bagal. Molecular structure: Property relationships. *Journal of Chemical Education*, 64(7):575, 1987. 4
- [51] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *CVPR*, 2021. 2, 3, 7, 8, 9, 14
- [52] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *IEEE TPAMI*, 2020. 2, 3, 7, 14
- [53] Weijing Shi and Raj Rajkumar. Point-gnn: Graph neural network for 3d object detection in a point cloud. In *CVPR*, 2020. 1, 14
- [54] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *arXiv:1907.05600*, 2019. 3
- [55] Yang Song and Diederik P Kingma. How to train your energy-based models. *arXiv:2101.03288*, 2021. 5, 17
- [56] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, et al. Score-based generative modeling through stochastic differential equations. *arXiv:2011.13456*, 2020. 3, 6
- [57] Daniel F Veber, Stephen R Johnson, Hung-Yuan Cheng, Brian R Smith, Keith W Ward, and Kenneth D Kopple. Molecular properties that influence the oral bioavailability of drug candidates. *Journal of medicinal chemistry*, 45(12):2615–2623, 2002. 2, 9, 21
- [58] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, et al. Dynamic graph cnn for learning on point clouds. *ACM ToG*, 2020. 1
- [59] Zichao Wang, Weili Nie, Zhuoran Qiao, Chaowei Xiao, Richard Baraniuk, and Anima Anandkumar. Retrieval-based controllable molecule generation. *arXiv preprint arXiv:2208.11126*, 2022. 14

- [60] Zongze Wu, Dani Lischinski, and Eli Shechtman. Stylespace analysis: Disentangled controls for stylegan image generation. In *CVPR*, 2021. 1
- [61] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, et al. Gan inversion: A survey. *arXiv:2101.05278*, 2021. 2
- [62] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, et al. Pointflow: 3d point cloud generation with continuous normalizing flows. In *ICCV*, 2019. 9, 14
- [63] Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley, Philipp Eiden, et al. Analyzing learned molecular representations for property prediction. *Journal of chemical information and modeling*, 59(8):3370–3388, 2019. 14
- [64] Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy network for goal-directed molecular graph generation. *Advances in neural information processing systems*, 31, 2018. 2, 14
- [65] Chengxi Zang and Fei Wang. Moflow: an invertible flow model for generating molecular graphs. In *KDD*, 2020. 7, 14