

Shape-conditioned 3D Molecule Generation via Equivariant Diffusion Models

Ziqi Chen¹, Bo Peng¹, Srinivasan Parthasarathy^{1,2}, Xia Ning^{1,2,3*},

¹Computer Science and Engineering, The Ohio State University, Columbus, OH 43210

²Translational Data Analytics Institute, The Ohio State University, Columbus, OH 43210

³Biomedical Informatics, The Ohio State University, Columbus, OH 43210

chen.8484@buckeyemail.osu.edu, peng.307@buckeyemail.osu.edu, srini@cse.ohio-state.edu, ning.104@osu.edu

Abstract

Ligand-based drug design aims to identify novel drug candidates of similar shapes with known active molecules. In this paper, we formulated an *in silico* shape-conditioned molecule generation problem to generate 3D molecule structures conditioned on the shape of a given molecule. To address this problem, we developed a translation- and rotation-equivariant shape-guided generative model ShapeMol. ShapeMol consists of an equivariant shape encoder that maps molecular surface shapes into latent embeddings, and an equivariant diffusion model that generates 3D molecules based on these embeddings. Experimental results show that ShapeMol can generate novel, diverse, drug-like molecules that retain 3D molecular shapes similar to the given shape condition. These results demonstrate the potential of ShapeMol in designing drug candidates of desired 3D shapes binding to protein target pockets.

Introduction

Generating novel drug candidates is a critical step in drug discovery to identify possible therapeutic solutions. Conventionally, this process is characterized based on knowledge and experience from medicinal chemists, and is resource- and time-consuming. Recently, computational approaches to molecule generation have been developed to accelerate the conventional paradigm. Existing molecular generative models largely focus on generating either molecule SMILES strings or molecular graphs (Gómez-Bombarelli et al. 2018; Jin, Barzilay, and Jaakkola 2018; Chen et al. 2021), with a recent shift towards 3D molecular structures. Several models (Luo et al. 2021; Peng et al. 2022; Guan et al. 2023) have been designed to generate 3D molecules conditioned on the protein targets, aiming to facilitate structure-based drug design (SBDD) (Batool, Ahmad, and Choi 2019), given that molecules exist in 3D space and the efficacy of drug molecules depends on their 3D structures fitting into protein pockets. However, SBDD relies on the availability of high-quality 3D structures of protein binding pockets, which are lacking for many targets (Zheng et al. 2013).

Different from SBDD, ligand-based drug design (LBDD) (Acharya et al. 2011) utilizes ligands known to interact with a protein target, and does not require knowledge of protein structures. In LBDD, shape-based virtual

screening tools such as ROCS (Hawkins, Skillman, and Nicholls 2006) have been widely used to identify molecules with similar shapes to known ligands by enumerating molecules in chemical libraries. However, virtual screen tools cannot probe the novel chemical space. Therefore, it is highly needed to develop generative methods to generate novel molecules with desired 3D shapes.

In this paper, we present a novel generative model for 3D molecule generation conditioned on given 3D shapes. Our method, denoted to as ShapeMol, employs an equivariant shape embedding module to map 3D molecule surface shapes into shape latent embeddings. It then uses a conditional diffusion generative model to generate molecules conditioned on the shape latent embeddings, by iteratively denoising atom positions and atom features (e.g., atom type and aromaticity). During molecule generation, ShapeMol can utilize additional shape guidance by pushing the predicted atoms far from the condition shapes to those shapes. ShapeMol with shape guidance is denoted as ShapeMol+g. The major contributions of this paper are as follows:

- To the best of our knowledge, ShapeMol is the first diffusion-based method for 3D molecule generation conditioned on 3D molecule shapes.
- ShapeMol leverages a new equivariant shape embedding module to learn 3D surface shape embeddings from cloud points sampled over molecule surfaces.
- ShapeMol uses a novel conditional diffusion model to generate 3D molecule structures. The diffusion model is equivariant to the translation and rotation of molecule shapes. A new weighting scheme over diffusion steps is developed to ensure accurate molecule shape prediction.
- ShapeMol utilizes new shape guidance to direct the generated molecules to better fit the shape condition.
- ShapeMol+g achieves the highest average 3D shape similarity between the generated molecules and condition molecules, compared to the state-of-the-art baseline.

For reproducibility purposes, detailed parameters in all the experiments, code and data are reported in Supplementary Section S1.

*Contact author

Related Work

Molecule Generation

A variety of deep generative models have been developed to generate molecules using various molecule representations, including generating SMILES string representations (Gómez-Bombarelli et al. 2018), or 2D molecular graph representations (Jin, Barzilay, and Jaakkola 2018; Chen et al. 2021). Recent efforts have been dedicated to the generation of 3D molecules. These 3D molecule generative models can be divided into two categories: autoregressive models and non-autoregressive models. Autoregressive models generate 3D molecules by sequentially adding atoms into the 3D space (Luo et al. 2021; Peng et al. 2022). While these models ensure the validity and connectivity of generated molecules, any errors made in sequential predictions could accumulate in subsequent predictions. Non-autoregressive models generate 3D molecules using flow-based methods (Garcia Satorras et al. 2021) or diffusion methods (Guan et al. 2023). In these models, atoms are generated or adjusted all together. For example, Hoogeboom *et al.* (2022) developed an equivariant diffusion model, in which an equivariant network is employed to jointly predict both the positions and features of all atoms.

Shape-Conditioned Molecule Generation

Following the idea of ligand-based drug design (LBDD) (Acharya et al. 2011), previous work has been focused on generating molecules with similar 3D shapes to those of efficacy, based on the observation that structurally similar molecules tend to have similar properties. Papadopoulos *et al.* (2021) developed a reinforcement learning method to generate SMILES strings of molecules that are similar to known antagonists of DRD2 receptors in 3D shapes and pharmacophores. Imrie *et al.* (2021) generated 2D molecular graphs conditioned on 3D pharmacophores using a graph-based autoencoder. However, there is limited work that generates 3D molecule structures conditioned on 3D shapes. Adams and Coley (2023) developed a shape-conditioned generative framework SQUID for 3D molecule generation. SQUID learns a variational autoencoder to generate fragments conditioned on given 3D shapes, and decodes molecules by sequentially attaching fragments with fixed bond lengths and angles. While LBDD plays a vital role in drug discovery, the problem of generating 3D molecule structures conditioned on 3D shapes is still under-addressed.

Definitions and Notations

Problem Definition

Following Adams and Coley (2023), we focus on the 3D molecule generation *conditioned on* the shape of a given molecule (e.g., a ligand). Specifically, we aim to generate a new molecule M_y , conditioned on the 3D shape of a given molecule M_x , such that 1) M_y is similar to M_x in their 3D shapes, measured by $\text{Sim}_s(s_x, s_y)$, where s is the 3D shape of M ; and 2) M_y is dissimilar to M_x in their 2D molecular graph structures, measured by $\text{Sim}_g(M_x, M_y)$. This conditional 3D

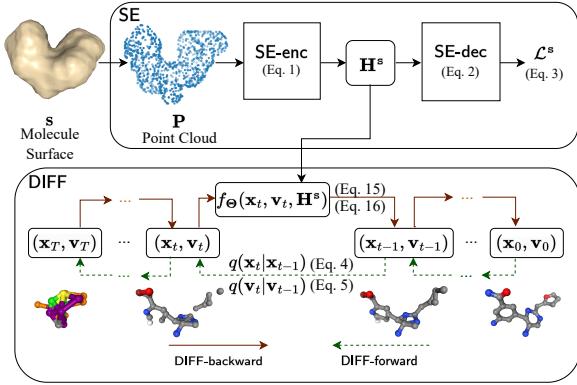


Figure 1: Model Architecture of ShapeMol

shape generation problem is motivated by the fact that in ligand-based drug design, it is desired to find chemically diverse and novel molecules that share similar shapes and similar activities with known active ligands (Ripphausen, Nisius, and Bajorath 2011). Such chemically diverse and novel molecules could expand the search space for drug candidates and potentially enhance the development of effective drugs.

Representations and Notations

We represent a molecule M as a set of atoms $M = \{a_1, a_2, \dots, a_{|M|} | a_i = (\mathbf{x}_i, \mathbf{v}_i)\}$, where $|M|$ is the number of atoms in M ; a_i is the i -th atom in M ; $\mathbf{x}_i \in \mathbb{R}^3$ represents the 3D coordinates of a_i ; and $\mathbf{v}_i \in \mathbb{R}^K$ is a_i 's one-hot atom feature vector indicating the atom type and its aromaticity. Following Guan *et al.* (2023), bonds between atoms can be uniquely determined by the atom types and the atomic distances among atoms. We represent the 3D surface shape s of a molecule M as a point cloud constructed by sampling points over the molecular surface. Details about the construction of point clouds from the surface of molecules are available in Supplementary Section S2. We denote the point cloud as $\mathcal{P} = \{z_1, z_2, \dots, z_{|\mathcal{P}|} | z_j = (\mathbf{z}_j)\}$, where $|\mathcal{P}|$ is the number of points in \mathcal{P} ; z_j is the j -th point; and $\mathbf{z}_j \in \mathbb{R}^3$ represents the 3D coordinates of z_j . We denote the latent embedding of \mathcal{P} as $\mathbf{H}^s \in \mathbb{R}^{d_p \times 3}$, where d_p is the dimension of the latent embedding.

Method

ShapeMol consists of an equivariant shape embedding module SE that maps 3D molecular surface shapes to latent embeddings, and an equivariant diffusion model DIFF that generates 3D molecules conditioned on these embeddings. Figure 1 presents the overall architecture of ShapeMol.

Equivariant Shape Embedding (SE)

ShapeMol pre-trains a shape embedding module SE to generate surface shape embeddings \mathbf{H}^s . SE uses an encoder SE-enc to map \mathcal{P} to the equivariant latent embedding \mathbf{H}^s . SE employs a decoder SE-dec to optimize \mathbf{H}^s by recovering the signed distances (Park et al. 2019) of sampled query points in 3D space to the molecule surface based on \mathbf{H}^s . ShapeMol uses \mathbf{H}^s to guide the diffusion process later.

Shape Encoder (SE-enc) SE-enc generates equivariant shape embeddings \mathbf{H}^s from the 3D surface shape \mathcal{P} of molecules, such that \mathbf{H}^s is equivariant to both translation and rotation of \mathcal{P} . That is, any translation and rotation applied to \mathcal{P} is reflected in \mathbf{H}^s accordingly. To ensure translation equivariance, SE-enc shifts the center of each \mathcal{P} to zero to eliminate all translations. To ensure rotation equivariance, SE-enc leverages Vector Neurons (VNs) (Deng et al. 2021) and Dynamic Graph Convolutional Neural Networks (DGCNNs) (Wang et al. 2019) as follows:

$$\{\mathbf{H}_1^p, \mathbf{H}_2^p, \dots, \mathbf{H}_{|\mathcal{P}|}^p\} = \text{VN-DGCNN}(\{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_{|\mathcal{P}|}\}),$$

$$\mathbf{H}^s = \sum_j \mathbf{H}_j^p / |\mathcal{P}|,$$

where $\text{VN-DGCNN}(\cdot)$ is a VN-based DGCNN network to generate equivariant embedding $\mathbf{H}_j^p \in \mathbb{R}^{d_p \times 3}$ for each point z_j in \mathcal{P} ; and $\mathbf{H}^s \in \mathbb{R}^{d_p \times 3}$ is the embedding of \mathcal{P} generated via a mean-pooling over the embeddings of all the points. Note that $\text{VN-DGCNN}(\cdot)$ generates a matrix as the embedding of each point (i.e., \mathbf{H}_j^p) to guarantee the equivariance.

Shape Decoder (SE-dec) To optimize \mathbf{H}^s , SE learns a decoder SE-dec to predict the signed distance of a query point z_q sampled from 3D space using Multilayer Perceptrons (MLPs) as follows:

$$\tilde{o}_q = \text{MLP}(\text{concat}(\langle \mathbf{z}_q, \mathbf{H}^s \rangle, \|\mathbf{z}_q\|^2, \text{VN-In}(\mathbf{H}^s))), \quad (1)$$

where \tilde{o}_q is the predicted signed distance of z_q , with positive and negative values indicating z_q is inside or outside the surface shape, respectively; $\langle \cdot, \cdot \rangle$ is the dot-product operator; $\|\mathbf{z}_q\|^2$ is the Euclidean norm of the coordinates of z_q ; $\text{VN-In}(\cdot)$ is an invariant VN network (Deng et al. 2021) that converts the equivariant shape embedding $\mathbf{H}^s \in \mathbb{R}^{d_p \times 3}$ into an invariant shape embedding. Thus, SE-dec predicts the signed distance between the query point and 3D surface by jointly considering the position of the query point ($\|\mathbf{z}_q\|^2$), the molecular surface shape ($\text{VN-In}(\cdot)$) and the interaction between the point and surface $\langle \cdot, \cdot \rangle$. The predicted signed distance \tilde{o}_q is used to calculate the loss for the optimization of \mathbf{H}^s (discussed below). As shown in the literature (Deng et al. 2021), \tilde{o}_q remains invariant to the rotation of the 3D molecule surface shapes (i.e., \mathcal{P}). We present the sampling process of z_q in the Supplementary Section S3.

SE Pre-training ShapeMol pre-trains SE by minimizing the squared-errors loss between the predicted and the ground-truth signed distances of query points as follows:

$$\mathcal{L}^s = \sum_{z_q \in \mathcal{Z}} \|o_q - \tilde{o}_q\|^2, \quad (2)$$

where \mathcal{Z} is the set of sampled query points and o_q is the ground-truth signed distance of query point z_q . By pretraining SE, ShapeMol learns \mathbf{H}^s that will be used as the condition in the following 3D molecule generation.

Shape-Conditioned Molecule Generation

In ShapeMol, a shape-conditioned molecule diffusion model, referred to as DIFF, is used to generate a 3D molecule structure (i.e., atom coordinates and features) conditioned on a given 3D surface shape that is represented by

the shape latent embedding \mathbf{H}^s (Eq.). Following the denoising diffusion probabilistic models (Ho, Jain, and Abbeel 2020), DIFF includes a forward diffusion process based on a Markov chain, denoted as DIFF-forward, which gradually adds noises step by step to the atom positions and features $\{(\mathbf{x}_i, \mathbf{v}_i)\}$ in the training molecules. The noisy atom positions and features at step t are represented as $\{(\mathbf{x}_{i,t}, \mathbf{v}_{i,t})\}$ ($t = 1, \dots, T$), and the molecules without any noise are represented as $\{(\mathbf{x}_{i,0}, \mathbf{v}_{i,0})\}$. At the final step T , $\{(\mathbf{x}_{i,T}, \mathbf{v}_{i,T})\}$ are completely unstructured and resemble a simple distribution like a Normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ or a uniform categorical distribution $\mathcal{C}(1/K)$, in which \mathbf{I} and $\mathbf{1}$ denotes the identity matrix and identity vector, respectively.

During training, DIFF is learned to reverse the forward diffusion process via another Markov chain, referred to as the backward generative process and denoted as DIFF-backward, to remove the noises in the noisy molecules. During inference, DIFF first samples noisy atom positions and features at step T from simple distributions and then generates a 3D molecule structure by removing the noises in the noisy molecules step by step until t reaches 1.

Forward Diffusion Process (DIFF-forward) Following the previous work (Guan et al. 2023), at step $t \in [1, T]$, a small Gaussian noise and a small categorical noise are added to the continuous atom positions and discrete atom features $\{(\mathbf{x}_{i,t-1}, \mathbf{v}_{i,t-1})\}$, respectively. When no ambiguity arises, we will eliminate subscript i in the notations and use $(\mathbf{x}_{t-1}, \mathbf{v}_{t-1})$ for brevity. The noise levels of the Gaussian and categorical noises are determined by two predefined variance schedules $(\beta_t^x, \beta_t^v) \in (0, 1)$, where β_t^x and β_t^v are selected to be sufficiently small to ensure the smoothness of DIFF-forward. The details about variance schedules are available in Supplementary Section S4.1. Formally, for atom positions, the probability of \mathbf{x}_t sampled given \mathbf{x}_{t-1} , denoted as $q(\mathbf{x}_t | \mathbf{x}_{t-1})$, is defined as follows,

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t | \sqrt{1 - \beta_t^x} \mathbf{x}_{t-1}, \beta_t^x \mathbf{I}), \quad (3)$$

where $\mathcal{N}(\cdot)$ is a Gaussian distribution of \mathbf{x}_t with mean $\sqrt{1 - \beta_t^x} \mathbf{x}_{t-1}$ and covariance $\beta_t^x \mathbf{I}$. Following Hoogeboom et al. (2021), for atom features, the probability of \mathbf{v}_t across K classes given \mathbf{v}_{t-1} is defined as follows,

$$q(\mathbf{v}_t | \mathbf{v}_{t-1}) = \mathcal{C}(\mathbf{v}_t | (1 - \beta_t^v) \mathbf{v}_{t-1} + \beta_t^v \mathbf{1}/K), \quad (4)$$

where \mathcal{C} is a categorical distribution of \mathbf{v}_t derived by noising \mathbf{v}_{t-1} with a uniform noise $\beta_t^v \mathbf{1}/K$ across K classes.

Since the above distributions form Markov chains, the probability of any \mathbf{x}_t or \mathbf{v}_t can be derived from \mathbf{x}_0 or \mathbf{v}_0 :

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t | \sqrt{\bar{\alpha}_t^x} \mathbf{x}_0, (1 - \bar{\alpha}_t^x) \mathbf{I}), \quad (5)$$

$$q(\mathbf{v}_t | \mathbf{v}_0) = \mathcal{C}(\mathbf{v}_t | \bar{\alpha}_t^v \mathbf{v}_0 + (1 - \bar{\alpha}_t^v) \mathbf{1}/K), \quad (6)$$

$$\text{where } \bar{\alpha}_t^u = \prod_{\tau=1}^t \alpha_\tau^u, \quad \alpha_\tau^u = 1 - \beta_\tau^u, \quad u = x \text{ or } v. \quad (7)$$

Note that $\bar{\alpha}_t^u$ ($u = x$ or v) is monotonically decreasing from 1 to 0 over $t = [1, T]$. As $t \rightarrow 1$, $\bar{\alpha}_t^x$ and $\bar{\alpha}_t^v$ are close to 1, leading to that \mathbf{x}_t or \mathbf{v}_t approximates \mathbf{x}_0 or \mathbf{v}_0 . Conversely, as $t \rightarrow T$, $\bar{\alpha}_t^x$ and $\bar{\alpha}_t^v$ are close to 0, leading to that $q(\mathbf{x}_T | \mathbf{x}_0)$ resembles $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and $q(\mathbf{v}_T | \mathbf{v}_0)$ resembles $\mathcal{C}(1/K)$.

Using Bayes theorem, the ground-truth Normal posterior of atom positions $p(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ can be calculated in a closed-form (Ho, Jain, and Abbeel 2020) as below,

$$p(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}|\mu(\mathbf{x}_t, \mathbf{x}_0), \tilde{\beta}_t^x \mathbf{I}), \quad (8)$$

$$\mu(\mathbf{x}_t, \mathbf{x}_0) = \frac{\sqrt{\bar{\alpha}_{t-1}^x} \beta_t^x}{1-\bar{\alpha}_t^x} \mathbf{x}_0 + \frac{\sqrt{\alpha_t^x(1-\bar{\alpha}_{t-1}^x)}}{1-\bar{\alpha}_t^x} \mathbf{x}_t, \tilde{\beta}_t^x = \frac{1-\bar{\alpha}_{t-1}^x}{1-\bar{\alpha}_t^x} \beta_t^x. \quad (9)$$

Similarly, the ground-truth categorical posterior of atom features $p(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathbf{v}_0)$ can be calculated (Hoogeboom et al. 2021) as below,

$$p(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathbf{v}_0) = \mathcal{C}(\mathbf{v}_{t-1}|\mathbf{c}(\mathbf{v}_t, \mathbf{v}_0)), \quad (10)$$

$$\mathbf{c}(\mathbf{v}_t, \mathbf{v}_0) = \tilde{\mathbf{c}} / \sum_{k=1}^K \tilde{c}_k, \quad (11)$$

$$\tilde{\mathbf{c}} = [\alpha_t^v \mathbf{v}_t + \frac{1-\alpha_t^v}{K}] \odot [\bar{\alpha}_{t-1}^v \mathbf{v}_0 + \frac{1-\bar{\alpha}_{t-1}^v}{K}], \quad (12)$$

where \tilde{c}_k denotes the likelihood of k -th class across K classes in $\tilde{\mathbf{c}}$; \odot denotes the element-wise product operation; $\tilde{\mathbf{c}}$ is calculated using \mathbf{v}_t and \mathbf{v}_0 and normalized so as to represent probabilities. The proof of the above equations is available in Supplementary Section S4.2.

Backward Generative Process (DIFF-backward) DIFF learns to reverse DIFF-forward by denoising from $(\mathbf{x}_t, \mathbf{v}_t)$ to $(\mathbf{x}_{t-1}, \mathbf{v}_{t-1})$ at $t \in [1, T]$, conditioned on the shape latent embedding \mathbf{H}^s . Specifically, the probabilities of $(\mathbf{x}_{t-1}, \mathbf{v}_{t-1})$ denoised from $(\mathbf{x}_t, \mathbf{v}_t)$ are estimated by the approximates of the ground-truth posteriors $p(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ (Eq. 8) and $p(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathbf{v}_0)$ (Eq. 10). Given that $(\mathbf{x}_0, \mathbf{v}_0)$ is unknown in the generative process, a predictor $f_\Theta(\mathbf{x}_t, \mathbf{v}_t, \mathbf{H}^s)$ is employed to predict at t the atom position and feature $(\mathbf{x}_0, \mathbf{v}_0)$ as below,

$$(\tilde{\mathbf{x}}_{0,t}, \tilde{\mathbf{v}}_{0,t}) = f_\Theta(\mathbf{x}_t, \mathbf{v}_t, \mathbf{H}^s), \quad (13)$$

where $\tilde{\mathbf{x}}_{0,t}$ and $\tilde{\mathbf{v}}_{0,t}$ are the predictions of \mathbf{x}_0 and \mathbf{v}_0 at t ; Θ is the learnable parameter. Following Ho *et al.* (2020), with $\tilde{\mathbf{x}}_{0,t}$, the probability of \mathbf{x}_{t-1} denoised from \mathbf{x}_t , denoted as $p(\mathbf{x}_{t-1}|\mathbf{x}_t)$, can be estimated by the approximated posterior $p_\Theta((\mathbf{x}_{t-1}|\mathbf{x}_t, \tilde{\mathbf{x}}_{0,t}))$ as below,

$$\begin{aligned} p(\mathbf{x}_{t-1}|\mathbf{x}_t) &\approx p_\Theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \tilde{\mathbf{x}}_{0,t}) \\ &= \mathcal{N}(\mathbf{x}_{t-1}|\mu_\Theta(\mathbf{x}_t, \tilde{\mathbf{x}}_{0,t}), \tilde{\beta}_t^x \mathbf{I}), \end{aligned} \quad (14)$$

where $\mu_\Theta(\mathbf{x}_t, \tilde{\mathbf{x}}_{0,t})$ is an estimate of $\mu(\mathbf{x}_t, \mathbf{x}_0)$ by replacing \mathbf{x}_0 with its estimate $\tilde{\mathbf{x}}_{0,t}$ in Eq. 8. Similarly, with $\tilde{\mathbf{v}}_{0,t}$, the probability of \mathbf{v}_{t-1} denoised from \mathbf{v}_t , denoted as $p(\mathbf{v}_{t-1}|\mathbf{v}_t)$, can be estimated by the approximated posterior $p_\Theta(\mathbf{v}_{t-1}|\mathbf{v}_t, \tilde{\mathbf{v}}_{0,t})$ as below,

$$p(\mathbf{v}_{t-1}|\mathbf{v}_t) \approx p_\Theta(\mathbf{v}_{t-1}|\mathbf{v}_t, \tilde{\mathbf{v}}_{0,t}) = \mathcal{C}(\mathbf{v}_{t-1}|\mathbf{c}_\Theta(\mathbf{v}_t, \tilde{\mathbf{v}}_{0,t})), \quad (15)$$

where $\mathbf{c}_\Theta(\mathbf{v}_t, \tilde{\mathbf{v}}_{0,t})$ is an estimate of $\mathbf{c}(\mathbf{v}_t, \mathbf{v}_0)$ by replacing \mathbf{v}_0 with its estimate $\tilde{\mathbf{v}}_{0,t}$ in Eq. 10.

Equivariant Shape-Conditioned Molecule Predictor In DIFF-backward, the predictor $f_\Theta(\mathbf{x}_t, \mathbf{v}_t, \mathbf{H}^s)$ (Eq. 13) predicts the atom positions and features $(\tilde{\mathbf{x}}_{0,t}, \tilde{\mathbf{v}}_{0,t})$ given the noisy data $(\mathbf{x}_t, \mathbf{v}_t)$ conditioned on \mathbf{H}^s . For brevity, in this subsection, we eliminate the subscript t in the notations when no ambiguity arises. $f_\Theta(\cdot)$ leverages two multi-layer graph neural networks: (1) an equivariant graph neural network, denoted as EQ-GNN, that equivariantly predicts atom

positions that change under transformations, and (2) an invariant graph neural network, denoted as INV-GNN, that invariantly predicts atom features that remain unchanged under transformations. Following the previous work (Guan et al. 2023; Hoogeboom et al. 2022), the translation equivariance of atom position prediction is achieved by shifting a fixed point (e.g., the center of point clouds \mathcal{P}) to zero, and therefore only rotation equivariance needs to be considered.

Atom Coordinate Prediction In EQ-GNN, the atom position $\mathbf{x}_i^{l+1} \in \mathbb{R}^3$ of a_i at the $(l+1)$ -th layer is calculated in an equivariant way as below,

$$\Delta \mathbf{x}_i^{l+1} = \sum_{j \in N(a_i), i \neq j} (\mathbf{x}_i^l - \mathbf{x}_j^l) \text{MHA}^x(d_{ij}^l, \mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, \text{VN-In}(\mathbf{H}^s)),$$

$$\mathbf{x}_i^{l+1} = \mathbf{x}_i^l + \text{Mean}(\Delta \mathbf{x}_i^{l+1}) + \text{VN-Lin}(\mathbf{x}_i^l, \Delta \mathbf{x}_i^{l+1}, \mathbf{H}^s), \quad (16)$$

where $N(a_i)$ is the set of N -nearest neighbors of a_i based on atomic distances; $\Delta \mathbf{x}_i^{l+1} \in \mathbb{R}^{n_h \times 3}$ aggregates the neighborhood information of a_i ; $\text{MHA}^x(\cdot)$ denotes the multi-head attention layer in EQ-GNN with n_h heads; d_{ij}^l is the distance between i -th and j -th atom positions \mathbf{x}_i^l and \mathbf{x}_j^l at the l -th layer; $\text{Mean}(\Delta \mathbf{x}_i^{l+1})$ converts $\Delta \mathbf{x}_i^{l+1}$ into a 3D vector via meaning pooling to adjust the atom position; $\text{VN-Lin}(\cdot) \in \mathbb{R}^3$ denotes the equivariant VN-based linear layer (Deng et al. 2021). $\text{VN-Lin}(\cdot)$ adjusts the atom positions to fit the shape condition represented by \mathbf{H}^s , by considering the current atom positions \mathbf{x}_i^l and the neighborhood information $\Delta \mathbf{x}_i^{l+1}$. The learned atom position \mathbf{x}_i^L at the last layer L of EQ-GNN is used as the prediction of $\tilde{\mathbf{x}}_{i,0}$, that is,

$$\tilde{\mathbf{x}}_{i,0} = \mathbf{x}_i^L. \quad (17)$$

Atom Feature Prediction In INV-GNN, inspired by the previous work (Guan et al. 2023) and VN-Layer (Deng et al. 2021), the atom feature embedding $\mathbf{h}_i^{l+1} \in \mathbb{R}^{d_h}$ of the i -th atom a_i at the $(l+1)$ -th layer of INV-GNN is updated in an invariant way as follows,

$$\mathbf{h}_i^{l+1} = \mathbf{h}_i^l + \sum_{j \in N(a_i), i \neq j} \text{MHA}^h(d_{ij}^l, \mathbf{h}_i^l, \mathbf{h}_j^l, \text{VN-In}(\mathbf{H}^s)), \mathbf{h}_i^0 = \mathbf{v}_i, \quad (18)$$

where $\text{MHA}^h(\cdot) \in \mathbb{R}^{d_h}$ denotes the multi-head attention layer in INV-GNN. The learned atom feature embedding \mathbf{h}_i^L at the last layer L encodes the neighborhood information of a_i and the conditioned molecular shape, and is used to predict the atom features as follows:

$$\tilde{\mathbf{v}}_{i,0} = \text{softmax}(\text{MLP}(\mathbf{h}_i^L)). \quad (19)$$

The proof of equivariance in Eq. 16 and invariance in Eq. 18 is available in Supplementary Section S5.1 and S5.2.

Model Training ShapeMol optimizes DIFF by minimizing the squared errors between the predicted positions $(\tilde{\mathbf{x}}_{0,t})$ and the ground-truth positions (\mathbf{x}_0) of atoms in molecules. Given a particular step t , the error is calculated as follows:

$$\mathcal{L}_t^x(\mathcal{M}) = w_t^x \sum_{\forall a \in \mathcal{M}} \|\tilde{\mathbf{x}}_{0,t} - \mathbf{x}_0\|^2, \quad (20)$$

$$\text{where } w_t^x = \min(\lambda_t, \delta), \lambda_t = \bar{\alpha}_t^x / (1 - \bar{\alpha}_t^x),$$

where w_t^x is a weight at step t , and is calculated by clipping the signal-to-noise ratio $\lambda_t > 0$ with a threshold $\delta > 0$. Note that because $\tilde{\alpha}_t^x$ decreases monotonically as t increases from 1 to T (Eq. 7), w_t^x decreases monotonically over t as well until it is clipped. Thus, w_t^x imposes lower weights on the loss when the noise level in x_t is higher (i.e., at later/larger step t). This encourages the model training to focus more on accurately recovering molecule structures when there are sufficient signals in the data, rather than being potentially confused by major noises in the data.

ShapeMol also minimizes the KL divergence (Kullback and Leibler 1951) between the ground-truth posterior $p(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathbf{v}_0)$ (Eq. 10) and its approximate $p_\Theta(\mathbf{v}_{t-1}|\mathbf{v}_t, \tilde{\mathbf{v}}_{0,t})$ (Eq. 15) for discrete atom features to optimize DIFF, following the literature (Hoogeboom et al. 2021). Particularly, the KL divergence at t for a given molecule is calculated as follows:

$$\begin{aligned} \mathcal{L}_t^v(\mathbf{M}) &= \sum_{\forall a \in \mathbf{M}} \text{KL}(p(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathbf{v}_0) | p_\Theta(\mathbf{v}_{t-1}|\mathbf{v}_t, \tilde{\mathbf{v}}_{0,t})), \\ &= \sum_{\forall a \in \mathbf{M}} \text{KL}(\mathbf{c}(\mathbf{v}_t, \mathbf{v}_0) | \mathbf{c}_\Theta(\mathbf{v}_t, \tilde{\mathbf{v}}_{0,t})), \end{aligned} \quad (21)$$

where $\mathbf{c}(\mathbf{v}_t, \mathbf{v}_0)$ is a categorical distribution of \mathbf{v}_{t-1} (Eq. 11); $\mathbf{c}_\Theta(\mathbf{v}_t, \tilde{\mathbf{v}}_{0,t})$ is an estimate of $\mathbf{c}(\mathbf{v}_t, \mathbf{v}_0)$ (Eq. 15). The overall ShapeMol loss function is defined as follows:

$$\mathcal{L} = \sum_{\forall \mathbf{M} \in \mathcal{M}} \sum_{\forall t \in \mathcal{T}} (\mathcal{L}_t^x(\mathbf{M}) + \xi \mathcal{L}_t^v(\mathbf{M})), \quad (22)$$

where \mathcal{M} is the set of all the molecules in training; \mathcal{T} is the set of sampled timesteps; $\xi > 0$ is a hyper-parameter to balance $\mathcal{L}_t^x(\mathbf{M})$ and $\mathcal{L}_t^v(\mathbf{M})$. During training, step t is uniformly sampled from $\{1, 2, \dots, 1000\}$. The derivation of the loss functions is available in Supplementary Section S6.

Molecule Generation During inference, ShapeMol generates novel molecules by gradually denoising $(\mathbf{x}_T, \mathbf{v}_T)$ to $(\mathbf{x}_0, \mathbf{v}_0)$ using the equivariant shape-conditioned molecule predictor. Specifically, ShapeMol samples \mathbf{x}_T and \mathbf{v}_T from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\mathcal{C}(1/K)$, respectively. After that, ShapeMol samples \mathbf{x}_{t-1} from \mathbf{x}_t using $p_\Theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \tilde{\mathbf{x}}_{0,t})$ (Eq. 14). Similarly, ShapeMol samples \mathbf{v}_{t-1} from \mathbf{v}_t using $p_\Theta(\mathbf{v}_{t-1}|\mathbf{v}_t, \tilde{\mathbf{v}}_{0,t})$ (Eq. 15) until t reaches 1.

ShapeMol with Shape Guidance During molecule generation, ShapeMol can also utilize additional shape guidance by pushing the predicted atoms to the shape of the given molecule \mathbf{M}_x . Particularly, following Adams and Coley *et al.* (2023), the shape used for guidance is defined as a set of points \mathcal{Q} sampled according to atom positions in \mathbf{M}_x . Particularly, for each atom a_i in \mathbf{M}_x , 20 points are randomly sampled into \mathcal{Q} from a Gaussian distribution centered at \mathbf{x}_i with variance ϕ . Given the predicted atom position $\tilde{\mathbf{x}}_{0,t}$ at step t , ShapeMol applies the shape guidance by adjusting the predicted positions to \mathbf{M}_x as follows:

$$\mathbf{x}_{0,t}^* = (1 - \sigma) \tilde{\mathbf{x}}_{0,t} + \sigma \sum_{\mathbf{z} \in n(\tilde{\mathbf{x}}_{0,t}; \mathcal{Q})} \mathbf{z} / n, \text{ when } \sum_{\mathbf{z} \in n(\tilde{\mathbf{x}}_{0,t}; \mathcal{Q})} d(\tilde{\mathbf{x}}_{0,t}, \mathbf{z}) / n > \gamma, \quad (23)$$

where $\sigma > 0$ is the weight used to balance the prediction $\tilde{\mathbf{x}}_{0,t}$ and the adjustment; $d(\tilde{\mathbf{x}}_{0,t}, \mathbf{z})$ is the Euclidean distance between $\tilde{\mathbf{x}}_{0,t}$ and \mathbf{z} ; $n(\tilde{\mathbf{x}}_{0,t}; \mathcal{Q})$ is the set of n -nearest neighbors of $\tilde{\mathbf{x}}_{0,t}$ in \mathcal{Q} based on $d(\cdot)$; $\gamma > 0$ is a distance threshold.

By doing the above adjustment, the predicted atom positions will be pushed to those of \mathbf{M}_x if they are sufficiently far away. Note that the shape guidance is applied exclusively for steps

$$t = T, T-1, \dots, S, \text{ where } S > 1, \quad (24)$$

not for all the steps, and thus it only adjusts predicted atom positions when there are a lot of noises and the prediction needs more guidance. ShapeMol with the shape guidance is referred to as ShapeMol+g.

Experiments

Data

Following SQUID (Adams and Coley 2023), we used molecules in the MOSES dataset (Polykovskiy et al. 2020), with their 3D conformers calculated by RDKit (Landrum et al. 2023). We used the same training and test split as in SQUID. Please note that SQUID further modifies the generated conformers into artificial ones, by adjusting acyclic bond distances to their empirical means and fixing acyclic bond angles using heuristic rules. Unlike SQUID, we did not make any additional adjustments to the calculated 3D conformers, as ShapeMol is designed with sufficient flexibility to accept any 3D conformers as input and generate 3D molecules without restrictions on fixed bond lengths or angles. Limited by the predefined fragment library, SQUID also removes molecules with fragments not present in its fragment library. In contrast, we kept all the molecules, as ShapeMol is not based on fragments. Our final training dataset contains 1,593,653 molecules, out of which a random set of 1,000 molecules was selected for validation. Both the ShapeMol-enc and DIFF models are trained using this training set. 1,000 test molecules (i.e., conditions) as used in SQUID are used to test ShapeMol.

Baselines

We compared ShapeMol and ShapeMol+g with the state-of-the-art baseline SQUID and a virtual screening method over the training dataset, denoted as VS. As far as we know, SQUID is the only generative baseline that generates 3D molecules conditioned on molecule shapes. SQUID consists of a fragment-based generative model based on variational autoencoder that sequentially decodes fragments from molecule latent embeddings and shape embeddings, and a rotatable bond scoring framework that adjusts the angles of rotatable bonds between fragments to maximize the 3D shape similarity with the condition molecule. VS aims to sift through the training set to identify molecules with high shape similarities with the condition molecule. For SQUID, we assessed two interpolation levels, $\lambda = 0.3$ and 1.0 (prior), following the original SQUID paper (Adams and Coley 2023). For SQUID, ShapeMol and ShapeMol+g, we generated 50 molecules for each testing molecule (i.e., condition) as the candidates for evaluation. For VS, we randomly sampled 500 training molecules for each testing molecule, and considered the top-50 molecules with the highest shape similarities as candidates for evaluation.

Evaluation Metrics We use shape similarity $\text{Sim}_s(s_x, s_y)$ and molecular graph similarity $\text{Sim}_g(M_x, M_y)$ to measure the generated new molecules M_y with respect to the condition M_x . Higher Sim_s and meanwhile lower Sim_g indicate better model performance. We also measure the diversity (div) of the generated molecules, calculated as 1 minus average pairwise Sim_g among all generated molecules. Higher div indicates better performance. Details about the evaluation metrics are available in Supplementary Section S7.

Performance Comparison

Overall Comparison Table 1 presents the overall comparison of shape-conditioned molecule generation among VS, SQUID, ShapeMol and ShapeMol+g. As shown in Table 1, ShapeMol+g achieves the highest average shape similarity 0.746 ± 0.036 , with 2.3% improvement from the best baseline VS (0.729 ± 0.039), although at the cost of a slightly higher graph similarity (0.241 ± 0.050 in ShapeMol+g vs 0.226 ± 0.038 in VS). This indicates that ShapeMol+g could generate molecules that align more closely with the shape conditions than those in the dataset. Furthermore, ShapeMol+g achieves the second-best performance in maximum shape similarity maxSim_s at 0.852 ± 0.034 among all the methods. While it underperforms the best baseline (0.904 ± 0.070 for SQUID with $\lambda=0.3$) on this metric, ShapeMol+g achieves substantially lower maximum graph similarity maxSim_g of 0.247 ± 0.068 compared with the best baseline (0.549 ± 0.243). This highlights the ability of ShapeMol+g in generating novel molecules that resemble the shape conditions. ShapeMol+g also achieves the lowest standard deviation values on both the average and maximum shape similarities (0.036 and 0.034 , respectively) among all the methods, further demonstrating its ability to consistently generate molecules with high shape similarities.

ShapeMol+g performs substantially better than ShapeMol on 3D shape similarity metrics (e.g., 0.746 ± 0.036 vs 0.689 ± 0.044 on avgSim_s). The superior performance of ShapeMol+g highlights the importance of shape guidance in the generative process. Although ShapeMol underperforms ShapeMol+g, it still outperforms SQUID with $\lambda=1.0$ in terms of the avgSim_s (i.e., 0.689 ± 0.044 vs 0.670 ± 0.069).

In terms of the quality of generated molecules, 98.7% of molecules from ShapeMol+g and 98.8% from ShapeMol are connected, and every connected molecule is unique. SQUID with λ values of 0.3 or 1.0 ensures the 100% connectivity among generated molecules by sequentially attaching fragments. However, out of these connected molecules, 94.2% and 95.0% are unique for SQUID with λ value of 0.3 or 1.0, respectively. In terms of the drug-likeness (QED), both ShapeMol+g and ShapeMol achieve QED values (e.g., 0.749 for ShapeMol+g) close to those of SQUID with λ as 0.3 and 1.0 (e.g., 0.760 for SQUID with $\lambda=0.3$). All the generative methods produce slightly inferior QED values to real molecules (0.795 for VS). In terms of diversity, ShapeMol+g and ShapeMol achieve higher diversity values (e.g., 0.703 ± 0.053 for ShapeMol+g) than SQUID with $\lambda=1.0$ (0.677 ± 0.065), though slightly lower than SQUID with $\lambda=0.3$ and VS. Overall, ShapeMol and ShapeMol+g are

able to generate connected, unique and diverse molecules with good drug-likeness scores.

Please note that unlike SQUID, which neglects distorted bonding geometries in real molecules and limits itself to generating molecules with fixed bond lengths and angles, both ShapeMol and ShapeMol+g are able to generate molecules without such limitations. Given the superior performance of ShapeMol+g in shape-conditioned molecule generation, it could serve as a promising tool for ligand-based drug design.

Comparison of Diffusion Weighting Schemes While previous work (Peng et al. 2023; Guan et al. 2023) applied uniform weights on different diffusion steps, ShapeMol uses different weights (i.e., w_t^x in Eq. 20). We conducted an ablation study to demonstrate the effectiveness of this new weighting scheme. Particularly, we trained two DIFF modules with the varying step weights w_t^x (with $\delta = 10$ in Eq. 20) and uniform weights, respectively, while fixing all the other hyper-parameters in ShapeMol and ShapeMol+g. Table 2 presents their performance comparison.

The results in Table 2 show that the different weights on different steps substantially improve the quality of the generated molecules. Specifically, ShapeMol with different weights ensures higher molecular connectivity and drug-likeness than that with uniform weights (98.8% vs 89.4% for connectivity; 0.748 vs 0.660 for QED). ShapeMol with different weights also produces molecules with bond length distributions closer to those of real molecules (i.e., lower Jensen-Shannon divergence), for example, the Jensen-Shannon (JS) divergence of bond lengths between real and generated molecules decreases from 0.115 to 0.095 when different weights are applied. The same trend can be observed for ShapeMol+g, for which the different weights also improve the generated molecule qualities. Since w_t^x increases as the noise level in the data decreases (See discussions earlier in “Model Training”), the results in Table 2 demonstrate the effectiveness of the new weighting scheme in promoting new molecules generated more similarly to real ones when the noise level in data is small.

Parameter Study We conducted a parameter study to evaluate the impact of the distance threshold γ (Eq. 23) and the step threshold S (Eq. 24) in the shape guidance. Particularly, using the same trained DIFF module, we sampled molecules with different values of γ and S and present the results in Table 3. As shown in Table 3, the average shape similarities avgSim_s and maximum shape similarities maxSim_s consistently decrease as γ and S increase. For example, when $S = 50$, avgSim_s and maxSim_s decreases from 0.794 to 0.763 and 0.890 to 0.861, respectively, as γ increases from 0.2 to 0.6. Similarly, when $\gamma = 0.2$, avgSim_s and maxSim_s decreases from 0.794 to 0.746 and 0.890 to 0.852, respectively, as S increases from 50 to 300. As presented in “ShapeMol with Shape Guidance”, larger γ and S indicate stronger shape guidance in ShapeMol+g. These results demonstrate that stronger shape guidance in ShapeMol+g could effectively induce higher shape similarities between the given molecule and generated molecules.

It is also noticed that as shown in Table 3, incorporating

Table 1: Overall Comparison on Shape-Conditioned Molecule Generation

| method | #c% | #u% | QED | avgSim _s (std) | avgSim _g (std) | maxSim _s (std) | maxSim _g (std) | div(std) |
|-------------------------|-------|-------|-------|---------------------------|---------------------------|---------------------------|---------------------------|---------------|
| VS | 100.0 | 100.0 | 0.795 | 0.729 (0.039) | 0.226 (0.038) | 0.807 (0.042) | 0.241 (0.087) | 0.759 (0.015) |
| SQUID ($\lambda=0.3$) | 100.0 | 94.2 | 0.766 | 0.717 (0.083) | 0.349 (0.088) | 0.904 (0.070) | 0.549 (0.243) | 0.677 (0.065) |
| SQUID ($\lambda=1.0$) | 100.0 | 95.0 | 0.760 | 0.670 (0.069) | 0.235 (0.045) | 0.842 (0.061) | 0.271 (0.096) | 0.744 (0.046) |
| ShapeMol | 98.8 | 100.0 | 0.748 | 0.689 (0.044) | 0.239 (0.049) | 0.803 (0.042) | 0.243 (0.068) | 0.712 (0.055) |
| ShapeMol+g | 98.7 | 100.0 | 0.749 | 0.746 (0.036) | 0.241 (0.050) | 0.852 (0.034) | 0.247 (0.068) | 0.703 (0.053) |

Columns represent: “#c%”: the percentage of connected molecules; “#u%”: the percentage of unique molecules; “QED”: the average drug-likeness of generated molecules; “avgSim_s/avgSim_g”: the average of shape or graph similarities between the condition molecules and generated molecules; “std”: the standard deviation; “maxSim_s”: the maximum of shape similarities between the condition molecules and generated molecules; “maxSim_g”: the graph similarities between the condition molecules and the molecules with the maximum shape similarities; “div”: the diversity among the generated molecules.

Table 2: Comparison of Diffusion Weighting Schemes

| method | weights | #c% | #u% | QED | JS divergence | |
|------------|---------|------|-------|-------|---------------|-------|
| | | | | | bond | C-C |
| ShapeMol | w_t^x | 98.8 | 100.0 | 0.748 | 0.095 | 0.321 |
| | uniform | 89.4 | 100.0 | 0.660 | 0.115 | 0.393 |
| ShapeMol+g | w_t^x | 98.7 | 100.0 | 0.749 | 0.093 | 0.317 |
| | uniform | 90.1 | 100.0 | 0.671 | 0.112 | 0.384 |

Columns represent: “weights”: different weighting schemes; “JS distance of bond/C-C”: the Jensen-Shannon (JS) divergence of bond length among all the bond types (“bond”)/carbon-carbon single bonds (“C-C”) between real molecules and generated molecules; All the others are identical to those in Table 1.

Table 3: Parameter Study in Shape Guidance

| γ | S | QED | JS. bond | avgSim _s | avgSim _g | maxSim _s | maxSim _g |
|----------|-----|-------|----------|---------------------|---------------------|---------------------|---------------------|
| - | - | 0.748 | 0.094 | 0.689 | 0.239 | 0.803 | 0.243 |
| 0.2 | 50 | 0.630 | 0.110 | 0.794 | 0.236 | 0.890 | 0.244 |
| 0.2 | 100 | 0.666 | 0.105 | 0.786 | 0.238 | 0.883 | 0.245 |
| 0.2 | 300 | 0.749 | 0.093 | 0.746 | 0.241 | 0.852 | 0.247 |
| 0.4 | 50 | 0.678 | 0.106 | 0.779 | 0.240 | 0.875 | 0.245 |
| 0.4 | 100 | 0.700 | 0.103 | 0.772 | 0.241 | 0.870 | 0.247 |
| 0.4 | 300 | 0.752 | 0.093 | 0.738 | 0.242 | 0.845 | 0.247 |
| 0.6 | 50 | 0.706 | 0.103 | 0.763 | 0.242 | 0.861 | 0.246 |
| 0.6 | 100 | 0.720 | 0.100 | 0.758 | 0.242 | 0.857 | 0.247 |
| 0.6 | 300 | 0.753 | 0.093 | 0.731 | 0.242 | 0.838 | 0.247 |

Columns represent: “ γ ”/“ S ”: distance threshold/step threshold in shape guidance; “JS. bond”: the JS divergence of bond length distributions of all the bond types between real molecules and generated molecules; All the others are identical to those in Table 1.

shape guidance enables a trade-off between the quality of the generated molecules (QED), and the shape similarities (avgSim_s and maxSim_s) between the given molecule and the generated ones. For example, when $\gamma = 0.2$, QED increases from 0.630 to 0.749 and avgSim_s decreases from 0.794 to 0.746 as S increases from 50 to 300. These results indicate the effects of γ and S in guiding molecule generation conditioned on given shapes.

Case Study Figure 2 presents three generated molecules from three methods given the same condition molecule. As shown in Figure 2, the molecule generated by ShapeMol

has higher shape similarity (0.835) with the condition molecule than those from the baseline methods (0.759 for VS and 0.749 for SQUID). Particularly, the molecule from ShapeMol has the surface shape (represented as blue shade in Figure 2d) most similar to that of the condition molecule. All three molecules have low graph similarities with the condition molecule and higher QED scores than the condition molecule. This example shows the ability of ShapeMol to generate novel molecules that are more similar in 3D shape to condition molecules than those from baseline methods.

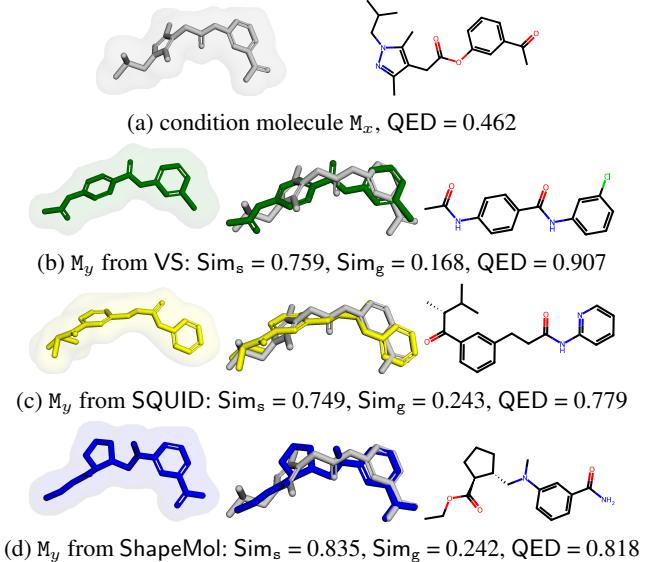


Figure 2: Generated 3D Molecules from Different Methods. Molecule 3D shapes are in shades; generated molecules are superpositioned with the condition molecule; and the molecular graphs of generated molecules are presented.

Discussions and Conclusions

In this paper, we develop a novel generative model ShapeMol, which generates 3D molecules conditioned on the 3D shape of given molecules. ShapeMol utilizes a pre-trained equivariant shape encoder to generate equivariant embeddings for 3D shapes of given molecules. Conditioned on the embeddings, ShapeMol learns an equivariant diffusion model to generate novel molecules. To improve the shape similarities between the given molecule and the generated ones, we develop ShapeMol+g, which incorporates

shape guidance to push the generated atom positions to the shape of the given molecule. We compare ShapeMol and ShapeMol+g against state-of-the-art baseline methods. Our experimental results demonstrate that ShapeMol and ShapeMol+g could generate molecules with higher shape similarities, and competitive qualities compared to the baseline methods. In future work, we will explore generating 3D molecules jointly conditioned on the shape and the electrostatic, considering that the electrostatic of molecules could also determine the binding activities of molecules.

Acknowledgements

This project was made possible, in part, by support from the National Science Foundation grant nos. IIS-2133650 (X.N.), and The Ohio State University President’s Research Excellence program (X.N.). Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agency.

References

Acharya, C.; Coop, A.; Polli, J. E.; and MacKerell, A. D. 2011. Recent Advances in Ligand-Based Drug Design: Relevance and Utility of the Conformationally Sampled Pharmacophore Approach. *Current Computer Aided-Drug Design*, 7(1): 10–22.

Adams, K.; and Coley, C. W. 2023. Equivariant Shape-Conditioned Generation of 3D Molecules for Ligand-Based Drug Design. In *The Eleventh International Conference on Learning Representations*.

Batool, M.; Ahmad, B.; and Choi, S. 2019. A Structure-Based Drug Discovery Paradigm. *International Journal of Molecular Sciences*, 20(11): 2783.

Chen, Y.; Fernando, B.; Bilen, H.; Nießner, M.; and Gavves, E. 2022. 3D Equivariant Graph Implicit Functions. In *Lecture Notes in Computer Science*, 485–502. Springer Nature Switzerland.

Chen, Z.; Min, M. R.; Parthasarathy, S.; and Ning, X. 2021. A deep generative model for molecule optimization via one fragment modification. *Nature Machine Intelligence*, 3(12): 1040–1049.

Deng, C.; Litany, O.; Duan, Y.; Poulenard, A.; Tagliasacchi, A.; and Guibas, L. J. 2021. Vector Neurons: A General Framework for $SO(3)$ -Equivariant Networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 12200–12209.

Garcia Satorras, V.; Hoogeboom, E.; Fuchs, F.; Posner, I.; and Welling, M. 2021. E(n) Equivariant Normalizing Flows. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 4181–4192. Curran Associates, Inc.

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. 2018. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science*, 4(2): 268–276.

Guan, J.; Qian, W. W.; Peng, X.; Su, Y.; Peng, J.; and Ma, J. 2023. 3D Equivariant Diffusion for Target-Aware Molecule Generation and Affinity Prediction. In *The Eleventh International Conference on Learning Representations*.

Hawkins, P. C. D.; Skillman, A. G.; and Nicholls, A. 2006. Comparison of Shape-Matching and Docking as Virtual Screening Tools. *Journal of Medicinal Chemistry*, 50(1): 74–82.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 6840–6851. Curran Associates, Inc.

Hoogeboom, E.; Nielsen, D.; Jaini, P.; Forré, P.; and Welling, M. 2021. Argmax Flows and Multinomial Diffusion: Learning Categorical Distributions. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 12454–12465. Curran Associates, Inc.

Hoogeboom, E.; Satorras, V. G.; Vignac, C.; and Welling, M. 2022. Equivariant Diffusion for Molecule Generation in 3D. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 8867–8887. PMLR.

Imrie, F.; Hadfield, T. E.; Bradley, A. R.; and Deane, C. M. 2021. Deep generative design with 3D pharmacophoric constraints. *Chemical Science*, 12(43): 14577–14589.

Jin, W.; Barzilay, R.; and Jaakkola, T. 2018. 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*, 2323–2332. PMLR.

Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA*, 2015.

Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; and Catanzaro, B. 2021. DiffWave: A Versatile Diffusion Model for Audio Synthesis. In *International Conference on Learning Representations*.

Kullback, S.; and Leibler, R. A. 1951. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1): 79–86.

Landrum, G.; Tosco, P.; Kelley, B.; Ric; Cosgrove, D.; Sriniker; Gedeck; Vianello, R.; NadineSchneider; Kawashima, E.; N, D.; Jones, G.; Dalke, A.; Cole, B.; Swain, M.; Turk, S.; AlexanderSavelyev; Vaucher, A.; Wójcikowski, M.; Ichiru Take; Probst, D.; Ujihara, K.; Scalfani, V. F.; Godin, G.; Lehtivarjo, J.; Pahl, A.; Walker, R.; Francois Berenger; Jasondbiggs; and Stretts123. 2023. rdkit/rdkit: 2023_03_2 (Q1 2023) Release.

Luo, S.; Guan, J.; Ma, J.; and Peng, J. 2021. A 3D Generative Model for Structure-Based Drug Design. In Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*.

Nichol, A. Q.; and Dhariwal, P. 2021. Improved Denoising Diffusion Probabilistic Models. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 8162–8171. PMLR.

Papadopoulos, K.; Giblin, K. A.; Janet, J. P.; Patronov, A.; and Engkvist, O. 2021. De novo design with deep generative models based on 3D similarity scoring. *Bioorganic & Medicinal Chemistry*, 44: 116308.

Park, J. J.; Florence, P.; Straub, J.; Newcombe, R.; and Lovegrove, S. 2019. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Peng, X.; Guan, J.; Liu, Q.; and Ma, J. 2023. MolDiff: Addressing the Atom-Bond Inconsistency Problem in 3D Molecule Diffusion Generation. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 27611–27629. PMLR.

Peng, X.; Luo, S.; Guan, J.; Xie, Q.; Peng, J.; and Ma, J. 2022. Pocket2Mol: Efficient Molecular Sampling Based on 3D Protein Pockets. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvari, C.; Niu, G.; and Sabato, S., eds., *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, 17644–17655. PMLR.

Polykovskiy, D.; Zhebrak, A.; Sanchez-Lengeling, B.; Golovanov, S.; Tatanov, O.; Belyaev, S.; Kurbanov, R.; Artyamonov, A.; Aladinskiy, V.; Veselov, M.; Kadurin, A.; Johansson, S.; Chen, H.; Nikolenko, S.; Aspuru-Guzik, A.; and Zhavoronkov, A. 2020. Molecular Sets (MOSES): A Benchmarking Platform for Molecular Generation Models. *Frontiers in Pharmacology*, 11.

Ravi, N.; Reizenstein, J.; Novotny, D.; Gordon, T.; Lo, W.-Y.; Johnson, J.; and Gkioxari, G. 2020. Accelerating 3D Deep Learning with PyTorch3D. *arXiv:2007.08501*.

Ripphausen, P.; Nisius, B.; and Bajorath, J. 2011. State-of-the-art in ligand-based virtual screening. *Drug Discovery Today*, 16(9-10): 372–376.

Vainio, M. J.; Puranen, J. S.; and Johnson, M. S. 2009. ShaEP: Molecular Overlay Based on Shape and Electrostatic Potential. *Journal of Chemical Information and Modeling*, 49(2): 492–502.

Wang, Y.; Sun, Y.; Liu, Z.; Sarma, S. E.; Bronstein, M. M.; and Solomon, J. M. 2019. Dynamic Graph CNN for Learning on Point Clouds. *ACM Trans. Graph.*, 38(5).

Wójcikowski, M.; Zielenkiewicz, P.; and Siedlecki, P. 2015. Open Drug Discovery Toolkit (ODDT): a new open-source player in the drug discovery field. *Journal of Cheminformatics*, 7(1).

Zheng, H.; Hou, J.; Zimmerman, M. D.; Wlodawer, A.; and Minor, W. 2013. The future of crystallography in drug discovery. *Expert Opinion on Drug Discovery*, 9(2): 125–137.

Shape-conditioned 3D Molecule Generation via Equivariant Diffusion Models (Supplementary Materials)

Ziqi Chen¹, Bo Peng¹, Srinivasan Parthasarathy^{1,2}, Xia Ning^{1,2,3*},

¹Computer Science and Engineering, The Ohio State University, Columbus, OH 43210

²Translational Data Analytics Institute, The Ohio State University, Columbus, OH 43210

³Biomedical Informatics, The Ohio State University, Columbus, OH 43210

chen.8484@buckeyemail.osu.edu, peng.307@buckeyemail.osu.edu, srini@cse.ohio-state.edu, ning.104@osu.edu

S1 Parameters for Reproducibility

We implemented both SE and DIFF using Python-3.7.16, PyTorch-1.11.0, PyTorch-scatter-2.0.9, Numpy-1.21.5, Scikit-learn-1.0.2. We trained the models using a Tesla V100 GPU with 32GB memory and a CPU with 80GB memory on Red Hat Enterprise 7.7. We released the code, data, and the trained model at Google Drive¹.

S1.1 Parameters of SE

In SE, we tuned the dimension of all the hidden layers including VN-DGCNN layers (Eq.), MLP layers (Eq. 1) and VN-In layer (Eq. 1), and the dimension d_p of generated shape latent embeddings \mathbf{H}^s with the grid-search algorithm in the parameter space presented in Table S1. We determined the optimal hyper-parameters according to the mean squared errors of the predictions of signed distances for 1,000 validation molecules that are selected as described in Section “Data” in the main manuscript. The optimal dimension of all the hidden layers is 128, and the optimal dimension d_p of shape latent embedding \mathbf{H}^s is 32. The optimal number of points $|\mathcal{P}|$ in the point cloud \mathcal{P} is 512. We sampled 1,024 query points in \mathcal{Z} for each molecule shape. We constructed graphs from point clouds, which are employed to learn \mathbf{H}^s with VN-DGCNN layer (Eq.), using the k -nearest neighbors based on Euclidean distance with $k = 20$. We set the number of VN-DGCNN layers as 4. We set the number of MLP layers in the decoder (Eq. 1) as 4. We set the number of VN-In layers as 1.

We optimized the SE model via Adam (Kingma and Ba 2015) with its parameters (0.950, 0.999), learning rate 0.001, and batch size 16. We evaluated the validation loss every 2,000 training steps. We scheduled to decay the learning rate with a factor of 0.6 and a minimum learning rate of 1e-6 if the validation loss does not decrease in 5 consecutive evaluations. The optimal SE model has 27.3K learnable parameters. We trained the SE model with \sim 172,600 training steps. The training took 80 hours with our GPUs. The trained SE model achieved the minimum validation loss at 162,000 steps.

*Contact author

¹<https://drive.google.com/drive/folders/146cpjuwenKGTd6Zh4sYBy-Wv6BMfGwe4?usp=sharing>

Table S1: Hyper-Parameter Space for SE Optimization

| Hyper-parameters | Space |
|-----------------------------------|-------------|
| hidden layer dimension | {64, 128} |
| dimension d_p of \mathbf{H}^s | {32, 64} |
| # points in \mathcal{P} | {512, 1024} |
| # query points in \mathcal{Z} | 1024 |
| # nearest neighbors | 20 |
| # VN-DGCNN layers (Eq) | 4 |
| # MLP layers in Eq 1 | 4 |

Table S2: Hyper-Parameter Space for DIFF Optimization

| Hyper-parameters | Space |
|--|-------|
| hidden layer dimension | 128 |
| weight of atom type loss ξ (Eq. 22) | 100 |
| threshold of step weight δ (Eq. 20) | 10 |
| # atom features K | 15 |
| # EQ-GNN/INV-GNN layers | 8 |
| # heads n_h in $\text{MHA}^x/\text{MHA}^v$ | 16 |
| # nearest neighbors N (Eq. 16 and 18) | 8 |
| # diffusion steps T | 1,000 |

S1.2 Parameters of DIFF

Table S2 presents the parameters used to train DIFF. In DIFF, we set the dimension of all the hidden layers including MHA^x and VN-Lin layer (Eq. 16), MHA^h and VN-In layer (Eq. 18) and MLP layer (Eq. 19) as 128. We set the number of layers L in EQ-GNN and INV-GNN as 8. Each layer is a multi-head attention layer (MHA^x or MHA^h) with 16 heads. We set the number of VN-Lin and VN-In layer as 1, and the number of MLP layer as 2.

We constructed graphs from atoms in molecules, which are employed to predict atom coordinates and features (Eq. 16 and 18), using the N -nearest neighbors based on Euclidean distance with $N = 8$. We used $K = 15$ atom features in total, indicating the atom types and its aromaticity. These atom features include 10 non-aromatic atoms (i.e., “H”, “C”, “N”, “O”, “F”, “P”, “S”, “Cl”, “Br”, “I”), and 5 aromatic atoms (i.e., “C”, “N”, “O”, “P”, “S”). We set the number of diffusion steps T as 1,000. We set the weight ξ of atom type loss (Eq. 22) as 100 to balance the values of atom type loss and atom coordinate loss. We set

the threshold δ (Eq. 20) as 10. The parameters β_t^x and β_t^y of variance scheduling in the forward diffusion process of DIFF are discussed in Supplementary Section S4.1. Following SQuID, we did not perform extensive hyperparameter tuning for DIFF given that the used hyperparameters have enabled good performance.

We optimized the DIFF model via Adam (Kingma and Ba 2015) with its parameters (0.950, 0.999), learning rate 0.001, and batch size 32. We evaluated the validation loss every 2,000 training steps. We scheduled to decay the learning rate with a factor of 0.6 and a minimum learning rate of 1e-5 if the validation loss does not decrease in 10 consecutive evaluations. The DIFF model has 2.7M learnable parameters. We trained the DIFF model with $\sim 900,000$ training steps. The training took 60 hours with our GPUs. The trained DIFF achieved the minimum validation loss at 746,000 steps.

During inference, given a condition molecule, we determined the number of atoms for the molecule to be generated, according to the shape size of the condition molecule. We assumed that molecules with similar shape sizes should have similar numbers of atoms. Therefore, we identified the training molecules with similar shape sizes to the condition molecule. From these molecules, we built an atom number distribution and sampled the number of atoms in the molecule to be generated from it. Particularly, to calculate the shape size of each molecule, we first converted it into an atomic density grid of side length 22 Å with 0.5 Å resolution and then counted the non-empty voxels in this grid. For a condition molecule with shape size n_s , we selected training molecules with shape sizes in the range $[n_s - 200, n_s + 200]$ to build the atom number distribution. Following Adams and Coley (Adams and Coley 2023), we set the variance ϕ of atom-centered Gaussians as 0.049, which is used to build a set of points for shape guidance in Section “ShapeMol with Shape Guidance” in the main manuscript. The optimal distance threshold γ is 0.2, and the optimal stop step S for shape guidance is 300. With shape guidance, each time we updated the atom position (Eq. 23), we randomly sampled the weight σ from [0.2, 0.8]. For each condition molecule, it took around 40 seconds on average to generate 50 molecule candidates with our GPUs.

S2 Point Cloud Construction

In ShapeMol, we represented molecular surface shapes using point clouds (\mathcal{P}). \mathcal{P} serves as input to ShapeMol-enc, from which we derive shape latent embeddings. To generate \mathcal{P} , we initially generated a molecular surface mesh using the algorithm from the Open Drug Discovery Toolkit (Wójcikowski, Zielenkiewicz, and Siedlecki 2015). Following this, we uniformly sampled points on the mesh surface with probability proportional to the face area, using the algorithm from PyTorch3D (Ravi et al. 2020). This point cloud \mathcal{P} is then centralized by setting the center of its points to zero.

S3 Query Point Sampling

As described in Section “Shape Decoder (SE-dec)”, the signed distances of query points z_q to molecule surface shape \mathcal{P} are used to optimize SE. In this sec-

tion, we present how to sample these points z_q in 3D space. Particularly, we first determined the bounding box around the molecular surface shape, using the maximum and minimum (x, y, z) -axis coordinates for points in our point cloud \mathcal{P} , denoted as $(x_{\min}, y_{\min}, z_{\min})$ and $(x_{\max}, y_{\max}, z_{\max})$. We extended this box slightly by defining its corners as $(x_{\min} - 1, y_{\min} - 1, z_{\min} - 1)$ and $(x_{\max} + 1, y_{\max} + 1, z_{\max} + 1)$. For sampling $|\mathcal{Z}|$ query points, we wanted an even distribution of points inside and outside the molecule surface shape. When a bounding box is defined around the molecule surface shape, there could be a lot of empty spaces within the box that the molecule does not occupy due to its complex and irregular shape. This could lead to that fewer points within the molecule surface shape could be sampled within the box. Therefore, we started by randomly sampling $3k$ points within our bounding box to ensure that there are sufficient points within the surface. We then determined whether each point lies within the molecular surface, using an algorithm from Trimesh ² based on the molecule surface mesh. If there are n_w points found within the surface, we selected $n = \min(n_w, k/2)$ points from these points, and randomly choose the remaining $k - n$ points from those outside the surface. For each query point, we determined its signed distance to the molecule surface by its closest distance to points in \mathcal{P} with a sign indicating whether it is inside the surface.

S4 Forward Diffusion (DIFF-forward)

S4.1 Variance Scheduling in DIFF-forward

Following Guan *et al.* (Guan et al. 2023), we used a sigmoid β schedule for the variance schedule β_t^x of atom coordinates as below,

$$\beta_t^x = \text{sigmoid}(w_1(2t/T - 1))(w_2 - w_3) + w_3 \quad (S1)$$

in which w_i ($i=1, 2$, or 3) are hyperparameters; T is the maximum step. We set $w_1 = 6$, $w_2 = 1.e - 7$ and $w_3 = 0.01$. For atom types, we used a cosine β schedule (Nichol and Dhariwal 2021) for β_t^y as below,

$$\begin{aligned} \bar{\alpha}_t^y &= \frac{f(t)}{f(0)}, f(t) = \cos\left(\frac{t/T + s}{1 + s} \cdot \frac{\pi}{2}\right)^2 \\ \beta_t^y &= 1 - \alpha_t^y = 1 - \frac{\bar{\alpha}_t^y}{\bar{\alpha}_{t-1}^y} \end{aligned} \quad (S2)$$

in which s is a hyperparameter and set as 0.01.

As shown in Section “Forward Diffusion Process”, the values of β_t^x and β_t^y should be sufficiently small to ensure the smoothness of forward diffusion process. In the meanwhile, their corresponding $\bar{\alpha}_t$ values should decrease from 1 to 0 over $t = [1, T]$. Figure S1 shows the values of β_t and $\bar{\alpha}_t$ for atom coordinates and atom types with our hyperparameters. Please note that the value of β_t^x is less than 0.1 for 990 out of 1,000 steps. This guarantees the smoothness of the forward diffusion process.

²<https://trimesh.org/>

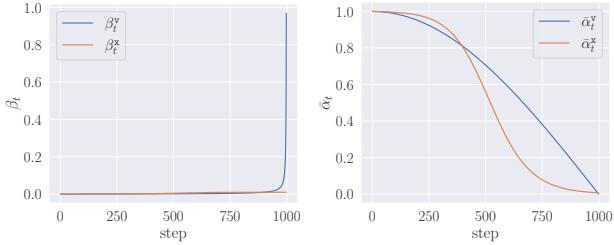


Figure S1: Schedule

S4.2 Derivation of Forward Diffusion Process

In ShapeMol, a Gaussian noise and a categorical noise are added to continuous atom position and discrete atom features, respectively. Here, we briefly describe the derivation of posterior equations (i.e., Eq. 8, and 10) for atom positions and atom types in our work. We refer readers to Ho *et al.* (2020) and Kong *et al.* (2021) for a detailed description of diffusion process for continuous variables and Hoogeboom *et al.* (Hoogeboom et al. 2021) for the description of diffusion process for discrete variables.

For continuous atom positions, as shown in Kong *et al.* (2021), according to Bayes theorem, given $q(\mathbf{x}_t|\mathbf{x}_{t-1})$ defined in Eq. 3 and $q(\mathbf{x}_t|\mathbf{x}_0)$ defined in Eq. 5, the posterior $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ is derived as below (superscript x is omitted for brevity),

$$\begin{aligned}
q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) &= \frac{q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{x}_0)q(\mathbf{x}_{t-1}|\mathbf{x}_0)}{q(\mathbf{x}_t|\mathbf{x}_0)} \\
&= \frac{\mathcal{N}(\mathbf{x}_t|\sqrt{1-\beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I})\mathcal{N}(\mathbf{x}_{t-1}|\sqrt{\bar{\alpha}_{t-1}}\mathbf{x}_0, (1-\bar{\alpha}_{t-1})\mathbf{I})}{\mathcal{N}(\mathbf{x}_t|\sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1-\bar{\alpha}_t)\mathbf{I})} \\
&= (2\pi\beta_t)^{-\frac{3}{2}}(2\pi(1-\bar{\alpha}_{t-1}))^{-\frac{3}{2}}(2\pi(1-\bar{\alpha}_t))^{\frac{3}{2}} \times \exp\left(-\frac{\|\mathbf{x}_t - \sqrt{\bar{\alpha}_t}\mathbf{x}_{t-1}\|^2}{2\beta_t} - \frac{\|\mathbf{x}_{t-1} - \sqrt{\bar{\alpha}_{t-1}}\mathbf{x}_0\|^2}{2(1-\bar{\alpha}_{t-1})}\right. \\
&\quad \left. + \frac{\|\mathbf{x}_t - \sqrt{\bar{\alpha}_t}\mathbf{x}_0\|^2}{2(1-\bar{\alpha}_t)}\right) \\
&= (2\pi\tilde{\beta}_t)^{-\frac{3}{2}} \exp\left(-\frac{1}{2\tilde{\beta}_t}\|\mathbf{x}_{t-1} - \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t}\mathbf{x}_0 - \frac{\sqrt{\bar{\alpha}_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}\mathbf{x}_t\|^2\right) \\
\text{where } \tilde{\beta}_t &= \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}\beta_t.
\end{aligned} \tag{S3}$$

Therefore, the posterior of atom positions is derived as below,

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}|\frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t}\mathbf{x}_0 + \frac{\sqrt{\bar{\alpha}_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}\mathbf{x}_t, \tilde{\beta}_t\mathbf{I}). \tag{S4}$$

For discrete atom features, as shown in Hoogeboom *et al.* (Hoogeboom et al. 2021) and Guan *et al.* (Guan et al. 2023), according to Bayes theorem, the posterior

$q(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathbf{v}_0)$ is derived as below (superscript v is omitted for brevity),

$$q(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathbf{v}_0) = \frac{q(\mathbf{v}_t|\mathbf{v}_{t-1}, \mathbf{v}_0)q(\mathbf{v}_{t-1}|\mathbf{v}_0)}{\sum_{\mathbf{v}_{t-1}} q(\mathbf{v}_t|\mathbf{v}_{t-1}, \mathbf{v}_0)q(\mathbf{v}_{t-1}|\mathbf{v}_0)} \tag{S5}$$

For $q(\mathbf{v}_t|\mathbf{v}_{t-1}, \mathbf{v}_0)$, we have

$$\begin{aligned}
q(\mathbf{v}_t|\mathbf{v}_{t-1}, \mathbf{v}_0) &= \mathcal{C}(\mathbf{v}_t|(1-\beta_t)\mathbf{v}_{t-1} + \beta_t/K) \\
&= \begin{cases} 1 - \beta_t + \beta_t/K, & \text{when } \mathbf{v}_t = \mathbf{v}_{t-1}, \\ \beta_t/K, & \text{when } \mathbf{v}_t \neq \mathbf{v}_{t-1}, \end{cases} \\
&= \mathcal{C}(\mathbf{v}_{t-1}|(1-\beta_t)\mathbf{v}_t + \beta_t/K).
\end{aligned} \tag{S6}$$

Therefore, we have

$$\begin{aligned}
&q(\mathbf{v}_t|\mathbf{v}_{t-1}, \mathbf{v}_0)q(\mathbf{v}_{t-1}|\mathbf{v}_0) \\
&= \mathcal{C}(\mathbf{v}_{t-1}|(1-\beta_t)\mathbf{v}_t + \beta_t\frac{1}{K})\mathcal{C}(\mathbf{v}_{t-1}|\bar{\alpha}_{t-1}\mathbf{v}_0 + (1-\bar{\alpha}_{t-1})\frac{1}{K}) \\
&= [\alpha_t\mathbf{v}_t + \frac{1-\alpha_t}{K}] \odot [\bar{\alpha}_{t-1}\mathbf{v}_0 + \frac{1-\bar{\alpha}_{t-1}}{K}].
\end{aligned} \tag{S7}$$

Therefore, with $q(\mathbf{v}_t|\mathbf{v}_{t-1}, \mathbf{v}_0)q(\mathbf{v}_{t-1}|\mathbf{v}_0) = \tilde{\mathbf{c}}$, the posterior is as below,

$$q(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathbf{v}_0) = \mathcal{C}(\mathbf{v}_{t-1}|\mathbf{c}(\mathbf{v}_t, \mathbf{v}_0)) = \frac{\tilde{\mathbf{c}}}{\sum_k^K \tilde{c}_k}. \tag{S8}$$

S5 Backward Generative Process (DIFF-backward)

S5.1 Proof of Equivariance in Atom Coordinate Prediction

In this section, we prove that our design in Eq. 16 in the main manuscript updates atom coordinates in an equivariant way. A function $f(\mathbf{x})$ is equivariant with respect to rotation, if given any rotation matrix $\mathbf{R} \in \mathbb{R}^{3 \times 3}$, it satisfies,

$$f(\mathbf{R}\mathbf{x}) = \mathbf{R}f(\mathbf{x}). \tag{S9}$$

In Eq. 16, the embedding $\Delta\mathbf{x}_i^{l+1} \in \mathbb{R}^{n_h \times 3}$, which aggregates the neighborhood information of a_i , is updated by atom coordinates \mathbf{x}_i^l from previous layer and embeddings from multi-head attention layer $\text{MHA}^x(\cdot)$ as below,

$$\begin{aligned}
\Delta\mathbf{x}_i^{l+1} &= \sum_{j \in N(a_i), i \neq j} (\mathbf{x}_i^l - \mathbf{x}_j^l) \text{MHA}^x(d_{ij}^l, \mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, \text{VN-In}(\mathbf{H}^s)), \\
\mathbf{x}_i^{l+1} &= \mathbf{x}_i^l + \text{Mean}(\Delta\mathbf{x}_i^{l+1}) + \text{VN-Lin}(\mathbf{x}_i^l, \Delta\mathbf{x}_i^{l+1}, \mathbf{H}^s).
\end{aligned}$$

The embeddings from $\text{MHA}^x(\cdot)$ and atom feature embeddings \mathbf{h}_i^l will be proved to be invariant in Supplementary Section S5.2 later. Therefore, $\text{MHA}^x(\cdot)$ and \mathbf{h}_i^{l+1} remain the same with rotation \mathbf{R} . Then, we can prove that $\Delta\mathbf{x}_i^{l+1}$ is updated in an equivariant way as below,

$$\begin{aligned}
&\sum_{j \in N(a_i), i \neq j} (\mathbf{R}\mathbf{x}_i^l - \mathbf{R}\mathbf{x}_j^l) \text{MHA}^x(d_{ij}^l, \mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, \text{VN-In}(\mathbf{H}^s)) \\
&= \mathbf{R} \sum_{j \in N(a_i), i \neq j} (\mathbf{x}_i^l - \mathbf{x}_j^l) \text{MHA}^x(d_{ij}^l, \mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, \text{VN-In}(\mathbf{H}^s)) \\
&= \mathbf{R} \Delta\mathbf{x}_i^{l+1}.
\end{aligned}$$

Given that mean pooling is a linear operation that respects rotation equivariance, $\text{Mean}(\Delta \mathbf{x}_i^{l+1})$ is equivariant to rotations. The vector-neuron-based linear layer $\text{VN-Lin}(\cdot)$ with leaky-ReLu activation function also produces rotational equivariant embeddings according to Deng *et al.* (Deng et al. 2021). Therefore, we prove that atom coordinates \mathbf{x}_i^{l+1} are updated in an equivariant way as below,

$$\begin{aligned}
& \mathbf{R}\mathbf{x}_i^l + \text{Mean}(\mathbf{R}\Delta\mathbf{x}_i^{l+1}) + \text{VN-Lin}(\mathbf{R}\mathbf{x}_i^l, \mathbf{R}\Delta\mathbf{x}_i^{l+1}, \mathbf{R}\mathbf{H}^s) \\
&= \mathbf{R}\mathbf{x}_i^l + \mathbf{R}\text{Mean}(\Delta\mathbf{x}_i^{l+1}) + \mathbf{R}\text{VN-Lin}(\mathbf{x}_i^l, \Delta\mathbf{x}_i^{l+1}, \mathbf{H}^s) \\
&= \mathbf{R}(\mathbf{x}_i^l + \text{Mean}(\Delta\mathbf{x}_i^{l+1}) + \text{VN-Lin}(\mathbf{x}_i^l, \Delta\mathbf{x}_i^{l+1}, \mathbf{H}^s)) \\
&= \mathbf{R}\mathbf{x}_i^{l+1}
\end{aligned} \tag{S10}$$

S5.2 Proof of Invariance in Atom Feature Prediction

In this section, we prove that our design in Eq. 18 in the main manuscript updates atom features in an invariant way. A function $f(\mathbf{x})$ is invariant with respect to rotation, if it satisfies,

$$f(\mathbf{R}\mathbf{x}) = f(\mathbf{x}). \tag{S11}$$

Recall that in Eq. 18, atom feature embeddings \mathbf{h}_i^{l+1} is updated according to atomic distances d_{ij}^l , embeddings \mathbf{h}_i from previous layer and invariant shape embeddings $\text{VN-In}(\mathbf{H}^s)$ from the vector-neuron based invariant layer adapted by (Chen et al. 2022) as follows,

$$\mathbf{h}_i^{l+1} = \mathbf{h}_i^l + \sum_{j \in N(a_i), i \neq j} \text{MHA}^h(d_{ij}^l, \mathbf{h}_i^l, \mathbf{h}_j^l, \text{VN-In}(\mathbf{H}^s)), \mathbf{h}_i^0 = \mathbf{v}_i.$$

Atomic distances are invariant to transformations. Atom feature embeddings are initialized with atom feature vectors, which are also invariant to transformations. The $\text{VN-In}(\cdot)$ layer maps the equivariant shape embedding $\mathbf{H}^s \in \mathbb{R}^{d_p \times 3}$ into the invariant shape embedding $\mathbf{h}^s \in \mathbb{R}^{d_p}$ as below,

$$\mathbf{h}_c^s = \text{MLP}(\langle \mathbf{H}_c^s, \frac{\bar{\mathbf{H}}^s}{\|\bar{\mathbf{H}}^s\|} \rangle).$$

where $\bar{\mathbf{H}}^s = \frac{1}{d_p} \sum_{c=1}^{d_p} \mathbf{H}_c^s$ and \mathbf{H}_c^s is the c -th channel of \mathbf{H}^s ; \mathbf{h}_c^s denotes the c -th channel of \mathbf{h}^s ; $\langle \cdot \rangle$ is the inner-product operation which the results will not be changed by any rotation vector. Given that all the variables in the right-hand side of Eq. 18 are invariant to rotations, we prove that the atom feature embeddings are updated in an invariant way. Therefore, the invariant atom feature embeddings lead to invariant atom feature predictions in Eq. 19. Similarly, we can prove that the embeddings from MHA^x layer (Eq. 16) are invariant to rotations.

S6 ShapeMol Loss Function Derivation

In this section, we demonstrate that a step weight w_t^x based on the signal-to-noise ratio λ_t should be included into the

atom position loss (Eq. 20). In the diffusion process for continuous variables, the optimization problem is defined as below (Ho, Jain, and Abbeel 2020),

$$\begin{aligned}
& \arg \min_{\Theta} KL(q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) | p_{\Theta}(\mathbf{x}_{t-1} | \mathbf{x}_t, \tilde{\mathbf{x}}_{0,t})) \\
&= \arg \min_{\Theta} \frac{\bar{\alpha}_{t-1}(1 - \alpha_t)}{2(1 - \bar{\alpha}_{t-1})(1 - \bar{\alpha}_t)} \|\tilde{\mathbf{x}}_{0,t} - \mathbf{x}_0\|^2 \\
&= \arg \min_{\Theta} \frac{1 - \alpha_t}{2(1 - \bar{\alpha}_{t-1})\alpha_t} \|\tilde{\epsilon}_{0,t} - \epsilon_0\|^2,
\end{aligned}$$

where $\epsilon_0 = \frac{\mathbf{x}_t - \sqrt{\bar{\alpha}_t} \mathbf{x}_0}{\sqrt{1 - \bar{\alpha}_t}}$ is the ground-truth noise variable sampled from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and is used to sample \mathbf{x}_t from $\mathcal{N}(\mathbf{x}_t | \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I})$ in Eq. 6; $\tilde{\epsilon}_0 = \frac{\mathbf{x}_t - \sqrt{\bar{\alpha}_t} \tilde{\mathbf{x}}_{0,t}}{\sqrt{1 - \bar{\alpha}_t}}$ is the predicted noise variable.

Ho *et al.* (Ho, Jain, and Abbeel 2020) further simplified the above objective as below and demonstrated that the simplified one can achieve better performance:

$$\begin{aligned}
& \arg \min_{\Theta} \|\tilde{\epsilon}_{0,t} - \epsilon_0\|^2 \\
&= \arg \min_{\Theta} \frac{\bar{\alpha}_t}{1 - \bar{\alpha}_t} \|\tilde{\mathbf{x}}_{0,t} - \mathbf{x}_0\|^2,
\end{aligned} \tag{S12}$$

where $\lambda_t = \frac{\bar{\alpha}_t}{1 - \bar{\alpha}_t}$ is the signal-to-noise ratio. While previous work (Guan et al. 2023) applies uniform step weights across different steps, we demonstrate that a step weight should be included into the atom position loss according to Eq. S12. However, the value of λ_t could be very large when $\bar{\alpha}_t$ is close to 1 as t approaches 1. Therefore, we clip the value of λ_t with threshold δ in Eq. 20.

S7 Evaluation Metrics

We calculated the shape similarity $\text{Sim}_s(s_x, s_y)$ via the overlap volumes between two aligned molecules as in the literature (Adams and Coley 2023). Each molecule candidate M_y for evaluation is aligned with the condition molecule M_x by the ShaEP tool (Vainio, Puranen, and Johnson 2009). For the molecular graph similarity $\text{Sim}_g(M_x, M_y)$, we used the Tanimoto similarity over Morgan fingerprints between M_x and M_y calculated by RDKit (Landrum et al. 2023).

References

Adams, K.; and Coley, C. W. 2023. Equivariant Shape-Conditioned Generation of 3D Molecules for Ligand-Based Drug Design. In *The Eleventh International Conference on Learning Representations*.

Chen, Y.; Fernando, B.; Bilen, H.; Nießner, M.; and Gavves, E. 2022. 3D Equivariant Graph Implicit Functions. In *Lecture Notes in Computer Science*, 485–502. Springer Nature Switzerland.

Deng, C.; Litany, O.; Duan, Y.; Poulenard, A.; Tagliasacchi, A.; and Guibas, L. J. 2021. Vector Neurons: A General Framework for $\text{SO}(3)$ -Equivariant Networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 12200–12209.

Guan, J.; Qian, W. W.; Peng, X.; Su, Y.; Peng, J.; and Ma, J. 2023. 3D Equivariant Diffusion for Target-Aware Molecule

Generation and Affinity Prediction. In *The Eleventh International Conference on Learning Representations*.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 6840–6851. Curran Associates, Inc.

Hoogeboom, E.; Nielsen, D.; Jaini, P.; Forré, P.; and Welling, M. 2021. Argmax Flows and Multinomial Diffusion: Learning Categorical Distributions. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 12454–12465. Curran Associates, Inc.

Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, 2015*.

Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; and Catanzaro, B. 2021. DiffWave: A Versatile Diffusion Model for Audio Synthesis. In *International Conference on Learning Representations*.

Landrum, G.; Tosco, P.; Kelley, B.; Ric; Cosgrove, D.; Sriniker; Gedeck; Vianello, R.; NadineSchneider; Kawashima, E.; N, D.; Jones, G.; Dalke, A.; Cole, B.; Swain, M.; Turk, S.; AlexanderSavelyev; Vaucher, A.; Wójcikowski, M.; Ichiru Take; Probst, D.; Ujihara, K.; Scalfani, V. F.; Godin, G.; Lehtivarjo, J.; Pahl, A.; Walker, R.; Francois Berenger; Jasonbiggs; and Strets123. 2023. rdkit/rdkit: 2023_03_2 (Q1 2023) Release.

Nichol, A. Q.; and Dhariwal, P. 2021. Improved Denoising Diffusion Probabilistic Models. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 8162–8171. PMLR.

Ravi, N.; Reizenstein, J.; Novotny, D.; Gordon, T.; Lo, W.-Y.; Johnson, J.; and Gkioxari, G. 2020. Accelerating 3D Deep Learning with PyTorch3D. *arXiv:2007.08501*.

Vainio, M. J.; Puranen, J. S.; and Johnson, M. S. 2009. ShaEP: Molecular Overlay Based on Shape and Electrostatic Potential. *Journal of Chemical Information and Modeling*, 49(2): 492–502.

Wójcikowski, M.; Zielenkiewicz, P.; and Siedlecki, P. 2015. Open Drug Discovery Toolkit (ODDT): a new open-source player in the drug discovery field. *Journal of Cheminformatics*, 7(1).