
Reprogramming Pretrained Target-Specific Diffusion Models for Dual-Target Drug Design

Xiangxin Zhou^{1,2}

Jiaqi Guan³

Yijia Zhang⁴

Xingang Peng⁵

Liang Wang^{1,2}

Jianzhu Ma^{4,6,*}

¹School of Artificial Intelligence, University of Chinese Academy of Sciences

²New Laboratory of Pattern Recognition (NLPR),

State Key Laboratory of Multimodal Artificial Intelligence Systems (MAIS),
Institute of Automation, Chinese Academy of Sciences (CASIA)

³Department of Computer Science, University of Illinois Urbana-Champaign

⁴Department of Electronic Engineering, Tsinghua University

⁵Institute for Artificial Intelligence, Peking University

⁶Institute for AI Industry Research, Tsinghua University

Abstract

Dual-target therapeutic strategies have become a compelling approach and attracted significant attention due to various benefits, such as their potential in overcoming drug resistance in cancer therapy. Considering the tremendous success that deep generative models have achieved in structure-based drug design in recent years, we formulate dual-target drug design as a generative task and curate a novel dataset of potential target pairs based on synergistic drug combinations. We propose to design dual-target drugs with diffusion models that are trained on single-target protein-ligand complex pairs. Specifically, we align two pockets in 3D space with protein-ligand binding priors and build two complex graphs with shared ligand nodes for $SE(3)$ -equivariant composed message passing, based on which we derive a composed drift in both 3D and categorical probability space in the generative process. Our algorithm can well transfer the knowledge gained in single-target pretraining to dual-target scenarios in a zero-shot manner. We also repurpose linker design methods as strong baselines for this task. Extensive experiments demonstrate the effectiveness of our method compared with various baselines.

1 Introduction

A promising paradigm of rational drug design is structure-based drug design (SBDD) [1], which uses computational chemistry tools in which the 3D structure of a protein target is used as the basis to identify or design new chemical entities. The foundation of structure-based drug design has been grounded in the lock-and-key hypothesis, positing that an optimal ligand molecule should possess a structure that is complementary to the target site. Recently, dual-target drug design, which aims to design “one key” for “two locks”, has attracted significant attention. Precisely, dual-target drug design [4] is a strategy in pharmaceutical research that aims to design single ligand molecules capable of interacting with two different biological targets simultaneously. A dual-target drug can potentially lower the odds of resistance developing [62] and effectively manage the disease which involve complex biological pathways with multiple proteins [49]. Recent years have witnessed a noticeable

*Correspondence to: Jianzhu Ma <majianzhu@tsinghua.edu.cn>.

increase in the FDA’s approval of dual-target drugs [32, 34]. Please refer to Appendix A for a more comprehensive understanding of motivation, significance, and current practices of dual-target drug design.

Deep learning, particularly deep generative models [63] and geometric deep learning [46], has been introduced to SBDD and achieved promising results. Peng et al. [45], Zhang et al. [65] proposed to sequentially generate atoms or fragments using auto-regressive generative models conditioned on a specific protein binding site. Guan et al. [19], Lin et al. [33], Schneuing et al. [53] proposed to generate ligand molecules with diffusion models and achieved high binding affinity. However, due to the scarcity of data resources and high computational complexity, there is limited progress on introducing powerful generative models into dual-target drug design. Besides, there also lacks a comprehensive benchmark and dataset for evaluating the dual-target drug design, which also hinders the community from developing AI-powered computational tools for dual-target drug design.

To overcome the aforementioned challenges, we first propose a dataset for dual-target drug design. The design of dual-target drugs for arbitrary target pairs lacks substantive purpose. Inspired by the concept of drug synergism [58], where the combined effect of two drugs surpasses the effects of each drug when used individually, we carefully select pairs of targets from combinations of drugs that demonstrate significant synergistic interactions. The effectiveness of such combination therapy [39, 50, 44] has demonstrated significant efficacy in tumor eradication at both cellular level and in vivo study. Designing dual-target drugs for the paired targets may further improve the efficacy and reduce side effects. Additionally, we also provide a reference ligand for each target and the 3D structure of each protein-ligand complex in our dataset. Besides, we formulate the dual-target drug design as a generative task, based on which we further propose to reprogram pretrained target-specific diffusion models as introduced by Guan et al. [19] for the dual-target setting in zero-shot manner. More specifically, we first align dual pockets in 3D space with protein-ligand interaction priors that encapsulate the intricate features of the pockets. We compose the predicted drift terms in both 3D and categorical probability space in the reverse generative process of the diffusion model to generate dual-target drugs. We name this method as COMPDIFF. We further improve this method by building two complex graphs with shared ligand nodes for SE(3)-equivariant composed message passing. In this method, we compose the SE(3)-equivariant message at each layer of the equivariant neural network instead of only on the output level. We name this method as DUALDIFF. Our approach effectively transfers the knowledge acquired from pretraining on single-target datasets, circumventing the challenging demand for extensive training data required for dual-target drug design. We also repurpose linker design methods [22, 17] as a strong baseline for this task. We outline strategies to identify potential fragments from the synergistic drug combinations, serving as input for these linker design methods. We highlight our main contributions as follows:

- We present a meticulously curated dataset derived from synergistic drug combinations for dual-target drug design, offering new opportunities for AI-driven drug discovery.
- We propose SE(3)-equivariant composed message for compositional generative sampling to reprogram pretrained single-target diffusion models for dual-target drug design in a zero-shot way.
- We propose fragment selection methods from synergistic drug combinations for repurposing linker design methods as strong baselines for dual-target drug design.
- Our method can be viewed as a general framework where any pretrained generative models for SBDD can be applied to dual-target drug design without any fine-tuning. We select TargetDiff as a demonstrative demo in our work.

2 Related Work

Structure-based Drug Design Structure-based drug design (SBDD) aims to design ligand molecules that can bind to specific protein targets. The introduction of deep generative models has marked a paradigm shift, yielding noteworthy outcomes. Ragoza et al. [48] utilized a variational autoencoder to generate 3D molecules within atomic density grids. Luo et al. [42], Peng et al. [45], Liu et al. [36] employed an autoregressive model to sequentially construct 3D molecules atom by atom, while Zhang et al. [65] introduced a method for generating 3D molecules by successively predicting molecular fragments in an auto-regressive way. Guan et al. [19], Schneuing et al. [53], Lin et al. [33] introduced diffusion models [21] to SBDD, which first generate the types and positions of atoms by iteratively denoising with an SE(3)-equivariant neural network [52, 18] and then determine bond types by post-processing. Some recent studies have endeavored to further improve

the aforementioned methods through the integration of biochemical prior knowledge. Guan et al. [20] proposed decomposed priors, bond diffusion and validity guidance to improve the quality of ligand molecules generated by diffusion models. Zhang and Liu [64] augmented molecule generation through global interaction between subpocket prototypes and molecular motifs. Huang et al. [23] incorporated protein-ligand interaction prior into both forward and reverse processes to improve the diffusion models. Zhou et al. [66] integrated conditional diffusion models with iterative optimization to optimize properties of generated molecules. The above works focus on structure-based single-target drug design, while our work aims at dual-target drug design.

Molecular Linker Design Molecular linker design, which enables the connection of molecular fragments to form potent compounds, is an effective approach in rational drug discovery. Approaches like DeLinker [26] and Develop [27] design linkers by utilizing molecular graphs with distance and angle information between anchor atoms, but they lack 3D structural information of molecules. More recent techniques, such as 3DLinker [22] and DiffLinker [25], generate linkers directly in 3D space using conditional VAEs and diffusion models, respectively, but they assume known fragment poses. LinkerNet [17] relaxes this assumption by co-designing molecular fragment poses and the linker, making it applicable in cases where fragment poses are unknown, such as in the linker design of PROTACs (PROteolysis TArgeting Chimeras). Since pharmacophore combination is a traditional strategy to design dual-target drugs, we repurpose linker design methods as strong baselines for dual-target drug design. Please refer to Appendix B for extended related works.

3 Method

In this section, we will present the pipeline of our work, from dataset curation to method. In Section 3.1, we will introduce how we curate the dual-target dataset based on synergistic drug combinations and how we derived the protein-ligand complex structures. In Section 3.2, we will show how we reprogram the pretrained target-specific diffusion models for dual-target drug design and introduce two methods, COMPDIFF and DUALDIFF. In Section 3.3, we will show how we repurpose linker design methods for dual-target drug design.




Figure 1: Overview of our method for dual-target drug design. (a) Illustration of COMPDIFF and DUALDIFF. We first align two pockets in 3D space with protein-ligand binding prior and build two complex graph with shared ligand nodes. We then compose the $SE(3)$ -equivariant message to derive the drift on output level (COMPDIFF) or at each layer of the equivariant neural network (DUALDIFF). Based on the composed drift, we can generate dual-target ligand molecules by compositional reverse sampling. (b) Illustration of repurposing linker design methods for dual-target drug design. We first identify binding-related fragments from the reference molecules for each of the dual targets and then apply linker design methods to link the fragments and derive a complete molecule that can bind to the dual targets separately.

3.1 Data Curation

Designing dual-target drugs for random pairs of targets lacks significant intent. However, by taking cues from drug synergy, where two drugs together deliver an impact greater than the sum of their separate effects [58], we carefully select target pairs to ensure the dataset holds practical significance for drug discovery.

Drug Synergy To collect drug combination pairs, we start from DrugCombDB² [35]. DrugCombDB is a comprehensive database devoted to the curation of drug combinations from various data sources including high-throughput screening (HTS) assays, manual curations from the literature, FDA Orange Book and external databases. DrugCombDB comprises a total of 448,555 combinations of drugs, encompassing 2,887 unique drugs and 124 human cancer cell lines. Particularly, DrugCombDB has more than 6,000,000 quantitative dose responses, from which we determine whether a drug combination is synergistic or not. Specifically, a drug combination with positive zero interaction potency (ZIP), Bliss, Loewe and the highest single agent (HSA) scores simultaneously in at least one cell line is supposed to be a synergistic one. Please refer to Appendix C for a comprehensive understanding of these scores.

Drug Information After collecting synergistic drug combinations, we need to collect other necessary information (e.g., SMILES and targets) according to their drug names provided by DrugCombDB. Before this procedure, we collect synonyms and cross-matching ID (e.g., CAS Number and ChEBI ID) mainly from DrugBank³ [30] and Therapeutic Target Database (TTD)⁴ [67]. This step facilitates comprehensive literature reviews, ensuring that all relevant data sources that may use alternate names for a drug is considered. We then collect SMILES or structures (if possible) also mainly from DrugBank and TTD. To identify drug targets, we also utilize DrugBank and TTD as the primary data sources, and supplement these with manual curation from the literature (e.g., [28]). For drugs for which we cannot find either SMILES or targets, we exclude them from our previously collected dataset of positive drug combinations.

Complex Structures For certain drug-target pairs, we incorporate their complex structures directly into our dataset if they are available in PDBBind [38], a repository of protein-ligand binding structures sourced from the Protein Data Bank (PDB) [2]. For drug-target pairs not present in PDBBind, we initially attempt to retrieve the target structures from PDB; if the structures are unavailable, we then source them from the AlphaFold Protein Structure Database (AlphaFold DB) [59] and exclude those whose confidence scores, referred to as pLDDT which provide an assessment of the structures predicted by AlphaFold 2 [29], are less than 70. For these protein targets with structures from PDB or AlphaFold DB, we first utilize P2Rank [31], a program that precisely predicts ligand-binding pockets from a protein structure, to find the most possible pocket given the target structure, and use AutoDock Vina [12] to obtain the protein-ligand complex structures. For each drug, there may exist more than one targets, in which case we use AutoDock Vina to measure the binding affinity and selected the target with the best binding affinity. Finally, we obtain 12,917 positive drug combinations with protein-ligand complex structures, among which there are 438 unique drugs. The 12,917 pairs of targets can be used to evaluate the ability of methods for dual-target drug design. And the binding ligands can be used for reference molecules.

3.2 Reprogramming Target-Specific Diffusion Models for Dual-Target Drug Design

Diffusion models [21, 8, 56, 54] have been introduced to structure-based drug design and achieved promising results [19, 53, 33, 20]. We will first revisit the background of diffusion models for SBDD [19] and then introduce how we apply diffusion models trained on single-target protein-ligand datasets to dual-target drug design in a zero-shot manner. Our method is illustrated in Figure 1 (a).

Diffusion Models for single-target SBDD In this following, we denote the type of an atom as $v \in \mathbb{R}^K$ and the coordinate of an atom as $x \in \mathbb{R}^3$, where K is the number of atom types of our interest. For single-target drug design, given a protein binding site denoted as a set of atoms

²<http://drugcombd.benglab.org/main>

³<https://go.drugbank.com/>

⁴<https://idrblab.net/ttd/>

$\mathcal{P} = \{(\mathbf{x}_P^{(i)}, \mathbf{v}_P^{(i)})\}_{i=1}^{N_P}$, where N_P is the number of protein atoms, our goal is to generate binding molecules $\mathcal{M} = \{(\mathbf{x}_L^{(i)}, \mathbf{v}_L^{(i)})\}_{i=1}^{N_M}$. For brevity, we denote the molecule as $M = [\mathbf{x}, \mathbf{v}]$, where $[\cdot, \cdot]$ denotes the concatenation operator and $\mathbf{x} \in \mathbb{R}^{N_M \times 3}$, $\mathbf{x} \in \mathbb{R}^{N_M \times K}$ denote the coordinates in 3D space and one-hot atom types, respectively. So we can use generative models to model the conditional distribution $p(M|\mathcal{P})$.

In the forward *diffusion* process of the diffusion model, noises are gradually injected into the data sample (i.e., small molecule $M_0 \sim p(M|\mathcal{P})$) and lead to a sequence of latent variable M_1, M_2, \dots, M_T . The final distribution $p(M_T|\mathcal{P})$, also known as prior distribution, is approximately standard normal distribution for atom positions and uniform distribution for atom types. The reverse *generative* process learns to recover data distribution from the noise distribution with a neural network parameterized by θ . The forward and reverse processes are both Markov chains defined as follows:

$$q(M_{1:T}|M_0, \mathcal{P}) = \prod_{t=1}^T q(M_t|M_{t-1}, \mathcal{P}) \quad \text{and} \quad p_{\theta}(M_{0:T-1}|M_T, \mathcal{P}) = \prod_{t=1}^T p_{\theta}(M_{t-1}|M_t, \mathcal{P}). \quad (1)$$

More specifically, the forward transition kernel in Guan et al. [19] are defined as follows:

$$q(M_t|M_{t-1}, \mathcal{P}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}) \cdot \mathcal{C}(\mathbf{v}_t|(1 - \beta_t) \mathbf{v}_{t-1} + \beta_t \mathbf{v}_t/K), \quad (2)$$

where $\{\beta_t\}_{t=1}^T$ are fixed noise schedule. The above diffusion process can be efficiently sampled directly from time step 0 to t as follows:

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t} \mathbf{x}_0, (1 - \bar{\alpha}_t) \mathbf{I}) \quad \text{and} \quad q(\mathbf{v}_t|\mathbf{v}_0) = \mathcal{C}(\mathbf{v}_t|\bar{\alpha}_t \mathbf{v}_0 + (1 - \bar{\alpha}_t)/K), \quad (3)$$

where $\alpha_t := 1 - \beta_t$ and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$. The posterior can be easily computed via Bayes theorem as:

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0), \tilde{\beta}_t \mathbf{I}) \quad \text{and} \quad q(\mathbf{v}_t|\mathbf{v}_0) = \mathcal{C}(\mathbf{v}_t|\bar{\alpha}_t \mathbf{v}_0 + (1 - \bar{\alpha}_t)/K), \quad (4)$$

where $\tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t$, $\tilde{\mu}_t(\mathbf{x}_t, \mathbf{x}_0) = \frac{\sqrt{\bar{\alpha}_{t-1}} \beta_t}{1 - \bar{\alpha}_t} \mathbf{x}_t + \frac{\sqrt{1 - \bar{\alpha}_t}}{1 - \bar{\alpha}_t} \mathbf{x}_0$, $\tilde{\mathbf{c}}_t(\mathbf{v}_t, \mathbf{v}_0) = \mathbf{c}^* / \sum_{k=1}^K c_k^*$ and $\mathbf{c}^*(\mathbf{v}_t, \mathbf{v}_0) = [\alpha_t \mathbf{v}_t + (1 - \alpha_t)/K] \odot [\bar{\alpha}_{t-1} \mathbf{v}_0 + (1 - \bar{\alpha}_{t-1})/K]$.

Accordingly, the reverse transition kernel are defined as follows:

$$p_{\theta}(M_{t-1}|M_t, \mathcal{P}) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P}), \sigma_t^2 \mathbf{I}) \cdot \mathcal{C}(\mathbf{v}_{t-1}|\mathbf{c}_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P})). \quad (5)$$

Guan et al. [19] use SE(3)-equivariant neural networks [52, 18] to parameterize $\mu_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P})$ and $\mathbf{c}_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P})$. More specifically, the $[\mathbf{x}_0, \mathbf{v}_0]$ are first predicted using neural network f_{θ} , i.e., $[\hat{\mathbf{x}}_0, \hat{\mathbf{v}}_0] = f_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P})$ and then substitute in the posterior as in Equation (4). At the l -th layer of f_{θ} , the hidden embedding \mathbf{h} and coordinates \mathbf{x} of each atom are updated alternately as follows:

$$\mathbf{h}_i^{l+1} = \mathbf{h}_i^l + \sum_{j \in \mathcal{V}, i \neq j} f_{\theta_h}(\mathbf{h}_i^l, \mathbf{h}_j^l, d_{ij}^l, \mathbf{e}_{ij}), \quad (6)$$

$$\mathbf{x}_i^{l+1} = \mathbf{x}_i^l + \sum_{j \in \mathcal{V}, i \neq j} (\mathbf{x}_i^l - \mathbf{x}_j^l) f_{\theta_x}(\mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, d_{ij}^l, \mathbf{e}_{ij}) \cdot \mathbf{1}_{\text{ligand}}, \quad (7)$$

where \mathcal{V} is a k-nearest neighbors (knn) graph, $d_{ij} = \|\mathbf{x}_i - \mathbf{x}_j\|$ is the Euclidean distance between two atoms i and j , \mathbf{e}_{ij} is an additional feature that indicates the connection is between protein atoms, ligand atoms or protein atom and ligand atom, and $\mathbf{1}_{\text{ligand}}$ is a mask for ligand nodes since only coordinates of ligand atoms are supposed to be updated.

The diffusion model is trained to minimize the KL-divergence between the ground-truth posterior $q(M_{t-1}|M_0, M_t, \mathcal{P})$ and the estimated posterior $p_{\theta}(M_{t-1}|M_t, \mathcal{P})$. After being trained, given a specific pocket, the ligand molecule can be generated by first sampling from prior distribution and sequentially applying the reverse generative process defined above.

Problem Definition of Dual-Target Drug Design The goal of dual-target drug design is to design a ligand molecule M that can bind to both given pocket \mathcal{P}_1 and pocket \mathcal{P}_2 . The problem can be also formulated as a generative task which models the conditional distribution $p(M|\mathcal{P}_1, \mathcal{T}\mathcal{P}_2)$. Notebly, we introduce a transformation operator \mathcal{T} here. This is because protein pockets exhibit a wide variety of shapes and chemical characteristics and it is necessary to achieve spatial alignment of the dual pockets when modeling the conditional distribution with both of them as conditions. To maintain a neat but significant setting, we restrict the transformation \mathcal{T} to encompass solely translations \mathcal{T}_T and rotations \mathcal{T}_R , i.e., $\mathcal{T} = \mathcal{T}_T \circ \mathcal{T}_R$. To be more precise, $(\mathcal{T}_T \circ \mathcal{T}_R)\mathcal{P}_2 = \{(\mathbf{R}\mathbf{x}_{P_2}^{(i)} + \mathbf{t}), \mathbf{v}_{P_2}^{(i)}\}_{i=1}^{N_{P_2}}$, where $\mathbf{R} \in \text{SO}(3)$ represents the rotation and $\mathbf{t} \in \mathbb{R}^3$ represents the translation.

Aligning Dual Targets with Protein-Ligand Binding Priors Protein pockets can have intricate 3D structures with varying depths, widths, and surface contours, and distinct chemical properties, including differences in surface electron potentials, hydrophobicity, and the distribution of functional groups. The complex nature of pocket characteristics hinder us from directly aligning two pockets. Nevertheless, the binding mode is determined by the complex nature of pocket information, thus the protein-binding priors can effectively summarize the essential information needed for aligning the two pockets. This motivates us to propose to use a ligand molecule as a probe to implicitly reflect the spatial arrangement of the two pockets and then align them with the protein-ligand binding priors. More specifically, we first dock a ligand molecule to pocket \mathcal{P}_1 and pocket \mathcal{P}_2 separately. We can then compute \mathbf{R} and \mathbf{t} by aligning the two docked poses of the ligand molecule. Experiments have demonstrated that even ligand molecules capable of approximate binding to the two pockets can effectively indicate a specific spatial alignment between them. Further details are provided in Section 4.

SE(3)-Equivariant Composed Message and Compositional Generative Sampling Inspired by compositional visual generation [9, 37], we can model the dual-target drug design with the following composed distribution:

$$p(M|\mathcal{P}_1, \mathcal{T}\mathcal{P}_2) \propto p_{\theta}(M|\mathcal{P}_1)p_{\theta}(M|\mathcal{T}\mathcal{P}_2). \quad (8)$$

Following Liu et al. [37], for the atom position prediction, we can reparameterize $\mu_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P})$ with $\mathbf{x}_t - \epsilon_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P})$, so that we can rewrite the transition kernel of the reverse generative process as follows:

$$p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathcal{P}) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P}), \sigma_t^2 \mathbf{I}) = \mathcal{N}(\mathbf{x}_{t-1}; \mathbf{x}_t - \epsilon_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P}), \sigma_t^2 \mathbf{I}). \quad (9)$$

The reversed transition kernel corresponds to a step as follows:

$$\mathbf{x}_{t-1} = \mathbf{x}_t - \epsilon_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P}) + \mathcal{N}(\mathbf{0}, \sigma_t^2 \mathbf{I}). \quad (10)$$

where $\epsilon_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P})$ can be viewed as a drift term and $\mathcal{N}(\mathbf{0}, \sigma_t^2 \mathbf{I})$ can be viewed as a diffusion term. As Liu et al. [37] points out, this is analogous to the Langevin dynamics [11] that used to sample from Energy-Based Models (EBMs) [10, 55], which can be formulated as follows:

$$\mathbf{x}_{t-1} = \mathbf{x}_t - \frac{\lambda}{2} \nabla_{\mathbf{x}} E_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P}) + \mathcal{N}(\mathbf{0}, \sigma_t^2 \mathbf{I}). \quad (11)$$

The sampling procedure produces samples from the probability density $p_{\theta}(\mathbf{x}|\mathcal{P}) \propto \exp(-E_{\theta}(\mathbf{x}|\mathcal{P}))$ where $E_{\theta}(\mathbf{x}|\mathcal{P})$ is an energy function parameterized by model θ . Thus, accordingly, the composed distribution of atom positions can be written as:

$$p(\mathbf{x}|\mathcal{P}_1, \mathcal{T}\mathcal{P}_2) \propto p_{\theta}(\mathbf{x}|\mathcal{P}_1)p_{\theta}(\mathbf{x}|\mathcal{T}\mathcal{P}_2) \propto \exp\left(-\left(E_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P}_1) + E_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{T}\mathcal{P}_2)\right)\right).$$

This composed distribution corresponds to Langevin dynamics as follows:

$$\mathbf{x}_{t-1} = \mathbf{x}_t - \frac{\lambda}{2} \nabla_{\mathbf{x}} \left(-\left(E_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P}_1) + E_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{T}\mathcal{P}_2)\right) \right). \quad (12)$$

Accordingly, each step in the compositional reverse generative sampling process can be defined as:

$$\mathbf{x}_{t-1} = \mathbf{x}_t - \eta \left(\left(\epsilon_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{P}_1) + \epsilon_{\theta}([\mathbf{x}_t, \mathbf{v}_t], t, \mathcal{T}\mathcal{P}_2) \right) \right), \quad (13)$$

where we additionally introduce a hyperparameter η here to control the strength of the drift. In theory, this is equivalent to making a more flexible assumption that $p(\mathbf{x}|\mathcal{P}_1, \mathcal{T}\mathcal{P}_2) \propto [p_{\theta}(\mathbf{x}|\mathcal{P}_1)p_{\theta}(\mathbf{x}|\mathcal{T}\mathcal{P}_2)]^{\eta}$. In this case the transition kernel is defined as $p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathcal{P}_1, \mathcal{T}\mathcal{P}_2) = [p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathcal{P}_1)p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathcal{T}\mathcal{P}_2)]^{\eta}$. In practice, we set $\eta = 1/2$ by default. This composition operation is equivalent to averaging two $\hat{\mathbf{x}}_0$ predicted on two complex graphs, i.e. \mathcal{V}_1 and \mathcal{V}_2 . Similarly, for the atom types which are discrete variables, we can also compose the transition kernel as follows: $p_{\theta}(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathcal{P}_1, \mathcal{T}\mathcal{P}_2) \propto p_{\theta}(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathcal{P}_1)p_{\theta}(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathcal{T}\mathcal{P}_2)$. Note that $p_{\theta}(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathcal{P}_1)$ is categorical distribution and its dimension (i.e., the number of atom types of our interest) is K , which is small in practice. So $p_{\theta}(\mathbf{v}_{t-1}|\mathbf{v}_t, \mathcal{P}_1, \mathcal{T}\mathcal{P}_2)$ can be computed analytically. We name the compositional reverse sampling with composed transition kernel (i.e., composed drift) as COMPDIFF.

We further improve the compositional reverse sampling by introducing the composition into each layer of the equivariant neural network in the pretrained diffusion model. For brevity, we denote the SE(3)-equivariant message at the l -th layer for the i -th atom of the complex graph \mathcal{V}_v ($v = 1, 2$) as introduced in Equation (7) as follows:

$$\Delta \mathbf{h}_i^l(\mathcal{V}_v) := \sum_{j \in \mathcal{V}_v, i \neq j} f_{\theta_h}(\mathbf{h}_i^l, \mathbf{h}_j^l, d_{ij}^l, \mathbf{e}_{ij}), \quad (14)$$

$$\Delta \mathbf{x}_i^l(\mathcal{V}_v) := \sum_{j \in \mathcal{V}_v, i \neq j} (\mathbf{x}_i^l - \mathbf{x}_j^l) f_{\theta_x}(\mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, d_{ij}^l, \mathbf{e}_{ij}) \cdot \mathbf{1}_{\text{ligand}}. \quad (15)$$

The above SE(3)-equivariant message can also be interpreted as drift in 3D and latent space. Thus we can also compose them as follows:

$$\mathbf{h}_i^{l+1} = \mathbf{h}_i^l + (\Delta \mathbf{h}_i^l(\mathcal{V}_1) + \Delta \mathbf{h}_i^l(\mathcal{V}_2))/2 \quad \text{and} \quad \mathbf{x}_i^{l+1} = \mathbf{x}_i^l + (\Delta \mathbf{x}_i^l(\mathcal{V}_1) + \Delta \mathbf{x}_i^l(\mathcal{V}_2))/2. \quad (16)$$

We name the compositional reverse sampling with the above SE(3)-equivariant message at each layer as DUALDIFF. The proof of SE(3)-equivariance can be found in Appendix G. This more meticulous composition is supposed to lead to higher-quality samples.

3.3 Repurposing Linker Design Methods for Dual-Target Drug Design

Pharmacophore combination is a prevalent strategy in traditional dual-target drug design, requiring the specialized knowledge of chemists. To automate this procedure, We design a strategy to identify crucial fragments from reference molecules of dual targets in our dataset and apply linker design methods [25, 17] to link the fragments and obtain complete molecules. Please refer to Appendix D for a more detailed justification for the choice of baselines.

Specifically, we break all rotatable bonds of reference molecule \mathcal{M}_1 (resp. \mathcal{M}_2) of target \mathcal{P}_1 (resp. \mathcal{P}_2) to obtain fragments. Since DiffLinker [25] requires relative positions of fragments as input, we dock all fragments derived from \mathcal{M}_1 and \mathcal{M}_2 to \mathcal{P}_2 (or \mathcal{P}_1) and select the pair of fragments which has the best sum of binding affinity and no physical conflicts (i.e., the minimum between atoms from the two fragments is large than 1.4Å). We then apply DiffLinker to link the two fragments, considering the existence of the pocket \mathcal{P}_2 (or \mathcal{P}_1). Since LinkerNet [17] models the translation and rotation of fragments by neural networks and does not require relative position of fragments as input, we directly dock all fragments derived from \mathcal{M}_1 (resp. \mathcal{M}_2) to \mathcal{P}_1 (resp. \mathcal{P}_2) and select their respective fragment with the best binding affinity. And we then link them using LinkerNet to obtain complete molecules.

4 Experiment

4.1 Experimental Setup

Dataset We use our dataset introduced in Section 3.1. All 12,917 pairs of targets (including 438 unique targets) are used for evaluation. For each target, there is an associated reference molecule that can be considered a benchmark for high-quality ligand molecules and utilized in linker design methods.

Baselines We compare our method with various baselines: **Pocket2Mol** [45] generates 3D molecules atom by atom in an autoregressive manner given a specific protein binding site. **TargetDiff** [19] is a diffusion-based method which generates atom coordinates and atom types in a non-autoregressive way. Note that Pocket2Mol and TargetDiff are both proposed for structure-based single-target drug design. **DiffLinker** [25] is a diffusion-based model for linker design with given fragment poses. **LinkerNet** [17] is a diffusion-based model for co-designing molecular fragment poses and the linker. DiffLinker and LinkerNet are repurposed for dual-target drug design as we introduced in Section 3.3. The code is available at <https://github.com/zhouxiangxin1998/DualDiff>.

Evaluation We evaluate generated ligand molecules from the perspectives of target binding affinity and molecular properties. We employ AutoDock Vina [12] to estimate the target binding affinity, following Peng et al. [45], Guan et al. [19]. We first evaluate Pocket2Mol and TargetDiff under the

Table 1: Summary of different properties of reference molecules and molecules generated by baselines and our methods under the **dual-target** setting. (\uparrow) / (\downarrow) denotes a larger / smaller number is better. Top 2 results are highlighted with **bold text** and underlined text, respectively.

Methods	P-1 Vina Dock (\downarrow)		P-2 Vina Dock (\downarrow)		Max Vina Dock (\downarrow)		Dual High Aff. (\uparrow)		QED (\uparrow)		SA (\uparrow)		Diversity (\uparrow)	
	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.
Reference	-7.60	-7.80	-6.02	-7.30	-5.46	-7.09	-	-	0.53	0.54	0.74	0.77	-	-
Pocket2Mol	-4.82	-4.76	-4.63	-4.64	-4.40	-4.42	0.2%	0.0%	0.49	0.48	0.88	0.90	0.82	0.83
TargetDiff	-8.62	-8.61	-6.89	-7.67	<u>-6.57</u>	<u>-7.39</u>	29.0%	20.0%	0.50	0.51	0.58	0.58	0.70	0.71
DiffLinker	-7.05	-7.87	-7.27	-7.92	-5.87	-7.18	24.6%	0.0%	0.43	0.42	0.30	0.29	0.52	0.54
LinkerNet	-8.20	-8.37	-8.13	-8.38	-7.17	-7.72	35.7%	0.0%	0.61	0.63	<u>0.72</u>	<u>0.73</u>	0.37	0.34
COMPDIFF	-8.32	-8.42	<u>-8.37</u>	<u>-8.47</u>	<u>-7.50</u>	<u>-7.78</u>	<u>35.9%</u>	<u>30.0%</u>	0.55	0.57	0.59	0.59	<u>0.72</u>	<u>0.72</u>
DUALDIFF	<u>-8.41</u>	<u>-8.51</u>	-8.48	-8.55	-7.66	-7.88	36.3%	30.2%	0.56	<u>0.58</u>	0.59	0.59	0.67	0.67

single-target setting as a preliminary verification (see Appendix E). For dual-target drug design, we use each method to design 10 molecules for each pair of targets, denoted as \mathcal{P}_1 and \mathcal{P}_2 . (For reference molecule, Pocet2Mol and TargetDiff, the ligand molecule is generated for \mathcal{P}_1 but the target binding affinity is evaluated on both \mathcal{P}_1 and \mathcal{P}_2 .) We then collect all generated molecules across 12,917 pairs of targets and report the mean and median (denoted as “Avg.” and “Med.” respectively) of affinity-related metrics (P-1 Vina Dock, P-2 Vina Dock, Max Vina Dock, and Dual High Affinity) and property-related metrics (drug-likeness QED [3], synthesizability SA [14], and diversity). Vina Dock incorporates a re-docking step to assess the highest binding affinity achievable. Here we introduce P-1 Vina Dock and P-2 Vina Dock to represent the Vina Dock score evaluated on \mathcal{P}_1 and \mathcal{P}_2 , respectively. Besides, we introduce Max Vina Dock, which represents the maximum Vina Dock of a given molecule towards \mathcal{P}_1 and \mathcal{P}_2 . The Vina Dock will be low if and only if the molecule can bind to both targets simultaneously, which is the goal of dual-target drug design. Additionally, we report Dual High Affinity (abbreviated as Dual High Aff.) which represents the proportion of generated molecules that exhibit binding affinity that exceeds that of the reference molecules on both the respective targets. This reflects the success rate in achieving higher binding affinities simultaneously on both targets in dual-target drug design. We also evaluate the RMSD between docked poses towards dual targets.

4.2 Main Results

We compare all methods under the dual-target setting. The results are reported in Table 1. Our methods, COMPDIFF and DUALDIFF, significantly outperforms other methods, especially in terms of binding affinity. Notably, DUALDIFF achieves the highest Dual High Affinity among all methods. In line with our expectations, for the single-target drug design methods (e.g., Pocket2Mol and TargetDiff), we observe a significant decline in performance according to P-2 Vina Dock compared to P-1 Vina Dock, which shows their inability in dual-target drug design. LinkerNet also achieves promising results except diversity. Note that DiffLinker and LinkerNet are provided with reference molecules while COMPDIFF and DUALDIFF are not. This indicates the strong generative abilities of our methods. Finally, DUALDIFF outperforms COMPDIFF, which shows that composition of SE(3)-equivariant message at each layer is more effective than only at the output level.

As shown in Figure 2, DUALDIFF performs better than LinkerNet the RMSD between docked poses towards dual targets. This indicates that the molecules generated by DUALDIFF can bind to dual targets with smaller conformation change. We also benchmark the inference time of baselines and our methods in Appendix H. Additionally, we provide visualization of examples of generated molecules in Figure 3. See more examples in Appendix F.




Figure 2: RMSD between docked poses towards dual targets of different methods.




Figure 3: Reference molecules and examples of ligand molecules by different methods generated for the dual targets (UniProt ID: P18507 (top) and Q9UBS5 (bottom)).

Table 2: Ablation on different ways of aligning dual targets.

Methods	P-1 Vina Dock (↓)		P-2 Vina Dock (↓)		Max Vina Dock (↓)		Dual High Aff. (↑)		QED (↑)		SA (↑)		Diversity (↑)	
	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.
COMPDIFF-Center	-8.10	-8.26	-8.08	-8.26	-7.23	-7.60	30.8%	22.0%	0.52	0.53	0.60	0.59	0.73	0.73
COMPDIFF-RMSD	-8.29	-8.44	-8.35	-8.46	-7.47	-7.78	35.6%	30.0%	0.55	0.56	0.59	0.59	0.72	0.72
COMPDIFF-Score	-8.32	-8.42	-8.37	-8.47	-7.50	-7.78	35.9%	30.0%	0.55	0.57	0.59	0.59	0.72	0.72
DUALDIFF-Center	-8.12	-8.29	-8.12	-8.28	-7.32	-7.66	30.0%	22.2%	0.52	0.54	0.59	0.59	0.69	0.69
DUALDIFF-RMSD	-8.40	-8.51	-8.45	-8.53	-7.63	-7.87	35.8%	30.0%	0.55	0.57	0.59	0.59	0.67	0.67
DUALDIFF-Score	-8.41	-8.51	-8.48	-8.55	-7.66	-7.88	36.3%	30.2%	0.56	0.58	0.59	0.59	0.67	0.67

Table 3: Ablation on different strategies of identifying fragments for linker design methods.

Methods	P-1 Vina Dock (↓)		P-2 Vina Dock (↓)		Max Vina Dock (↓)		Dual High Aff. (↑)		QED (↑)		SA (↑)		Diversity (↑)	
	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.
DiffLinker-5	-6.74	-6.74	-6.85	-6.81	-6.21	-6.32	9.4%	0.0%	0.58	0.60	0.32	0.31	0.57	0.62
DiffLinker-pocket-5	-7.35	-7.74	-7.20	-7.75	-6.25	-7.14	20.2%	0.0%	0.45	0.45	0.31	0.30	0.58	0.62
DiffLinker-8	-7.22	-7.32	-7.49	-7.49	-6.65	-6.86	15.7%	0.0%	0.61	0.63	0.32	0.32	0.49	0.51
DiffLinker-pocket-8	-7.05	-7.87	-7.27	-7.92	-5.87	-7.18	24.6%	0.0%	0.43	0.42	0.30	0.29	0.52	0.54
LinkerNet-5	-7.54	-7.54	-7.56	-7.58	-6.98	-7.05	19.9%	0.0%	0.69	0.72	0.77	0.79	0.46	0.46
LinkerNet-self-5	-7.54	-7.55	-7.55	-7.56	-6.98	-7.05	20.0%	0.0%	0.69	0.72	0.77	0.79	0.46	0.46
LinkerNet-8	-8.09	-8.23	-8.23	-8.31	-7.29	-7.63	34.0%	0.0%	0.62	0.65	0.72	0.74	0.38	0.35
LinkerNet-self-8	-8.20	-8.37	-8.13	-8.38	-7.17	-7.72	35.7%	0.0%	0.61	0.63	0.72	0.73	0.37	0.34

4.3 Ablation Studies

Effects of Alignment of Dual Targets We perform different methods to align the pockets of dual targets for COMPDIFF and DUALDIFF. See the results in Table 2. Naively, we can align two pockets by their geometric centers (denoted as “-Center”). In our method, we propose to align pockets with protein-ligand binding priors. We select the ligand with minimum RMSD between docked poses (resp. minimum sum of Vina Dock scores) towards dual targets as the anchor to align the dual targets, which is denoted as “-RMSD” (resp. “-Score”). DUALDIFF-Score achieved the best performance among all variants, demonstrating the effectiveness of the alignment method.

Different Strategies of Identifying Fragments for Linker Design Methods We conduct ablation on different strategies of identifying fragments for DiffLinker and LinkerNet for dual-target drug design. Since we use Vina Dock to select fragments, we have tried different box sizes for docking, i.e., 5Å and 8Å. For DiffLinker, since the relative poses of fragments are required as input, we dock fragments derived from \mathcal{M}_1 and \mathcal{M}_2 to target \mathcal{P}_1 (or \mathcal{P}_2). We then apply DiffLinker both with and without considering the pocket. The corresponding methods are denoted as “DiffLinker-5/8-pocket-5/8-pocket-5/8-pocket-8”. For LinkerNet, the relative poses of fragments are not required. So we try two setting: dock fragments from \mathcal{M}_1 and \mathcal{M}_2 to target \mathcal{P}_1 (or \mathcal{P}_2); dock fragments from \mathcal{M}_1 (resp. \mathcal{M}_2) to target \mathcal{P}_1 (resp. \mathcal{P}_2). The difference is that all fragments are docked to the same pocket in the former setting while the fragments from two ligand molecules are docked to their respective pockets. The

corresponding methods are denoted as “LinkerNet-5/-8/-self-5/-self-8”. The results are reported in Table 3. The results show that a larger docking box size allows for better selection of fragments. As we expected, DiffLinker with consideration of pockets achieves better performance. Among all variants, “LinkerNet-self-8” achieves best performance, which implies that adjusting relative poses of fragments may play a crucial role in linker design. This feature of LinkerNet allows for selecting the most important fragments for each pockets of the dual targets.

5 Conclusion

In this work, we introduced a novel dataset for dual-target drug design. We formulate this problem as a generative task and propose compositional reverse sampling to reprogram pretrained target-specific model for dual-target drug design, which can successfully generate dual-target ligands and outperform all baselines, including repurposed linker design methods. Our research lays the groundwork for dual-target drug design using generative methods, with future progress expected.

Acknowledgments

This work is supported by the National Science and Technology Major Project (2023ZD0120901) and the National Natural Science Foundation of China No. 62377030. We also thank the reviewers for their valuable feedback.

References

- [1] Amy C Anderson. 2003. The process of structure-based drug design. *Chemistry & biology*, 10(9):787–797.
- [2] Helen Berman, Kim Henrick, and Haruki Nakamura. 2003. Announcing the worldwide protein data bank. *Nature structural & molecular biology*, 10(12):980–980.
- [3] G Richard Bickerton, Gaia V Paolini, Jérémie Besnard, Sorel Muresan, and Andrew L Hopkins. 2012. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2):90–98.
- [4] Maria Laura Bolognesi and Andrea Cavalli. 2016. Multitarget drug discovery and polypharmacology.
- [5] Aislyn DW Boran and Ravi Iyengar. 2010. Systems approaches to polypharmacology and drug discovery. *Current opinion in drug discovery & development*, 13(3):297.
- [6] Yosef Buganim, Dina A Faddah, Albert W Cheng, Elena Itsikovich, Styliani Markoulaki, Kibibi Ganz, Sandy L Klemm, Alexander Van Oudenaarden, and Rudolf Jaenisch. 2012. Single-cell expression analyses during cellular reprogramming reveal an early stochastic and a late hierachic phase. *Cell*, 150(6):1209–1222.
- [7] Di Chen, Xi Liu, Yiping Yang, Hongjun Yang, and Peng Lu. 2015. Systematic synergy modeling: understanding drug synergy from a systems biology perspective. *BMC Systems Biology*.
- [8] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794.
- [9] Yilun Du, Shuang Li, and Igor Mordatch. 2020. Compositional visual generation with energy based models. *Advances in Neural Information Processing Systems*, 33:6637–6647.
- [10] Yilun Du, Shuang Li, Joshua Tenenbaum, and Igor Mordatch. 2020. Improved contrastive divergence training of energy based models. *arXiv preprint arXiv:2012.01316*.
- [11] Yilun Du and Igor Mordatch. 2019. Implicit generation and modeling with energy based models. *Advances in Neural Information Processing Systems*, 32.
- [12] Jerome Eberhardt, Diogo Santos-Martins, Andreas F Tillack, and Stefano Forli. 2021. AutoDock Vina 1.2. 0: New docking methods, expanded force field, and python bindings. *Journal of chemical information and modeling*, 61(8):3891–3898.
- [13] Gamaleldin F. Elsayed, Ian Goodfellow, and Jascha Sohl-Dickstein. 2019. Adversarial Reprogramming of Neural Networks. In *International Conference on Learning Representations*.
- [14] Peter Ertl and Ansgar Schuffenhauer. 2009. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of cheminformatics*, 1(1):1–11.
- [15] Michael A Fischbach. 2011. Combination therapies for combating antimicrobial resistance. *Current opinion in microbiology*, 14(5):519–523.
- [16] Silvia Giordano and Alessio Petrelli. 2008. From single-to multi-target drugs in cancer therapy: when aspecificity becomes an advantage. *Current medicinal chemistry*, 15(5):422–432.
- [17] Jiaqi Guan, Xingang Peng, PeiQi Jiang, Yunan Luo, Jian Peng, and Jianzhu Ma. 2023. LinkerNet: Fragment Poses and Linker Co-Design with 3D Equivariant Diffusion. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- [18] Jiaqi Guan, Wesley Wei Qian, Wei-Ying Ma, Jianzhu Ma, and Jian Peng. 2021. Energy-inspired molecular conformation optimization. In *international conference on learning representations*.
- [19] Jiaqi Guan, Wesley Wei Qian, Xingang Peng, Yufeng Su, Jian Peng, and Jianzhu Ma. 2023. 3D Equivariant Diffusion for Target-Aware Molecule Generation and Affinity Prediction. In *International Conference on Learning Representations*.

[20] Jiaqi Guan, Xiangxin Zhou, Yuwei Yang, Yu Bao, Jian Peng, Jianzhu Ma, Qiang Liu, Liang Wang, and Quanquan Gu. 2023. DecomDiff: Diffusion Models with Decomposed Priors for Structure-Based Drug Design. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 11827–11846. PMLR.

[21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851.

[22] Yinan Huang, Xingang Peng, Jianzhu Ma, and Muhan Zhang. 2022. 3DLinker: an E (3) equivariant variational autoencoder for molecular linker design. *arXiv preprint arXiv:2205.07309*.

[23] Zhilin Huang, Ling Yang, Xiangxin Zhou, Zhilong Zhang, Wentao Zhang, Xiawu Zheng, Jie Chen, Yu Wang, CUI Bin, and Wenming Yang. 2023. Protein-ligand interaction prior for binding-aware 3d molecule diffusion models. In *The Twelfth International Conference on Learning Representations*.

[24] Bliss Chester I. 1939. The toxicity of poisons applied jointly 1. *Annals of applied biology*, 26(3):585–615.

[25] Ilia Igashov, Hannes Stärk, Clément Vignac, Victor Garcia Satorras, Pascal Frossard, Max Welling, Michael Bronstein, and Bruno Correia. 2022. Equivariant 3d-conditional diffusion models for molecular linker design. *arXiv preprint arXiv:2210.05274*.

[26] Fergus Imrie, Anthony R Bradley, Mihaela van der Schaar, and Charlotte M Deane. 2020. Deep generative models for 3D linker design. *Journal of chemical information and modeling*, 60(4):1983–1995.

[27] Fergus Imrie, Thomas E Hadfield, Anthony R Bradley, and Charlotte M Deane. 2021. Deep generative design with 3D pharmacophoric constraints. *Chemical science*, 12(43):14577–14589.

[28] Francesco Iorio, Theo A Knijnenburg, Daniel J Vis, Graham R Bignell, Michael P Menden, Michael Schubert, Nanne Aben, Emanuel Gonçalves, Syd Barthorpe, Howard Lightfoot, et al. 2016. A landscape of pharmacogenomic interactions in cancer. *Cell*, 166(3):740–754.

[29] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873):583–589.

[30] Craig Knox, Mike Wilson, Christen M Klinger, Mark Franklin, Eponine Oler, Alex Wilson, Allison Pon, Jordan Cox, Na Eun Chin, Seth A Strawbridge, et al. 2024. Drugbank 6.0: the drugbank knowledgebase for 2024. *Nucleic Acids Research*, 52(D1):D1265–D1275.

[31] Radoslav Krivák and David Hoksza. 2018. P2Rank: machine learning based tool for rapid and accurate prediction of ligand binding sites from protein structure. *Journal of cheminformatics*, 10:1–12.

[32] Ying Hong Li, Pan Pan Wang, Xiao Xu Li, Chun Yan Yu, Hong Yang, Jin Zhou, Wei Wei Xue, Jun Tan, and Feng Zhu. 2016. The human kinome targeted by FDA approved multi-target drugs and combination products: A comparative study from the drug-target interaction network perspective. *PloS one*, 11(11):e0165737.

[33] Haitao Lin, Yufei Huang, Meng Liu, Xuanjing Li, Shuiwang Ji, and Stan Z Li. 2022. Diffbp: Generative diffusion of 3d molecules for target protein binding. *arXiv preprint arXiv:2211.11214*.

[34] Hui-Heng Lin, Le-Le Zhang, Ru Yan, Jin-Jian Lu, and Yuanjia Hu. 2017. Network analysis of drug–target interactions: a study on FDA-approved new molecular entities between 2000 to 2015. *Scientific reports*, 7(1):12230.

[35] Hui Liu, Wenhao Zhang, Bo Zou, Jinxian Wang, Yuanyuan Deng, and Lei Deng. 2020. Drug-CombDB: a comprehensive database of drug combinations toward the discovery of combinatorial therapy. *Nucleic acids research*, 48(D1):D871–D881.

[36] Meng Liu, Youzhi Luo, Kanji Uchino, Koji Maruhashi, and Shuiwang Ji. 2022. Generating 3d molecules for target protein binding. *arXiv preprint arXiv:2204.09410*.

[37] Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B Tenenbaum. 2022. Compositional visual generation with composable diffusion models. In *European Conference on Computer Vision*, pages 423–439. Springer.

[38] Zhihai Liu, Minyi Su, Li Han, Jie Liu, Qifan Yang, Yan Li, and Renxiao Wang. 2017. Forging the basis for developing protein–ligand interaction scoring functions. *Accounts of chemical research*, 50(2):302–309.

[39] Georgina V Long, Daniil Stroyakovskiy, Helen Gogas, Evgeny Levchenko, Filippo de Braud, James Larkin, Claus Garbe, Thomas Jouary, Axel Hauschild, Jean Jacques Grob, et al. 2014. Combined BRAF and MEK inhibition versus BRAF inhibition alone in melanoma. *New England Journal of Medicine*, 371(20):1877–1888.

[40] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35:5775–5787.

[41] Jin-Jian Lu, Wei Pan, Yuan-Jia Hu, and Yi-Tao Wang. 2012. Multi-target drugs: the trend of drug research and development. *PloS one*, 7(6):e40262.

[42] Shitong Luo, Jiaqi Guan, Jianzhu Ma, and Jian Peng. 2021. A 3D generative model for structure-based drug design. *Advances in Neural Information Processing Systems*, 34:6229–6239.

[43] Igor Melnyk, Vijil Chenthamarakshan, Pin-Yu Chen, Payel Das, Amit Dhurandhar, Inkit Padhi, and Devleena Das. 2023. Reprogramming pretrained language models for antibody sequence infilling. In *International Conference on Machine Learning*, pages 24398–24419. PMLR.

[44] Reza Bayat Mokhtari, Tina S Homayouni, Narges Baluch, Evgeniya Morgatskaya, Sushil Kumar, Bikul Das, and Herman Yeger. 2017. Combination therapy in combating cancer. *Oncotarget*, 8(23):38022.

[45] Xingang Peng, Shitong Luo, Jiaqi Guan, Qi Xie, Jian Peng, and Jianzhu Ma. 2022. Pocket2mol: Efficient molecular sampling based on 3d protein pockets. In *International Conference on Machine Learning*, pages 17644–17655. PMLR.

[46] Alexander S Powers, Helen H Yu, Patricia Suriana, Rohan V Koodli, Tianyu Lu, Joseph M Paggi, and Ron O Dror. 2023. Geometric deep learning for structure-based ligand design. *ACS Central Science*, 9(12):2257–2267.

[47] Lauren N Puls, Matthew Eadens, and Wells Messersmith. 2011. Current status of SRC inhibitors in solid tumor malignancies. *The oncologist*, 16(5):566–578.

[48] Matthew Ragoza, Tomohide Masuda, and David Ryan Koes. 2022. Generating 3D molecules conditional on receptor binding sites with deep generative models. *Chemical science*, 13(9):2701–2713.

[49] Rona R Ramsay, Marija R Popovic-Nikolic, Katarina Nikolic, Elisa Uliassi, and Maria Laura Bolognesi. 2018. A perspective on multi-target drug discovery and design for complex diseases. *Clinical and translational medicine*, 7:1–14.

[50] Caroline Robert, Boguslawa Karaszewska, Jacob Schachter, Piotr Rutkowski, Andrzej Mackiewicz, Daniil Stroyakovskiy, Michael Lichinitser, Reinhard Dummer, Florent Grange, Laurent Mortier, et al. 2015. Improved overall survival in melanoma with combined dabrafenib and trametinib. *New England Journal of Medicine*, 372(1):30–39.

[51] Loewe S. 1953. The problem of synergism and antagonism of combined drugs. *Arzneimittelforschung*, 3(6):285–90.

[52] Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. 2021. E (n) equivariant graph neural networks. In *International conference on machine learning*, pages 9323–9332. PMLR.

[53] Arne Schneuing, Yuanqi Du, Charles Harris, Arian Jamasb, Ilia Igashov, Weitao Du, Tom Blundell, Pietro Lió, Carla Gomes, Max Welling, Michael Bronstein, and Bruno Correia. 2022. Structure-based Drug Design with Equivariant Diffusion Models. *arXiv preprint arXiv:2210.13695*.

[54] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR.

[55] Yang Song and Diederik P Kingma. 2021. How to train your energy-based models. *arXiv preprint arXiv:2101.03288*.

[56] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2020. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.

[57] Dejuan Sun, Yuqian Zhao, Shouyue Zhang, Lan Zhang, Bo Liu, and Liang Ouyang. 2020. Dual-target kinase drug design: Current strategies and future directions in cancer therapy. *European Journal of Medicinal Chemistry*, 188:112025.

[58] Ronald J Tallarida. 2001. Drug synergism: its detection and applications. *Journal of Pharmacology and Experimental Therapeutics*, 298(3):865–872.

[59] Mihaly Varadi, Stephen Anyango, Mandar Deshpande, Sreenath Nair, Cindy Natassia, Galabina Yordanova, David Yuan, Oana Stroe, Gemma Wood, Agata Laydon, et al. 2022. AlphaFold Protein Structure Database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic acids research*, 50(D1):D439–D444.

[60] Bhagwan Yadav, Krister Wennerberg, Tero Aittokallio, and Jing Tang. 2015. Searching for Drug Synergy in Complex Dose–Response Landscapes Using an Interaction Potency Model. *BMC Systems Biology*.

[61] Chao-Han Huck Yang, Yun-Yun Tsai, and Pin-Yu Chen. 2021. Voice2series: Reprogramming acoustic models for time series classification. In *International conference on machine learning*, pages 11808–11819. PMLR.

[62] Jing Ye, Junhao Wu, and Bo Liu. 2023. Therapeutic strategies of dual-target small molecules to overcome drug resistance in cancer therapy. *Biochimica et Biophysica Acta (BBA)-Reviews on Cancer*, 1878(3):188866.

[63] Xiangxiang Zeng, Fei Wang, Yuan Luo, Seung-gu Kang, Jian Tang, Felice C Lightstone, Evandro F Fang, Wendy Cornell, Ruth Nussinov, and Feixiong Cheng. 2022. Deep generative molecular design reshapes drug discovery. *Cell Reports Medicine*, 3(12).

[64] Zaixi Zhang and Qi Liu. 2023. Learning Subpocket Prototypes for Generalizable Structure-based Drug Design. *arXiv preprint arXiv:2305.13997*.

[65] Zaixi Zhang, Yaosen Min, Shuxin Zheng, and Qi Liu. 2022. Molecule generation for target protein binding with structural motifs. In *The Eleventh International Conference on Learning Representations*.

[66] Xiangxin Zhou, Xiwei Cheng, Yuwei Yang, Yu Bao, Liang Wang, and Quanquan Gu. 2024. DecompOpt: Controllable and Decomposed Diffusion Models for Structure-based Molecular Optimization. In *The Twelfth International Conference on Learning Representations*.

[67] Ying Zhou, Yintao Zhang, Donghai Zhao, Xinyuan Yu, Xinyi Shen, Yuan Zhou, Shanshan Wang, Yunqing Qiu, Yuzong Chen, and Feng Zhu. 2024. TTD: Therapeutic Target Database describing target druggability information. *Nucleic acids research*, 52(D1):D1465–D1477.

A Motivation, Significance, and Current Practices of Dual-Target Drug Design

Motivation and significance of dual-target drug design The reviewers and audience can refer to Giordano and Petrelli [16], Boran and Iyengar [5] and [41] for the concept of dual-target (or multi-target) drugs. Dual-target drugs have various advantages, especially in aspects such as anti-tumor and overcoming drug resistance.

It has been revealed that single-target drugs may not always induce the desired effect on the entire biological system even if they successfully inhibit or activate a specific target [47]. One reason is that organisms can affect effectiveness through compensatory ways. Another reason is the development of resistance either by self-modification of the target through mutation or by the adoption of new pathways by a cancer cell, for the growth and multiplication [47]. Cancer and those nervous and cardiovascular system diseases are complicated, thereby promoting the multi-targeted therapies as a better pathway to achieve the desired treatment. For example, the multiple tyrosine kinase inhibitor imatinib⁵ induces better anti-cancer effects compared with that of gefitinib⁶, which involves a single target, further indicating that drugs with multiple targets may exhibit a better chance of affecting the complex equilibrium of whole cellular networks than drugs that act on a single target. Refer to Lu et al. [41] for more examples. Notably, combination therapy also performs well in overcoming drug resistance [15]. However, under some circumstances, the combination therapy might have additive and even synergistic effects in theory; however, it often leads to some unpredictable side effects, such as increased toxicity. Thus, multi-target drugs are a better choice.

Current practices for dual-target drug design Refer to Sun et al. [57] for a comprehensive review. The mainstream strategies are drug repurposing, skeleton modification, and linked/merged/fused pharmacophores, which involve complicated pipelines and require domain knowledge from experts. To the best of our knowledge, we are the first one to formulate dual-target drug design as a generative task and provide simple yet effective solutions.

B Extended Related Works

Our work is also relevant to neural model reprogramming⁷. We will discuss the connections as follows:

Elsayed et al. [13] proposed adversarial reprogramming. Specifically, given a model trained on a task which maps x to $f(x)$, the adversary aims to repurpose the model for a new task which maps $g(\tilde{x})$ for inputs \tilde{x} , by learning mapping functions $h_f(\cdot; \theta)$ and $h_g(\cdot; \theta)$ to make $h_g(f(h_f(\tilde{x})))$ approximate $g(\tilde{x})$. They successfully demonstrated their methods on image classification tasks. For example, a trained adversarial program can cause a classifier trained on ImageNet to an MNIST classifier.

Yang et al. [61] reprograms acoustic models for time series classification, through input transformation learning and output label mapping, which are similar to $h_f(\cdot)$ and $h_g(\cdot)$, respectively. The motivation is that data scarcity hinders researchers from using large-scale deep learning models for time-series tasks while many large-scale pre-trained speech processing models are available. Melnyk et al. [43] also follows this paradigm to reprogram pretrained language models, BERT trained on English corpus, for the antibody sequence infilling task.

The above works all aim to reprogram a model trained on a task to perform a new task. Our work also falls into this paradigm, since we focus on reprogramming models trained for single-target drug design to design dual-target drugs. And we also have a similar motivation of Yang et al. [61] and Melnyk et al. [43], because there is no training data for dual-target drug design. Differently, we focus on reprogramming a diffusion model equipped with SE(3)-equivariance which has not been widely explored in previous works due to its complex sampling process during inference. And, notably, all the above works require additional training to learn the additional parameters of input transformation (i.e., $h_f(\cdot; \theta)$ in Elsayed et al. [13]) for reprogramming. However, our framework works in a zero-shot manner, which means we only need to modify the sampling process of the diffusion model to perform

⁵<https://en.wikipedia.org/wiki/Imatinib>

⁶<https://en.wikipedia.org/wiki/Gefitinib>

⁷Note that the reprogramming here is not relevant to the biological concept “cell reprogramming” [6].

the novel task without any additional training. Specifically, our proposed pocket alignment resembles the “input transformation”, i.e., $h_f(\cdot)$, and the compositional sampling resembles the “output label mapping”, i.e., $h_g(\cdot)$, but no learnable parameters are introduced for reprogramming. Yang et al. [61] provided theoretical results on selecting a pre-trained model for reprogramming, but it is not easy to directly apply the results to our case due to the intricate properties of diffusion models.

C Term Definitions in Drug Synergy

Drug Synergy Drug synergy refers to the phenomenon where two or more drugs used in combination, produce a more therapeutic effect than the sum of their individual effects. When drugs with distinct binding targets are strategically paired, they can leverage each other’s strengths and compensate for respective weaknesses, which enables lower individual drug doses, thereby reducing the risk of adverse effects. This characteristic is usually used in cancer and HIV treatments, and has been proven to be a new but promising way to combat complex diseases [7].

Zero Interaction Potency (ZIP) Yadav et al. [60] proposed ZIP score to describe the drug interaction by comparing the alteration in the potency of dose-response curves in the context of single drug administration versus their concurrent use in combinations. In the two-drug combination scenario, we will refer to drug A and drug B, respectively. The effects of these drugs are defined as E_{AB} for combination, E_A and E_B for individual situations ($0 \leq E \leq 1$). The ZIP score can be defined as follows:

$$S_{ZIP} = \bar{E}_{AB} - (\bar{E}_A + \bar{E}_B - \bar{E}_A \cdot \bar{E}_B). \quad (17)$$

The \bar{E}_{AB} in Equation (17) is the average response values obtained by fitting dose-response curves independently in each dimension of the measured combinatorial data sub-tensor, the explanation for the other variables are the same.

Bliss The Bliss independence model [24] assumes a stochastic process where the two drugs elicit their effects independently. In this model, the expected combination effect can be calculated based on the probability of the independent events occurring:

$$S_{Bliss} = E_{AB} - (1 - (1 - E_A)(1 - E_B)). \quad (18)$$

Equation (18) is similar to Equation (17), the difference between which is that Equation (17) employs fitted drug responses instead of observed ones.

Loewe The Loewe score [51] forecasts the dose combination that will produce a specific effect, it calculates the expected response as if both drugs are the same. Assume drug A can produce effect E_A at dose x_A , and drug B can produce effect E_B at dose x_B , then the loewe affinity states that the expected effect E_{Loewe} can be determined by:

$$\frac{x_A}{X_A} + \frac{x_B}{X_B} = 1, \quad (19)$$

where X_A and X_B are the doses drug A or B alone that produces effect E_{Loewe} . The Loewe score is then defined as follows:

$$S_{Loewe} = E_{AB} - E_{Loewe}. \quad (20)$$

Highest Single Agent (HSA) The HSA is a straightforward scoring system for estimating drug synergy, which determines the incremental effect of combining drugs by comparing the enhanced combined effect to their individual effects [60]. The HSA score can be calculated as follows:

$$S_{HSA} = E_{AB} - \max(E_A, E_B) \quad (21)$$

The ZIP score is the most prevalently utilized metric in the assessment of drug synergy. Furthermore, these scores can be collectively analyzed to identify optimal drug combinations for targeted therapy.

D Justification for the Choice of Baselines

One of the traditional strategies for dual-target drug design (refer to Sun et al. [57] as a comprehensive review) is linking pharmacophores by domain experts, such as chemists or pharmaceutical scientists.

We use SOTA deep-learning-based linker design methods (DiffLinker [25] and LinkerNet [19]) to mimic this procedure. Since each target pocket has a corresponding ligand molecule in our curated dataset, we can break the ligand molecule into fragments and select the ones critical for binding as pharmacophores. In this case, given respective critical fragments of dual pockets, we can design linkers and reassemble them into complete molecules. Ideally, the designed molecules contain critical pharmacophores that account for binding to dual targets and serve as a potential candidate for dual-target drugs.

E Preliminary Verification on Single-Target Setting

We evaluate Pocket2Mol and TargetDiff under the setting of single-target drug design on 438 unique targets in our dataset as a preliminary verification. We use each method to generate 10 molecules for each target and collect all generated molecules across 438 proteins and report the mean, trimmed mean (i.e., averaging that removes 10% of the largest and smallest values before calculating the mean) and median (denoted as “Avg.”, “T-Avg.” and “Med.” respectively) of affinity-related metrics (Vina Score, Vina Min, Vina Dock, and High Affinity) and property-related metrics (drug-likeness QED [3], synthesizability SA [14], and diversity). Vina Score assesses binding affinity directly based on the generated 3D molecules. Vina Min carries out a energy minimization over the local structure before estimation. Vina Dock incorporates a re-docking step to assess the highest binding affinity achievable. Meanwhile, High Affinity evaluates the proportion of generated molecules that have stronger binding affinity than the reference molecule for each protein tested.

The results are shown in Table 4, where molecules generated by TargetDiff exhibits binding affinities that are either comparable to or slightly greater than those of the reference molecules. Pocket2Mol achieves strong performance in terms of QED, SA, and diversity but fails in Vina-related metrics. Therefore, TargetDiff can be regarded as an effective molecular generative model on our dataset.

Table 4: Summary of different properties of reference molecules and molecules generated by baselines under the **single-target** setting. (↑) / (↓) denotes a larger / smaller number is better. Considering outliers significantly affect the mean, we ignore some abnormal values of Vina Score and Vina Min.

Methods	Vina Score (↓)			Vina Min (↓)			Vina Dock (↓)			High Affinity (↑)		QED (↑)		SA (↑)		Diversity (↑)	
	Avg.	T-Avg.	Med.	Avg.	T-Avg.	Med.	Avg.	T-Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.	Avg.	Med.
Reference	-	-8.02	-7.96	-	-8.07	-8.04	-8.09	-8.37	-8.25	-	-	0.54	0.55	0.74	0.78	-	-
Pocket2Mol	-3.14	-3.19	-3.16	-3.81	-3.77	-3.76	-4.87	-4.86	-4.85	2.1%	0.0%	0.49	0.49	0.87	0.89	0.83	0.85
TargetDiff	-	-7.38	-7.38	-	-7.89	-7.86	-8.77	-8.78	-8.75	56.0%	55.6%	0.51	0.53	0.58	0.58	0.69	0.70

F Visualization of More Examples

Here we provide visualization of more examples as shown in Figure 4.




Figure 4: Visualization of more reference molecules and examples designed by TargetDiff, LinkerNet and DUALDIFF.

G Proof of SE(3)-Equivariance

We denote the global SE(3) transformation as T_g , and which means the transformation as $T_g(\mathbf{x}_i) = \mathbf{R}_g \mathbf{x}_i + \mathbf{b}$, where $\mathbf{R}_g \in \mathbb{R}^{3 \times 3}$ is the rotation matrix and $\mathbf{b} \in \mathbb{R}^3$ is the translation vector. In line with Section 3.2, we define the composed message passing as follows:

$$\Delta \mathbf{h}_i^l(\mathcal{V}_v) := \sum_{j \in \mathcal{V}_v, i \neq j} f_{\theta_h}(\mathbf{h}_i^l, \mathbf{h}_j^l, d_{ij}^l, \mathbf{e}_{ij}) \quad (22)$$

$$\Delta \mathbf{x}_i^l(\mathcal{V}_v) := \sum_{j \in \mathcal{V}_v, i \neq j} (\mathbf{x}_i^l - \mathbf{x}_j^l) f_{\theta_x}(\mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, d_{ij}^l, \mathbf{e}_{ij}) \cdot \mathbf{1}_{\text{ligand}} \quad (23)$$

$$\mathbf{h}_i^{l+1} = \mathbf{h}_i^l + \frac{1}{2} (\Delta \mathbf{h}_i^l(\mathcal{V}_1) + \Delta \mathbf{h}_i^l(\mathcal{V}_2)) \quad (24)$$

$$\mathbf{x}_i^{l+1} := \phi(\mathbf{x}_i^l) = \mathbf{x}_i^l + \frac{1}{2} (\Delta \mathbf{x}_i^l(\mathcal{V}_1) + \Delta \mathbf{x}_i^l(\mathcal{V}_2)) \quad (25)$$

It is easy to see that the atomic distance $d_{ij}^l = \|\mathbf{x}_i - \mathbf{x}_j\|$ and \mathbf{e}_{ij} feature are both invariant to SE(3) transformation. The hidden embedding \mathbf{h}_i^l is also invariant since the its updates (as shown in Equation (22) and Equation (24)) are only related to invariant features.

We define $\Delta T_g(\mathbf{x}_i^l(\mathcal{V}_v))$ as

$$\begin{aligned}\Delta T_g(\mathbf{x}_i^l(\mathcal{V}_v)) &:= \sum_{j \in \mathcal{V}_v, i \neq j} (T_g(\mathbf{x}_i^l) - T_g(\mathbf{x}_i^l)) f_{\theta_x}(\mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, d_{ij}^l, \mathbf{e}_{ij}) \cdot \mathbf{1}_{\text{ligand}} \\ &= \sum_{j \in \mathcal{V}_v, i \neq j} (\mathbf{R}_g \mathbf{x}_i^l + \mathbf{b} - \mathbf{R}_g \mathbf{x}_j^l - \mathbf{b}) f_{\theta_x}(\mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, d_{ij}^l, \mathbf{e}_{ij}) \cdot \mathbf{1}_{\text{ligand}} \\ &= \sum_{j \in \mathcal{V}_v, i \neq j} \mathbf{R}_g(\mathbf{x}_i^l - \mathbf{x}_j^l) f_{\theta_x}(\mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, d_{ij}^l, \mathbf{e}_{ij}) \cdot \mathbf{1}_{\text{ligand}}\end{aligned}$$

After applying T_g to \mathbf{x}_i^l , the updated position $\mathbf{x}_i^{l+1} = \phi(\mathbf{x}_i^l)$ can be written as (using the above results):

$$\begin{aligned}\phi(T_g(\mathbf{x}_i^l)) &= T_g(\mathbf{x}_i^l) + \frac{1}{2}(\Delta T_g(\mathbf{x}_i^l(\mathcal{V}_1)) + \Delta T_g(\mathbf{x}_i^l(\mathcal{V}_2))) \\ &= \mathbf{R}_g \mathbf{x}_i^l + \mathbf{b} + \frac{1}{2} \left(\sum_{v \in 1, 2} \sum_{j \in \mathcal{V}_v, i \neq j} \mathbf{R}_g(\mathbf{x}_i^l - \mathbf{x}_j^l) f_{\theta_x}(\mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, d_{ij}^l, \mathbf{e}_{ij}) \cdot \mathbf{1}_{\text{ligand}} \right) \\ &= \mathbf{R}_g \mathbf{x}_i^l + \frac{1}{2} \mathbf{R}_g \left(\sum_{v \in 1, 2} \sum_{j \in \mathcal{V}_v, i \neq j} (\mathbf{x}_i^l - \mathbf{x}_j^l) f_{\theta_x}(\mathbf{h}_i^{l+1}, \mathbf{h}_j^{l+1}, d_{ij}^l, \mathbf{e}_{ij}) \cdot \mathbf{1}_{\text{ligand}} \right) + \mathbf{b} \\ &= \mathbf{R}_g \mathbf{x}_i^l + \frac{1}{2} \mathbf{R}_g (\Delta \mathbf{x}_i^l(\mathcal{V}_1) + \Delta \mathbf{x}_i^l(\mathcal{V}_2)) + \mathbf{b} \\ &= \mathbf{R}_g \left(\mathbf{x}_i^l + \frac{1}{2} (\Delta \mathbf{x}_i^l(\mathcal{V}_1) + \Delta \mathbf{x}_i^l(\mathcal{V}_2)) \right) + \mathbf{b} \\ &= T_g \left(\mathbf{x}_i^l + \frac{1}{2} (\Delta \mathbf{x}_i^l(\mathcal{V}_1) + \Delta \mathbf{x}_i^l(\mathcal{V}_2)) \right) \\ &= T_g(\phi(\mathbf{x}_i^l))\end{aligned}$$

The above equation shows the SE(3)-equivariance of the atom position update formula Equation (25). Based on the fact that Equation (24) is SE(3)-invariant and Equation (25) is SE(3)-equivariant, we can say that the composition operation of our method is SE(3)-equivariant.

H Computational Efficiency

We test the time for generating 10 molecules by TargetDiff (single-target drug design) and DiffLinker, LinkerNet, CompDiff, and DualDiff (dual-target drug design). The results are shown in Table 5.

Table 5: Time Cost for generating 10 molecules by TargetDiff (single-target drug design) and DiffLinker, LinkerNet, CompDiff, and DualDiff (dual-target drug design).

Setting	Method	Time (seconds)
Single-Target	TargetDiff	220
Dual-Target	DiffLinker	31
	LinkerNet	30
	COMPDIFF	395
	DUALDIFF	493

Linker design methods (DiffLinker and LinkerNet) are faster than de novo design methods. The reasons are twofold: 1. For linker design, parts of the molecules are provided; thus, it is only necessary to generate the remaining part (i.e., the linker). 2. For DiffLinker, only one of the dual targets can serve as the input of the model. For linkerNet, it cannot use pockets as conditions. Thus, the computational cost is lower.

As the table shows, the inference time of CompDiff and DualDiff (dual-target drug design) is about twice as that of TargetDiff (single-target drug design). This is as expected because there are two heterogeneous graphs and the messages on the two graphs are processed and aggregated separately.

(sequentially in our implementation) and then composed as described in our paper. There is space for optimizing the inference speed of CompDiff and DualDiff by parallelling the message-passing operations over the two graphs (when GPU memory is sufficient). This code optimization will only speed up the inference and not change the generated results. Ideally, this will make the inference time of CompDiff and DualDiff almost the same as that of TargetDiff.

Though the computational cost of our method is higher than the baselines, it is still acceptable in practice.

General acceleration methods (e.g., pruning, quantization, and distillation) for deep learning models can be applied to the models used in our framework. Besides, fast diffusion sampling solvers can also be applied to our framework, with DPM-Solver [40] as one of the most representative works, which can significantly reduce the number of sampling steps without sacrificing the quality of generated samples.

I Discussion, Limitation, and Future Work

Our work provides a novel dataset and a general framework for dual-target drug design to the community. And our method can be easily adapted to the multi-target scenario. Our work is the first step towards generative dual-target drug design. There are still limitations in our work. For example, we do not consider flexibility of proteins in our work, which is a more practical setting, though this is also an issue for most works in SBDD. Another limitation of our work is the lack of validation through wet-lab experiments. We will leave these as future work.

J Societal Impacts

Our research holds the promise of significantly advancing the pharmaceutical industry by aiding in the development of potent dual-target drugs. This could potentially streamline the path to new treatments, making efficient drug discovery a more attainable goal. Moreover, emphasizing the ethical implementation of our methods is of paramount importance. It is also needed to ensure that these scientific achievements are utilized for social good, safeguarding against any misuse that could lead to negative consequences for society.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The main claims made in the abstract and introduction precisely reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Limitations are discussed in Appendix I.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: The paper provide the full set of assumptions and a complete and correct proof.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The paper fully discloses all the information needed to reproduce the main experimental results of the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We require authorization to release the code. Once we obtain the necessary approval, we will proceed with the code release.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The paper has provided the procedures of processing the dataset and the chosen hyperparameters.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: For RMSD between docked poses towards dual targets, the paper has reported error bars. The main results are reported in mean and average over a large test dataset. The experimental results are stable and significant.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: **[No]**

Justification: The experiments are done on several clusters. The overall accurate computer resources are not well recorded.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: **[Yes]**

Justification: The research conducted in the paper conforms, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: **[Yes]**

Justification: The broader impacts are discussed in Appendix J.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The creators or original owners of assets (e.g., code, data, models), used in the paper, are properly credited.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.