D^{2}G-TO: Task-aware and OOD-guided Discrete Graph Diffusion for Robust CNS Drug Discovery

Xue Zhai[1], **Chu-An Yang**[1,2], **Minghao Liu**[3], **Xu Dong**[1], **Han Wang**[1,2], **Weiwei Han**[3], **Ting Gao**[1*] and **LiHong Hu**[1*]

[1]School of Information Science and Technology, Northeast Normal University, Changchun, China
[2]Institute of Computational Biology, School of Information Science and Technology, Northeast Normal University, Changchun, China
[3]Key Laboratory for Molecular Enzymology and Engineering of Ministry of Education, School of Life Sciences, Jilin University, Changchun, China
lhhu@nenu.edu.cn, gaot080@nenu.edu.cn

Abstract

Central nervous system (CNS) drug discovery is constrained by an immense and sparse chemical search space. Meanwhile, molecules that simultaneously achieve brain penetration, target efficacy, and synthesizability are extremely scarce. However, existing generative models rarely couple rigorous multi-objective control with robustness to distribution shift, limiting their reliability in realistic CNS design. We introduce D^{2}G-TO, a task-aware and out-of-distribution (OOD)-guided discrete graph diffusion framework that unifies multi-pharmacological properties with structural distribution guidance. A novel Structural Similarity Guidance mechanism steers generation toward in-distribution regions while repelling OOD modes, maintaining structural distributional consistency in realistic scenarios. Across BBBP, BACE, and QM9 benchmarks, D^{2}G-TO achieves strong validity, diversity, and other metrics. In an Alzheimer's disease case study, we subject the generated molecules to cross-property pharmacological prediction and systematic ADMET profiling, followed by structure-based molecular docking against BACE-1 to assess binding-mode plausibility. D^{2}G-TO identifies candidates that jointly satisfy blood–brain barrier permeability, β-site amyloid precursor protein-cleaving enzyme 1 inhibition, and synthetic accessibility. Thus, D^{2}G-TO has the potential to serve as an efficient *in silico* engine for early-stage CNS drug design. The code is available at https://github.com/zhaix922/DDG_TO.

1 Introduction

The therapeutic landscape of central nervous system (CNS) disorders remains one of the most intractable frontiers in biomedical research. CNS drug discovery success rates are markedly lower than in most other disease areas, while clinical demand vastly exceeds the available treatment options. This persistent gap reflects three tightly coupled challenges:

- The chemical space is vast while truly viable structures are extremely sparse [Paul *et al.*, 2010; Scannell *et al.*, 2012; Stumpfe and Bajorath, 2012];
- CNS drugs must finely balance blood–brain barrier permeability and overall efficacy, where even slight deviations can readily erode clinical benefit [Pardridge, 2005; Cummings *et al.*, 2014; DiChiara *et al.*, 2024];
- Complex molecular scaffolds markedly amplify risks in synthesis and development, forcing synthetic accessibility to become a hard constraint already in the early design stage [Dichiara *et al.*, 2009];

Jointly optimizing chemical space exploration, CNS-critical pharmacological properties, and synthetic accessibility at the hit and lead discovery stages has thus become a central challenge that ultimately shapes clinical translation success.

In recent years, computer-aided drug design (CADD) and deep learning have advanced rapidly, with graph neural networks (GNNs) and related representation learning methods achieving strong performance in property prediction and virtual screening, thereby helping to shorten development cycles and improve success rates [Cheng *et al.*, 2021]. Building on this, different generative models have been introduced for molecular design, aiming to directly propose molecules with desired properties in vast chemical spaces [Yang *et al.*, 2023; Du *et al.*, 2024]. However, in practical drug discovery, property control and molecular generation remain weakly integrated, and the combined challenge of multi-objective design and distributional robustness has yet to be systematically addressed [Cheng *et al.*, 2021]. Most existing frameworks still implicitly assume that training and deployment share similar data distributions, making it difficult to explicitly handle distribution shift and out-of-distribution (OOD) structures in realistic CNS scenarios [Liu *et al.*, 2023; Wu *et al.*, 2022].

To address these challenges, we propose D^{2}G-TO (Task-aware and OOD-guided Discrete Graph Diffusion), a task-aware and distribution-consistency discrete diffusion framework for CNS-oriented molecular graph generation. D^{2}G-TO injects multiple CNS-relevant pharmacological objectives—such as target inhibition, brain penetration, and synthetic

accessibility—directly into the diffusion process. While the structural guidance pulls samples toward in-distribution (ID) graphs and pushes them away from potential OOD modes. D^2G-TO achieves strong performance in molecular validity, diversity, and distributional robustness across CNS benchmarks. In an Alzheimer's disease (AD) case study, D^2G-TO prioritizes early-stage candidate compounds that jointly satisfy efficacy and developability requirements, followed by ADMET profiling and structure-based docking to assess binding-mode plausibility. Thus, D^2G-TO is practical *in silico* design to accelerate CNS drug discovery.

2 Related Work

2.1 Multiple Challenges in CNS Drug Discovery

Drug discovery remains one of the most formidable challenges in CNS disorders. Despite decades of research and investment, progress in clinical translation has been slow—such as AD, where repeated Phase III failures and success rate reduced to below 0.4% [Cummings *et al.*, 2014]. This stagnation arises from three intertwined challenges: chemical space, pharmacological property, and synthetic accessibility.

The vastness of chemical space provides a fundamental challenge for CNS drug discovery. With an estimated $> 10^{60}$ drug-like molecules, exploration through traditional screening is infeasible [Paul *et al.*, 2010; Scannell *et al.*, 2012; Stumpfe and Bajorath, 2012; Vamathevan *et al.*, 2019; Cheng *et al.*, 2021; Mullard, 2017; Berdigaliyev and Aljofan, 2020; Zhang *et al.*, 2025]. And pharmacological properties impose intrinsic constraints on molecular optimization. CNS drug design must simultaneously balance blood–brain barrier (BBB) permeability and precise target selectivity, both serving as key determinants in diseases such as AD [Pardridge, 2005; Cummings *et al.*, 2014]. However, these requirements often conflict, narrowing the drug design space and hindering balanced optimization. In addition, synthetic accessibility further increases these challenges, as complex or unstable scaffolds demand costly synthesis cost and validation, prolonging development and limiting scalability [DiChiara *et al.*, 2024; Dichiara *et al.*, 2009].

These challenges need for computational frameworks capable of navigating complex drug design trade-offs while jointly optimizing pharmacological, synthetic, and structural objectives. CADD has shown substantial potential in supporting multiple stages of drug development, significantly enhancing the efficiency of drug discovery [Button *et al.*, 2019]. With the rapid advancement of artificial intelligence, deep learning models have emerged—particularly GNNs—demonstrating strong capabilities in encoding molecular topology and predicting key pharmacological properties[Cheng *et al.*, 2021; Choung *et al.*, 2023; Huang *et al.*, 2021; Jing *et al.*, 2025]. However, most related studies on CNS remain focused on property prediction and optimization. These models effectively evaluate existing compounds, but they limit capability in generating novel molecules that simultaneously satisfy pharmacological, synthetic, and structural constraints. To address these limitations, generative modeling offers new possibilities for advancing CNS drug design.

2.2 Generative Diffusion and OOD Generation

Generative modeling has emerged as a transformative paradigm for drug discovery, providing an approach to explore the vast and uncharted chemical space and to design novel, diverse, pharmacological compounds beyond existing libraries [Cheng *et al.*, 2021; Zhavoronkov *et al.*, 2019; Du *et al.*, 2024; Ren *et al.*, 2023].

Early based sequence and graph generative frameworks, such as SMILES-based RNNs and LSTMs [Moret *et al.*, 2023], as well as variational autoencoders including JT-VAE [Jin *et al.*, 2018] and its Bayesian variant JTVAE-BO [Gómez-Bombarelli *et al.*, 2018], demonstrated the feasibility of learning molecular distributions for de novo design. However, these models often depend on heuristic latent sampling or post hoc property optimization. Therefore, they limit the generation capacity about multi-objective constraints, particularly those related to pharmacological property, synthetic accessibility, and structural rationality. Subsequent graph-based genetic algorithms GB-GA expanded chemical diversity through stochastic graph mutations, but lacked mechanisms to integrate domain priors for guided optimization [Jensen, 2019]. Additionally, fragment-based methods like MARS improved sampling efficiency though adaptive graph edits and MCMC exploration, yet remained limited in jointly handling complex, multi-objective constraints[Xie *et al.*, 2021].

Diffusion models, based in probabilistic formulation, have recently emerged as a dominant approach for generation. Graph diffusion models further extend the capability to capture complex structural distributions [Liu *et al.*, 2023; Yang *et al.*, 2023]. Continuous diffusion models, such as EDP-GNN and GDSS, generate molecular graphs through iterative denoising with score-based Langevin dynamics, while discrete approaches like DiGress introduces a discrete denoising diffusion framework that reformulates graph generation as node and edge classification tasks [Lee *et al.*, 2023; Niu *et al.*, 2020; Vignac *et al.*, 2023; Jo *et al.*, 2022]. Furthermore, GraphDiT extends discrete diffusion framework with a denoiser based on transformer and a graph-dependent noise model, and DiffMC-Gen introduces a dual denoising architecture that combines discrete graph diffusion with SE(3)-equivariant denoising in continuous 3D space [Liu *et al.*, 2024; Yang *et al.*, 2025].

Despite these advances, most diffusion models assume a consistent distribution between training and testing graphs. This condition is rarely satisfied in the real world for CNS drug design. Newly generated molecules often deviate from training distributions, resulting in OOD samples that affect chemical validity and pharmacological reliability. Addressing this distributional robustness challenge has become crucial for trustworthy generation models [Wu *et al.*, 2022; Liu *et al.*, 2024]. Recent studies such as MOOD and PGR-MOOD introduce OOD generation in continuous diffusion process. But these methods remain external to the generative process and absent from discrete diffusion models [Lee *et al.*, 2023; Shen *et al.*, 2024]. Therefore, developing distribution-aware diffusion frameworks that integrate property control, structural rationality, and distributional robustness constitutes an essential step toward reliable CNS drug design.

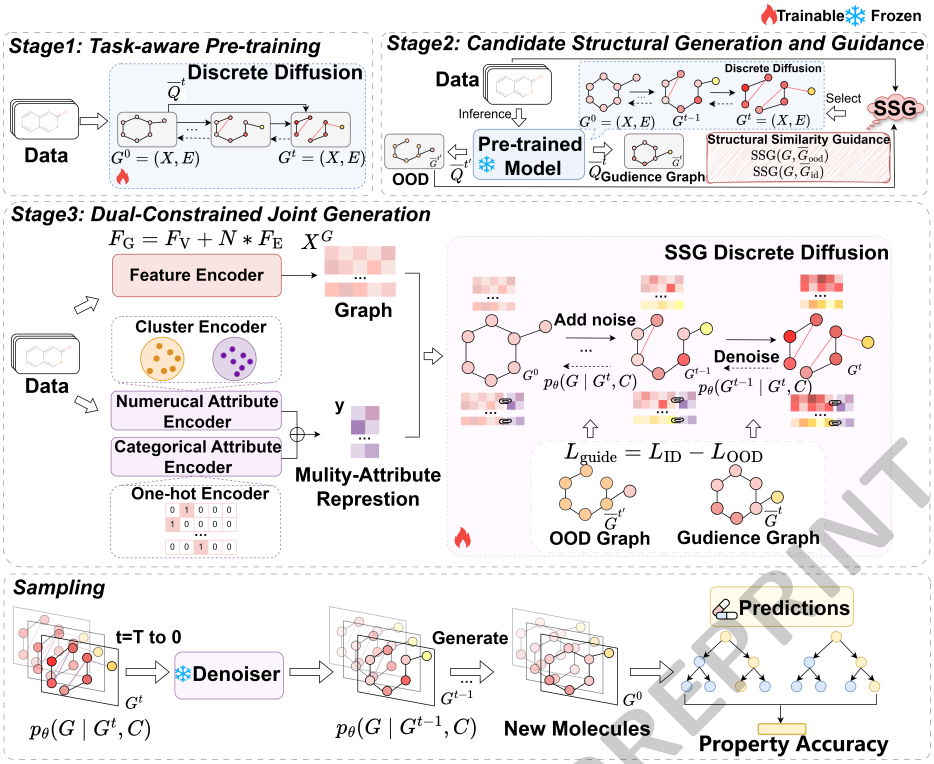

Figure 1: Overview of the proposed D^{2}G-TO framework.

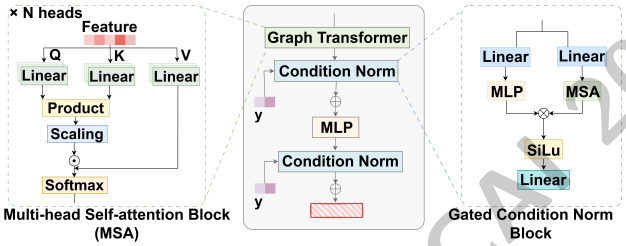

Figure 2: Denoising Network Architecture based on Graph Transformer.

3 Method

We propose Task-aware and OOD-guided Discrete Graph Diffusion (D^{2}G-TO) method, a three-stage discrete diffusion framework for molecular graph generation. As illustrated in Figure 1, the framework unifies task-aware modeling and distribution-guided diffusion within a single generative pipeline. In Stage 1, we pre-train a discrete diffusion model with multi-pharmacological conditioning to learn property-aware molecular graph representations. In Stage 2, we generate candidate ID/OOD graphs and introduce a Structural Similarity Guidance (SSG) module to select guidance structures that remain close to ID regions while repelling OOD regions. In Stage 3, we train denoising network based on Graph Transformer for noise molecular graphs to generate new molecules. The reverse diffusion process is governed by dual constraints that couple property-aware task optimization with distribution-consistent structural guidance, enabling task- and structure-guided joint molecular generation for precise and robust CNS drug discovery.

3.1 Task-Aware Pre-training

In Stage 1, we introduce multi-pharmacological properties (e.g., permeability, synthetic accessibility) into a discrete diffusion framework to learn molecular graph representations conditioned on these properties, providing a stable generative prior for subsequent stages (Figure 1, Stage 1).

Given a molecular graph $G = (X, E)$, with X the node/atom feature and E the edge/bond feature tensorthe, the forward diffusion process incrementally perturbs node and edge states:

$$q\left(G^t \mid G^{t-1}\right) = \prod_i q\left(x_i^t \mid x_i^{t-1}\right) \prod_{(i,j)} q\left(e_{ij}^t \mid e_{ij}^{t-1}\right), \quad (1)$$

where $x_i^t$and e_{ij}^tdenote node and edge states at timestep t. The reverse process reconstructs clean molecules under multi-pharmacological conditions $\{y_m\}_{m=1}^M$:

$$p_\theta(G^{t-1} \mid G^t, y_1, \ldots, y_M) = f_\theta(G^t, y_1, \ldots, y_M), \quad (2)$$

where $f_\theta(\cdot)$is a denoising network conditioned on multi-pharmacological properties. These property-aware representations serve as the basis for distribution-guided diffusion in later stages.

3.2 Candidate Structural Generation and Guidance

In Stage 2, we simulate pseudo-OOD structures using training set and construct guidance graphs in the discrete diffusion

space. The stage learns explicit structural distribution boundaries, thereby avoids anomalies and enhances structural robustness during the generation.

In Figure 1 (Stage 2), we use the pretrained model (Stage 1) to generate candidate guidance structures and inject targeted perturbations at node and edge levels to construct a set of candidate OOD graphs G_{OOD}. And they are used only as repulsive structural references rather than desired generation targets. These graphs differ markedly from the training distribution (ID) in topological patterns, thereby approximating potential OOD regions.

To distinguish ID and OOD samples, we introduce the Structural Similarity Guidance (SSG) module, which evaluates, for each candidate graph, its structural deviation from both ID and OOD reference sets. SSG is defined as a joint topology–feature distance. This design preserves global structural information, making it amenable to end-to-end optimization within diffusion training [Shen *et al.*, 2024; Vayer *et al.*, 2019]:

$$d^{SSG}(G_1, G_2) = \alpha \|\mathcal{A}_1 - \mathcal{A}_2\|_F^2 + \frac{1-\alpha}{N_1 N_2} \sum_{i,j} \|x_1(i) - x_2(j)\|_2^2, \quad (3)$$

where $\mathcal{A}_1, \mathcal{A}_2$ are graph adjacency matrices, $i = 1, \ldots, N_1$, $j = 1, \ldots, N_2$, and $x_k(i)$ denotes node features in graph G_k. The first term captures global topological divergence, while the second term measures cross-graph feature deviation.

Based on SSG distances to ID and OOD sets, we define a bilateral selection criterion:

$$G_{ID} = \arg\min_{G} \left(d_{ID}^{SSG}(G) - \gamma d_{OOD}^{SSG}(G) \right), \quad (4)$$

where $\gamma > 0$ is a weighting hyperparameter. The selected graphs G_{ID} and G_{OOD} are used as structural guidance in the subsequent generation stage.

3.3 Dual-Constrained Joint Generation

As illustrated in Figure 1 (Stage 3), the reverse diffusion process is governed by dual constraints that couple task-guided property optimization and distribution-guided structural constraint.

Forward diffusion

The forward process follows the same Markovian noising scheme in discrete graph space:

$$q\left(G^t \mid G^{t-1}\right) = \prod_i q\left(x_i^t \mid x_i^{t-1}\right) \prod_{(i,j)} q\left(e_{ij}^t \mid e_{ij}^{t-1}\right), \quad (5)$$

which gradually corrupts a clean molecular graph G^0 into a noisy graph G^T that serves as the start for denoising. Before reverse diffusion, we encode M pharmacological attributes $(y_1, \ldots, y_M)$ into a property-aware vector:

$$y = [\, \phi(y_1) \oplus \phi(y_2) \oplus \cdots \oplus \phi(y_M) \,], \quad (6)$$

where cluster encoder maps numerical attributes (e.g., synthetic accessibility) and one-hot encoder maps categorical attributes (e.g., permeability). The resulting vector y is used as a property-aware control signal for the denoising network.

Reverse denoising with joint constraints

The reverse diffusion starts from G^T and iteratively reconstructs molecular structures under the joint control of multi-pharmacological properties y and structural guidance (G_{ID}, G_{OOD}):

$$p_\theta\left(G^{t-1} \mid G^t, G_{ID}, G_{OOD}, y_1, \ldots, y_M\right) = f_\theta\left(G^t, y, G_{ID}, G_{OOD}\right), \quad (7)$$

where $f_\theta(\cdot)$ is a denoising network based on Graph Transformer with joint constraints.

Task-guided property optimization. At each step t to $t-1$, the denoiser employs a Condition Norm module to incorporate the multi-property vector y into feature updates:

$$\hat{h}_l = \mathrm{CondNorm}\left(h_l; y\right), \quad (8)$$

where h_l denotes hidden features at layer l. This process steers the denoising towards molecules to satisfy target design, such as high permeability and low synthetic accessibility. The task loss is defined as:

$$\mathcal{L}_{task} = \sum_{m=1}^{M} L_m\left(G^t, y_m\right), \quad (9)$$

where L_m denotes the supervised loss associated with each pharmacological property.

Distribution-consistency constraint. We train distribution consistency using the SSG module d^{SSG} introduced in Eq. 3. And we define the structural distribution consistency loss as:

$$\mathcal{L}_{SSG} = \beta_1 d_{ID}^{SSG}(G^t) - \beta_2\, d_{OOD}^{SSG}(G^t), \quad (10)$$

where $d_{ID}^{SSG}(\cdot)$ and $d_{OOD}^{SSG}(\cdot)$ denote SSG diversity to ID and OOD guidance sets, and $\beta_1, \beta_2 > 0$ control the balance between attraction to ID and repulsion from OOD.

Dual-Constraint Training Objective. The overall loss in Stage 3 is:

$$\mathcal{L}_{total} = \mathcal{L}_{diff} + \lambda_1 \mathcal{L}_{task} + \lambda_2 \mathcal{L}_{SSG}, \quad (11)$$

where $\mathcal{L}_{diff}$ denotes the diffusion reconstruction loss and $\lambda_1, \lambda_2 > 0$ are the weights.

From Noise to Molecule Generation. After training, molecule generation corresponds to the Sampling stage in Figure 1. Starting from a randomly initialized noisy molecular graph G^T, we use the denoising network $f_\theta(\cdot)$, and simulate the reverse Markov chain $p_\theta(G^{t-1} \mid G^t, C), C = (y, G_{ID}, G_{OOD})$. From $t = T$ down to $t = 0$ to obtain the final sample G^0.

Denoising Network Architecture. The denoising network is implemented as a Graph Transformer that updates node and edge features. In Figure 2, it combines Multi-head Self-attention Block (MSA) [Peebles and Xie, 2023; Dwivedi and Bresson, 2021] with Gated condition Norm Block to support task-aware, distribution-guided generation. The MSA is:

$$Attention(Q, K, V) = Softmax(\frac{QK^\top}{\sqrt{d_k}})V. \quad (12)$$

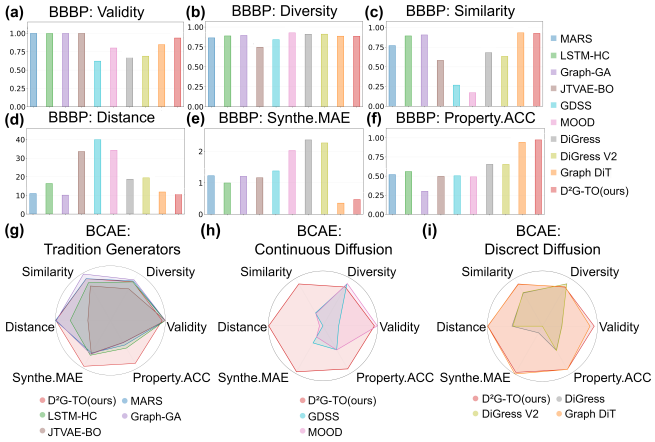

Figure 3: Overall performance comparison with baseline models on CNS benchmarks. (a-f) Bar plots on BBBP dataset. (g-i) Radar chart on BACE dataset.

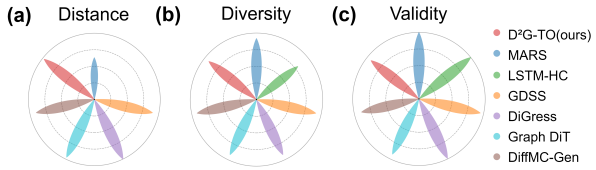

Figure 4: (a-c) Performance comparison with baseline models on QM9 dataset.

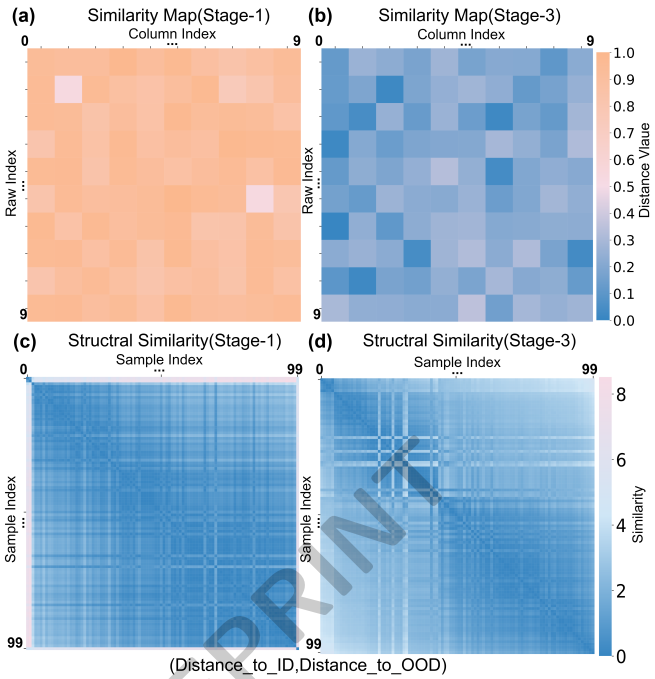

Figure 5: Analysis of structural distribution consistency. (a) (b) Similarity maps between generated molecules and guidance graphs . (c) (d) Sample×Sample structural similarity matrices.

Multi-pharmacological properties are incorporated via the Condition Norm module in Eq. 8, keeping latent representations aligned with the CNS drug design objectives throughout denoising. To fuse property-aware signals with structure-distribution information, we employ a gated fusion block:

$$h' = \sigma(MLP(h)) \otimes MSA(h), \tag{13}$$

where $\sigma(MLP(h))$ acts as an element-wise gate applied to the attention output $MSA(h)$, enabling an adaptive balance between property condition and structure features along the denoising trajectory.

Consequently, we enforce a dual constraint regime of property-aware task optimization and distribution-consistent structural guidance, enabling the model to generate molecules that satisfy multi-pharmacological property requirements while remaining structurally plausible distributions.

4 Experiments

4.1 Experiment Setup

Datasets. We evaluate on two CNS-related benchmarks and one large-scale small molecule dataset. For CNS drug design, we adopt the BBBP and BACE subsets of MoleculeNet [Wu *et al.*, 2018], which provide categorical property labels for BBB permeability and β-site amyloid precursor protein-cleaving enzyme 1 (BACE-1) inhibition, and we further assign each molecule a synthetic accessibility (SAS) score as a numerical property [Dichiara *et al.*, 2009]. To assess generation over a broader region of chemical space, we additionally include the QM9 dataset comprising approximately 130k small organic molecules [Ramakrishnan *et al.*, 2014].

Evaluation metrics. For CNS benchmarks (BBBP, BACE), we evaluate the following metrics: **Validity** (fraction of chemically valid molecules), **Diversity** (internal structural diversity among generated molecules), **Similarity** (fragment-based similarity to the reference set), **Distance** (Fréchet ChemNet Distance to the reference distribution), **Synthe.MAE** (mean absolute error between generated and target synthetic accessibility scores), and **Property.Acc** (classification accuracy on task-related categorical conditions such as BBBP/BACE labels, computed by an oracle). For QM9, we focus on Validity, Diversity, and Distance to assess basic chemical correctness, structural variability, and distributional differences in a large-scale small molecule setting. The evaluation oracle is implemented as a random forest trained on task-related molecules; therefore, Property.Acc should be interpreted as a surrogate indicator of conditional consistency rather than direct experimental evidence of pharmacological activity. Lower Distance and Synthe.MAE, and higher Validity, Diversity, Similarity, and Property.Acc indicate better performance.

Baselines. We compare with a broad set of molecular generative models, including traditional methods GraphGA [Jensen, 2019], MARS [Xie *et al.*, 2021], JTVAE with Bayesian optimization (JTVAE-BO) [Jin *et al.*, 2018], and LSTM-HC on SMILES [Brown *et al.*, 2019]; continuous diffusion models GDSS [Jo *et al.*, 2022] and MOOD [Lee *et al.*, 2023]; and discrete diffusion models DiGress [Vignac *et al.*, 2023] and GraphDiT [Liu *et al.*, 2024]. On QM9 exper-

	Validity	Coverage	Diversity	Similarity	Distance↓	Synthe.MAE↓	Property.ACC	Rank
MARS	**1.0000**	8/9	0.8637	0.7696	10.9791	1.2250	0.5189	5.5000
LSTM-HC	0.9990	8/9	0.8883	0.8932	16.3904	0.9969	0.5590	4.0000
Graph-GA	**1.0000**	9/9	0.8950	0.9059	**10.1659**	1.2082	0.3015	3.0000
JTVAE-BO	**1.0000**	5/9	0.7458	0.5821	33.5746	1.1619	0.4958	9.0000
GDSS	0.6218	3/9	0.8415	0.2679	39.9440	1.3788	0.5037	10.0000
MOOD	0.8008	9/9	**0.9273**	0.1715	34.2506	2.0284	0.4903	8.0000
DiGress	0.6960	9/9	0.9098	0.6805	18.6921	2.3658	0.6536	5.5000
DiGress V2	0.6892	9/9	0.9107	0.6336	19.4498	2.2694	0.6531	7.0000
Graph DiT	0.8468	9/9	0.8856	**0.9329**	11.8519	**0.3551**	0.9417	2.0000
D^2G-TO(ours)	0.9368	**9/9**	0.8843	0.9259	10.7091	0.4690	**0.9733**	**1.0000**

Table 1: Comparison with baseline models on the BBBP dataset. Lower is better (↓) for Distance and Synthe.MAE.

	Validity	Coverage	Diversity	Similarity	Distance↓	Synthe.MAE↓	Property.ACC	Rank
MARS	**1.0000**	8/8	0.8338	0.8827	**6.7923**	1.0123	0.5184	**1.0000**
LSTM-HC	0.9972	8/8	0.8146	0.7982	17.5585	0.9207	0.5816	5.0000
Graph-GA	**1.0000**	8/8	0.8585	**0.9805**	7.4104	0.9633	0.4690	4.0000
JTVAE-BO	**1.0000**	6/8	0.6682	0.7281	30.3696	0.9923	0.4628	9.0000
GDSS	0.2879	4/8	0.8756	0.2708	46.7539	1.6422	0.5036	10.0000
MOOD	0.9947	8/8	**0.8902**	0.2587	44.2394	1.8853	0.5062	6.0000
DiGress	0.3511	8/8	0.8862	0.6942	24.6560	2.0681	0.5061	8.0000
DiGress V2	0.3546	8/8	0.8812	0.7027	25.3270	2.3365	0.5113	7.0000
Graph DiT	0.8646	8/8	0.8240	0.8757	6.9836	**0.4053**	0.9050	2.5000
D^2G-TO(ours)	0.9412	**8/8**	0.8174	0.8813	7.1309	0.4708	**0.9051**	2.5000

Table 2: Comparison with baseline models on the BACE dataset. Lower is better (↓) for Distance and Synthe.MAE.

	Validity	Diversity	Distance↓	Rank
MARS	**1.0000**	0.92	7.34	3.0000
LSTM-HC	**1.0000**	0.80	18.43	6.5000
GDSS	0.9513	0.92	2.54	4.0000
DiGress	0.9939	0.92	**0.90**	**1.0000**
Graph DiT	0.9221	0.92	1.43	5.0000
DiffMC-Gen	0.8900	0.92	2.57	6.5000
D^2G-TO(ours)	0.9260	**0.92**	1.34	2.0000

Table 3: Comparison with baseline models on the QM9 dataset.

iments, we further include DiffMC-Gen [Yang *et al.*, 2025]. All baselines are implemented following their original settings and evaluated under the same protocol as ours.

4.2 Comparison with Baseline Methods

Comparisons between D^2G-TO and baseline models on the CNS benchmarks and the QM9 dataset are reported in Tables 1–3 and Figure 3–4. On BBBP and BACE, D^2G-TO delivers strong performance and improves Property Acc. by 3% over the best competing method. On the large-scale QM9 dataset, D^2G-TO also achieves the best overall rank among all baselines. The bar plots, radar and petal charts in Figure 3–4 further show that D^2G-TO attains a more favorable balance among Validity, Diversity, and distributional Distance metrics than traditional generators, continuous diffusion models, and discrete diffusion models.

4.3 Model Analysis

Ablation Studies

To verify the effectiveness of the components in our D^2G-TO model, we design ablation studies on BBBP and BACE benchmarks. In Table 4, introducing SSG guidance consistently improves Property.ACC and reduces both Distance and Synthe MAE. These changes indicate that the generated molecules more closely follow the target chemical-space distribution and exhibit enhanced synthetic accessibility. While Validity, Coverage and Diversity are at least maintained compared with the w/o SSG guidence. These results show that the proposed guidance improves distributional quality without harming the basic generative capability.

Analysis of Structural Distribution Consistency

To further analyze structural distribution consistency, we project each generated molecule into the Structural Distribution State Space (SDSS) as a 2D descriptor $[d_{\mathrm{ID}}, d_{\mathrm{OOD}}]$. This SDSS effectively filters out semantic redundancy and high dimensional interference, and we then compute pairwise Euclidean distances in this space. On BBBP dataset, we visualize (i) similarity maps between generated molecules and guidance graphs, and (ii) sample×sample structural similarity matrices based on $\Delta = d_{\mathrm{OOD}} - d_{\mathrm{ID}}$ for Stage 1 (with-

Dataset	Method	Validity	Coverage	Diversity	Similarity	Distance↓	Synthe.MAE↓	Property.ACC
BBBP	w/o SSG	0.8691	8/9	0.8764	0.8991	13.4613	0.6442	0.9506
	D^{2}G-TO	**0.9368**	**9/9**	**0.8843**	**0.9259**	**10.7091**	**0.4690**	**0.9733**
BACE	w/o SSG	0.7031	7/8	0.8133	0.8599	9.1003	0.8057	0.8861
	D^{2}G-TO	**0.9412**	**8/8**	**0.8174**	**0.8813**	**7.1309**	**0.4708**	**0.9051**

Table 4: Ablation of Structural Similarity Guidance (SSG) on BBBP and BACE.

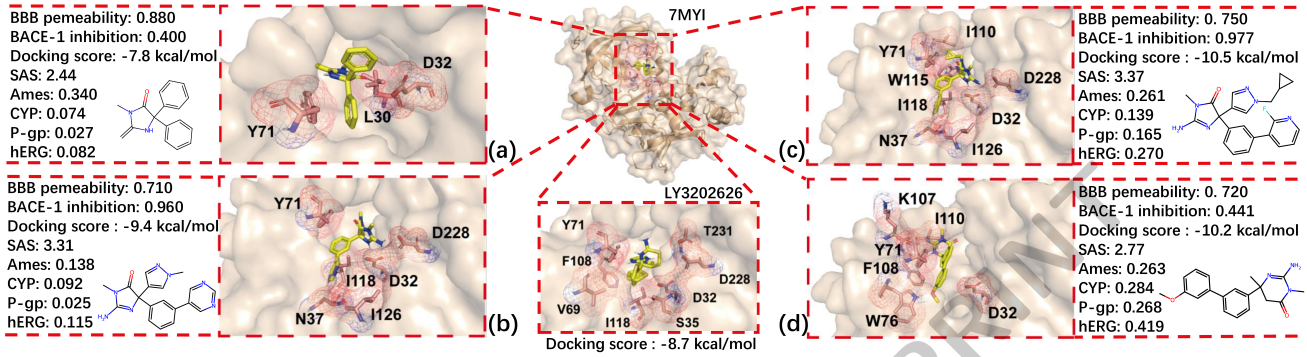

Figure 6: Case study results on AD: four selected CNS-oriented candidates.

out SSG guidance) and Stage 3 (D^{2}G-TO with SSG guidance). These plots characterize the distribution of generated molecules under an ID/OOD structural reference.

In Figures. 5, the Stage 1 (a) similarity map is dominated by warm colors and indicates that many molecules violate the "close to ID, far from OOD" requirement. But the Stage 3 (b) map is cooler color and shows that most samples now satisfy this constraint under SSG guidance. For the sample×sample matrices, molecules are first sorted in descending order of Δ. The Stage 1 (c) matrix appears nearly homogeneous blue with weak clustering, making it difficult to distinguish structurally good and bad samples. In contrast, the Stage 3 (d) matrix exhibits clear block structures that distinctly separate higher quality samples like guidance from lower quality samples like OOD. These observations demonstrate that SSG guidance substantially improves distributional consistency.

4.4 Case Study

We conduct a case study on AD by targeting BACE-1. D^{2}G-TO generates 10,000 molecules per design space from two complementary spaces (BACE-focused and BBBP-focused). After threshold filtering (1,689), ADMET screening (118), and multi-objective Pareto selection [Swanson *et al.*, 2024; Deb *et al.*, 2002], we obtain seven candidates. To assess structural plausibility, we dock these molecules into human BACE-1 receptor (PDB: 7MYI). And we additionally docked the clinical-stage BACE inhibitor LY3202626 as a reference [McKinzie *et al.*, 2021].

Docking results suggest that the seven generated candidates are predicted to adopt plausible binding poses in the BACE-1 pocket, but due to limited manuscript space, we show only four representative binding poses in Figure 6. LY3202626 shows a docking score of $-8.7\,\mathrm{kcal/mol}$, featuring hydrogen bonds with THR231/ASP228/ASP32 and TYR71-centered aromatic interactions, with additional contacts from VAL69/ILE118 and SER35. Molecule (a) adopts a compact pose stabilized by key interactions with TYR71 with additional hydrophobic stabilization from LEU30. Molecule (b) retains the key ASP32-TYR71-ILE118 interaction pattern and appears more favorable from a developability perspective. Molecule (c) shows the best docking score ($-10.5\,\mathrm{kcal/mol}$), featuring an extensive polar contact network with ASP32/ASN37/ASP228 and aromatic anchoring near TYR71. Molecule (d) provides an alternative scaffold, characterized by TYR71-centered aromatic interactions and supporting contacts from TRP76/PHE108. Overall, these candidates span diverse chemotypes with competitive docking scores and better developability, supporting D^{2}G-TO as a computational prioritization pipeline for early-stage virtual screening.

5 Conclusion

In this paper, we propose D^{2}G-TO, a task-aware and OOD-guided discrete diffusion framework for CNS-oriented molecular graph generation. We introduce Structural Similarity Guidance (SSG) module that explicitly steers the generation process toward ID regions and away from OOD structures. By jointly coupling multi-pharmacological properties with distribution-aware structural guidance, D^{2}G-TO enforces a dual task-structure guidance scheme within a unified generative model. The AD case study further shows that D^{2}G-TO can prioritize candidates predicted to satisfy key CNS constraints, supporting its utility for early-stage in silico drug discovery.

Acknowledgments

This work was supported by National Natural Science Foundation of China (22273010).

References

[Berdigaliyev and Aljofan, 2020] Nurken Berdigaliyev and Mohamad Aljofan. An overview of drug discovery and development. *Future Medicinal Chemistry*, 12(10):939–947, 2020.

[Brown *et al.*, 2019] Nathan Brown, Marco Fiscato, Marwin H. S. Segler, and Alain C. Vaucher. Guacamol: Benchmarking models for de novo molecular design. *Journal of Chemical Information and Modeling*, 59(3):1096–1108, 2019.

[Button *et al.*, 2019] Alexander Button, Daniel Merk, Jan A. Hiss, and Gisbert Schneider. Automated de novo molecular design by hybrid machine intelligence and rule-driven chemical synthesis. *Nature Machine Intelligence*, 1(7):307–315, 2019.

[Cheng *et al.*, 2021] Yu Cheng, Yongshun Gong, Yuansheng Liu, Bosheng Song, and Quan Zou. Molecular design in drug discovery: a comprehensive review of deep generative models. *Briefings in Bioinformatics*, 22(6):bbab344, 2021.

[Choung *et al.*, 2023] Oh-Hyeon Choung, Riccardo Vianello, Marwin Segler, Nikolaus Stiefl, and José Jiménez-Luna. Extracting medicinal chemistry intuition via preference machine learning. *Nature Communications*, 14(1):6651, 2023.

[Cummings *et al.*, 2014] Jeffrey L. Cummings, Travis Morstorf, and Kate Zhong. Alzheimer's disease drug-development pipeline: few candidates, frequent failures. *Alzheimer's Research & Therapy*, 6:37, 2014.

[Deb *et al.*, 2002] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.*, 6(2):182–197, 2002.

[Dichiara *et al.*, 2009] Maria Dichiara, Giuseppe Cosentino, Giorgia Giordano, Lorella Pasquinucci, Agostino Marrazzo, Giuliana Costanzo, et al. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of Cheminformatics*, 1(1):8, 2009.

[DiChiara *et al.*, 2024] M. DiChiara, G. Cosentino, G. Giordano, et al. Designing drugs optimized for both blood-brain barrier permeation and intra-cerebral partition. *Expert Opinion on Drug Discovery*, 19(3):317–329, 2024.

[Du *et al.*, 2024] Yuanqi Du, Arian R. Jamasb, Jeff Guo, Tianfan Fu, Charles Harris, Yingheng Wang, et al. Machine learning-aided generative molecular design. *Nature Machine Intelligence*, 6(6):589–604, 2024.

[Dwivedi and Bresson, 2021] Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs. In *AAAI Workshop on Deep Learning on Graphs: Methods and Applications*, 2021.

[Gómez-Bombarelli *et al.*, 2018] Rafael Gómez-Bombarelli, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, et al. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2):268–276, 2018.

[Huang *et al.*, 2021] Kexin Huang, Cao Xiao, Lucas M. Glass, and Jimeng Sun. Moltrans: Molecular interaction transformer for drug–target interaction prediction. *Bioinformatics*, 37(6):830–836, 2021.

[Jensen, 2019] Jan H. Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. *Chemical Science*, 10(12):3567–3572, 2019.

[Jin *et al.*, 2018] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of Machine Learning Research*, pages 2323–2332. PMLR, 2018.

[Jing *et al.*, 2025] Yankang Jing, Guangyi Zhao, Yuanyuan Xu, Terence McGuire, Ganqian Hou, Jack Zhao, et al. Gcn-bbb: Deep learning blood–brain barrier (bbb) permeability pharmacoanalytics with graph convolutional neural (gcn) network. *The AAPS Journal*, 27(3):73, 2025.

[Jo *et al.*, 2022] Jaehyeong Jo, Seul Lee, and Sung Ju Hwang. Score-based generative modeling of graphs via the system of stochastic differential equations. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, Proceedings of Machine Learning Research. PMLR, 2022.

[Lee *et al.*, 2023] Seul Lee, Jaehyeong Jo, and Sung Ju Hwang. Exploring chemical space with score-based out-of-distribution generation. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, volume 202 of *Proceedings of Machine Learning Research*, pages 18872–18892. PMLR, 2023.

[Liu *et al.*, 2023] Chengyi Liu, Wenqi Fan, Yunqing Liu, Jiatong Li, Hang Li, Hui Liu, et al. Generative diffusion models on graphs: Methods and applications. *arXiv preprint*, 2023. arXiv:02591.

[Liu *et al.*, 2024] Gang Liu, Jiaxin Xu, Tengfei Luo, and Meng Jiang. Graph diffusion transformers for multi-conditional molecular generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

[McKinzie *et al.*, 2021] David L. McKinzie, Leonard L. Winneroski, Steven J. Green, Erik J. Hembre, Jon A. Erickson, Brian A. Willis, et al. Discovery and early clinical development of ly3202626, a low-dose, cns-penetrant bace inhibitor. *J. Med. Chem.*, 64(12):8076–8100, 2021.

[Moret *et al.*, 2023] Michael Moret, Irene Pachon Angona, Leandro Cotos, Shen Yan, Kenneth Atz, Cyrill Brunner, et al. Leveraging molecular structure and bioactivity with chemical language models for de novo drug design. *Nature Communications*, 14(1):114, 2023.

[Mullard, 2017] Asher Mullard. The drug-maker's guide to the galaxy. *Nature*, 549(7673):445–447, 2017.

[Niu *et al.*, 2020] Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. Permutation invariant graph generation via score-based generative modeling. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, Proceedings of Machine Learning Research. PMLR, 2020.

[Pardridge, 2005] William M. Pardridge. The blood-brain barrier: bottleneck in brain drug development. *NeuroRx*, 2:3–14, 2005.

[Paul *et al.*, 2010] Steven M. Paul, Daniel S. Mytelka, Christopher T. Dunwiddie, Charles C. Persinger, Bernard H. Munos, Stacy R. Lindborg, et al. How to improve r&d productivity: the pharmaceutical industry's grand challenge. *Nature Reviews Drug Discovery*, 9(3):203–214, 2010.

[Peebles and Xie, 2023] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4195–4205, 2023.

[Ramakrishnan *et al.*, 2014] Raghunathan Ramakrishnan, Pavlo O. Dral, Matthias Rupp, and O. Anatole von Lilienfeld. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1:140022, 2014.

[Ren *et al.*, 2023] Feng Ren, Xiao Ding, Min Zheng, Mikhail Korzinkin, Xin Cai, Wei Zhu, et al. Alphafold accelerates artificial intelligence powered drug discovery: efficient discovery of a novel CDK20 small molecule inhibitor. *Chemical Science*, 14(6):1443–1452, 2023.

[Scannell *et al.*, 2012] Jack W. Scannell, Alex Blanckley, Helen Boldon, and Brian Warrington. Diagnosing the decline in pharmaceutical R&D efficiency. *Nature Reviews Drug Discovery*, 11(3):191–200, 2012.

[Shen *et al.*, 2024] Xu Shen, Yili Wang, Kaixiong Zhou, Shirui Pan, and Xin Wang. Optimizing ood detection in molecular graphs: A novel approach with diffusion models. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2024. Also available as arXiv:2404.15625.

[Stumpfe and Bajorath, 2012] Dagmar Stumpfe and Jürgen Bajorath. Exploring activity cliffs in medicinal chemistry: miniperspective. *Journal of Medicinal Chemistry*, 55(7):2932–2942, 2012.

[Swanson *et al.*, 2024] Kyle Swanson, Parker Walther, Jeremy Leitz, Souhrid Mukherjee, Joseph C. Wu, Rabindra V. Shivnaraine, et al. Admet-ai: a machine learning admet platform for evaluation of large-scale chemical libraries. *Bioinformatics*, 40(7):btae416, 2024.

[Vamathevan *et al.*, 2019] Jessica Vamathevan, Dominic Clark, Paul Czodrowski, Ian Dunham, Edgardo Ferran, George Lee, et al. Applications of machine learning in drug discovery and development. *Nature Reviews Drug Discovery*, 18(6):463–477, 2019.

[Vayer *et al.*, 2019] Titouan Vayer, Nicolas Courty, Romain Tavenard, Laetitia Chapel, and Rémi Flamary. Optimal transport for structured data with application on graphs. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 6275–6284, 2019.

[Vignac *et al.*, 2023] Clément Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. DiGress: Discrete denoising diffusion for graph generation. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.

[Wu *et al.*, 2018] Zhenqin Wu, Bharath Ramsundar, Evan N. Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S. Pappu, et al. Moleculenet: A benchmark for molecular machine learning. *Chemical Science*, 9(2):513–530, 2018.

[Wu *et al.*, 2022] Bingzhe Wu, Jintang Li, Junchi Yu, Yatao Bian, Hengtong Zhang, Chaochao Chen, et al. Trustworthy graph learning: Reliability, explainability, and privacy protection. *arXiv preprint arXiv:2205.10014*, 2022.

[Xie *et al.*, 2021] Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, et al. Mars: Markov molecular sampling for multi-objective drug discovery. In *International Conference on Learning Representations (ICLR)*, 2021. Preprint: arXiv:2103.10432.

[Yang *et al.*, 2023] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, et al. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.

[Yang *et al.*, 2025] Yuwei Yang, Shukai Gu, Bo Liu, Xiaoqing Gong, Ruiqiang Lu, Jiayue Qiu, et al. Diffmc-gen: A dual denoising diffusion model for multi-conditional molecular generation. *Advanced Science*, page 2417726, 2025.

[Zhang *et al.*, 2025] Kang Zhang, Xin Yang, Yifei Wang, Yunfang Yu, Niu Huang, Gen Li, et al. Artificial intelligence in drug development. *Nature Medicine*, 31(1):45–59, 2025.

[Zhavoronkov *et al.*, 2019] Alex Zhavoronkov, Yan A. Ivanenkov, Alex Aliper, Mark S. Veselov, Vladimir A. Aladinskiy, Anastasiya V. Aladinskaya, et al. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology*, 37(9):1038–1040, 2019.