bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

LATENT DIFFUSION FOR CONDITIONAL GENERATION OF MOLECULES

A PREPRINT

Benjamin Kaufman **Edward C. Williams** **Ryan Pederson** **Carl Underkoffler** **Zahid Panjwani**

Miles Wang-Henderson **Narbe Mardirossian** **Matthew H. Katcher** **Zack Strater**

Jean-Marc Grandjean **Bryan Lee** **John Parkhill**

Terray Therapeutics

ABSTRACT

Designing a small molecule therapeutic is a challenging multi-parameter optimization problem. Key properties, such as potency, selectivity, bioavailability, and safety must be jointly optimized to deliver an effective clinical candidate. We present COATI-LDM, a novel application of latent diffusion models to the conditional generation of property-optimized, drug-like small molecules. Diffusive generation of latent molecular encodings, rather than direct diffusive generation of molecular structures, offers an appealing way to handle the small and mismatched datasets that are common for molecular properties. We benchmark various diffusion guidance schemes and sampling methods against a pre-trained autoregressive transformer and genetic algorithms to evaluate control over potency, expert preference, and various physicochemical properties. We show that conditional diffusion allows control over the properties of generated molecules, with practical and performance advantages over competing methods. We also apply the recently introduced idea of particle guidance to enhance sample diversity. We prospectively survey a panel of medicinal chemists and determine that we can conditionally generate molecules that align with their preferences via a learned preference score. Finally, we present a partial diffusion method for the local optimization of molecular properties starting from a seed molecule. Conditional generation of small molecules using latent diffusion models on molecular encodings provides a highly practical and flexible alternative to prior molecular generation schemes.

1 Introduction

Robust generation of molecules subject to design constraints (i.e., inverse molecular design) has great promise to accelerate pharmaceutical and materials research. Models that generate novel chemical structures with desired properties are a subject of furious development, with methodologies such as genetic algorithms applied to molecular graph structures [Kwon et al., 2021], fine-tuning of autoregressive generators via reinforcement learning [Blaschke et al., 2020, You et al., 2019], sampling from generative models such as variational autoencoders or generative adversarial networks [Jin et al., 2018a, Bilodeau et al., 2022], and sampling from multimodal transformer-based models [Kaufman et al., 2024, Reidenbach et al., 2023, Liu et al., 2023].

To provide practical utility for drug discovery, generative design algorithms need to be highly controllable, computationally efficient, and able to leverage datasets of varying sizes to produce optimized molecular designs in chemical space both near hits of interest and distinct from well-explored regions. Drug design is a challenging multi-parameter problem, and some critical properties, like pharmacokinetic features determining bioavailability, are significantly more difficult to gather data for than others. In this work, we propose the use of latent diffusion models (LDMs) [Rombach et al., 2022] and large pre-trained molecular encoder-decoders [Kaufman et al., 2024] for generative molecule design. While diffusive methods have been applied in molecular graph generation [Vignac et al., 2022] and 3D conformer generation [Satorras et al., 2021, Schneuing et al., 2023, Pinheiro et al., 2023, Huang et al., 2023, Xu et al., 2023], we

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

abstract out the problem of encoding and decoding molecular structures to a large, pre-trained model. This work uses a pre-trained transformer-based contrastive (2D/3D) encoder-decoder trained on a large (∼250M) unconditioned sample of chemical space, which has learned a fixed-dimensional molecular embedding. This architecture is an extension of the COATI model as defined in [Kaufman et al., 2024]. We refer to this application of latent diffusion models to COATI embeddings as COATI-LDM.

Many properties must be optimized to produce a clinical molecule, and in many cases, limited data makes it a few-shot generation problem. Assays that assess absorption, distribution, metabolism, excretion, and toxicity (ADMET) are often costly and therefore the datasets are smaller. Public sources such as the ADMET datasets contained in the Therapeutic Data Commons [Huang et al., 2021] or for human feedback on molecule designs [Choung et al., 2023a] only contain a few thousand samples. Analogous to the work in [Sinha et al., 2021], which describes the application of latent diffusion for few-shot image generation, this work applies the same concept to drug-like chemical space, where data limitations are particularly severe and decoders are less standardized. Latent diffusion methods provide a straightforward and computationally efficient path toward applying a variety of physicochemical and activity constraints to generated molecules.

First, we show that latent diffusion methods can effectively capture the embedding distribution of a large dataset of empirical binding affinity measurements against Human Carbonic Anhydrase II (hCAII) [Sly and Hu, 1995]. Utilizing this dataset, we evaluate the ability of several conditional diffusion approaches to match *conditional* property distributions, i.e., the distribution of embedding vectors conditioned on binding affinity or physicochemical descriptors. We additionally evaluate each method's level of control over molecular properties via the correlation between conditioning values and the resulting properties of the generated molecules.

We then perform a set of practical benchmarks using empirical binding affinity data, first demonstrating that conditional latent diffusion models provide better coverage of hit-like chemical space compared to highly competitive evolutionary algorithm-based approaches. Additionally, we show that this method is competitive with transformer-based conditioning, despite the latter being trained on a much larger dataset. Furthermore, in analogy to image generation or large language models (LLMs), these conditional diffusion models can be aligned to human preferences (in this case, expert medicinal chemists) for producing molecules with high desirability by using a small dataset. We validate this alignment by prospectively surveying a small panel of medicinal chemists.

We also provide a straightforward method for conditionally sampling near a seed molecule which we call "partial diffusion": partially noising the seed molecules latent representation and then performing conditional denoising. We show that this method can jointly control molecule similarity, potency, and other properties, and obtain state-of-the-art results on a common constrained molecular optimization benchmark. Finally, we show that particle guidance [Corso et al., 2023a] can be used to generate diverse sets of molecules, raising the possibility of prospectively generating "coverage sets" of chemical space without resorting to clustering on pre-generated chemical sets.

Overall, we find conditional latent diffusion to be a versatile tool for molecular generation under various conditions and disparate datasets. A generative method that separates the data required to condition a molecular property from the data required to learn the structure of admissible chemical space is often vitally important because the data for the former is much scarcer than the data for the latter. Latent diffusion is an amenable approach to this problem, and we present the first application of latent diffusion models to conditional optimization of drug-like molecules using modular conditioning models, trained on separate data sources. We combine a separately-trained invertible continuous embedding of chemical space, derived from [Kaufman et al., 2024], with a straightforward application of conditional latent diffusion methods, and produce a highly performant generative molecular design method.

2 Related Works

Herein we offer a summary of generative approaches to drug-like molecule generation, but refer to reviews on this rapidly evolving topic for further detail [Bilodeau et al., 2022, Du et al., 2022]. The vast majority of data-driven generative molecular models focus on direct modification of a molecular structure representation in a stepwise fashion. Evolutionary algorithms [Kwon et al., 2021], which long predate data-driven machine learning but often outperform newer methods [Jensen, 2019], iteratively mutate a molecular structure using operators on graphs or SMILES strings. These methods require the manual curation of a library of chemically reasonable mutation and crossover operators and need to query (possibly costly) oracle functions – such as a docking program or other predictive model – at every step of generation. Methods based on reinforcement learning suffer from the same requirements, performing approximate gradient steps to fine-tune a generative model [Zhou et al., 2019, Blaschke et al., 2020] using feedback from an oracle function.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

More recent methods for conditional generative design perform local modifications to a molecular structure via a gradient ascent-like process [Fu et al., 2022] or via a diffusion process on graph structures themselves [Liu et al., 2024]. Generative models relying on 3D structural data generate molecules atom-by-atom [Hoogeboom et al., 2022, Schneuing et al., 2023] or by modifying torsion and bond angles [Corso et al., 2023b, Brocidiacono et al., 2023] of a complete molecular structure to produce realistic binding poses, using datasets such as PDBbind [Wang et al., 2004]. Others have used crystal structures of a target protein binding pocket as conditioning information [Schneuing et al., 2023]. In this work, we focus on the diffusive optimization of a continuous latent vector representation derived from a foundation encoding model, which allows the straightforward application of diffusion modeling without many domain-specific changes.

Methods that leverage a latent molecular representation to perform generative optimization have also been explored in the literature. Early attempts at using variational autoencoders (VAEs) to traverse a latent molecular space [Lim et al., 2018, Gómez-Bombarelli et al., 2018] often suffered from low validity (i.e., generating physically implausible molecule structures) and were often trained on small molecular datasets. More recent VAE methods use molecule graph-specific modifications [Jin et al., 2019a] to the encoder and decoder networks to achieve higher validity. Posterior collapse [Wang et al., 2023] remains a challenge with this technique. GeoLDM [Xu et al., 2023] trains an equivariant VAE on small molecule confirmations, and then performs diffusion on geometrically structured latent embeddings with an equivariant score model. However, the geometric structure of the latent space limits the applicability of well-known diffusion guidance approaches and adds additional constraints on the types of allowable guide functions. 3M-Diffusion [Zhu et al., 2024] performs diffusion on the latent embeddings of a VAE trained on molecular graphs, where the latent space has been aligned to textual descriptions of the molecules using contrastive learning. However, the model is designed to condition on text rather than directly on molecular properties.

More recent methods train autoencoding models from commercial catalogs containing hundreds of millions to billions of synthetically available molecules, which are excellent proxies for chemical realism and accessibility. Recurrent neural networks trained to jointly autoencode molecules and properties [Winter et al., 2019, Hoffman et al., 2022a] or multimodal transformer-based models [Liu et al., 2023, Kaufman et al., 2024] produce a latent code representing a molecule structure. Latent optimization algorithms are advantageous as they decouple the problem of representing chemical data from the construction of generative and predictive models.

Statistically significant experimental measures of performance for generative molecular design are costly to produce, which has frustrated fair comparisons between these methods. Simple benchmarks that estimate chemical validity [Brown et al., 2019] or the ability to produce novel molecules are useful toy examples, but are often unrepresentative of realistic drug design processes [Walters, 2023]. In this work, we focus on modeling an experimental dataset of binding affinities (see Section 4), with a large enough size to calculate Frechet distances. We also survey expert medicinal chemists (Section 5.3) to evaluate our ability to align models to their preferences.

3 Methods

Latent diffusion [Rombach et al., 2022] supposes that it is advantageous to generatively model a low-dimensional latent manifold rather than a structured high-dimensional one, given a high-performance pre-trained encoder-decoder. Image diffusion models, for example, perform diffusion in a vector space that can be decoded into an image. In the context of molecule diffusion, we require an invertible encoding for molecular structures, which we describe in Section 3.1. We then pair it with diffusion and flow matching models (Section 3.2) to generate molecule structures with optimized properties. All experiments were trained using a distributed cluster of A100 GPUs on the Nvidia DGX-C compute cloud [Nvidia, online]. Development and training of the encoder-decoder architecture on a 250-million molecule dataset (Section 4) took roughly 20,000 GPU hours, while the top-performing FID diffusion models trained on the same dataset took around 4,000 GPU hours. The associated code repository (https://github.com/terraytherapeutics/COATI-LDM) provides further details.

3.1 Molecular Encoder-Decoder

In this work, we rely on a large-scale pre-trained encoder-decoder that maps chemical space (in either 2D graph structures or 3D point clouds) to fixed-length latent vector representations. We adapt the general framework of COATI[Kaufman et al., 2024], in which contrastive learning is used to align 2D and 3D molecular encodings with decoding capability, and extend it by employing a chirality-aware 3D encoder based on Allegro [Musaelian et al., 2023], using DirectCLR [Jing et al., 2022] to control contrastive mode collapse, increasing the latent vector size to 512, and expanding the training dataset. More details on our encoder-decoder are available in Appendix B. We note that the general framework of molecular latent diffusion would function with any sufficiently expressive decodeable molecular representation.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

The molecular decoder we employ can complete prompts that contain a property condition, such as `[PROP] [LOGP] [3] [SMILES] ...`, which would generate a molecule likely to have a logP of 3. This is another possible modality for conditioned generation of molecules, which we compare against in this work. See Table 6 in Appendix B for a complete list of such available properties. We find in Section 5.4 that the transformer's property control is quite good; however, applying such a model to new properties requires a costly and laborious fine-tuning and vocabulary injection step that is not required in the latent diffusion framework.

3.2 Latent Diffusion

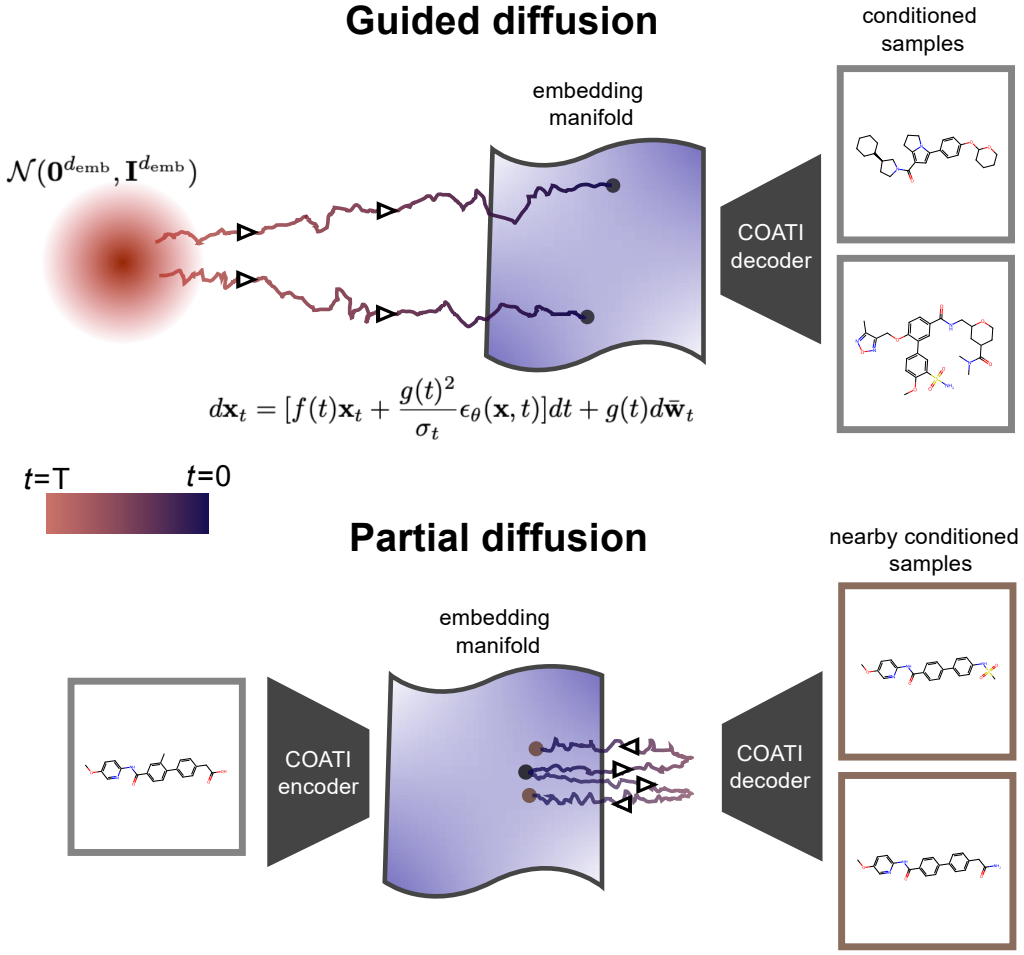

Figure 1: Latent diffusion for molecule generation allows models trained on scarce or non-overlapping datasets to condition generations on a large data manifold which represents a larger and more flexible chemical space. *Partial* diffusion allows one to start with a given molecule and perform a partial diffusion propagation ($t < T$) to obtain conditioned samples that are nearby in chemical space.

Latent diffusion [Rombach et al., 2022] offers a solution to the problem of generating samples in high-dimensional spaces by learning mappings between lower-dimensional embeddings and Gaussian noise. In this work, we focus on 512-dimensional molecular embeddings $\{\mathbf{v}_i \in \mathbb{R}^{512}\}$ and normally distributed noise $\mathcal{N}(\mathbf{0}^{512}, \mathbf{I}^{512})$. This converts the high dimension chemical generation problem into a mathematically much tamer problem of learning a *score network*, $\epsilon_\theta(\mathbf{v}, t, \mathbf{c}) \approx -\sigma_t \nabla_{\mathbf{v}_t} \log p(\mathbf{v}_t|\mathbf{c})$, which separates Gaussian noise ϵ, from a samples ($\mathbf{v}_i$) along a diffusion trajectory (possibly conditioned by conditions, $\mathbf{c}$) which maps data samples at time $t = 0$, onto normal uncorrelated samples at

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

$t = T$. [Ho et al., 2020] showed such processes can be learned from a straightforward denoising loss,

$$L_{\text{DDPM}} = ||\epsilon - \epsilon_\theta(\mathbf{v}, t, \mathbf{c})||^2 . \tag{1}$$

Our results are based on the mathematically convenient formalisms of denoising diffusion probabilistic models (DDPM) and denoising diffusion implicit models (DDIM) [Ho et al., 2020, Song et al., 2022] which are trained with the same loss function.

Flow-Matching (FM) models [Lipman et al., 2023, Ben-Hamu et al., 2024a] generalize the diffusion generative process to a more flexible family of integrated vector fields with adjustable priors and have shown superior performance to other diffusion methods in certain settings [Lipman et al., 2023, Stärk et al., 2023]. We will show some results from optimal transport conditional flow matching, which builds a linear inductive bias into the dynamics mapping data and prior distributions that could lead to more stable training. Besides minor changes to the training process these flow methods can be used to train joint-conditioned generators in exactly the same fashion as DDPM/DDIM, using the same score models as vector fields, and yielding similar performance.

The generative process can be conditioned to produce samples with a desired property in several ways. These techniques were originally developed for use with image classifiers, but are trivially compatible with arbitrary regression objectives. The simplest scheme, joint conditioning, appends variables of interest onto score network inputs during training and sampling. Classifier guidance (CG) [Dhariwal and Nichol, 2021], perturbs the trajectories with a force derived from some desired conditioning classification or regression objective: $-\nabla_{\mathbf{v}_t} \log p(c|\mathbf{v}_t)$. Better results are obtained by training the regression model or classifier with noised samples and joint knowledge of the point in the noise schedule (t) from which a sample is drawn. Throughout this work we use DUE regressors [van Amersfoort et al., 2022] with the same dimension as our latent space whenever we are discussing CG. These are trained using noised samples drawn from the diffusion schedule. The guidance force is: $-\nabla_{v_t}(\langle \text{DUE}(\mathbf{v}_t, \mathbf{t})\rangle - c)^2$. Another possible scheme is classifier-free guidance (CFG) [Ho and Salimans, 2022], which uses the difference between a joint-conditioned and unconditioned score model to guide the diffusion trajectories. D-Flow [Ben-Hamu et al., 2024b] utilizes the deterministic mapping flow matching learns between prior and samples to optimize the condition of a differentiable sample with respect to the initial point on the prior. This makes D-Flow good at achieving strong control while remaining on the data manifold, but the use of a Quasi-Newton solver renders the process much more costly and less predictable because local minima can be found. We experiment with all of these methodologies in the context of molecular latent diffusion, and demonstrate the relative strengths and weaknesses of each method.

Often within a real molecular design project, it is important to generate 'near' a previously known molecule due to accessibility of starting materials and increased confidence in structure-activity relationship (SAR) around a known hit. This is another source of strength for latent diffusion over transformer models. Integrating the diffusive generation process forwards and backwards allows a user to explore nearby chemical space while enforcing desired conditions (Fig. 1) and sampling properly on the manifold. Rather than simply adding noise to a latent embedding and decoding, this method traverses the learned manifold of valid molecular structures to find close analogs.

Score model network architecture has a strong effect on diffusion model performance [Heusel et al., 2017a]. The latent vectors of our encoder/decoder have no translation invariance, and so the U-Net architecture [Ronneberger et al., 2015] usually applied to diffusive image generation isn't a natural choice. However other components of the architecture are well tuned for diffusion tasks. Given this, we chose to use an architecture in the style of U-Net without convolutions, which is described in Appendix C. The optimization of a score model architecture for this setup is an area for future work.

3.3 Preference Alignment

The problem of "aligning" a generative model to human preferences has been widely studied in the context of Large Language Models (LLMs) [Ouyang et al., 2022, Sun et al., 2023, Rafailov et al., 2023] and image generation methods [Wallace et al., 2023, Zhang et al., 2024, Lin et al., 2023]. In the context of small molecular generative models, alignment entails producing molecules that are desirable from a property perspective (i.e., potency) as well as "drug-like", indicating a high likelihood of oral bioavailability, as well as synthetic accessibility. Medicinal chemists gain a strong intuition for the interplay between these properties. This 'chemical intuition' [Kutchukian et al., 2012, Pedreira et al., 2019] allows a medicinal chemist to estimate whether or not a molecule will be worth progressing forward in a program. Historical work into cheminformatics has attempted to quantify this intuition into a series of cutoffs, such as Lipinski's Rule of 5 or the Quantiative Estimate of Drug-Likeness (QED) [Lipinski, 2004, Bickerton et al., 2012]. More recent work models 'chemical intuition' as a preference learning problem, and trains score models on datasets of human-labeled molecule preferences [Takaoka et al., 2003, Sheridan et al., 2014, Choung et al., 2023b].

To align generations to medicinal chemists' preferences, we adopt the method of [Choung et al., 2023b], which learns a score for a dataset of molecular pairs, where one molecule was picked as "preferable" to the other. Such a model is

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

compact, simple to train, and, most importantly, *improvable* – the scores can easily be fine-tuned by collecting additional preference pairs. This model assumes that pairwise rankings between molecules m_i and m_j is explained by a score function $s(m)$ that provides an ordering of molecules [Burges et al., 2005]. The learning objective, then, is simply to maximize the likelihood that the chosen molecule has a higher score via a (regularized) binary cross-entropy objective. We fit this model using the latent molecular encoding described in Section 3.1. We then incorporate this information into diffusive generation via simple conditional generation – fitting a smaller conditional diffusion model and using classifier-free guidance, or using a preference score model to condition directly using classifier guidance.

4 Datasets

We utilize an ensemble of datasets to probe various properties of latent molecular diffusion models. Outside of Section 5.1, for classifier and classifier-free guidance experiments we use unconditional latent generators trained using an ensemble dataset of over 250 million molecules gathered from publicly available sources of characterized and/or purchasable compounds. This dataset primarily consists of combinatorially enumerated molecule structures using building blocks from Enamine [Grygorenko et al., 2020] as well as publicly available compounds from sources such as ChEMBL [Gaulton et al., 2012], with the objective of providing a close-to-comprehensive survey of practically synthetically accessible chemical space. This dataset was also used to train the encoder-decoder model described above; see Appendix B for more details on its construction.

To examine the ability of latent diffusion to satisfy realistic chemical objectives we use 500K samples of experimental binding affinity data of molecules to a target, hCAII [Sly and Hu, 1995], gathered from a high-throughput binding affinity assay. We also calculate physicochemical properties for these molecules using the RDKit [RDKit, online] to provide examples of conditioning on molecular properties. This dataset is primarily utilized in Section 5.1.

The preference alignment experiments in Section 5.3 use a small dataset (5K molecule pairs, 10K unique molecules) of human-labeled data. In that work, a rating panel of experienced medicinal chemists was assembled and asked to rank pairs of molecule structures based on their desirability. This dataset was first used to train a latent ranking model using the procedure described in [Choung et al., 2023b], which was then used to label data for conditional diffusion model training. Internally, we formed a panel of 3 practicing medicinal chemistry experts to validate that model, finding good agreement. This panel also produced additional preference data and was used to assess the success of conditional generation methods.

See Appendix A for links to these datasets.

5 Results

5.1 Sampling and Conditional Generation

Generative models are designed to provide an accurate model of their training data distribution. However, in the general case, defining the training data distribution (especially in high-dimensional or discrete cases) is difficult. In the case of generative molecular models, we want to evaluate how well the set of generated molecules represents the distribution of molecules used to train it. A molecular structure, in this setting, is a discrete graph object, computing a parametric distribution of which is difficult. Image generation models are typically evaluated via a comparison between the distributions of low-dimensional embedded samples and training or validation data embeddings. These metrics, such as Frechet Inception Distance (FID) and CLIP Maximum Mean Discrepancy (CMMD) [Heusel et al., 2017a, Jayasumana et al., 2023], compute statistics of lower-dimensional vectors and compare to a reference set. We adopt an analogous approach for evaluating latent molecular diffusion - evaluating coverage using the Frechet distance (FD) of the molecular embeddings generated by the encoder-decoder model described in Section 3.1. While forms of FD are widely used for measuring how generative models fit the distribution of their data, it is typically used in a comparative fashion, as a "good" or "bad" value can be embedding dependent. To contextualise, the current SOTA unconditional FD for image diffusion on the CIFAR-10 dataset (measured using Inception distance) is 1.78 [Kim et al., 2023], but has been progressively decreased from a value of 24.8 from its introduction in [Heusel et al., 2017b]. We use these metrics to compare the relative performance of 3 different diffusion-based conditional generation methods – conditional diffusion, classifier-free guidance, and classifier guidance [Ho and Salimans, 2022, Dhariwal and Nichol, 2021] and additionally evaluate a conditional flow matching model and another flow based method D-Flow [Ben-Hamu et al., 2024a]. We report these distances on the a held-out partition of the hCAII dataset (described in Section 4) in Figure 2. To evaluate the conditional FD, samples are drawn from the model matching the conditions of the samples in the holdout partition.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

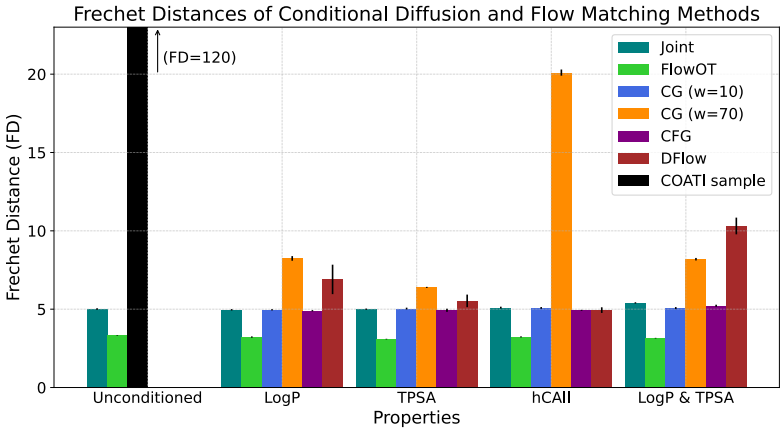

Figure 2: Average FD of different diffusion and flow approaches. Frechet Distances are measured against a 20,000 sample holdout set. The means and standard deviations are computed over five runs of 20,000 samples for each entry, with the exception of the direct sample from COATI which was just run once. w refers to the classifier guidance weight used during generation. For each set of the conditions, the conditions of the samples are matched to those in the holdout set. Optimal Transport Flow Matching outperforms all diffusion-based methods, indicating that the learned vector fields induce a better approximation of the training distribution for this dataset.

Separately, we evaluate conditional models' abilities to produce molecules with desired conditions. We produce a uniform sample of 4000 possible condition values (using the empirical conditional distribution function of the training set values), sample corresponding molecules, and compute Pearson correlations with the desired property. For every condition but hCAII potency, the molecules are decoded into SMILES strings and the properties are computed directly using the RDKit [RDKit, online]. To evaluate hCAII potency conditioning, we predict the value with an MLP trained on the entire hCAII dataset.

We report Pearson correlations between target conditions and the measured properties of the corresponding generations in Figure 3. This and 2 give insights into the generative perfomance of these approaches. We observe that the FDs are fairly low, indicating that they effectively capture the data distribution on this validation set. The three diffusion methods have roughly equivalent FDs across all properties, and the flow matching model performs better across the board. However, we find that better FDs do not lead to better property obedience, with the joint and CFG models having comparable or better correlations on all properties. Notably, CG does not perform as well on property obedience at this weight ($w = 10$), but improved at higher weights. We anticipate that future work will improve upon these metrics, which serve as a baseline for additional developments in generative molecule design on practical datasets.

To further explore this we investigate the impact that changing the weight parameters of the guidance methods (CFG and CG) has on FD and condition obedience. While there is no obvious trend when varying classifier-free guidance weights (see Appendix Figure 12) there is a clear trade-off between FD and MAE for classifier guidance, as shown in Figure 4. This type of trade-off between sample fidelity and condition is well suited to a user-adjustable strength of guidance.

5.2 Comparison to Genetic Algorithms

Genetic algorithms are a strong baseline for molecular generation methods, often outperforming data-driven generative methods on public benchmark tasks [Tripp and Hernández-Lobato, 2023]. Instead of directly modeling the manifold of chemical space, genetic algorithms build up a library of transformations assumed to span molecular space from fragments, and perform a population-based random search in the space of discrete chemical structures to maximize an objective function.

We construct a genetic generative baseline using a pool of fragments furnished by making successive cuts to molecules in ChEMBL 32 [Mendez et al., 2018] and collecting fragments with ≤ 12 heavy atoms. During each generation, molecules are produced by joining two or more fragments along the scissile atom positions, as determined by a SMARTS match, and the top-scoring percentile is disassembled into an enriched fragment pool for the next generation.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

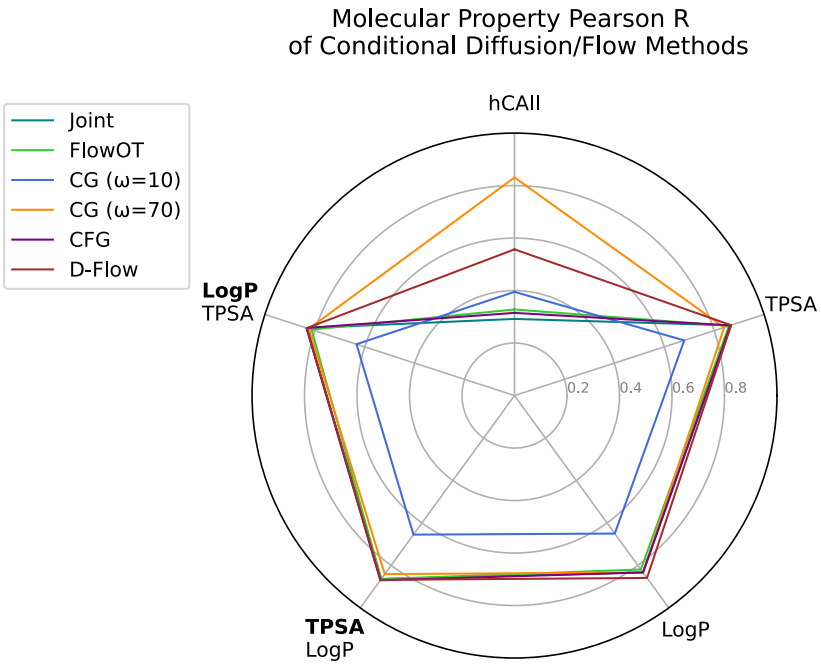

Figure 3: Visualization of Molecule Property/Target Condition Correlations. Values are computed against 4000 condition values that were uniformly sampled over the CDF of the properties, with the exception of the multi-property case, where the samples are treated as independent. For the multi-property case the property the metric is measured on is bolded. We observe that conditioning on more than one property does not degrade the level of control on individual properties. There is limited discrepancy on almost all properties outside of low weight CG. The one exception is hCAII binding, where the classifier guided approaches outperform.

We perform a head-to-head comparison between this fragment-based evolutionary approach and a diffusive generative model to elucidate the strengths and weaknesses of each generative method. To better evaluate the methods in a hit-finding context, generations were assessed by their coverage of a set of held-out "binding hits" from the hCAII data. This set was generated by sorting the top 10% of binding values and putting every other molecule into the holdout set. Additionally, the models were evaluated on 544 public hCAII active compounds (pIC50 $\geq$ 4) that were both not in training data and not constrained by the building block chemistry of the training data. Coverage was determined by the proportion of each evaluation set which passed a threshold of Tanimoto similarity [Bajusz et al., 2015] on Morgan ECFP6-2048 fingerprints [Morgan, 1965] to fifty thousand generated molecules for each approach. For example, a coverage of 15% means that 15% of the evaluation set had close analogs produced by the generative method.

We benchmark against three genetic algorithm variants – the first uses only the predicted hCAII binding affinity as the objective function, while the second utilizes a quantitative estimate of drug-likeness (QED) [Bickerton et al., 2012] cutoff (QED $\geq$0.5) as an extra condition when selecting parents, and the third equally weights predicted hCAII binding affinity and QED in its objective function. After 10 generations where the top 10% of molecules are the parents to further generations, a final selection is made from the top 50K molecules of any generation.

To compare against the genetic methods, which are not constrained to sample within the data distribution, we trained an unconditioned model on the large ensemble dataset (Section 4). We used the unconditional model within the CG and CFG setups, training a hCAII regressor for CG and a hCAII conditional diffuser for CFG. For comparison against the genetic algorithm's use of QED, one CG variation included a conditional QED diffuser with equal weight to the hCAII diffuser. We note that the "public hits" are in ChEMBL – i.e., 500 out of 1M sources of fragments used to produce the genetic fragment set - which could be seen as an unfair advantage for the genetic approach. However, we believe this comparison is still the most interesting, most closely reflecting industrial practice.

In the domain closest to the hCAII training data - the held out Terray binding affinity data, CFG outperformed CG and genetic approaches (Table 1) by generating more somewhat-similar (Tanimoto $\geq$0.4) and similar (Tanimoto $\geq$0.7) compounds. We hypothesize that the conditional diffuser utilized in CFG directly models the data distribution rather

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

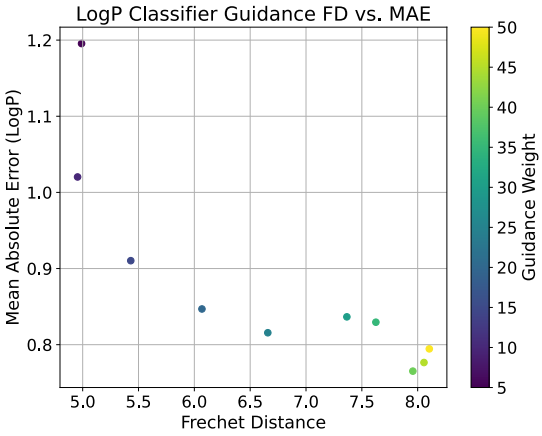

Figure 4: Mean Absolute Error and Frechet Distances of different guidance weights when conditioned on LogP. Increasing guidance weight improves conditioning MAE (i.e., the resulting generations better obey the conditions) but have a higher FD, indicating that the samples are dissimilar to their training set distribution. Frechet Distances are computed with 20,000 samples against the holdout set, while the MAEs are computed on a sample of 4000 conditions uniformly sampled from the LogP CDF of the hCAII training data.

Partition	**Tanimoto Cutoff**	**hCAII Train**	**CFG**	**CG**	**CG w/ Joint QED**	**Genetic**	**Genetic w/ QED Cutoff**	**Genetic w/ Joint QED**
Binding Hits	0.5	99.6	**42.7**	5.6	10.9	0.9	1.4	1.5
	0.6	75.7	**15.3**	0.9	2.1	0.1	0.2	0.2
	0.7	28.6	**4.3**	0.1	0.3	0	0	0
Public Hits	0.5	0.6	1.1	6.1	**15.1**	7.2	7.2	5.0
	0.6	0	0	2.4	**5.7**	1.8	2.6	1.3
	0.7	0	0	0.6	**2.0**	0	0.2	0

Table 1: Percent of known-potent held-out molecules with neighbors to 50K generations from diffusion methods and genetic algorithms. We find that the classifier-free guidance-conditioned diffuser is best able to recover hits in the hCAII data, which is unsurprising given that its conditioning model was trained on similar data. Classifier guidance, on the other hand, is best able to recover molecules on a public dataset.

than simply providing a gradient signal to guide a diffusive generator. This provides an interesting companion to the results observed in Section 5.1, where there is a clear trade-off between distribution matching and conditioning for CG but no such trade-off for CFG.

For the small set of non-Terray hCAII hits, CG generations showed better coverage than CFG and genetic approaches for "somewhat-similar" and "similar" Tanimoto cutoffs. Molecules from the public hCAII hits are not confined to building blocks and chemistries of the binding affinity assay. Besides the fact that they are known potent binders of hCAII, there is no reason to expect that 50K samples of a generative model would cover the space of these molecules. Even so, we maintain that coverage of known binders is a useful and quantifiable measure of generative generalization. We observe that the the coverage gap between genetic generation and diffusion narrows on these molecules far outside the classifier's dataset. Diffusion-based methods still outperformed, with substantially less wall-time.

The stark performance difference between CFG, genetic and CG on the "in-sample" but held out HCAII data vs. public test samples highlights how difficult it is to quantitatively compare molecular generation approaches because statistics do not converge well even with larger-than-average datasets. In practical applications training data is the limiting factor, and methods like CG and genetic algorithms which are less expressive, and therefore less overfit, easily outperform.

5.3 Medicinal Chemistry Preference Alignment

Table 2 shows correlations between conditioned preference scores and inferred preference values using conditional diffusion models trained on a RankNet-labeled hCAII dataset. To produce these labels, we used the trained preference model outlined in Section 3.3 to label the hCAII dataset described in Section 4, and trained a conditioned diffusion

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

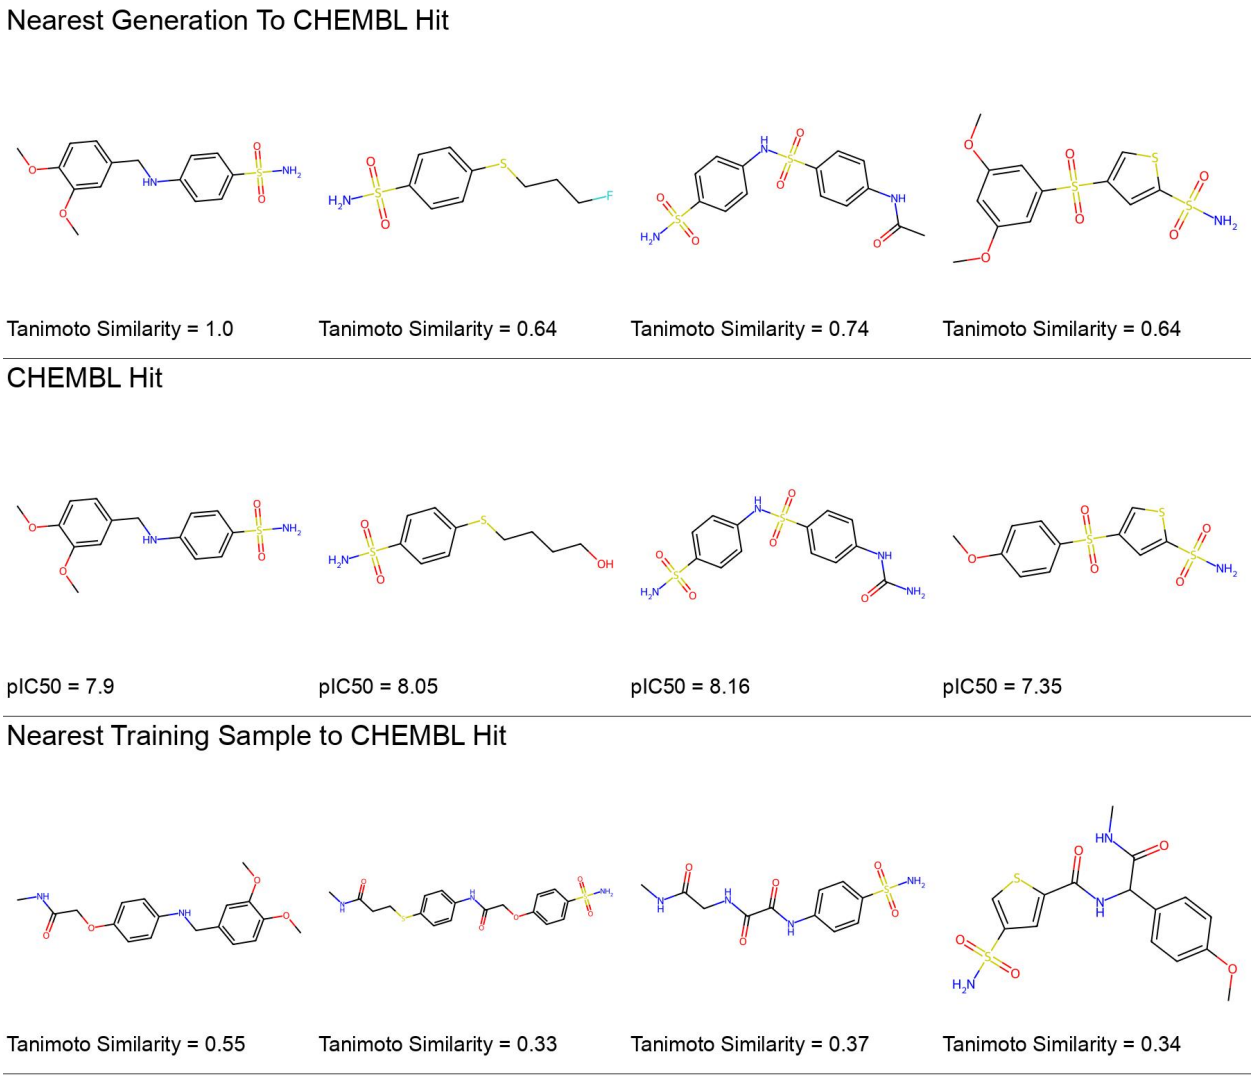

Figure 5: Notable hCAII Generations from the CG with QED model during the genetic comparison experiment. The most similar (ECFP6 Tanimoto) molecule from the public hits dataset is shown for each generation, as well as the most similar molecule from the regressor training data to the public hit.

model on the scores. These were then used to perform classifier-free guidance as in Section 3.2. We observe decent control over preference, indicating that the learned ranking scores are easily incorporated into a diffusion model. We also find that conditioning ability is not significantly degraded when conditioning on preference in addition to another property of interest.

We performed a qualitative analysis of preference conditioning - sampling molecules with high and low desirability as well as jointly conditioning on hCAII potency. Figure 6 contains a t-SNE plot demonstrating that the conditoning produces dissimilar chemical spaces. We see, primarily, that low-desirability molecules appear very large and synthetically complex, with unstable substructures.

To produce a statistical assessment of our ability to condition on preference, we applied a simplified version of the survey methodology found in [Choung et al., 2023b] to gather data from a group of 3 medicinal chemists. A Streamlit [Snowflake, 2024] app (see Figure 14 for an interface example) presented chemists with a random pair of molecules, who were prompted to select which molecule would be the most desirable. We first performed a survey to confirm that the preference model learned from public data matched the preferences of our medicinal chemistry team. We created a composite dataset consisting of samples from ChEMBL, randomly enumerated combinatorial molecules, and random decodes created by sampling from the molecular decoder described in Section 3.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

Conditioning Method	Preference r	hCAII & Preference r
Joint	0.91	0.32 & 0.91
Classifier-Free Guidance	0.92	0.37 & 0.92
Classifier Guidance	0.80	0.25 & 0.83

Table 2: Conditional control over the learned molecule preference score described in Section 3.3. 2048 samples from a uniform distribution were used to condition the diffusion methods, and the resulting vector was used to infer a learned preference score. We see that classifier guidance and direct joint conditioning have better control over preference, and, like with our other properties, conditioning on preference does not degrade conditioning on another property.

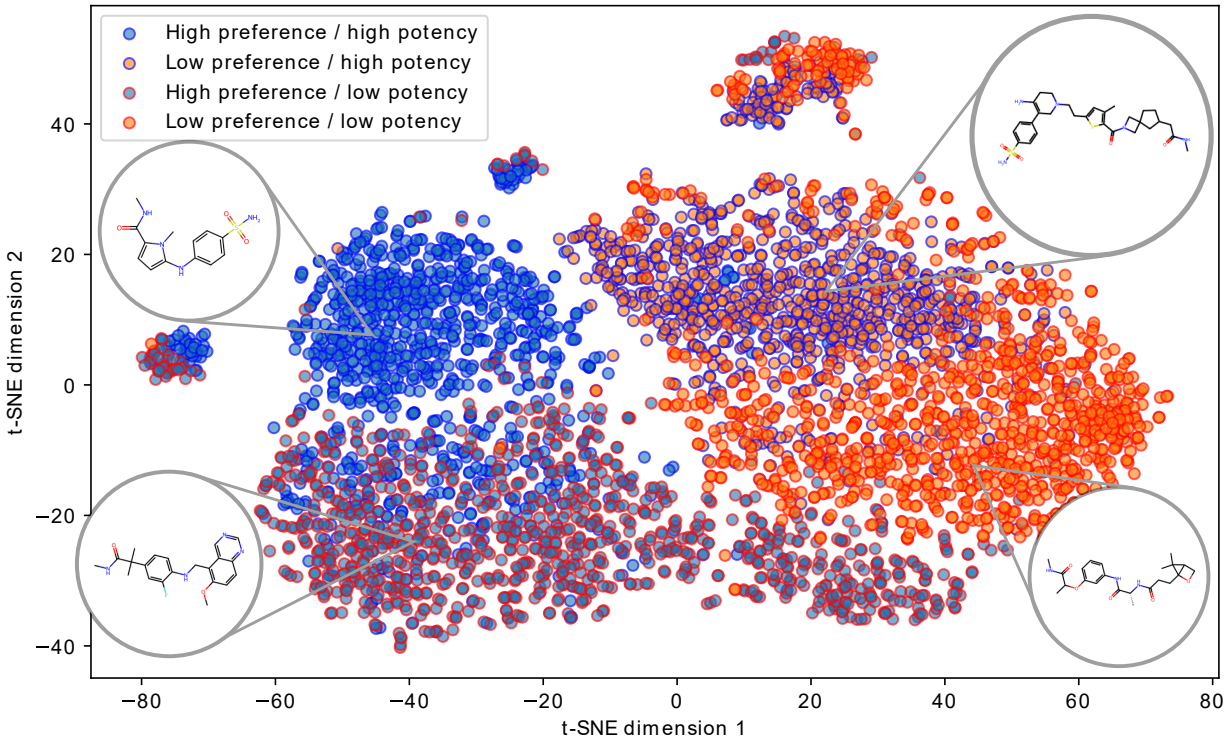

Figure 6: A t-SNE plot of 4096 samples from a CFG process, conditioned jointly on high and low preference as well as high and low hCAII potency. We see that the four conditions produce a fairly disjoint chemical space, with low preference samples appearing larger and more lipophilic, and high potency samples possessing chemical motifs that are present in known binders.

After we had confirmed that a preference model trained on public data agreed with our teams' preference (see Figure 13), we designed a subsequent experiment to prospectively test generative preference alignment. Using the low-data conditioned diffusion model described in Section 5.3, we used classifier-free guidance to generate two sets of molecules: (1) a set of molecules with high preference (preference scores 2 standard deviations above the mean) and (2) low preference (preference score between 1 and 2 standard deviations below the mean). We then presented our survey participants with 150 (randomized) pairs using the same methodology as our first-round experiment, without revealing that each pair contained one low-preference and one high-preference molecule. We noticed qualitatively that conditioning for extremely low preference (i.e., producing molecules with the minimum possible preference values) was too obvious – molecules tended to be either macrocylic or very large and lipophilic. To get a fairer survey of preference alignment, we weakened the low-preference condition to better represent drug-like chemical space.

As a simple evaluation of whether or not we can practically enrich generated molecular sets for medicinal chemist-preferred molecules, we then computed the fraction of reported pairs where the high-preference-score molecule was preferred. Across all observations, the high-preference generation was preferred over the low-preference generation in 76.8% of responses, without any fine-tuning using data from our chemistry team. We found that pairs where the

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

Method	Training Set Size	TPSA	fsp3	TPSA & fsp3
Autoregressive generation	263M	22.28	0.0592	26.32 & 0.0661
Joint	0.48M	18.62	0.0840	18.42 & 0.0879
Classifier-Free Guidance	0.48M	18.72	0.0860	18.33 & 0.0864
Classifier Guidance	0.48M	25.69	0.0815	24.11 & 0.0806

Table 3: RMSE of conditioned properties in latent diffusion models vs. the parent transformer. We see that diffusion models perform comparably to prompted autoregressive generation with a fraction of the training data.

chemists disagreed with the conditioning tended to have strange or unstable "high-preference" molecules (See Appendix D, indicating that the preference model can benefit from additional feedback data.

An off-the-shelf preference dataset, alongside this conditioning framework, allows us to produce molecules that are more in line with medicinal chemists' preference with a minimal amount of tuning. In future work, we anticipate incorporating active feedback from medicinal chemists to improve the quality of generated molecular structures.

5.4 Comparison to Transformer Conditioning

The frozen molecular encoder-decoder used in our latent diffusion scheme was trained to jointly encode and decode several molecular properties (the full list of properties can be found in Appendix B) and can be used as a tool for conditioned generation. Latent diffusion is still extremely practical, because of its low data demands, more efficient training, and incremental improvability. Still, the generation quality between these methods merits comparison.

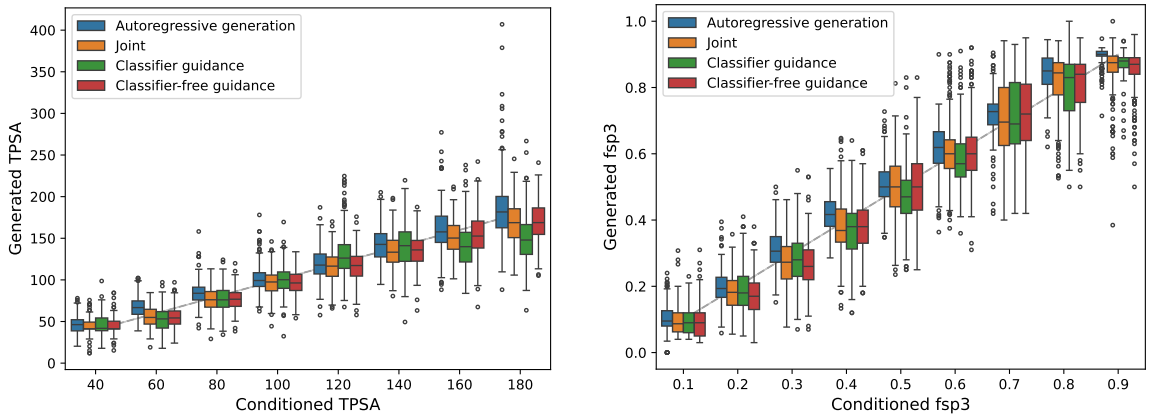

Figure 7: Box and whisker plots (with outliers marked as circles) of generated molecule properties versus the conditioned value for fsp3 and TPSA. Each distribution is over 256 condition-generated samples. The dashed gray line represents a perfect agreement. RMSEs are given in Table 3. Both methods perform well here, although the difficulty of retraining a large autoregressive model on new properties obviously precludes its utility for arbitrary generative tasks.

In Figure 7, we compare the distributions of property-conditioned generations for the following properties: topological polar surface area (TPSA) and the fraction of C atoms that are sp3 hybridized (fsp3). The values at which each property was conditioned are given in the x-axis and we generate 256 samples with each method at each value. In Table 3, we report the overall root mean squared error (RMSE) across these specific property conditions, where the multi-property conditioning (TPSA & fsp3) considers all such property value combinations and the RMSE is evaluated separately for each property. The total training set size is also given for context, where the COATI autoregressive generation model was trained on over 500x more molecules than the diffusion models. Despite the much more limited chemical training space, we find that diffusion-based models can sometimes outperform the corresponding autoregressive generator on these chemical properties. Additionally, for the COATI autoregressive generator, we find that performance degrades for the multi-property conditioned generations, whereas in the diffusion models, the performance is better maintained.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

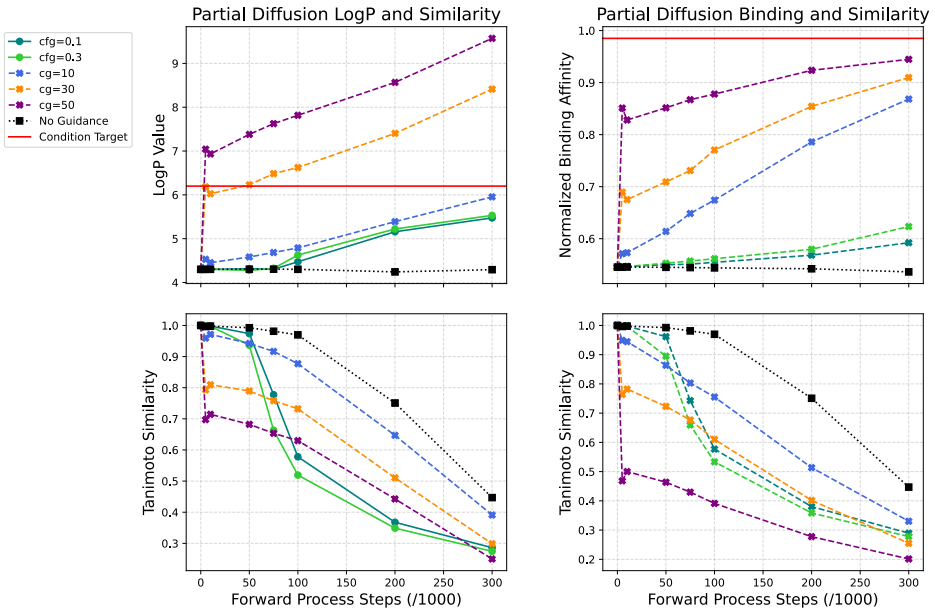

Figure 8: One use case of partial diffusion is to control the similarity of generations relative to a starting point, while also optimizing a desired property. We show this for the maximization of LogP and hCAII binding affinity. These plots show the average Tanimoto Similarity of 1000 diffusion trajectories to their randomly sampled starting molecules and the average property values of the molecules generated by the trajectories. Increasing the number of reverse diffusion steps reduces the similarity to the starting molecule while increasing the conditioned value. Lower weightings of classifier guidance offer a reasonable trade-off of property optimization and similarity to starting point. Classifier free guidance is less effective in this framework.

5.5 Nearby Sampling

Small molecule drug design is typically pursued as an iterative cycle of small modifications to an existing "lead" compound. The data domain of predictive molecular models is highly dependent on the chemical space used to fit them, so generating new, in-domain molecules with property improvements is highly relevant to progressing a drug development program. Typically, these modifications are made in attempts to modify suboptimal properties of the lead compound - e.g., improving activity or solubility, or removing toxicophores without reducing activity.

We approach this task as a partial diffusion process – adding limited noise to a starting point using the diffusive forward process, and then using the conditional score model to denoise back into the structured latent space. The number of noise steps, as shown in Figure 8, allows control over the resulting molecules' similarity to the start point. Classifier-free guidance performed poorly on LogP optimization, failing to effectively improve LogP while maintaining similarity to the starting point. Classifier guidance better achieved this desired behavior. hCAII binding affinity follows a similar pattern. An example of varying the number of forward diffusion steps can be found in Figure 9. This toy example shows the essential trade-off between adding noise and modifying properties - adding additional noise allows for optimization of the property in question (in this case, lipophilicity), but traverses further away from the seed molecule.

This setup was evaluated on a generative benchmark from [Jin et al., 2019b] where, given a set of starting molecules with QED $\in [0.6, 0.8]$, the goal is to generate a molecule with QED ≥ 0.9 with a Morgan fingerprint Tanimoto similarity to the starting molecule ≥ 0.4. Our methodology for the benchmark is detailed in appendix E. As shown in Table 4, latent molecule partial diffusion with COATI-LDM achieves state-of-the-art performance. Notably, this is achieved without directing the process with Tanimoto similarity at all, which differs from MolMIM and QMO which both incorporate it into gradient-free optimization routines, which is a handicap for the diffusion approach. In partial latent diffusion, the relative distance on the latent manifold resembles a Tanimoto similarity, and restricting the number of noise time steps will naturally tend to limit the corresponding Tanimoto similarity deviation of resulting generated molecules.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

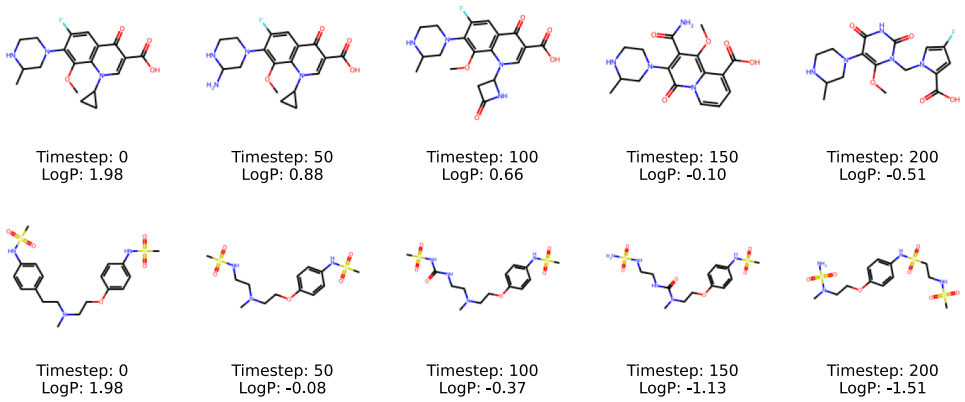

Figure 9: Varying the number of reverse diffusion steps and performing reverse diffusion with classifier guidance allows control over how similar output molecules are to a starting point, trading off property optimization with similarity. Here are examples of molecules generated by partial diffusion on two FDA approved drugs gatifloxacin and dofetilide. Classifier guidance was used to reduce LogP.

Method	Success (%)
JT-VAE[Jin et al., 2018a]	8.8
GCPN[You et al., 2018]	9.4
MMPA[Dalke et al., 2018]	32.9
VSeq2Seq[Bahdanau et al., 2015]	58.5
VJTN+GAN[Jin et al., 2018b]	60.6
AtomG2G[Jin et al., 2020]	73.6
HierG2G[Jin et al., 2020]	76.9
DESMILES[Maragakis et al., 2020]	77.8
QMO[Hoffman et al., 2022b]	92.8
MolMIM[Reidenbach et al., 2023]	94.6
COATI-LDM	**95.6**

Table 4: Success percentages for the conditional nearby generation task defined in [Jin et al., 2019b]. Given a set of starting molecules with QED $\in [0.6, 0.8]$, the objective is to generate a molecule with QED ≥ 0.9 with a Morgan fingerprint Tanimoto similarity to the starting molecule ≥ 0.4. Partial diffusion using COATI vectors achieves state-of-the-art performance.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

5.6 Diverse Generations via Particle Guidance

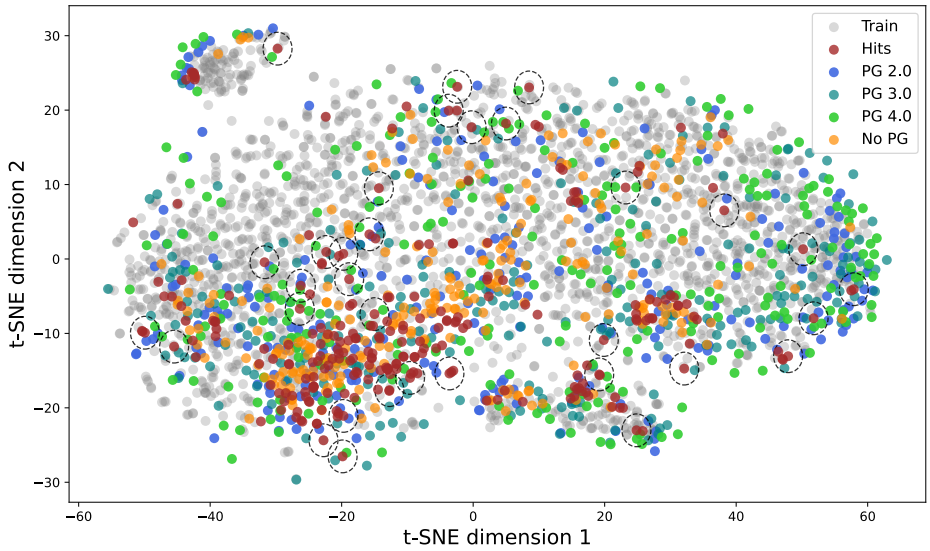

Figure 10: A t-SNE of dataset samples, hits, samples conditioned for potency, and samples conditioned for potency with particle guidance. Particle guidance improves sampling of the smaller hit clusters, without sacrificing focus on potency which is a useful generalization of the classic and often invoked sphere exclusion clustering technique.

Encouraging the generation of *diverse* molecular sets is an important task in molecular design, in particular during the early stages of drug discovery when lead series have not been solidified. Users may want to avoid the intellectual property space of a known drug, or rapidly assess SAR around a scaffold with as few measurements as possible. Practitioners of cheminformatics have used a method called Directed Sphere Exclusion [Gobbi and Lee, 2003] for decades to diversity filter ensembles of molecules. This method defines a sphere around a seed compound (typically via a Tanimoto distance radius), and assigns all molecules within that sphere to that compound's cluster. This procedure occurs iteratively until all molecules are clustered, producing a set with enhanced pairwise Tanimoto distances.

Latent diffusion also provides a convenient mechanism to build diversity into generated molecular samples. We employ the concept of particle guidance [Corso et al., 2023a] which simply adds a repulsive guidance force between samples at each timestep while sampling. Figure 10 is a t-SNE plot of 256 generations at varying levels of particle guidance sampled from a jointly conditioned DDPM model for hCAII potency and LogP values between 1.3 and 1.8. The plot depicts training samples, generated samples, and hits from the hCAII interleaved holdout set. Several known potent held-out molecules, especially those that are "isolated" from other hits, are only covered nearby with a sample when particle guidance is applied. These are shown in the figure with the t-SNE plotted points encompassed by dashed circles.

6 Conclusions

In this work, we present COATI-LDM, a novel latent diffusion framework for small molecule design and optimization. We show that latent diffusion can be applied directly and conveniently using conditioned diffusion, classifier-free guidance, and classifier guidance model schemes. We demonstrate that such schemes can accurately model molecular data distributions and provide control over molecular properties of interest, such as physicochemical properties and binding affinities, with competitive quality to available alternatives. Furthermore, we demonstrate the practical utility of this framework, such as the ability to simultaneously generate close analogs that improve upon molecular properties and diverse generation of chemical sets using a particle guidance approach. Additionally, we utilize an open-source medicinal chemistry preference dataset to align generations to medicinal chemists' interests and validate these generations using a prospective survey of trained chemists. In sum, this latent diffusion methodology offers a practical and effective tool for conditional molecular generation, with clear applications in the hit-to-lead and lead optimization stages of small molecule drug discovery.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

References

Youngchun Kwon, Seokho Kang, Youn-Suk Choi, and Inkoo Kim. Evolutionary design of molecules based on deep learning and a genetic algorithm. *Scientific Reports*, 11(1):17304, 2021. doi: 10.1038/s41598-021-96812-8. URL https://doi.org/10.1038/s41598-021-96812-8.

Thomas Blaschke, Josep Arús-Pous, Hongming Chen, Christian Margreitter, Christian Tyrchan, Ola Engkvist, Kostas Papadopoulos, and Atanas Patronov. Reinvent 2.0: An ai tool for de novo drug design. *Journal of Chemical Information and Modeling*, 60(12):5918–5922, 2020. doi: 10.1021/acs.jcim.0c00915. URL https://doi.org/10.1021/acs.jcim.0c00915. PMID: 33118816.

Jiaxuan You, Bowen Liu, Rex Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy network for goal-directed molecular graph generation, 2019.

Wengong Jin, Regina Barzilay, and Tommi S. Jaakkola. Junction tree variational autoencoder for molecular graph generation. *CoRR*, abs/1802.04364, 2018a. URL http://arxiv.org/abs/1802.04364.

Camille Bilodeau, Wengong Jin, Tommi Jaakkola, Regina Barzilay, and Klavs F. Jensen. Generative models for molecular discovery: Recent advances and challenges. *WIREs Computational Molecular Science*, 12(5):e1608, 2022. doi: https://doi.org/10.1002/wcms.1608. URL https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wcms.1608.

Benjamin Kaufman, Edward Williams, Carl Underkoffler, Ryan Pederson, Narbe Mardirossian, Ian Watson, and John Parkhill. Coati: multi-modal contrastive pre-training for representing and traversing chemical space. *Journal of Chemical Information*, (accepted), 2024.

Danny Reidenbach, Micha Livne, Rajesh K. Ilango, Michelle Gill, and Johnny Israeli. Improving small molecule generation using mutual information machine, 2023.

Shengchao Liu, Weili Nie, Chengpeng Wang, Jiarui Lu, Zhuoran Qiao, Ling Liu, Jian Tang, Chaowei Xiao, and Anima Anandkumar. Multi-modal molecule structure-text model for text-based retrieval and editing, 2023.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2022.

Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation, 2022. URL https://arxiv.org/abs/2209.14734.

Victor Garcia Satorras, Emiel Hoogeboom, Fabian B. Fuchs, Ingmar Posner, and Max Welling. E(n) equivariant normalizing flows, 2021. URL https://arxiv.org/abs/2105.09016.

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. Structure-based drug design with equivariant diffusion models, 2023.

Pedro O. Pinheiro, Joshua Rackers, Joseph Kleinhenz, Michael Maser, Omar Mahmood, Andrew Martin Watkins, Stephen Ra, Vishnu Sresht, and Saeed Saremi. 3d molecule generation by denoising voxel grids, 2023.

Han Huang, Leilei Sun, Bowen Du, and Weifeng Lv. Learning joint 2d and 3d diffusion models for complete molecule generation, 2023.

Minkai Xu, Alexander S Powers, Ron O Dror, Stefano Ermon, and Jure Leskovec. Geometric latent diffusion models for 3d molecule generation. In *International Conference on Machine Learning*, pages 38592–38610. PMLR, 2023.

Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Therapeutics data commons: Machine learning datasets and tasks for drug discovery and development. *arXiv preprint arXiv:2102.09548*, 2021.

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, 2023a. doi: 10.1038/s41467-023-42242-1. URL https://doi.org/10.1038/s41467-023-42242-1.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

Abhishek Sinha, Jiaming Song, Chenlin Meng, and Stefano Ermon. D2c: Diffusion-decoding models for few-shot conditional generation. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 12533–12548. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/682e0e796084e163c5ca053dd8573b0c-Paper.pdf.

William S. Sly and Peiyi Y. Hu. Human carbonic anhydrases and carbonic anhydrase deficiencies. *Annual Review of Biochemistry*, 64(1):375–401, 1995. doi: 10.1146/annurev.bi.64.070195.002111. URL https://doi.org/10.1146/annurev.bi.64.070195.002111. PMID: 7574487.

Gabriele Corso, Yilun Xu, Valentin de Bortoli, Regina Barzilay, and Tommi Jaakkola. Particle guidance: non-i.i.d. diverse sampling with diffusion models, 2023a.

Yuanqi Du, Tianfan Fu, Jimeng Sun, and Shengchao Liu. Molgensurvey: A systematic survey in machine learning models for molecule design, 2022.

Jan H. Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. *Chem. Sci.*, 10:3567–3572, 2019. doi: 10.1039/C8SC05372C. URL http://dx.doi.org/10.1039/C8SC05372C.

Zhenpeng Zhou, Steven Kearnes, Li Li, Richard N. Zare, and Patrick Riley. Optimization of molecules via deep reinforcement learning. *Scientific Reports*, 9(1):10752, 2019. doi: 10.1038/s41598-019-47148-x. URL https://doi.org/10.1038/s41598-019-47148-x.

Tianfan Fu, Wenhao Gao, Cao Xiao, Jacob Yasonik, Connor W. Coley, and Jimeng Sun. Differentiable scaffolding tree for molecular optimization, 2022.

Gang Liu, Jiaxin Xu, Tengfei Luo, and Meng Jiang. Inverse molecular design with multi-conditional diffusion guidance, 2024.

Emiel Hoogeboom, Victor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3d, 2022.

Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Diffdock: Diffusion steps, twists, and turns for molecular docking, 2023b.

Michael Brocidiacono, Konstantin I. Popov, David Ryan Koes, and Alexander Tropsha. Plantain: Diffusion-inspired pose score minimization for fast and accurate molecular docking, 2023.

Renxiao Wang, Xueliang Fang, Yipin Lu, and Shaomeng Wang. The pdbbind database: Collection of binding affinities for protein-ligand complexes with known three-dimensional structures. *Journal of Medicinal Chemistry*, 47(12):2977–2980, 06 2004. doi: 10.1021/jm030580l. URL https://doi.org/10.1021/jm030580l.

Jaechang Lim, Seongok Ryu, Jin Woo Kim, and Woo Youn Kim. Molecular generative model based on conditional variational autoencoder for de novo molecular design. *Journal of Cheminformatics*, 10(1):31, 2018. doi: 10.1186/s13321-018-0286-7. URL https://doi.org/10.1186/s13321-018-0286-7.

Rafael Gómez-Bombarelli, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2):268–276, 2018. doi: 10.1021/acscentsci.7b00572. URL https://doi.org/10.1021/acscentsci.7b00572. PMID: 29532027.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation, 2019a.

Yixin Wang, David M. Blei, and John P. Cunningham. Posterior collapse and latent variable non-identifiability, 2023.

Huaisheng Zhu, Teng Xiao, and Vasant G Honavar. 3m-diffusion: Latent multi-modal diffusion for text-guided generation of molecular graphs. *arXiv preprint arXiv:2403.07179*, 2024.

Robin Winter, Floriane Montanari, Frank Noé, and Djork-Arné Clevert. Learning continuous and data-driven molecular descriptors by translating equivalent chemical representations. *Chem. Sci.*, 10:1692–1701, 2019. doi: 10.1039/C8SC04175J. URL http://dx.doi.org/10.1039/C8SC04175J.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

Samuel C Hoffman, Vijil Chenthamarakshan, Kahini Wadhawan, Pin-Yu Chen, and Payel Das. Optimizing molecules using efficient queries from property evaluations. *Nature Machine Intelligence*, 4(1):21–31, 2022a.

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, 03 2019. doi: 10.1021/acs.jcim.8b00839. URL https://doi.org/10.1021/acs.jcim.8b00839.

Patrick Walters. Generative molecular design - we need to raise the bar, 2023. URL https://practicalcheminformatics.blogspot.com/2023/02/generative-molecular-design-we-need-to.html.

Nvidia, online. Dgx cloud user guide, 2024. URL https://docs.nvidia.com/dgx-cloud/cluster-user-guide/latest/content.html. Accessed: 2024-08-02.

Albert Musaelian, Simon Batzner, Anders Johansson, Lixin Sun, Cameron J. Owen, Mordechai Kornbluth, and Boris Kozinsky. Learning local equivariant representations for large-scale atomistic dynamics. *Nature Communications*, 14(1):579, 2023. doi: 10.1038/s41467-023-36329-y. URL https://doi.org/10.1038/s41467-023-36329-y.

Li Jing, Pascal Vincent, Yann LeCun, and Yuandong Tian. Understanding dimensional collapse in contrastive self-supervised learning, 2022.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models, 2022.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling, 2023.

Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, and Yaron Lipman. D-flow: Differentiating through flows for controlled generation, 2024a.

Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi Jaakkola. Harmonic self-conditioned flow matching for multi-ligand docking and binding site design. *arXiv preprint arXiv:2310.05764*, 2023.

Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis, 2021.

Joost van Amersfoort, Lewis Smith, Andrew Jesson, Oscar Key, and Yarin Gal. On feature collapse and deep kernel learning for single forward pass uncertainty, 2022.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance, 2022.

Heli Ben-Hamu, Omri Puny, Itai Gat, Brian Karrer, Uriel Singer, and Yaron Lipman. D-flow: Differentiating through flows for controlled generation. *arXiv preprint arXiv:2402.14017*, 2024b.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017a. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/8a1d694707eb0fefe65871369074926d-Paper.pdf.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback, 2022.

Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, Kurt Keutzer, and Trevor Darrell. Aligning large multimodal models with factually augmented rlhf, 2023.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2023.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam, Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using direct preference optimization, 2023.

Yinan Zhang, Eric Tzeng, Yilun Du, and Dmitry Kislyuk. Large-scale reinforcement learning for diffusion models, 2024.

Bill Yuchen Lin, Abhilasha Ravichander, Ximing Lu, Nouha Dziri, Melanie Sclar, Khyathi Chandu, Chandra Bhagavatula, and Yejin Choi. The unlocking spell on base llms: Rethinking alignment via in-context learning. *ArXiv preprint*, 2023.

Peter S Kutchukian, Nadya Y Vasilyeva, Jordan Xu, Mika K Lindvall, Michael P Dillon, Meir Glick, John D Coley, and Natasja Brooijmans. Inside the mind of a medicinal chemist: the role of human bias in compound prioritization during drug discovery. *PLoS One*, 7(11):e48476, November 2012.

Júlia G B Pedreira, Lucas S Franco, and Eliezer J Barreiro. Chemical intuition in drug design and discovery. *Curr. Top. Med. Chem.*, 19(19):1679–1693, 2019.

Christopher A Lipinski. Lead- and drug-like compounds: the rule-of-five revolution. *Drug Discov. Today Technol.*, 1 (4):337–341, December 2004.

G Richard Bickerton, Gaia V Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L Hopkins. Quantifying the chemical beauty of drugs. *Nat. Chem.*, 4(2):90–98, January 2012.

Yuji Takaoka, Yutaka Endo, Susumu Yamanobe, Hiroyuki Kakinuma, Taketoshi Okubo, Youichi Shimazaki, Tomomi Ota, Shigeyuki Sumiya, and Kensei Yoshikawa. Development of a method for evaluating drug-likeness and ease of synthesis using a data set in which compounds are assigned scores based on chemists' intuition. *Journal of chemical information and computer sciences*, 43(4):1269–1275, 2003.

Robert P Sheridan, Nicolas Zorn, Edward C Sherer, Louis-Charles Campeau, Charlie Chang, Jared Cumming, Matthew L Maddess, Philippe G Nantermet, Christopher J Sinz, and Paul D O'Shea. Modeling a crowdsourced definition of molecular complexity. *Journal of chemical information and modeling*, 54(6):1604–1616, 2014.

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, 2023b. doi: 10.1038/s41467-023-42242-1. URL `https://doi.org/10.1038/s41467-023-42242-1`.

Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, pages 89–96, 2005.

Oleksandr O Grygorenko, Dmytro S Radchenko, Igor Dziuba, Alexander Chuprina, Kateryna E Gubina, and Yurii S Moroz. Generating multibillion chemical space of readily accessible screening compounds. *iScience*, 23(11):101681, Nov 2020. ISSN 2589-0042 (Electronic); 2589-0042 (Linking). doi: 10.1016/j.isci.2020.101681.

Anna Gaulton, Louisa J Bellis, A Patricia Bento, Jon Chambers, Mark Davies, Anne Hersey, Yvonne Light, Shaun McGlinchey, David Michalovich, Bissan Al-Lazikani, and John P Overington. Chembl: a large-scale bioactivity database for drug discovery. *Nucleic Acids Res*, 40(Database issue):D1100–7, Jan 2012. ISSN 1362-4962 (Electronic); 0305-1048 (Print); 0305-1048 (Linking). doi: 10.1093/nar/gkr777.

RDKit, online. RDKit: Open-source cheminformatics. `http://www.rdkit.org`. [Online; accessed 11-April-2013].

Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and Sanjiv Kumar. Rethinking fid: Towards a better evaluation metric for image generation. *arXiv preprint arXiv:2401.09603*, 2023.

Dongjun Kim, Chieh-Hsin Lai, Wei-Hsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Uesaka, Yutong He, Yuki Mitsufuji, and Stefano Ermon. Consistency trajectory models: Learning probability flow ode trajectory of diffusion, 2023.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017b.

Austin Tripp and José Miguel Hernández-Lobato. Genetic algorithms are strong baselines for molecule generation, 2023.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

David Mendez, Anna Gaulton, A Patrícia Bento, Jon Chambers, Marleen De Veij, Eloy Félix, María Paula Magariños, Juan F Mosquera, Prudence Mutowo, Michał Nowotka, María Gordillo-Marañón, Fiona Hunter, Laura Junco, Grace Mugumbate, Milagros Rodriguez-Lopez, Francis Atkinson, Nicolas Bosc, Chris J Radoux, Aldo Segura-Cabrera, Anne Hersey, and Andrew R Leach. ChEMBL: towards direct deposition of bioassay data. *Nucleic Acids Research*, 47(D1):D930–D940, 11 2018. ISSN 0305-1048. doi: 10.1093/nar/gky1075. URL https://doi.org/10.1093/nar/gky1075.

Dávid Bajusz, Anita Rácz, and Károly Héberger. Why is tanimoto index an appropriate choice for fingerprint-based similarity calculations? *Journal of Cheminformatics*, 7(1):20, 2015. doi: 10.1186/s13321-015-0069-3. URL https://doi.org/10.1186/s13321-015-0069-3.

Harry L Morgan. The generation of a unique machine description for chemical structures-a technique developed at chemical abstracts service. *Journal of chemical documentation*, 5(2):107–113, 1965.

Snowflake. Streamlit, 2024. URL https://streamlit.io/.

Wengong Jin, Kevin Yang, Regina Barzilay, and Tommi Jaakkola. Learning multimodal graph-to-graph translation for molecule optimization. In *International Conference on Learning Representations*, 2019b.

Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy network for goal-directed molecular graph generation. *Advances in neural information processing systems*, 31, 2018.

Andrew Dalke, Jerome Hert, and Christian Kramer. mmpdb: An open-source matched molecular pair platform for large multiproperty data sets. *Journal of chemical information and modeling*, 58(5):902–910, 2018.

Dzmitry Bahdanau, Kyung Hyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*, 2015.

Wengong Jin, Kevin Yang, Regina Barzilay, and Tommi Jaakkola. Learning multimodal graph-to-graph translation for molecule optimization. In *International Conference on Learning Representations*, 2018b.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Hierarchical generation of molecular graphs using structural motifs. In *International conference on machine learning*, pages 4839–4848. PMLR, 2020.

Paul Maragakis, Hunter Nisonoff, Brian Cole, and David E Shaw. A deep-learning view of chemical space designed to facilitate drug discovery. *Journal of Chemical Information and Modeling*, 60(10):4487–4496, 2020.

Samuel C. Hoffman, Vijil Chenthamarakshan, Kahini Wadhawan, Pin-Yu Chen, and Payel Das. Optimizing molecules using efficient queries from property evaluations. *Nature Machine Intelligence*, 4(1):21–31, 2022b. doi: 10.1038/s42256-021-00422-y. URL https://doi.org/10.1038/s42256-021-00422-y.

Alberto Gobbi and Man-Ling Lee. Dise: Directed sphere exclusion. *Journal of Chemical Information and Computer Sciences*, 43(1):317–323, 01 2003. doi: 10.1021/ci025554v. URL https://doi.org/10.1021/ci025554v.

Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.

Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction, 2021.

Simon Axelrod and Rafael Gómez-Bombarelli. Geom, energy-annotated molecular conformations for property prediction and molecular generation. *Scientific Data*, 9(1):185, 2022. doi: 10.1038/s41597-022-01288-4. URL https://doi.org/10.1038/s41597-022-01288-4.

Kun Yao, John E. Herr, David W. Toth, Ryker Mcintyre, and John Parkhill. The tensormol-0.1 model chemistry: a neural network augmented with long-range physics, 2017. URL https://arxiv.org/abs/1711.06385.

Pavel Polishchuk. Crem: chemically reasonable mutations framework for structure generation. *Journal of Cheminformatics*, 12(1):28, 2020. ISSN 1758-2946. doi: 10.1186/s13321-020-00431-w. URL https://doi.org/10.1186/s13321-020-00431-w.

Patrick Ropp, Jesse Kaminsky, Sara Yablonski, and Jacob Durrant. Dimorphite-dl: An open-source program for enumerating the ionization states of drug-like small molecules. *Journal of Cheminformatics*, 11, 02 2019. doi: 10.1186/s13321-019-0336-9.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models, 2023.

Noam Shazeer. Glu variants improve transformer, 2020.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

A Human Carbonic Anhydrase II Data

The hCAII binding affinity data is produced using our in-house, proprietary tArray platform, which measures quantitative binding affinities of millions of small molecules to a biological target of interest (protein, RNA, etc.) through an image-based, fluorescent readout. The small molecule libraries are built combinatorially using commercial and custom building blocks - with the small molecules tethered to beads that are immobilized to a 32 million-well silicon chip that is the size of a nickel. The targets are either directly labeled with a fluorophore or detected through a fluorescently-labeled secondary detection modality (i.e., Anti-6X His Tag Antibody) that binds the complementary tag on the target. Each target-molecule interaction is measured 15 times per chip on average (governed by a Poisson distribution). 500k rows of data were randomly selected from this data for our experiments. Binding is regressed as a function of the cumulative distribution function (CDF) in our experiments. This data will be made available at https://github.com/terraytherapeutics/COATI-LDM at a later date.

B Encoder-Decoder Details

Our molecular encoder-decoder is derived from the architecture of [Kaufman et al., 2024], with a few architectural changes and an expanded training set. The E(3)-GNN architecture used for the 3D conformer encoder was replaced with an Allegro model [Musaelian et al., 2023], which provides resolution of chirality. The size of the transformer encoder was increased up to 62M parameters. We used 512-dimensional encodings for all experiments in this work (up from 256 in the previous architecture). We replace the InfoNCE [Oord et al., 2018] and Barlow Twins [Zbontar et al., 2021] losses used in the original work with DirectCLR [Jing et al., 2022], which we find empirically to reduce dimensional collapse in the latent representation.

The sources included in the encoder-decoder dataset (Table 5) were mostly synthetically reasonable catalogs (Enamine, Mcule, Wuxi, Zinc22, Terray Libraries, and Enamine building block space) but supplemented with compounds from in-vitro & in-vivo studies (Gostar, CHEMBL), conformer ensembles from GEOM-Drugs [Axelrod and Gómez-Bombarelli, 2022] & TensorMol [Yao et al., 2017], and permuted compounds (Protonation Enumeration, ChEMBL 33 CREM, Difficult CREM, Multi-Fragment) via CREM [Polishchuk, 2020] or Dimorphite-DL [Ropp et al., 2019].

As described in the main text, the molecular decoder we employ can complete prompts that contain a property condition. Table 6 contains a complete list of available properties which were obtained via RDKit [RDKit, online].

Name	Description	Samples
Enamine Real	Classical database of enumerated structures	59,999,028
Mcule	Purchasable Mcule supplier & ULTIMATE catalogs	39,999,889
Virtual Terray Libraries	Virtually enumerated molecules available to Terray's screening assay	39,999,708
Wuxi	Virtual screening catalog	39,953,341
Zinc22	Database of commercially-available compounds for virtual screening	35,262,038
Terray Libraries	Sample of combinatorial libraries used in Terray's screening assay	13,999,674
Gostar	Curated SAR database	9,088,018
Protonation Enumeration	ChEMBL and Terray library samples modified using Dimorphite-DL	8,147,175
GEOM-Drugs	Molecules with multiple energy-annotated conformations	6,999,873
ChEMBL 33 CREM	ChEMBL molecules modified using the CREM package	5,235,347
Difficult CREM	Low COATI[Kaufman et al., 2024] likelihood molecules modified using CREM	2,308,381
ChEMBL 33	Curated database of bioactive molecules with drug-like properties	1,325,368
TensorMol	Molecules with multiple DFT-optimized conformers	399,946
Enamine Building-Blocks	Combinatorial library building blocks from Enamine	299,953
Multi-Fragment	Two fragment molecules from ChEMBL	58,118

Table 5: Sources sampled to train the encoder-decoder.

C Score Network Architecture

This work replaces the convolution layers of a U-Net with 2-layer weight normalized [Karras et al., 2023] SwiGLU [Shazeer, 2020] modules which downsample and upsample. Three downsampling stages, followed by three upsampling stages gave the best results in our experiments. Any scalar variables made available to a score network are sinusoidally encoded. Unless otherwise mentioned 1000 timesteps are used in all sampling experiments with a linear noise schedule with β in range $[0.0001, .02]$.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

RDKit property	Description
MolLogP	LogP of the molecule.
FormalCharge	Formal charge of the molecule.
RingCount	Number of rings in the molecule.
NumHDonors	Number of hydrogen bond donors.
NHOHCount	Number of NHs or OHs.
ExactMolWt	Molecular weight of the molecule.
NumAliphaticCarbocycles	Number of aliphatic carbocycles.
NumHAcceptors	Number of hydrogen bond acceptors.
NumRotatableBonds	Number of rotatable bonds.
TPSA	Topological polar surface area.
NumAromaticCarbocycles	Number of aromatic carbocycles.
FractionCSP3	Fraction of SP3 carbons.
HeavyAtomCount	Number of heavy atoms.
QED	Quantitative Estimate of Drug-likeness.
NumAromaticRings	Number of aromatic rings.
NOCount	Number of nitro and nitroso groups.

Table 6: The list of specific RDKit properties included as tokens in the COATI encoder-decoder transformer model.

D Preference Modeling

To test our ability to control the "desirability" of a molecule, we first train an MLP RankNet using the methodology of [Choung et al., 2023b]. We observe in Figure 11 that these latent embeddings reproduce similar performance as the original fingerprint-based model.

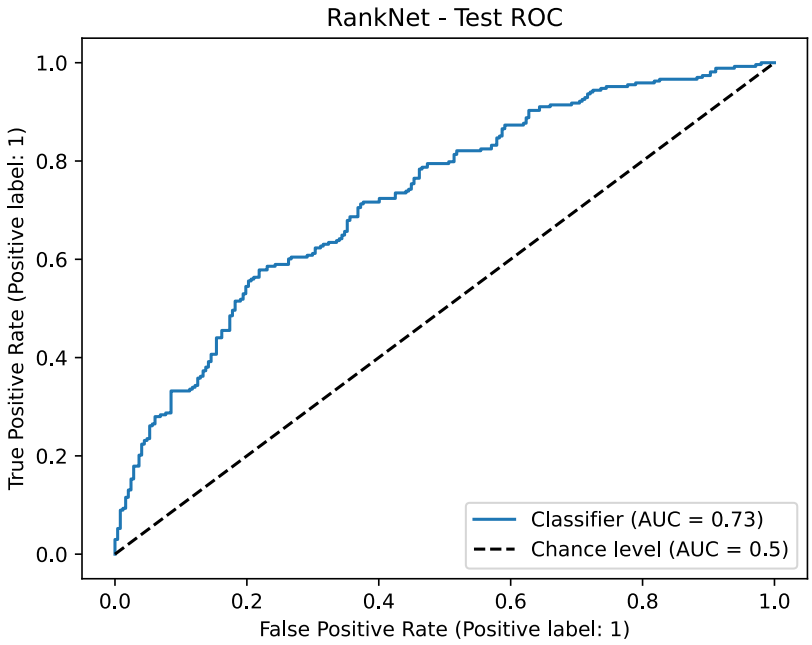

Figure 11: Receiver operating characteristic (ROC) curve of held-out molecule ranking pairs. This reproduces the results from [Choung et al., 2023b]

We can then relabel the RankNet training molecules using the model, and fit a conditional diffusion model using the procedure described in Section 3.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

Method	Preference
Joint	.923
Classifier-Free Guidance	.938
Classifier Guidance	.770

Table 7: Property/Target Correlations

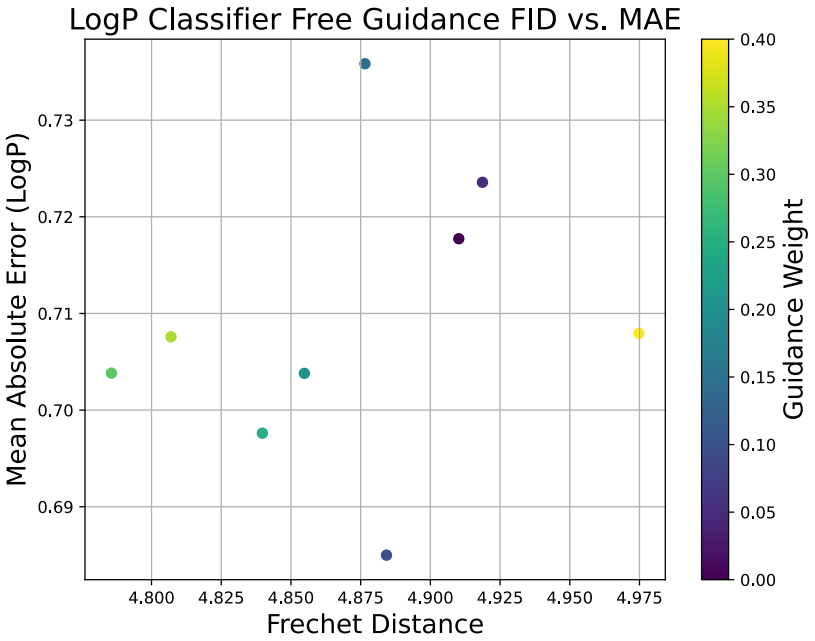

Figure 12: Mean Absolute Error and Frechet Distances of different Classifier-free guidance weights when conditioned on LogP. We observe no obvious correlation between FD and MAE, with a limited range of both FDs and MAEs across weights.

We see in Table 7 that control over the preference score, as measured by RankNet predictions on the resulting samples, is generally good. The conditional diffusion model can easily learn the target mapping ("copying" the RankNet scoring model) and perform conditioning of a large diffusion model trained on a very distinct dataset.

Given that one of the major benefits of this preference modeling is its improvability, we investigated potential deficiencies in the scoring or generation. Feedback from participants indicated that some pairs were ambiguous, and a qualitative analysis of the pairs where chemists disagreed with the model show strange or otherwise unstable groups. Figure 15 shows an example pair where the respondent chose the low-preference generation over the high-preference generation. We observe that the high-preference generation includes a highly strained ring system, indicating that the learned preference score has an incomplete understanding of chemical stability. However, this feedback can be easily integrated into the learned preference framework.

We hypothesize that the preference model may be incompletely aligned, or conditioning using such a small training dataset could lead to degenerate samples. In aggregate, we find these proof-of-concept results encouraging, in particular due to the improvability of our preference model. We anticipate exploring the adaptation of chemical preference alignment to the needs of specific teams or projects in future work.

E QED benchmark methodology

For the nearby diffusion on the QED benchmark task we used the same unconditional latent generator trained on the data in table 5 and trained a DUE Classifier Guide on a random 500k subset to predict QED. we embedded the 800

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

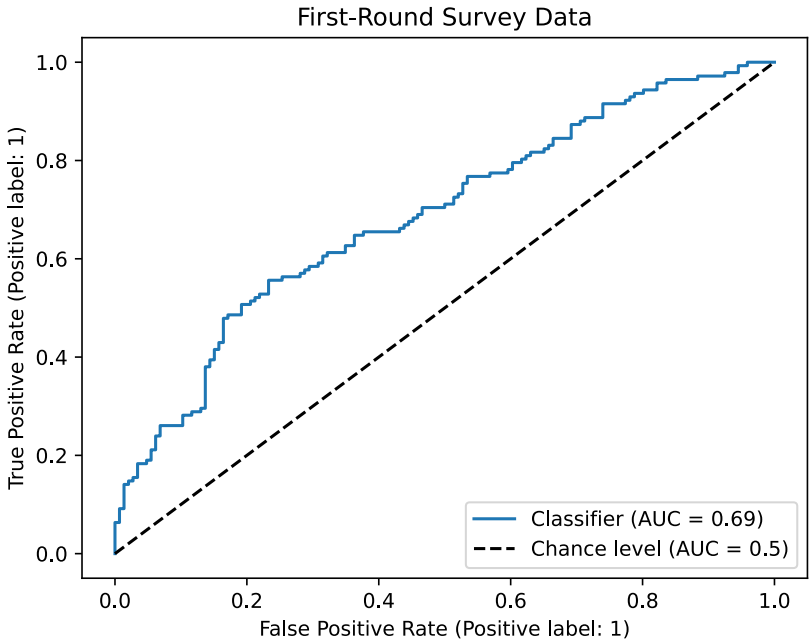

Figure 13: Pre-trained RankNet model inference on first round survey data. We find the preference model to have moderate agreement with our survey participants, on par with the held-out pairs from the model's data source.

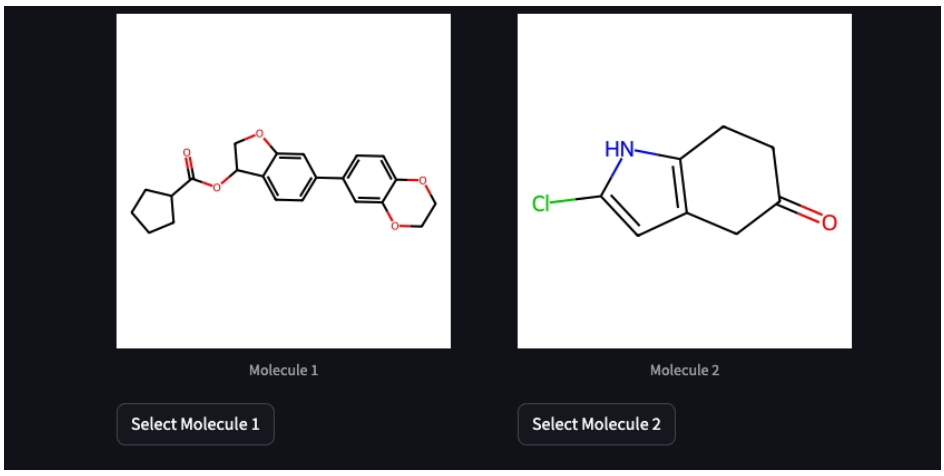

Figure 14: Streamlit molecule selection interface.

bioRxiv preprint doi: https://doi.org/10.1101/2024.08.22.609169; this version posted August 22, 2024. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY-NC-ND 4.0 International license.

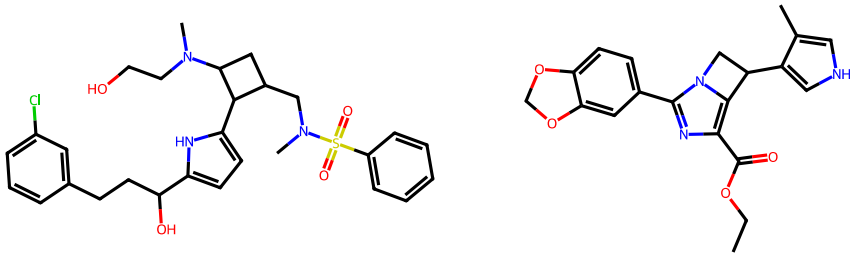

selected (low preference) not selected (high preference)

Figure 15: A pair of molecules where the low-preference generation was chosen. We observe that the high-preference generation includes a highly strained ring system, indicating that the learned preference score has an incomplete understanding of chemical stability.

molecules in the benchmark set with our encoder and then applied a grid-search like methodology to each of them up to the maximum 50,000 oracle calls specified by the benchmark. The pseudocode for this approach is in algorithm 1.

Algorithm 1 QED Benchmark Approach

```
function QEDBENCHMARK(inputMol)
    emb ← ENCODE(inputMol)
    guideWeights ← [10, 30, 50]
    noiseSteps ← [50, 100, 150, 200]
    batchSize ← 100
    oracleCalls ← 0
    for each w in guideWeights do
        for each s in noiseSteps do
            outputEmbs ← PARTIALDIFFUSION(emb, w, s, batchSize)
            outputMols ← DECODE(outputEmbs)
            for each mol in outputMols do
                if MEETSCRITERION(mol) then
                    return mol
                else
                    oracleCalls ← oracleCalls + 1
                end if
                if oracleCalls == 50000 then
                    return false
                end if
            end for
        end for
    end for
end function
```