

# Patterns

## Multi-objective latent space optimization of generative molecular design models

### Highlights

- Multi-objective optimization enhances the capability to design Pareto optimal molecules
- Molecules are ranked based on Pareto optimality, guiding the model optimization process
- Evaluations show that the method effectively pushes the Pareto front for multiple properties
- Predicted DRD2 inhibitors show superior performance to known drugs in *in silico* evaluations

### Authors

A N M Nafiz Abeer, Nathan M. Urban, M. Ryan Weil, Francis J. Alexander, Byung-Jun Yoon

### Correspondence

[bjyoon@ece.tamu.edu](mailto:bjyoon@ece.tamu.edu)

### In brief

Abeer et al. propose a multi-objective latent space optimization approach to make a generative model biased toward suggesting molecules that jointly meet multiple design criteria. The concept of Pareto optimality incorporated into a practical weighted retraining scheme enables the approach to handle the challenges that emerge in the generative design of novel molecules that have multiple optimized properties of interest. The efficacy of the proposed approach is demonstrated based on a JT-VAE (junction-tree variational autoencoder), a well-known deep generative model for molecular design, in various multi-objective molecular design tasks.



## Article

# Multi-objective latent space optimization of generative molecular design models

A N M Nafiz Abeer,<sup>1</sup> Nathan M. Urban,<sup>2</sup> M. Ryan Weil,<sup>3</sup> Francis J. Alexander,<sup>4</sup> and Byung-Jun Yoon<sup>1,2,5,\*</sup><sup>1</sup>Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843, USA<sup>2</sup>Computational Science Initiative, Brookhaven National Laboratory, Upton, NY 11973, USA<sup>3</sup>Strategic and Data Science Initiatives, Frederick National Laboratory, Frederick, MD 21702, USA<sup>4</sup>Computing, Environment and Life Sciences, Argonne National Laboratory, Lemont, IL 60439, USA<sup>5</sup>Lead contact\*Correspondence: [bjyoon@ece.tamu.edu](mailto:bjyoon@ece.tamu.edu)<https://doi.org/10.1016/j.patter.2024.101042>

**THE BIGGER PICTURE** A generative model is a type of machine learning method that generates new information from learning patterns from existing data. Generative models are accelerating the search for drug compounds with desired properties. However, the efficiency of these models in exploring the vast chemical space is typically limited by the molecules used during model training. In addition, identifying molecules that simultaneously have ideal properties is also a challenge. We propose a “multi-objective optimization approach,” a method that effectively biases a given generative model toward optimized molecules given an arbitrary number of properties. Improving molecule-generation models is essential to the development of drug discovery and molecule optimization schemes.

## SUMMARY

Molecular design based on generative models, such as variational autoencoders (VAEs), has become increasingly popular in recent years due to its efficiency for exploring high-dimensional molecular space to identify molecules with desired properties. While the efficacy of the initial model strongly depends on the training data, the sampling efficiency of the model for suggesting novel molecules with enhanced properties can be further enhanced via latent space optimization (LSO). In this paper, we propose a multi-objective LSO method that can significantly enhance the performance of generative molecular design (GMD). The proposed method adopts an iterative weighted retraining approach, where the respective weights of the molecules in the training data are determined by their Pareto efficiency. We demonstrate that our multi-objective GMD LSO method can significantly improve the performance of GMD for jointly optimizing multiple molecular properties.

## INTRODUCTION

The development of quantitative structure-activity relationship (QSAR)<sup>1</sup> models has accelerated the drug design process. However, designing molecules with the desired drug properties through direct optimization over the chemical space remains challenging due to the high dimensionality of the domain. While drug discovery based on high-throughput screening (HTS) systems<sup>2</sup> has been shown to be highly useful, the computational cost needed for screening a huge candidate pool is formidable. In addition, the design and operation of computational HTS pipelines have traditionally relied on expert intuition and various heuristics, resulting in suboptimal performance.<sup>3,4</sup> Furthermore, should one wish to consider drug candidates beyond the pool of known drugs and drug-like molecules, expanding the pool

faces the same challenges of molecular design in high-dimensional chemical space. As drug discovery involves consideration and optimization of multiple properties, which may conflict with one another, this multi-objective optimization aspect further exacerbates the aforementioned design challenges.

Recent advances in deep generative models provide promising alternatives to conventional computational approaches for drug discovery, which may be able to effectively address many of these challenges. A representative example is the work by Gómez-Bombarelli et al.,<sup>5</sup> in which they propose the use of a variational autoencoder (VAE) to convert the input molecules, originally represented by simplified molecular-input line-entry system (SMILES) strings, into a continuous lower-dimensional representation in a latent space. This approach effectively maps molecules in the original chemical space, which is high dimensional and



discrete, to a latent space, which is low dimensional and continuous, thereby enabling efficient numerical optimization in the latent space in pursuit of molecules with enhanced target attributes. In this study,<sup>5</sup> a Gaussian process (GP) was used to model and optimize the property predictor in the latent space, which was shown to significantly outperform molecular optimization in the original chemical space using a genetic algorithm (GA) as well as a randomized search in the latent space. Winter et al.<sup>6</sup> adopted particle swarm optimization (PSO), instead of Bayesian optimization (BO), aiming at further improving the computational efficiency for multi-objective molecular optimization, also in the latent space. In this work, multiple properties of interest were jointly optimized by defining a single-objective function through scalarization via weighted combination of multiple property scores. As noted in Gómez-Bombarelli et al.,<sup>5</sup> the generative model may not always suggest molecules with valid molecular structures, which may degrade the overall efficiency of the generative molecular design (GMD) approach. Empirically, this phenomenon has been shown to occur when data points representing the molecules are sampled in regions of the latent space that are far away from the region where the original training data were located. To deal with this shortcoming, the search for an optimized molecule with desirable attributes can be formulated as a constrained BO problem,<sup>7</sup> which has been shown to improve the validity of the novel molecules produced by the generative model. The junction-tree VAE (JT-VAE)<sup>8</sup> tackles this issue by taking a two-phase approach. In the first phase, the JT-VAE generates a junction tree that represents the overall scaffold for a molecular graph, which specifies the relative arrangement of valid subgraph structures learned from the training data. During the second phase, subgraphs corresponding to chemical substructures are combined according to the junction tree to obtain the final molecular graph. As a result, JT-VAE is capable of suggesting novel molecules in the latent space that can be decoded into legitimate molecules with a high chance. To explore the latent space to produce novel molecules with targeted attributes, the VAE may also be conditioned by the desired property values. For example, Kang and Cho<sup>9</sup> proposed a semisupervised VAE (SSVAE), which simultaneously performs property prediction and molecular generation, resulting in a conditioning of the model such that it suggests molecules in the latent space that are centered around a desired range of properties.

In addition to the aforementioned schemes that perform molecular optimization in the latent space, another popular strategy is to first train a generative network to model the input data distribution, which is followed by fine-tuning the model via reinforcement learning (RL) to meet the design criteria. The work by Olivecrona et al.<sup>10</sup> proposed a fine-tuning approach that facilitates the generation of high-scoring molecules without deviating away from the original input data distribution. Shi et al.<sup>11</sup> applied an RL-based policy to fine-tune a generative network model for molecular graphs. In the objective-reinforced generative adversarial network (ORGAN)<sup>12</sup> framework, the reward for generating molecules with better properties is assimilated into the loss function to guide the latent space distribution during model training. To reduce the potential bias in the generative network that may arise from the training dataset, Zhou et al.<sup>13</sup> proposed molecule deep Q networks (MoIQNs). In this work, the molecular generation problem was formulated as a Markov decision process

(MDP), and a deep Q network (DQN)<sup>14</sup> was used to find the optimal design policy for the given MDP. The allowable actions in the MDP were dictated by relevant domain knowledge (i.e., chemical reactions) to ensure the validity of the generated molecule. To jointly optimize multiple molecular attributes, MoIQN also resorted to scalarization by defining a single-objective function based on a weighted combination of multiple property scores.

Like other data-driven models, the initial capability of the generative molecular models to suggest novel molecules will be determined—at least to a certain extent—by the training data. In fact, the generated molecules are likely to reside in a chemical space similar to that of the molecules in the original training set, which may make it challenging to design new candidate molecules whose target attributes significantly exceed those of the original molecules—regardless of whether we perform molecular optimization in the latent space<sup>5,8,9</sup> or adopt a fine-tuning strategy.<sup>10–13</sup> To mitigate this issue, a number of recent efforts aimed to improve the generative models by incorporating additional training data, generated either from experiments or from simulations, where the goal was to ensure that these models were primed to suggest novel molecules with enhanced target properties that went beyond the initial molecules. Yang et al.<sup>15</sup> proposed an iterative retraining approach to improve the quality of the molecules sampled from the latent space of the generative model. In this approach, they pre-train the generative model jointly with a property predictor. The trained model is used to generate novel molecules from which a small batch of molecules with the best properties is selected using the predictor and added to the training dataset. The extended dataset is subsequently used for retraining to update the latent space to make it more amenable to producing better molecules with improved properties. A similar effort from Iovanac et al.<sup>16</sup> utilizes the grammar VAE,<sup>17</sup> jointly trained with a linear predictor network to sample new molecules along the latent dimension corresponding to the targeted property region. The new molecules go through further screening to be used with the training dataset to retrain the generative model iteratively. In Liu et al.,<sup>18</sup> the generator in a generative adversarial network (GAN) is iteratively updated under a chance-constrained optimization framework. They employ a validity function to guide the optimization of the property value within the region of grammatically correct input sequences or legitimate structures. A recent work by Tripp et al.<sup>19</sup> proposed a weighted retraining approach to reshape the latent space of a VAE to make it more sampling efficient for producing novel molecules with improved properties. For this purpose, the weights assigned to the data points in the training dataset are determined by the rank of the corresponding molecules according to the objective function that evaluates the property of interest. Another recent iterative approach<sup>20</sup> applies a GA along with domain knowledge to generate a new set of potentially improved candidates from the initial training data. The property values of these candidates are validated by high-throughput experiments or simulations. A deep neural network (DNN) property predictor is then retrained using a larger dataset augmented by the candidates generated by a GA. In contrast to the generative approaches,<sup>15,16,18,19</sup> the candidates produced by a GA generally remain closer to the training data. As a result, the predictions from the DNN remain

relatively reliable, but the novelty of the produced candidates tends to be limited. In addition, the (computational) cost of the experiments or simulations to assess the property of the novel candidates can be a significant burden when the candidate pool gets larger.

In drug discovery applications, potential drug candidates always have to simultaneously meet multiple design criteria. For example, in addition to their capability to intervene in a specific biomolecular target or mechanism, potential drug molecules are assessed based on various physiochemical properties that contribute to their bioavailability. In the generative molecule design schemes discussed above, multi-objective optimization is typically handled via scalarization, which turns the problem into a simpler single-objective optimization problem. However, the weights for combining the multiple objective functions are often selected by the designer in an *ad hoc* manner, despite their importance in guiding the multi-objective optimization process. This may potentially lead to suboptimal results, as the optimization process may be unintentionally dominated by a few objective functions, resulting in emphasizing certain attributes while ignoring others. Various other approaches have also been proposed to address multiple design criteria, a notable example of which is RationaleRL, recently proposed by Jin et al.<sup>21</sup> In this work, a policy gradient is applied to fine-tune a VAE for generating molecules from incomplete subgraphs that correspond to multi-property rationales. RationaleRL utilizes a property predictor to extract the subgraph rationale for each of the multiple target properties. From the set of extracted single-property rationales, RationaleRL identifies the combined subgraph that meets the multiple property constraints and then uses the fine-tuned generative model to complete the molecular graph for the combined subgraph. However, a practical limitation of this approach is that the number of training samples that satisfy all properties of interest may be scarce, which is exacerbated as the number of target design criteria increases. Markov molecular sampling (MARS)<sup>22</sup> was proposed to address the multi-objective drug-discovery problem by formulating the molecular design process as an iterative graph-editing process. For this purpose, MARS defines the target distribution by combining the scoring functions of multiple properties and adopts Markov chain Monte Carlo (MCMC) sampling to identify high-scoring molecular candidates. However, as the target distribution is obtained by taking either the sum or the product of the multiple scoring functions, it faces similar shortcomings compared with other scalarization-based approaches discussed earlier. To generate molecules with multiple target properties, Feng et al.<sup>23</sup> utilized a Langevin diffusion process as a stochastic generator of latent embeddings within a pre-trained autoencoder model. This stochastic generator, primarily governed by reference molecules selected based on multiple properties of interest, can be seen as an alternative to other optimization strategies in latent space.

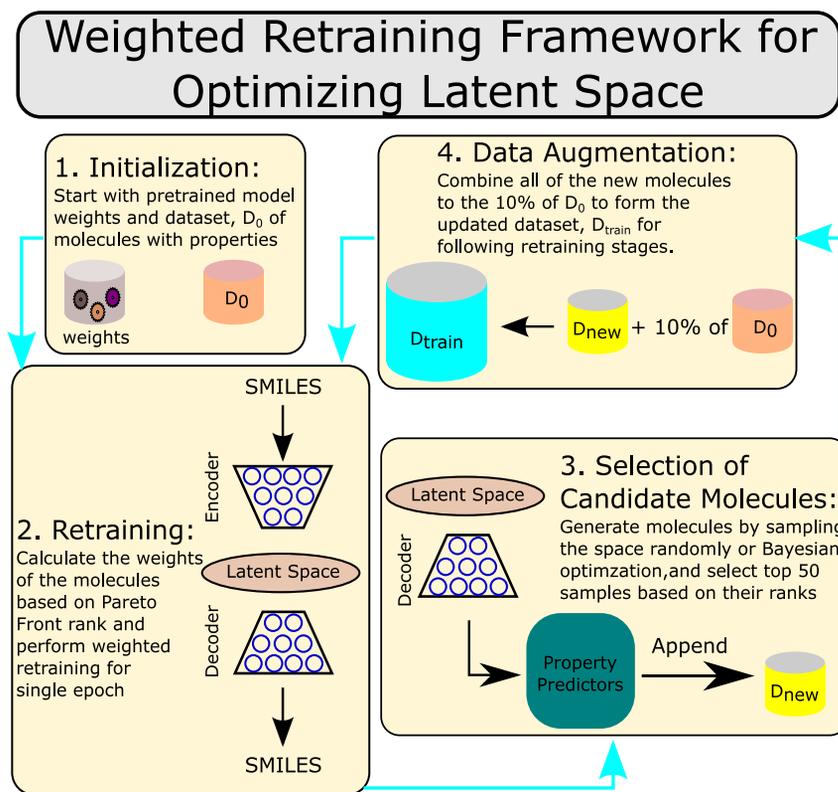
In this paper, we propose a novel multi-objective latent space optimization (MO-LSO) scheme that can effectively address the aforementioned limitations of existing generative models for molecular design—specifically, their limited capability for extrapolation in a multi-objective fashion beyond the molecular property space seen during training. We extend the weighted retraining framework recently proposed in Tripp et al.<sup>19</sup> to equip it with the inherent capability to enhance the efficiency of sampling

novel molecules in the latent space that simultaneously improves multiple target properties. This is achieved by ranking the molecules based on their Pareto optimality through non-dominated sorting (NDS), where the rankings are used both for generating improved molecules based on multiple design criteria to augment the training data and for determining the weights of the molecules in the (augmented) training set based on the relative importance. Our proposed MO-LSO scheme can naturally balance the trade-offs among multiple properties without any *ad hoc* scalarization that may potentially bias the optimization results. Furthermore, as the molecules are assessed based on their Pareto efficiency, our MO-LSO scheme scales very well computationally as the number of design criteria increases (empirically demonstrated up to three objectives in this study), and it is naturally equipped with the capability to handle the optimization of properties that may be highly correlated (or even redundant) or conflict with one another. We show that our proposed MO-LSO scheme can effectively shift the latent space representation of the molecules based on multiple design criteria, thereby substantially enhancing the sampling efficiency of the generative model for suggesting novel molecules that simultaneously improve multiple properties. Furthermore, by applying it to the design of dopamine receptor D2 (DRD2) inhibitors, we also demonstrate through *in silico* analysis that the generative model optimized by the proposed MO-LSO scheme is able to produce highly promising molecules that outperform known DRD2-inhibitory molecules.

## RESULTS

### Overview of the proposed method and the experimental setup

Figure 1 provides an overview of our proposed MO-LSO scheme for GMD. Given an initial training dataset of molecules, we rank the molecules based on the multiple molecular properties of interest using the Pareto ranking scheme described in Equation 1. A detailed description of the algorithm for multi-objective ranking of the molecules can be found in the [experimental methods](#). To fine-tune the generative model's latent space to make it more sampling efficient for desirable molecules, the molecules in the dataset are weighted according to Equation 2, where higher-ranked molecules are assigned with larger weights, while lower-ranked molecules are assigned with smaller weights. Retraining the generative model based on this weighted dataset biases the model toward higher-ranked molecules with more desirable properties, thereby making the latent space of the retrained model more amenable to suggesting novel molecules that are likely to be highly ranked based on the Pareto ranking scheme. After retraining the baseline model using the weighted dataset for a single epoch, we explore the latent space to search for new molecules that can potentially improve upon the molecules in the dataset in terms of the multiple target properties. While various multi-objective optimization schemes may be adopted for this purpose, we considered two potential approaches in this study: (1) random sampling and selection of the top-ranked molecules and (2) single-objective BO (SOBO). In the first approach, we generated 250 random molecules, of which we selected the top 50 (based on Pareto ranking). In the second approach, we used SOBO<sup>24</sup> to generate 50 molecules,



**Figure 1. Overview of the proposed multi-objective latent space optimization scheme**

The initial JT-VAE model is trained based on the original training dataset (step 1). The weights of the molecules in the dataset are adjusted according to their Pareto front ranking based on the properties of interest. Desirable molecules with a higher ranking are assigned larger weights, while molecules with a lower ranking are assigned smaller weights. The JT-VAE is retrained based on the reweighted dataset (step 2). The retrained model is used to suggest novel molecules with enhanced properties by sampling or optimization in the latent space (step 3). Top molecules are selected and used to augment the current training dataset (step 4). Steps 2–4 may be repeated for iterative retraining of the generative model.

cient (logP), (2) synthetic accessibility score (SAS), (3) natural product-likeness score (NP score),<sup>25</sup> and (4) the probability of inhibition against DRD2.<sup>26</sup> We aimed to maximize logP, NP score, and DRD2 inhibition property. On the other hand, we aimed to minimize SAS, as a lower SAS indicates better synthesizability of a given molecule. For the computation of logP and SAS, the RDKit package<sup>27</sup> was used.

We adopted the method in Ertl et al.<sup>25</sup> for assessing the NP score. The probability of inhibition against DRD2 was estimated by using a machine learning (ML) surrogate model, whose details are given in the [experimental methods](#).

### Weighted retraining via Pareto front rank effectively shifts the latent space for multiple objectives

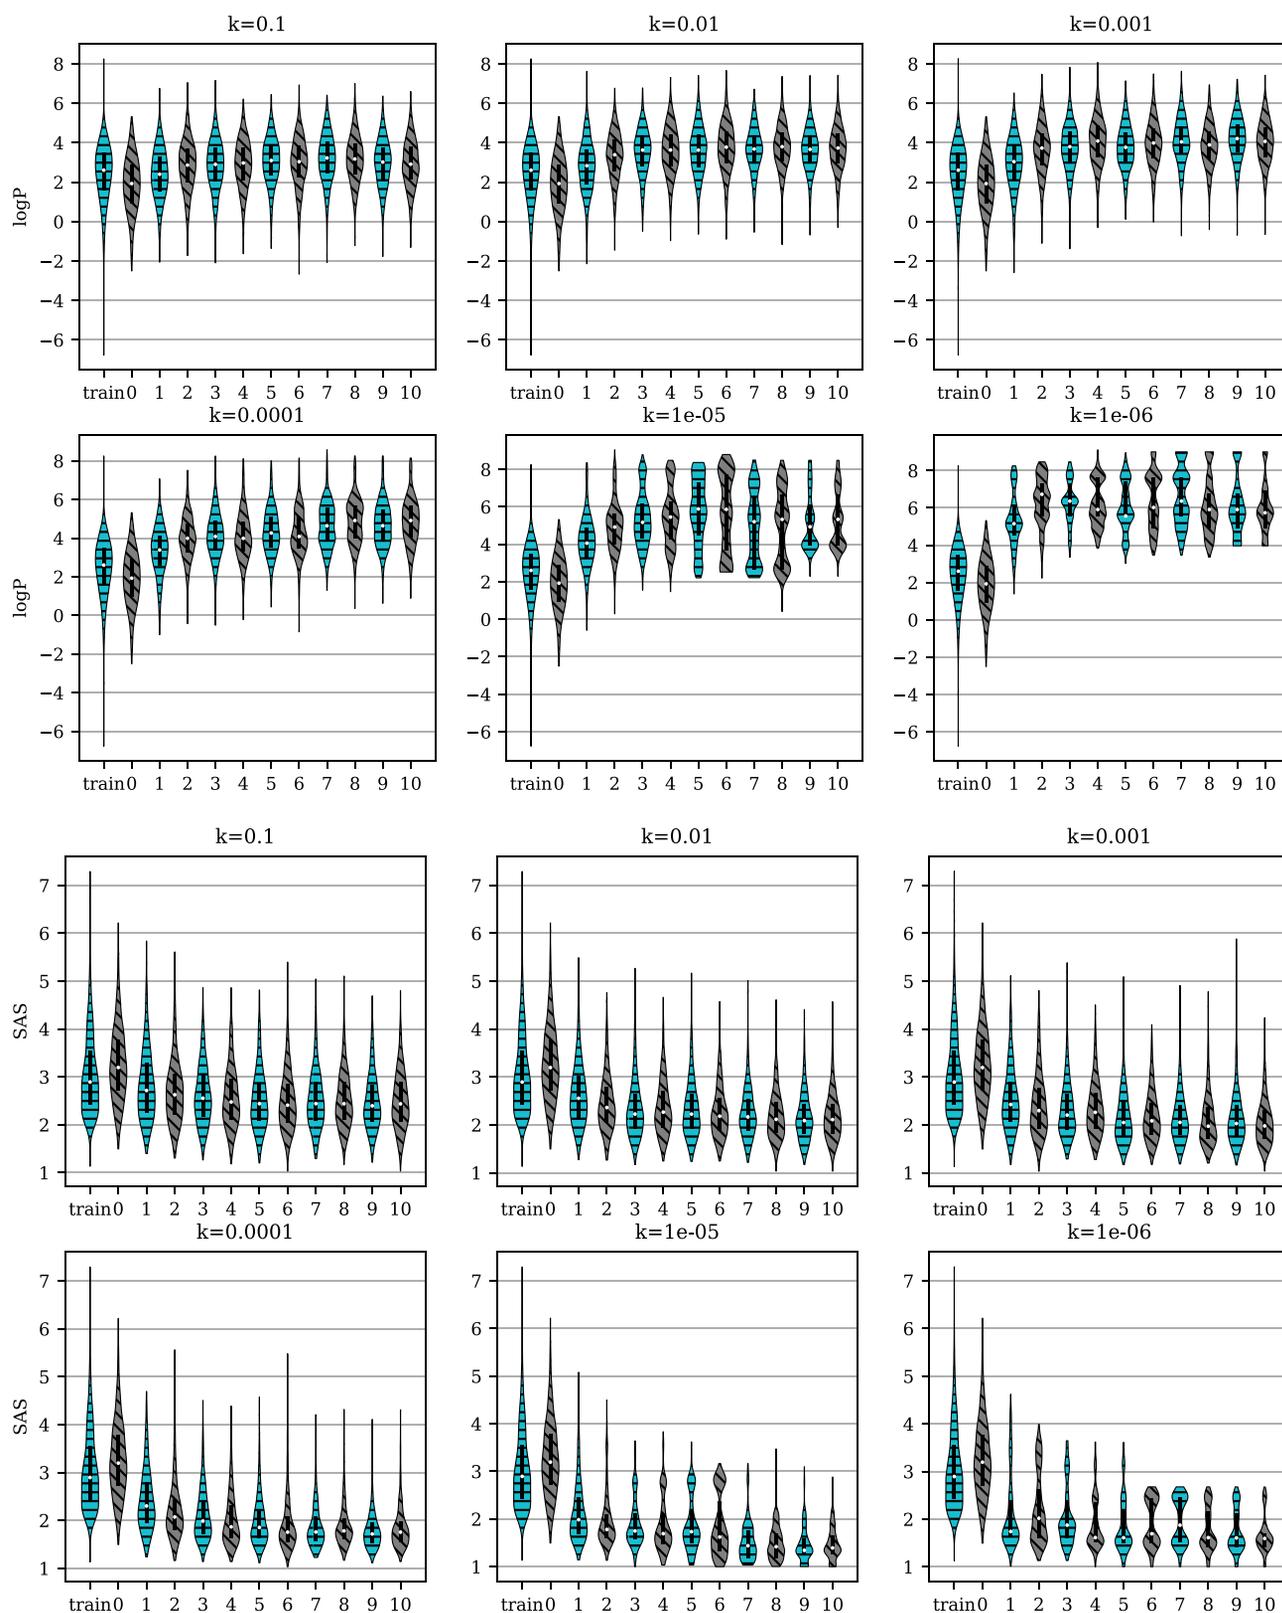
As the baseline model, we used the pre-trained JT-VAE shared by Tripp et al.<sup>19</sup> and applied the proposed multi-objective weighted retraining scheme. Initially, we used the complete ZINC dataset<sup>28</sup> for weighted training of the baseline model, where the dataset was split into training (218,969 molecules) and validation (24,333 molecules) sets as in Tripp et al.<sup>19</sup> After each weighted retraining step, 250 new molecules were randomly sampled, of which the top  $r = 50$  candidates were selected based on the Pareto front rank. The selected top candidates were used in the subsequent retraining stages. For details, please see the [experimental methods](#).

Figure 2 shows the evolution of the property distribution of the generated molecules as we iterate the weighted retraining cycles. Each plot shows the distributional changes over multiple iterations for a specific value of  $k$ , the hyperparameter that determines how the rankings translate into the weights. Results are shown for optimizing the latent space of the JT-VAE for enhancing the property pair (logP, SAS). Here we retrained the baseline model 10 times for several different values of  $k$  ranging from 0.1 to  $10^{-6}$ . Based on a given model at a specific iteration in the retraining cycle, we collected 1,000 molecules randomly sampled from the latent space to plot the distribution of the molecular property of interest. The x axis refers to the property distribution after the  $i$ -th iterative weighted retraining, where “train”

of which the unique molecules were selected. To evaluate and rank the novel molecules, the latent points were first decoded into the original molecular space, where their properties were assessed. The selected top novel molecules were then used to form a “candidate dataset” that could be used to augment the training dataset at hand, potentially pushing the current Pareto frontier and further improving the sampling efficiency of the current model for improved molecules.

Before the next retraining iteration, we create an updated training set by combining the selected top candidates with 10% of the initial dataset, which consisted of randomly selected molecules. The random down-selection of the molecules in the initial training data mainly aims at reducing the computational cost needed for shifting the latent space toward desirable directions within fewer iterations. The updated training set is then used for the weighted retraining of the generative model, which is subsequently used for identifying a new set of desirable molecules to be appended to the candidate dataset. This iterative retraining cycle—new candidate generation based on the current model, augmentation of the candidate set and creation of a new training set that integrates the additional candidates, re-ranking of the molecules, and performing another weighted retraining of the model—can be repeated until either the generated molecules meet the desired multi-objective criteria and converge in terms of the molecular properties or the total training cost (computation or time) reaches a pre-specified budget. We further describe the iterative retraining procedure in later subsections based on specific molecular optimization scenarios.

In this study, we considered the pairwise optimization of the following molecular properties: (1) water-octanol partition coefficient



**Figure 2. Evolution of the property distribution of the generated molecules due to latent space optimization via iterative weighted retraining**

The plots show how the property distribution changes as a result of weighted retraining of the JT-VAE based on the proposed multi-objective latent space optimization scheme. The latent space of the JT-VAE was optimized to suggest molecules with larger logP and smaller SAS. Results are shown for different values

(legend continued on next page)

corresponds to the property distribution of the molecules in the initial ZINC dataset. Furthermore, iteration 0 corresponds to the property distribution of the molecules sampled from the latent space of the baseline model (without any weighted retraining). These distributions are shown as a reference to show the relative improvement of the molecular property of interest as a result of the weighted retraining procedure.

As can be seen in Figure 2, the logP distribution of the molecules sampled from the latent space tends to shift upward, as desired, as the retraining cycles proceed. Similarly, the SAS values tend to decrease, as desired, indicating that iterative retraining generally improves the overall synthesizability of the molecules suggested by the JT-VAE. As discussed in the [experimental methods](#), a smaller  $k$  places a greater emphasis on higher-scoring molecules. We can see its impact on the weighted retraining results in Figure 2, where using a smaller  $k$  leads to a more rapid and more pronounced shift of the property distribution. However, the use of a smaller  $k$  value makes the overall retraining process dominated by a smaller set of high-scoring molecules, which may have an impact on the overall diversity of the generated molecules and skew the property distribution of the molecules. For example, using  $k = 10^{-5}$  or  $k = 10^{-6}$  results in bimodal (or multi-modal) property distributions after several iterations of weighted retraining, reflecting this phenomenon. Here, the model learns the latent space mainly based on a limited number of high-scoring molecules at the Pareto front. Consequently, the molecules sampled in the learned latent space become clustered around those high-ranking molecules, which may limit the diversity of the molecules generated by the retrained model.

Table S1 shows the average property values for the molecules in the training data, molecules generated using the initial model (based on 1,000 randomly generated molecules), and molecules generated by the retrained model (again based on 1,000 randomly generated molecules) for different values of  $k$ . In addition, we evaluated the structural diversity (Table S1) of the molecules generated by different versions of JT-VAE models, to assess the ability of a given model to learn, represent, and sample from a wider chemical space.

The diversity was measured in terms of the average structural distance (based on extended connectivity fingerprint [ECFC4]) over all pairs in a given set of molecules. For smaller values of  $k$ , we observe that the structural diversity is reduced as expected. As mentioned earlier, using a smaller  $k$  assigns relatively higher weights to a smaller group of high-ranking molecules, which has the effect of making the model “see” this small group of molecules more frequently while retraining. Consequently, the model learns the latent space representation of the chemical space mainly based on these select molecules, which may make the molecules sampled from the learned latent space bear higher similarity to one another.

To further demonstrate the effectiveness of our proposed approach MO-LSO, we performed additional experiments for simultaneous optimization of two or three molecular properties

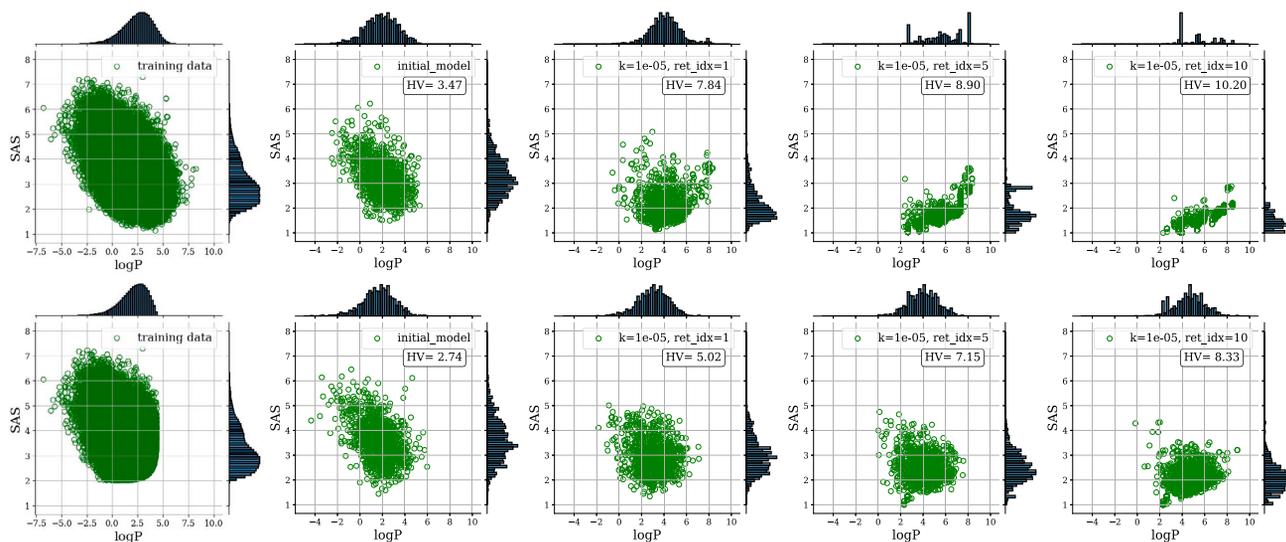
and compared the results against the scalarization baseline of Tripp et al.<sup>19</sup> and a Markov molecular sampling scheme for multi-objective drug discovery, called MARS.<sup>22</sup> Specifically, the latter approach involves training a molecular generative model from scratch targeting multiple properties of interest, whereas our approach starts with a generative model that is trained in a self-supervised fashion without the molecular property values of the training samples. However, through iterative retraining via our approach, the generative model succeeds in outperforming MARS for quite a few property combinations (Table S6). Furthermore, we investigated the impact of retraining on the JT-VAE's molecule reconstruction performance. The results obtained from these experiments are summarized and discussed in the [supplemental information](#).

### Multi-objective latent space optimization effectively recovers sampling efficiency for incomplete dataset

Next, we investigated the ability of the weighted retraining scheme to propose high-scoring molecules when such molecules are absent in the training dataset. For this purpose, we first removed the top 20% molecules—selected based on their Pareto front rank—from the training data that are used to train the baseline model.<sup>19</sup> Since the original baseline model has seen the complete dataset during its training, for each property pair, we trained a separate baseline model based on the reduced dataset that does not contain the top 20% molecules with the highest Pareto front rank for the given property pair. In all cases, a learning rate of 0.0007 was used with a batch size of 32 for 30 epochs for model training. The scatterplots in Figure 3 show the progression of the two properties logP and SAS, where the latent space of the JT-VAE is optimized for the given property pair with  $k = 10^{-5}$ . In each plot, we also show the hypervolume of the property space that is dominated by the Pareto front. The hypervolume is computed with respect to the average property values of the molecules in the complete training dataset. The top row of Figure 3 shows the evolution of the Pareto front when the complete dataset was used. On the other hand, the bottom row in Figure 3 shows the trends for the reduced dataset, which does not contain the top 20% molecules. As can be seen in the bottom row of Figure 3, although the initial dataset does not contain many molecules with logP greater than 4 and SAS lower than 2, the iterative multi-objective weighted retraining still manages to effectively push the latent space toward a desirable region that contains high-scoring molecules, where the trends are similar to the case when the complete dataset is used. This is also illustrated by the larger hypervolume achieved by the optimized model compared to the initial pre-trained model in both cases.

The capability of the proposed MO-LSO method to recover high-performance molecules—despite the absence of such molecules in the training data—is demonstrated even more clearly when applied to the optimization of inhibitory molecules for DRD2. In this experiment, we considered pairwise property optimization of DRD2 inhibition along with one of the properties

of  $k$ , which determines the sensitivity of the weight to ranking. In each graph, the first violin plot (labeled “train”) shows the property distribution of all molecules in the initial training dataset. The subsequent violin plots show the property distribution of 1,000 randomly sampled molecules after  $i$ -th iterative retraining ( $i = 0$  corresponds to the original JT-VAE without any retraining). The results clearly show that the distribution of logP is shifted upward, while that of SAS is shifted downward during the iterative retraining process, as desired.



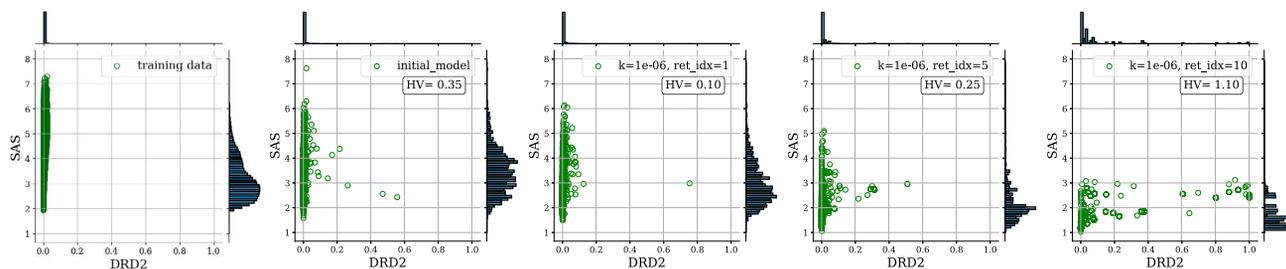
**Figure 3. Evolution of the Pareto front via latent space optimization of the generative model**

The latent space of the JT-VAE has been jointly optimized to maximize logP and minimize SAS of the molecules suggested by the generative model. The scatterplots show the (logP, SAS) distribution of the molecules in the initial training dataset (column 1), molecules sampled in the latent space of the baseline model (column 2), and molecules suggested by the optimized model after iteration 1 (column 3), iteration 5 (column 4), and iteration 10 (column 5). The plots in the top row show the trends for the case when the complete training dataset was used, while the bottom row shows the trend when a reduced dataset was used. The results show that the Pareto front gradually shifts toward the desired direction (i.e., bottom right for larger logP and smaller SAS) resulting in a larger hypervolume (HV) of the Pareto-front-dominated property space in both cases.

among logP, SAS, and NP score. As before, in each experiment for a given property pair, we removed the top 20% molecules from the training data based on the Pareto front rank. What makes this experiment especially interesting is the fact that the training dataset is highly imbalanced and contains a relatively small number of inhibitory molecules against DRD2. As a result, the removal of the top 20% molecules leaves virtually no active DRD2 inhibitors in the training data. Moreover, since we train a new baseline model with the reduced dataset, the trained model does not initially possess any knowledge of inhibitory molecules against DRD2. Consequently, the MO-LSO method needs to guide the optimization of the latent space of the generative model toward a completely unexplored region, making the optimization task more challenging.

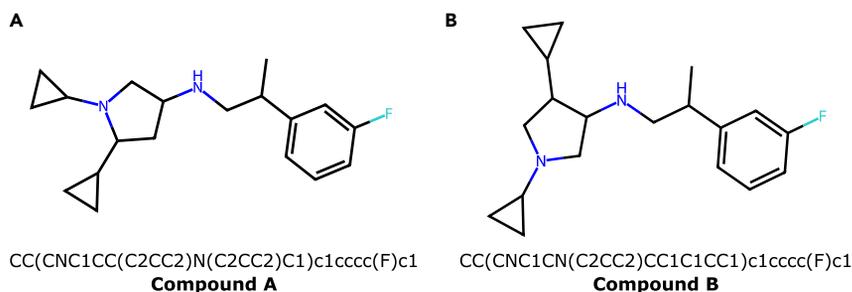
Figure 4 shows the results for the multi-objective weighted retraining with  $k = 10^{-6}$  for the property pair DRD2 and SAS. The

leftmost scatterplot shows the distribution of the training data after removing the top 20% molecules. The horizontal axis shows the probability of inhibition (the higher the better) and the vertical axis shows the SAS (a lower score corresponds to better synthesizability). The second subplot (from the left) shows the property distribution of 1,000 molecules randomly sampled in the latent space of the baseline model. Next, the third, fourth, and fifth (i.e., rightmost) subplots depict the property distribution of the molecules after the first, fifth, and tenth weighted retraining cycle, respectively. While the initial training data are lacking DRD2-inhibitory molecules, we can see a relatively large number of inhibitory molecules after the tenth weighted retraining. Moreover, the SAS distribution of the sampled molecules becomes more skewed toward smaller values, as expected. We have repeated similar experiments for two other property pairs, (DRD2, logP) and (DRD2, NP score), which all showed similar



**Figure 4. Transition of DRD2 and SAS toward optimum direction while starting with no DRD2-active training samples**

The first two scatterplots represent the training data and the molecules generated from the baseline model, respectively. In the rest of the jointplots, the molecules from the learned model after the first, fifth, and tenth weighted retrains for  $k = 10^{-6}$  are shown in the objective space. The hypervolume (HV) in each plot indicates the volume of the property space that is dominated by the Pareto front. As we iterate the weighted retraining, the resulting hypervolume tends to increase.



**Figure 5. Molecular structure and SMILES of two generated compounds with predicted DRD2 activity**

Both compounds A and B show the lowest docking energy among the pool of molecules that we generated by considering the DRD2 activity as one of the objectives. The weighted retraining framework starts from the initial training dataset that does not contain any active molecules.

trends. Details of these simulation results can be found in [Table S2](#).

### **In silico analysis of the designed DRD2 inhibitors**

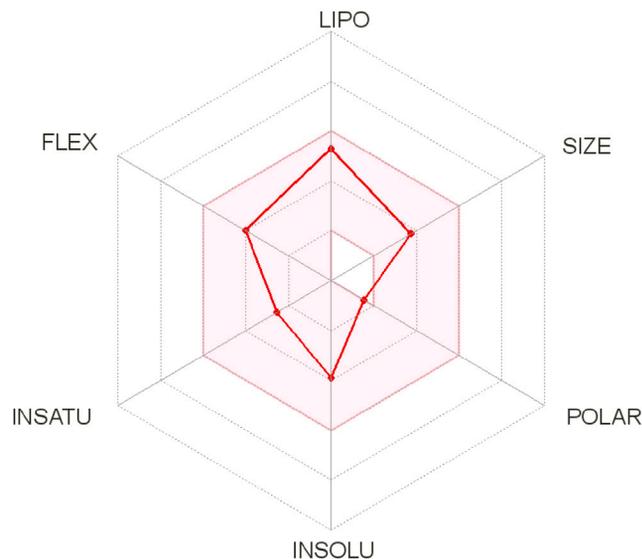
Given the long history of structure-based rational design, the use of methods that are structurally unaware to design “active” molecules can be quite disconcerting. However, the use of support vector machine (SVM) model training with thousands of active and inactive compounds that sample the available contacts in the pocket seems to have encoded the space well enough that, while the method is not structurally aware, the information required is implicitly available. Potentially of more value is removing the inherent bias in what an “active should look like” seen when working with computational and medicinal chemists. In this study, tossing out the concept of what a hit should look like and taking a data-driven generative approach for molecular design led to some strange-looking molecules that have attractive properties and, by multiple measures of *in silico* modeling (docking, molecular mechanics generalized Born surface area [MM-GBSA], and long-duration molecular dynamics [MD]), were superior to known DRD2 inhibitors. For example, one of the designed compounds, CC(CNC1CC(C2CC2)N(C2CC2)C1)c1cccc(F)c1 (which we refer to as “compound A”), has energies  $-3.43$  kcal/mol lower than the crystal ligand by docking and  $-3.78$  with MM-GBSA. Furthermore, MD simulations showed stable, unstrained interactions over the length of the simulation.

An MD system was created using Dibenzy 3,5-pyridinedicarboxylate (DPPC) for the membrane, TIP4PEW water, and physiological levels of NaCl. Following 100 ps of minimization, the system was run using Desmond and the OPLS4 force field at 310.15K for 15 and 100 ns to evaluate the movement and ligand pose stability. For docking, Glide grids were created from the system at time 14.98 ns as well as from the original crystal structure, and all residues within 5 Å of the ligand were set to allow rotation. Glide XP was used for docking since, while computationally very expensive, the sampling and pose minimization seem to best replicate the crystal position of the native ligand.

The generated novel molecules with predicted DRD2 activity were prepared using Ligprep in Schrodinger 2021.3, with the molecules enumerated around chiral centers in a pH range of  $7.4 \pm 2.0$ . The crystal structure of 6CM4 was prepared with the protein preparation tool in Schrodinger 2021.3 at a pH of 7.4. The generated ligands were docked along with 667 inactive bait compounds as well as two non-reverse agonist DRD2 ligands as controls due to the possibility that the risperidone bound structure would not be favorable to the docking of tradi-

tional antagonists. This concern was unfounded, since both domperidone and L-741626 were able to achieve low-energy poses. From the docking, two compounds—compound A (shown in [Figure 5A](#)) and another compound, which we refer to as compound B (shown in [Figure 5B](#))—achieved the lowest energies of any compound, including the crystal ligand. MD simulations using the same parameters as before were run for 15 ns on compounds A and B to evaluate whether the confirmations predicted were stable (was the compound ejected from the pocket) and if relevant contacts were maintained. In addition, the Schrodinger prime MM-GBSA calculations were carried out on all compounds/poses with docking scores the same as or equal to the lowest scoring pose of the crystal ligand.

While *in silico* results do not often translate into activity *in vivo*, they are often used to select what molecules get made and their priority for testing, so the results of this effort and the computational chemistry evaluation produced “high-priority compounds” that are very different from the native ligand and the molecules on the market. This ability to explore *in silico* vast amounts of novel chemistry space while still generating compounds with desirable properties is likely to lead to the types of cost and time reductions in discovery that have been promised by the AI/ML community. The actual compounds predicted have desirable properties ([Figure 6](#) shows the properties for compound A as predicted by SwissADME<sup>29</sup>). In addition, the binding of compound A ([Figure 7A](#)) shows the typical aromatic interaction between W386 and the ligand, in this case, the six-membered ring on the ligand, as well as the potential for salt bridges from the ligand Ns to the Os on D114. Of more interest is the two cyclopropane groups, which allow the compounds to have hydrophobic interactions with the largely hydrophobic mouth of the pocket, which includes residues 389, 392, 184, 189, 416, 412, etc. ([Figure 7B](#)), without having undesirable steric effects that bulkier ring systems may present. The potential desolvation effects from these residues, while still allowing for the protein to have a closed confirmation, have the potential to allow the compounds to be selective and active at low concentrations. As previously mentioned, this molecular series would not have been ranked high by computational or medicinal chemists based on how they looked, but once the poses were reviewed and the ease of the synthetic routes evaluated, it is generally agreed that the series is a high priority. For illustrative purpose, we performed the ADMET analysis for the top candidate for DRD2. In practice, if one desires to select the optimal candidate leads for the properties of interest, a more comprehensive screening like the one performed in Feng et al.<sup>23</sup> may be required.



**Figure 6. Properties of a top molecule predicted by the optimized generative model**

The properties of a top molecule (compound A) suggested by the JT-VAE, whose latent space was optimized by the proposed method. Six parameters—POLAR (polarity), INSOLU (insolubility), INSATU (instauration), FLEX (rotatable bond flexibility), LIPO (lipophilicity), and SIZE (molecular weight)—are shown. We can see that the suggested compound is within the colored zone, which corresponds to the physicochemical space suitable for oral bioavailability.<sup>29</sup>

### Molecule optimization and selection via Bayesian optimization can enhance latent space optimization results

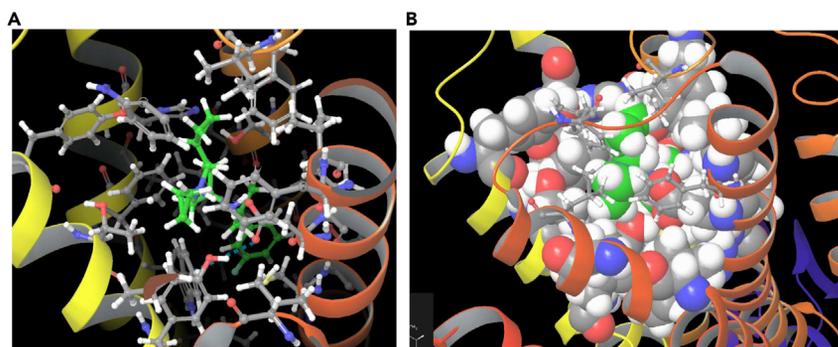
In previous sections, we generated novel molecules through random sampling in the latent space of the trained generative model, from which the top molecules were selected to augment the training data for the next cycle of weighted retraining. Although we adopted random sampling to demonstrate how the proposed multi-objective weighted retraining can enhance the sampling efficiency of the generative model for multiple target molecular properties, the overall efficacy of the GMD can be further improved by leveraging more sophisticated optimization techniques—such as BO, GAs, and PSO<sup>6</sup>—to optimize the molecules in the latent space. To demonstrate this, we utilized BO to optimize molecules in the latent space of the generative model for logP and SAS. A single-objective function was defined through scalarization, where logP (to be maximized) was penalized by SAS (to be minimized) after standardizing both values based on the mean and standard deviation of the respective values in the initial dataset. Based on the expected improvement acquisition function, we generated 50 molecules, and only unique samples were added to the training dataset. Compared to the random generation strategy, the BO approach is computationally expensive, since it requires training the GP surrogate model as well as optimizing the acquisition function. However, the BO approach is more sample efficient, as it requires fewer property predictions to find the best data points to augment the training dataset. On the other hand, the random-generation approach requires a larger number of evaluations (250 in our case) to select the best 50 candidates for data augmentation.

Figure 8 shows the results for  $k = 10^{-3}$ . As shown in the top row, iterative weighted retraining with data augmentation through random sampling continues to shift the latent space distribution toward the desired direction, although not significantly, even after the 10th iteration. On the other hand, data augmentation through BO shifts the latent space distribution much more effectively, as can be seen in the bottom row of Figure 8. This is especially dramatic for logP, as after the 10th iterative retraining, the retrained generative model is capable of generating molecules with remarkably higher logP compared to those in the original training data. The candidate molecules suggested by the generative model tend to be somewhat biased toward the higher logP region with slightly higher SAS, which is likely an artifact due to the use of SOBO in this experiment, where the objective function was defined as a linear combination of logP and SAS. This may be addressed through different scalarization or the use of multi-objective BO (MOBO). Nonetheless, the example in Figure 8 clearly shows the potential advantage of utilizing BO (or other advanced optimization schemes) for effective molecular optimization in the latent space and thereby more effectively enhance the sampling efficiency of the generative model through the proposed MO-LSO approach.

## DISCUSSION

The presented framework for multi-objective optimization illustrates the potential of weighted retraining in bringing the molecule-generation model into the expected multi-dimensional-objective region. The JT-VAE model's latent space is optimized for different pairs of molecular properties, and the effects of ranks are studied. We have found that the weight formulation from the ranks dictates a trade-off between the diversity and the shifts of the property distribution in the latent space. To verify the strength of our approach, similar experiments were repeated with relatively poor training data. Even starting with no DRD2-active molecules, weighted retrained models still managed to configure its latent space for the active region. These outcomes are more pronounced given the fact that random selection from the latent space is used to propose the candidate molecules for retraining stages. To speed up the reshaping of the latent space, the BO is shown to be promising.

The ranking scheme of our framework is contingent on the robust property predictors. How well these surrogate models can tackle the unexplored chemical space is critical. Even if these models are less accurate, the retrained latent space can still be the exploration field for further screening. Given a reliable and fast approximate mapping from a molecule to its property value, the weighted retraining approach can optimize the latent space jointly for more practical properties that are responsible for a higher attrition rate of proposed drugs. With the availability of surrogate models like protein-ligand binding score<sup>30</sup> and inhibition of bile salt export pump,<sup>31</sup> our approach can optimize the latent space in producing candidate drugs that are most likely to be active against a specific target without causing possible damage to patients. In terms of computational cost, retraining the generative network multiple times may be slightly expensive for a larger network compared to the one we have used. However, with the application of distributed training,<sup>32</sup> the training time



**Figure 7. Binding pose of compound A in the DRD2 structure**

(A) The ball-and-stick representation shown highlights how the compound fits the broader pocket. (B) The space-fill model highlights the interaction between cyclopropanes and the hydrophobic mouth of the receptor.

can be significantly mitigated. Moreover, we have seen a decreasing trend in the diversity of the molecules of the latent space with an increase in the shift of the property distribution. A diversity-oriented candidate selection strategy could be a better answer to this issue.

As shown in Figure 8, the latent space optimization outcomes can be enhanced by adopting more effective sampling strategies in the latent space that can identify and suggest novel molecules with more desirable properties. In this study, this was demonstrated by replacing the random sampling scheme with BO. In fact, it may be possible to further improve the GMD through the proposed MO-LSO scheme by incorporating more sophisticated optimization techniques that can effectively explore the unknown landscape of the multiple-objective functions under immense uncertainties. Optimal experimental design (OED)<sup>33–36</sup> and active learning<sup>37–39</sup> techniques that build on objective-based uncertainty quantification (objective-UQ) based on MOCU (mean objective cost of uncertainty)<sup>40,41</sup> may provide practical solutions for such “uncertainty-aware” sampling in the latent space. These are topics of our ongoing investigation.

## EXPERIMENTAL PROCEDURES

### Resource availability

#### Lead contact

Byung-Jun Yoon is the lead contact for this study and can be reached at [bjyoon@ece.tamu.edu](mailto:bjyoon@ece.tamu.edu).

#### Materials availability

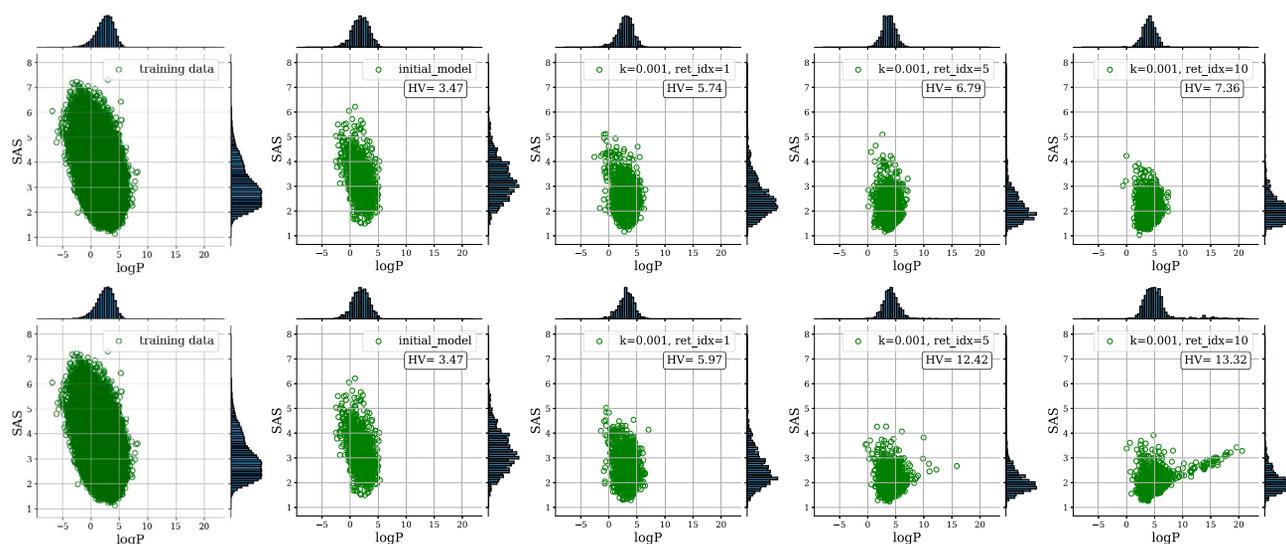
This study did not generate any physical materials.

#### Data and code availability

The specific training and validation split of the ZINC dataset provided in <https://github.com/cambridge-mlg/weighted-retraining> was used in this work for retraining JT-VAE. The source code of the proposed MO-LSO method can be downloaded from <https://github.com/nafizabeer/GMD-MO-LSO> and Zenodo.<sup>42</sup>

### Generative molecular design using the junction-tree variational autoencoder

In this study, we have used the JT-VAE for investigating the MO-LSO of deep generative models for GMD. While various deep generative models with latent space molecular representation have been proposed to date, JT-VAE is widely known for its high reconstruction accuracy when decoding latent samples into the original molecular space. Compared to other VAE models, where novel molecules sampled in the latent space often fail to decode into legitimate molecules, JT-VAE effectively addresses this issue by decomposing molecular



**Figure 8. Effect of Bayesian optimization on the progression of the model in the objective space**

The first and second rows represent the scatterplots for random selection and Bayesian optimization strategy, respectively. In each row, the objective space is for the training data and the molecules from the baseline model and retrained model after the first, fifth, and tenth weighted retrains for  $k = 10^{-3}$ . Higher increase in hypervolume (HV) for the Bayesian optimization strategy shows its effectiveness over random selection.

**Algorithm 1. Find the Pareto front by identifying the set of non-dominated data points**

```

Require: N data points with their objective scores
Initialize  $P' = \{1, 2, 3, \dots, N\}$   $\triangleright$  Set of non-dominated points,  $P'$ 
and  $i \leftarrow 1$ 
while  $i \leq |P'|$  do
  Initialize  $k \leftarrow 0$ 
  for (each  $j \in P' \wedge j \neq i$ ) do
    if  $(x^{(j)})$  does not dominate  $x^{(P'(i))}$  in any objective score) then
       $P' \leftarrow P' \setminus \{j\}$ 
    else if  $j < i$  then
       $k \leftarrow k + 1$ 
    end if
   $i \leftarrow k + 1$ 
end for
end while

```

graphs into a junction tree of chemical substructures. To suggest novel molecules, JT-VAE reconstructs the junction tree from the sampled latent embedding and assembles the chemical substructures into a molecular graph.<sup>8</sup> The decoded tree structure serves as a scaffold guiding the generation of the molecular graph to reconstruct the molecule, resulting in a high fraction of valid molecules. However, we note that the proposed MO-LSO scheme can be applied to various other types of VAEs<sup>5,17</sup> in a straightforward manner without any modification.

**Non-dominated sorting and Pareto ranking**

In a multi-objective optimization problem, there may be no single solution that is optimal in terms of every objective. In practice, different objectives may conflict with one another, where optimizing one objective may result in a suboptimal solution for one or more other objectives. The concept of Pareto optimality provides an effective way of addressing this issue and is widely utilized in the context of multi-objective optimization. Instead of finding the solution that optimizes a single objective, Pareto optimization aims to identify the Pareto optimal set, which is defined as the collection of all solutions that are not dominated by any other solution in the feasible solution space.<sup>43</sup> Consider the problem of jointly optimizing  $K$  objective functions  $f_1(\mathbf{x}), \dots, f_k(\mathbf{x})$ . Without loss of generality, we assume that the goal is to maximize all  $K$  objective functions. Let  $\mathbf{x}^i$  and  $\mathbf{x}^j$  be two points in the solution space.  $\mathbf{x}^i$  is said to dominate  $\mathbf{x}^j$  if  $\mathbf{x}^i$  is as good as  $\mathbf{x}^j$  in terms of all  $K$  objectives (i.e.,  $f_k(\mathbf{x}^i) \geq f_k(\mathbf{x}^j), \forall k = 1, \dots, K$ ) and if there is at least one objective such that  $\mathbf{x}^i$  outperforms  $\mathbf{x}^j$  (i.e.,  $\exists k$  s.t.  $f_k(\mathbf{x}^i) > f_k(\mathbf{x}^j)$ ). The Pareto optimal points  $\mathbf{x}$  in the solution space form the Pareto front (or Pareto frontier) in the objective space.

Considering all objectives simultaneously, all points in the Pareto front are equivalent to one another, as no point is either more preferable or less preferable than the others. As no point dominates any other point in the Pareto front, these points in the Pareto front may be assigned the same ranking. The process of finding the Pareto front is summarized in Algorithm 1. Once the Pareto front is identified, we may remove these Pareto optimal points  $\mathcal{P}_1$  from the dataset—i.e., “peel off” the first Pareto front—and move on to identify the next Pareto front  $\mathcal{P}_2$  in the remaining dataset. This process may be repeated until all points in the dataset  $D$  are exhausted to find all possible Pareto optimal sets  $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3, \dots, \mathcal{P}_S$ . Note that the sets  $\mathcal{P}_1, \dots, \mathcal{P}_S$  form a partition of  $D$  such that  $\mathcal{P}_i \cap \mathcal{P}_j = \emptyset$  for  $i \neq j$  and  $D = \bigcup_j \mathcal{P}_j$ . This Pareto ranking process is

summarized in Algorithm 2 and illustrated in Figure 9. Based on these results, we may rank all data points such that all points in the  $j$ -th Pareto front are assigned the following ranking:

$$rank_D(\mathbf{x}) = \sum_{i=1}^{j-1} |\mathcal{P}_i| \quad \forall \mathbf{x} \in \mathcal{P}_j. \quad (\text{Equation 1})$$

Despite practical differences from Equation 1, it is worth noting that similar multi-objective ranking schemes based on the concept of Pareto optimality have been previously explored in different contexts. For example, the work by Obayashi et al.<sup>44</sup> adopted a similar ranking scheme for a multi-objective GA (MOGA) to identify solutions within the population of GA solutions that are nearly Pareto optimal.

**Weighted retraining of generative models based on Pareto ranking**

We adopt the ranking scheme in Equation 1 for weighted retraining of the generative model (i.e., the JT-VAE in this study) to steer its latent space toward a region that is more sampling-efficient for enhanced molecules with multiple target properties. More specifically, we extend the weighted retraining scheme in Tripp et al.<sup>19</sup>—originally designed for latent space optimization based on a single objective—to enable flexible MO-LSO regardless of the number of objectives without any *ad hoc* scalarization of the objective function. We calculate the weight for every data point  $\mathbf{x}$  in the dataset  $D$  as follows:

$$w(\mathbf{x}, k, D) = \frac{1}{kN + rank_D(\mathbf{x})}, \quad (\text{Equation 2})$$

where this weight  $w(\mathbf{x}, k, D)$  determines the influence of a given data point  $\mathbf{x}$  on the training loss. As a consequence, a more desirable molecule with a higher Pareto ranking (i.e., a smaller  $rank_D(\mathbf{x})$ ) is assigned a larger weight  $w(\mathbf{x}, k, D)$ , thereby playing a more important role in retraining the generative model.  $k$  is a hyperparameter that adjusts the influence of the rank on the computed weight. A larger  $k$  makes the weight distribution more uniform, while a smaller  $k$  assigns large weights to relatively fewer high-rank data points.  $N = |D|$  is the cardinality of the training set  $D$ .

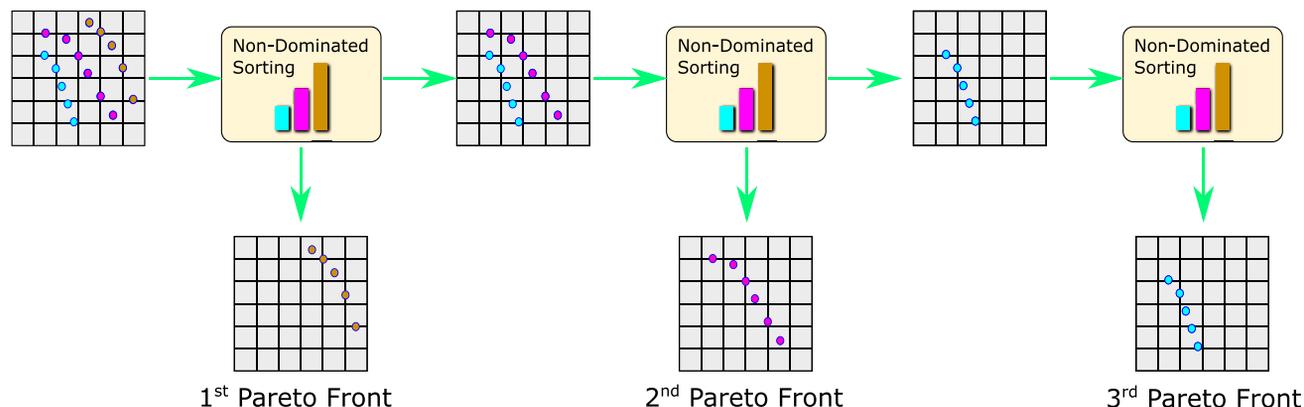
The iterative weighted retraining is performed as follows. We first start with the initial training dataset,  $D_{train} = D_0$ , and a null set,  $D_{new} = \emptyset$ . We retrain the model based on the given  $D_{train}$  after reweighting every data point (i.e.,

**Algorithm 2. Pareto front ranking of data points**

```

Require: N data points with their objective scores
Initialize  $P = \{1, 2, 3, \dots, N\}, j = 1$ 
while  $|P| \neq 0$  do
  Step 1: Find the non-dominated set,  $P'$  from  $P$  using Algorithm 1
  Step 2:  $P_j \leftarrow P', P \leftarrow P \setminus P'$  and  $j \leftarrow j + 1 \triangleright P_j$ :  $j^{\text{th}}$  Pareto front
end while

```



**Figure 9. Illustration of the Pareto front ranking process**

Suppose the main objective is to design novel molecules such that two target properties are jointly maximized. For weighted retraining of the generative model, we first identify the Pareto optimal molecules that are located on the Pareto front of the current dataset. All molecules in this first Pareto front are ranked 1. Next, these Pareto optimal molecules are removed from the dataset, after which we identify the second Pareto front among the remaining molecules. The Pareto optimal molecules in the second Pareto front are ranked 2. The process of removing the Pareto optimal molecules in the current dataset and the identification of the next Pareto front in the reduced dataset may be repeated until all molecules are ranked.

molecule)  $\mathbf{x} \in D_{train}$  by  $w(\mathbf{x}, k, D_{train})$  given by Equation 2. After the first weighted retraining with  $D_{train} = D_0$ , we randomly generate 250 molecules from the latent space of this model and select the top 20% among the generated molecules, which are then added to the set  $D_{new}$ . Next, we update the training dataset by augmenting it with the selected new molecules:

$$D_{train} \leftarrow \bar{D}_0 \cup D_{new}, \quad (\text{Equation 3})$$

where  $\bar{D}_0$  may be either the initial training dataset  $D_0$  in its entirety or a randomly sampled subset to reduce computation. In this study, we used  $\bar{D}_0$  by randomly sampling 10% of the molecules in the initial training data  $D_0$ . In the next iteration, we can repeat the process by reweighting the molecules in  $D_{train}$ , retraining the generative model, sampling new molecules in the latent space of the retrained model, and then selecting the top 20% molecules and adding them to  $D_{new}$ . The weighted retraining cycle may be repeated a pre-determined number of times or until a stopping criterion is reached. In this study, we repeated the cycle for 10 iterations to investigate the overall impact on enhancing the sampling efficiency of the generative model for suggesting novel molecules that simultaneously improve multiple target properties.

### Molecular property predictors

In this study, we considered four different molecular properties for validating the capability of the proposed MO-LSO scheme and assessing its performance. The tested properties included the logP, synthesizability (SAS), NP score, and inhibition of DRD2.

#### Partition coefficient

As a quantitative measure of lipophilicity, the octanol-water logP is one of the standard properties for selecting potential drugs according to Lipinski's "rule of 5."<sup>45</sup> We used the RDKit's `rdkit.Chem.Crippen`<sup>46</sup> module to get the calculated logP values from SMILES representation.

#### Synthetic accessibility score

The SAS of a molecule serves as a surrogate for quantifying the degree of difficulty in developing it. Although SAS does not account for the additional constraints the medicinal chemist may have in particular laboratory settings, e.g., restriction of using particular reagents, it is still useful in screening from a large number of molecules for further evaluation. We obtained SASs of molecules in our work from the RDKit-based implementation of SAS estimation method. In addition to the fragment score and the complexity penalty as in the original work,<sup>47</sup> this implementation includes the score based on the molecular symmetry.

#### Natural product-likeness score

Through the evolutionary selection process, natural products often contain bioactive substructures that can be utilized in drugs.<sup>48</sup> The NP score for a

molecule quantifies how much similarity its substructure has with the natural products. A higher score indicates that the molecular structure is more likely within the natural product space. We used the implementation from Ertl et al.,<sup>25</sup> which aggregates the individual scores for all the fragments of a molecule.

#### Inhibition of DRD2

DRD2 has a long history of being used as the target protein for antipsychotic drugs.<sup>49</sup> More recent research findings<sup>50–53</sup> demonstrate the effectiveness of DRD2-targeting drugs against a wide range of cancer cells. Thus, searching for DRD2-inhibiting molecules provides us with an opportunity to showcase our proposed approach in a practical drug-discovery-like scenario.

We used an ML surrogate model in Olivecrona et al.<sup>10</sup> to predict the efficacy of a given molecule in inhibiting the activity of DRD2. The surrogate model for predicting the activity against the dopamine type 2 receptor DRD2 is built as a binary SVM classifier with a radial basis function ( $\gamma = 2^{-6}$ ). The Morgan fingerprint, with radius 3 (FCFC6) computed by RDKit<sup>27</sup> is used as the input feature for the SVM classifier. The probability of being active predicted by this model for each molecule is treated as a property to be maximized.

Since there is a class imbalance issue due to the smaller number of active molecules, Olivecrona et al.<sup>10</sup> split the active samples in such a way that the structural similarity among the samples from the train and test/validation dataset is less. We have used the same split dataset consisting of 7,218 active and 100,000 inactive molecules to train, test, and validate the SVM model in Scikit-learn<sup>54</sup> (version 0.23.2) with a regularization parameter  $C = 2^7$ . The overall performance of the DRD2 activity classifier used in this study is summarized in Table 1 for the training, validation, and test sets. Accuracy is defined as the fraction of samples correctly classified to be active or inactive. AUC denotes the area under the ROC (receiver operating characteristic) curve. Precision shows the ratio between correctly classified active molecules and all samples that are predicted to be active. Recall is the fraction of all active samples that are correctly classified.

**Table 1. Performance of the DRD2 activity classifier used in this study**

Dataset	Accuracy	AUC	Precision	Recall
Train	0.9998	0.9999	0.9969	1.0
Validation	0.9807	0.8745	0.9747	0.7498
Test	0.9842	0.9074	0.9770	0.8178

The classifier predicts whether a given molecule may be an effective inhibitor of DRD2 ("active") or not ("inactive").

## SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.patter.2024.101042>.

## ACKNOWLEDGMENTS

This work represents a multi-institutional effort. Funding sources include the following: federal funds from the National Cancer Institute, National Institutes of Health, and Department of Health and Human Services; Leidos Biomedical Research Contract no. 75N91019D00024; and Task Order 75N91019F00134 through the Accelerating Therapeutics for Opportunities in Medicine (ATOM) Consortium under CRADA TC02349. This work was supported in part by Brookhaven National Laboratory (BNL) LDRD no. 21-044. Portions of this research were conducted with the advanced computing resources provided by Texas A&M High Performance Research Computing (HPRC).

## AUTHOR CONTRIBUTIONS

F.J.A., B.-J.Y., and N.M.U. initiated the project. B.-J.Y. and N.M.U. proposed the idea. A.N.M.N.A. developed, refined, and evaluated the method. M.R.W. evaluated the DRD2 activity of the predicted molecules. A.N.M.N.A. and B.-J.Y. analyzed the results and wrote the paper. All authors have edited and verified the paper.

## DECLARATION OF INTERESTS

B.-J.Y. is a member of the advisory board of *Patterns*.

Received: January 18, 2024

Revised: June 27, 2024

Accepted: July 18, 2024

Published: August 12, 2024

## REFERENCES

- Muratov, E.N., Bajorath, J., Sheridan, R.P., Tetko, I.V., Filimonov, D., Poroikov, V., Oprea, T.I., Baskin, I.I., Varnek, A., Roitberg, A., et al. (2020). Qsar without borders. *Chem. Soc. Rev.* 49, 3525–3564. <https://doi.org/10.1039/D0CS00098A>.
- Kerns, E.H. (2001). High throughput physicochemical profiling for drug discovery. *J. Pharmaceut. Sci.* 90, 1838–1858. <https://doi.org/10.1002/jps.1134>.
- Woo, H.-M., Qian, X., Tan, L., Jha, S., Alexander, F.J., Dougherty, E.R., and Yoon, B.-J. (2021). Optimal decision making in high-throughput virtual screening pipelines. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2109.11683>.
- Woo, H.-M., Allam, O., Chen, J., Jang, S.S., and Yoon, B.-J. (2023). Optimal high-throughput virtual screening pipeline for efficient selection of redox-active organic materials. *iScience* 26, 105735. <https://doi.org/10.1016/j.isci.2022.105735>.
- Gómez-Bombarelli, R., Wei, J.N., Duvenaud, D., Hernández-Lobato, J.M., Sánchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T.D., Adams, R.P., and Aspuru-Guzik, A. (2018). Automatic chemical design using a data-driven continuous representation of molecules. *ACS Cent. Sci.* 4, 268–276. <https://doi.org/10.1021/acscentsci.7b00572>.
- Winter, R., Montanari, F., Steffen, A., Briem, H., Noé, F., and Clevert, D.-A. (2019). Efficient multi-objective molecular optimization in a continuous latent space. *Chem. Sci.* 10, 8016–8024. <https://doi.org/10.1039/C9SC01928F>.
- Griffiths, R.-R., and Hernández-Lobato, J.M. (2020). Constrained bayesian optimization for automatic chemical design using variational autoencoders. *Chem. Sci.* 11, 577–586. <https://doi.org/10.1039/C9SC04026A>.
- Jin, W., Barzilay, R., and Jaakkola, T.S. (2018). Junction tree variational autoencoder for molecular graph generation. In Proceedings of the 35th International Conference on Machine Learning, ICML 2018. Proceedings of Machine Learning Research PMLR, J.G. Dy and A. Krause, eds., pp. 2328–2337. <http://proceedings.mlr.press/v80/jin18a.html>.
- Kang, S., and Cho, K. (2019). Conditional molecular design with deep generative models. *J. Chem. Inf. Model.* 59, 43–52. <https://doi.org/10.1021/acs.jcim.8b00263>.
- Olivecrona, M., Blaschke, T., Engkvist, O., and Chen, H. (2017). Molecular de-novo design through deep reinforcement learning. *J. Cheminf.* 9, 48. <https://doi.org/10.1186/s13321-017-0235-x>.
- Shi, C., Xu, M., Zhu, Z., Zhang, W., Zhang, M., and Tang, J. (2020). Graphaf: a flow-based autoregressive model for molecular graph generation. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2001.09382>.
- Sanchez-Lengeling, B., Outeiral, C., Guimaraes, G., and Aspuru-Guzik, A. (2017). Optimizing distributions over molecular space. an objective-reinforced generative adversarial network for inverse-design chemistry (organic). Preprint at ChemRxiv. <https://doi.org/10.26434/chemrxiv.5309668.v2>.
- Zhou, Z., Kearnes, S., Li, L., Zare, R.N., and Riley, P. (2019). Optimization of molecules via deep reinforcement learning. *Sci. Rep.* 9, 10752. <https://doi.org/10.1038/s41598-019-47148-x>.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature* 518, 529–533. <https://doi.org/10.1038/nature14236>.
- Yang, K., Jin, W., Swanson, K., Barzilay, D., and Jaakkola, T. (2020). Improving molecular design by stochastic iterative target augmentation. In Proceedings of the 37th International Conference on Machine Learning vol. 119 of *Proceedings of Machine Learning Research*, H.D. III and A. Singh, eds. (PMLR), pp. 10716–10726. <https://proceedings.mlr.press/v119/yang20e.html>.
- Iovanac, N.C., MacKnight, R., and Savoie, B.M. (2022). Actively searching: Inverse design of novel molecules with simultaneously optimized properties. *J. Phys. Chem.* 126, 333–340. <https://doi.org/10.1021/acs.jpca.1c08191>.
- Kusner, M.J., Paige, B., and Hernández-Lobato, J.M. (2017). Grammar variational autoencoder. In *International conference on machine learning* (PMLR), pp. 1945–1954.
- Liu, X., Liu, Q., Song, S., and Peng, J. (2020). A chance-constrained generative framework for sequence optimization. In Proceedings of the 37th International Conference on Machine Learning, 119, H.D. III and A. Singh, eds. (PMLR), pp. 6271–6281. <https://proceedings.mlr.press/v119/liu20i.html>.
- Tripp, A., Daxberger, E., and Hernández-Lobato, J.M. (2020). Sample-efficient optimization in the latent space of deep generative models via weighted retraining. *Adv. Neural Inf. Process. Syst.* 33, 11259–11272.
- Kim, Y., Kim, Y., Yang, C., Park, K., Gu, G.X., and Ryu, S. (2021). Deep learning framework for material design space exploration using active transfer learning and data augmentation. *npj Comput. Mater.* 7, 140. <https://doi.org/10.1038/s41524-021-00609-2>.
- Jin, W., Barzilay, D., and Jaakkola, T. (2020). Multi-objective molecule generation using interpretable substructures. In Proceedings of the 37th International Conference on Machine Learning, 119, H.D. III and A. Singh, eds. (PMLR), pp. 4849–4859. <https://proceedings.mlr.press/v119/jin20b.html>.
- Xie, Y., Shi, C., Zhou, H., Yang, Y., Zhang, W., Yu, Y., and Li, L. (2021). Mars: Markov molecular sampling for multi-objective drug discovery. In *International Conference on Learning Representations*.
- Feng, H., Wang, R., Zhan, C.-G., and Wei, G.-W. (2023). Multiobjective molecular optimization for opioid use disorder treatment using generative network complex. *J. Med. Chem.* 66, 12479–12498. <https://doi.org/10.1021/acs.jmedchem.3c01053>.
- Frazier, P.I. (2018). A tutorial on bayesian optimization. Preprint at arXiv. <https://doi.org/10.48550/arXiv.1807.02811>.

25. Ertl, P., Roggo, S., and Schuffenhauer, A. (2008). Natural product-likeness score and its application for prioritization of compound libraries. *J. Chem. Inf. Model.* *48*, 68–74. <https://doi.org/10.1021/ci700286x>.
26. Creese, I., Burt, D.R., and Snyder, S.H. (1976). Dopamine receptor binding predicts clinical and pharmacological potencies of antischizophrenic drugs. *Science* *192*, 481–483. <https://doi.org/10.1126/science.3854>.
27. Landrum, G. (2013). Rdkit: A software suite for cheminformatics, computational chemistry, and predictive modeling, *8*, p. 31. <https://www.rdkit.org>.
28. Irwin, J.J., Sterling, T., Mysinger, M.M., Bolstad, E.S., and Coleman, R.G. (2012). Zinc: A free tool to discover chemistry for biology. *J. Chem. Inf. Model.* *52*, 1757–1768. <https://doi.org/10.1021/ci3001277>.
29. Daina, A., Michielin, O., and Zoete, V. (2017). Swissadme: a free web tool to evaluate pharmacokinetics, drug-likeness and medicinal chemistry friendliness of small molecules. *Sci. Rep.* *7*, 42717. <https://doi.org/10.1038/srep42717>.
30. Jones, D., Kim, H., Zhang, X., Zemla, A., Stevenson, G., Bennett, W.F.D., Kirshner, D., Wong, S.E., Lightstone, F.C., and Allen, J.E. (2021). Improved protein–ligand binding affinity prediction with structure-based deep fusion inference. *J. Chem. Inf. Model.* *61*, 1583–1592. <https://doi.org/10.1021/acs.jcim.0c01306>.
31. McLoughlin, K.S., Jeong, C.G., Sweitzer, T.D., Minnich, A.J., Tse, M.J., Bennion, B.J., Allen, J.E., Calad-Thomson, S., Rush, T.S., and Brase, J.M. (2021). Machine learning models to predict inhibition of the bile salt export pump. *J. Chem. Inf. Model.* *61*, 587–602. <https://doi.org/10.1021/acs.jcim.0c00950>.
32. Jacobs, S.A., Moon, T., McLoughlin, K., Jones, D., Hysom, D., Ahn, D.H., Gyllenhaal, J., Watson, P., Lightstone, F.C., Allen, J.E., et al. (2021). Enabling rapid covid-19 small molecule drug design through scalable deep learning of generative models. *Int. J. High Perform. Comput. Appl.* *35*, 469–482. <https://doi.org/10.1177/10943420211010930>.
33. Dehghannasiri, R., Yoon, B.-J., and Dougherty, E.R. (2015a). Optimal experimental design for gene regulatory networks in the presence of uncertainty. *IEEE ACM Trans. Comput. Biol. Bioinf* *12*, 938–950. <https://doi.org/10.1109/TCBB.2014.2377733>.
34. Dehghannasiri, R., Yoon, B.-J., and Dougherty, E.R. (2015b). Efficient experimental design for uncertainty reduction in gene regulatory networks. In *BMC Bioinformatics*, *16* (Springer), pp. 1–18. <https://doi.org/10.1186/1471-2105-16-S13-S2>.
35. Hong, Y., Kwon, B., and Yoon, B.-J. (2021). Optimal experimental design for uncertain systems based on coupled differential equations. *IEEE Access* *9*, 53804–53810. <https://doi.org/10.1109/ACCESS.2021.3071038>.
36. Woo, H.-M., Hong, Y., Kwon, B., and Yoon, B.-J. (2021b). Accelerating optimal experimental design for robust synchronization of uncertain kuramoto oscillator model using machine learning. *IEEE Trans. Signal Process.* *69*, 6473–6487. <https://doi.org/10.1109/TSP.2021.3130967>.
37. Zhao, G., Dougherty, E., Yoon, B.-J., Alexander, F.J., and Qian, X. (2021a). Efficient active learning for gaussian process classification by error reduction. In *Thirty-fifth Conference on Neural Information Processing Systems*.
38. Zhao, G., Dougherty, E., Yoon, B.-J., Alexander, F., and Qian, X. (2021b). Uncertainty-aware active learning for optimal bayesian classifier. In *International Conference on Learning Representations*.
39. Zhao, G., Dougherty, E., Yoon, B.-J., Alexander, F.J., and Qian, X. (2021c). Bayesian active learning by soft mean objective cost of uncertainty. In *International Conference on Artificial Intelligence and Statistics (PMLR)*, pp. 3970–3978.
40. Yoon, B.-J., Qian, X., and Dougherty, E.R. (2013). Quantifying the objective cost of uncertainty in complex dynamical systems. *IEEE Trans. Signal Process.* *61*, 2256–2266. <https://doi.org/10.1109/TSP.2013.2251336>.
41. Yoon, B.-J., Qian, X., and Dougherty, E.R. (2021). Quantifying the multi-objective cost of uncertainty. *IEEE Access* *9*, 80351–80359. <https://doi.org/10.1109/ACCESS.2021.3085486>.
42. Abeer, A.N.M.N. (2024). GMD-MO-LSO: Multi-Objective Latent Space Optimization of Generative Molecular Design Models. Zenodo. <https://doi.org/10.5281/zenodo.12730304>.
43. Deb, K. (2004). *Multi-objective Optimization Using Evolutionary Algorithms* (John Wiley & Sons).
44. Obayashi, S., Takahashi, S., and Takeguchi, Y. (1998). Niching and elitist models for mogas. In *Parallel Problem Solving from Nature—PPSN V: 5th International Conference, 1998 Proceedings 5* (Springer), pp. 260–269. <https://doi.org/10.1007/BFb0056869>.
45. Lipinski, C.A., Lombardo, F., Dominy, B.W., and Feeney, P.J. (1997). Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv. Drug Deliv. Rev.* *23*, 3–25. [https://doi.org/10.1016/S0169-409X\(96\)00423-1](https://doi.org/10.1016/S0169-409X(96)00423-1).
46. Wildman, S.A., and Crippen, G.M. (1999). Prediction of physicochemical parameters by atomic contributions. *J. Chem. Inf. Comput. Sci.* *39*, 868–873. <https://doi.org/10.1021/ci9903071>.
47. Ertl, P., and Schuffenhauer, A. (2009). Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *J. Cheminf.* *1*, 1–11. <https://doi.org/10.1186/1758-2946-1-8>.
48. Harvey, A.L. (2008). Natural products in drug discovery. *Drug Discov. Today* *13*, 894–901. <https://doi.org/10.1016/j.drudis.2008.07.004>.
49. Wang, S., Che, T., Levit, A., Shoichet, B.K., Wacker, D., and Roth, B.L. (2018). Structure of the d2 dopamine receptor bound to the atypical antipsychotic drug risperidone. *Nature* *555*, 269–273. <https://doi.org/10.1038/nature25758>.
50. Bakadlag, R., Jandaghi, P., Hoheisel, J.D., and Riazalhosseini, Y. (2019). The potential of dopamine receptor d2 (drd2) as a therapeutic target for tackling pancreatic cancer. *Expert Opin. Ther. Targets* *23*, 365–367. <https://doi.org/10.1080/14728222.2019.1606904>.
51. Yeh, C.-T., Wu, A.T., Chang, P.M.-H., Chen, K.-Y., Yang, C.-N., Yang, S.-C., Ho, C.-C., Chen, C.-C., Kuo, Y.-L., Lee, P.-Y., et al. (2012). Trifluoperazine, an antipsychotic agent, inhibits cancer stem cell growth and overcomes drug resistance of lung cancer. *Am. J. Respir. Crit. Care Med.* *186*, 1180–1188. <https://doi.org/10.1164/rccm.201207-1180OC>.
52. Yong, M., Yu, T., Tian, S., Liu, S., Xu, J., Hu, J., and Hu, L. (2017). Dr2 blocker thioridazine: A promising drug for ovarian cancer therapy corrigendum in/10.3892/ol. 2020.11285. *Oncol. Lett.* *14*, 8171–8177. <https://doi.org/10.3892/ol.2017.7184>.
53. Tung, M.-C., Lin, Y.-W., Lee, W.-J., Wen, Y.-C., Liu, Y.-C., Chen, J.-Q., Hsiao, M., Yang, Y.-C., and Chien, M.-H. (2022). Targeting drd2 by the antipsychotic drug, penfluridol, retards growth of renal cell carcinoma via inducing stemness inhibition and autophagy-mediated apoptosis. *Cell Death Dis.* *13*, 400. <https://doi.org/10.1038/s41419-022-04828-3>.
54. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* *12*, 2825–2830.