

MIMOSA: Multi-constraint Molecule Sampling for Molecule Optimization

Tianfan Fu,¹ Cao Xiao,² Xinhao Li,³ Lucas Glass,^{2,4} Jimeng Sun⁴

¹College of Computing, Georgia Institute of Technology

²Analytics Center of Excellence, IQVIA

³Department of Chemistry, North Carolina State University

⁴Department of Statistics, Temple University

⁵Department of Computer Science, University of Illinois, Urbana-Champaign

Abstract

Molecule optimization is a fundamental task for accelerating drug discovery, with the goal of generating new valid molecules that maximize multiple drug properties while maintaining similarity to the input molecule. Existing generative models and reinforcement learning approaches made initial success, but still face difficulties in simultaneously optimizing multiple drug properties. To address such challenges, we propose the Multi-constraint MOleculE SAMpling (MIMOSA) approach, a sampling framework to use input molecule as an initial guess and sample molecules from the target distribution. MIMOSA first pretrains two property-agnostic graph neural networks (GNNs) for molecule topology and substructure-type prediction, where a *substructure* can be either atom or single ring. For each iteration, MIMOSA uses the GNNs' prediction and employs three basic substructure operations (*add*, *replace*, *delete*) to generate new molecules and associated weights. The weights can encode multiple constraints including similarity and drug property constraints, upon which we select promising molecules for next iteration. MIMOSA enables flexible encoding of multiple property- and similarity-constraints and can efficiently generate new molecules that satisfy various property constraints and achieved up to 49.1% relative improvement over the best baseline in terms of success rate.

1 Introduction

Designing molecules with desirable properties is a fundamental task in drug discovery. Traditional methods such as high throughput screening (HTS) tests large compound libraries to identify molecules with desirable properties, which are inefficient and costly (Polishchuk, Madzhidov, and Varnek 2013; Huang et al. 2020a,b). Two important machine learning tasks have been studied in this context:

- Molecule generation aims at creating new and diverse molecule graphs with some desirable properties (Jin, Barzilay, and Jaakkola 2018; You et al. 2018);
- Molecule optimization takes a more targeted approach to find molecule Y with improved drug properties such as drug likeness and biological activity given an input molecule X (Jin et al. 2019; Zhou et al. 2019).

Existing works on molecule optimization and molecule generation tasks can be categorized as generative models (Kusner, Paige, and Hernández-Lobato 2017; Dai et al. 2018; Gómez-Bombarelli et al. 2018) and reinforcement learning (RL) methods (You et al. 2018; Zhou et al. 2019). Most existing works only optimize a single property, while multiple properties need to be optimized in order to develop viable drug candidates. Recently, (Jin, Barzilay, and Jaakkola 2020b) proposed a molecule generation algorithm that can optimize multiple properties which is a related but different task than molecule optimization since they do not take any specific input molecule as the anchor. Nigam et al. (2020) proposed a genetic algorithm (GA) for molecule generation and optimization. In this work, we propose a *sampling based strategy* to tackle the molecule optimization for multi-properties.

To allow for flexible and efficient molecule optimization on multiple properties, we propose a new sampling based molecule optimization framework named Multi-constraint MOleculE SAMpling (MIMOSA). MIMOSA uses the input molecule as an initial guess and pretrains two graph neural networks (GNNs) on molecule topology and substructure-type predictions to produce better molecule embedding for sampling, where *substructure* can be either an atom or a ring. In each iteration, MIMOSA uses the prediction and employs three basic substructure operations (*add*, *replace*, *delete*) to generate new molecule candidates and associated weights. The weights thus effectively encode multiple constraints including similarity to the input molecule and various drug properties, upon which we accept promising molecules for next iteration sampling. MIMOSA iteratively produces new molecule candidates and can efficiently draw molecules that satisfy all constraints. The main contributions of our paper are listed below.

- A new sampling framework for flexible encoding of multiple constraints. We reformulate molecule optimization task in a sampling framework to draw molecules from the target distribution (Eq. (1)). The framework provides flexible and efficient encoding of multi-property and similarity constraints as a target distribution (Section 3.1).
- Efficient sampling augmented by GNN pretraining. With the help of two pretrained GNN models, we designed a Markov Chain Monte Carlo (MCMC) based molecule

sampling method that enables efficient sampling from a target distribution (Section 3.2). This enables MIMOSA to leverage vast amount molecule data in an unsupervised manner without the need of any knowledge of molecule pairs (i.e., an input molecule and an enhanced molecule) as many existing methods do.

- **Guaranteed unbiased sampling.** We provide theoretical analysis to show that the proposed MCMC method draws unbiased samples from the target distribution, i.e., exhibiting ergodicity and convergence (Section 3.3).

We compare MIMOSA with state-of-the-art baselines on optimizing several important properties across multiple settings, MIMOSA achieves 43.7% success rate (49.1% relative improvement over the best baseline GA (Nigam et al. 2020)) when optimizing DRD and PLogP jointly.

2 Related Work

Generative models for molecule optimization project an input molecule to a latent space, then search in the latent space for new and better molecules. For example, Gómez-Bombarelli et al. (2018), Blaschke et al. (2018) utilized SMILES strings as molecule representations to generate molecules. Since string-based approaches often create many invalid molecules, Kusner, Paige, and Hernández-Lobato (2017) and Dai et al. (2018) designed grammar constraints to improve the chemical validity. Recently, Nigam et al. (2020) proposed to explore molecule generation using a genetic algorithm. Another line of works focus on graph representations of molecules, e.g., CGVAE (Constrained Graph VAE) (Liu et al. 2018), JTVAE (Junction Tree VAE) based approaches (Jin, Barzilay, and Jaakkola 2018; Jin et al. 2019; Jin, Barzilay, and Jaakkola 2020a; Fu, Xiao, and Sun 2020; Fu et al. 2020a, 2021). Although almost perfect on generating valid molecules, most of them rely on paired data as training data.

Reinforcement learning for molecule optimization are also developed on top of molecule generators for achieving desirable properties. For example, Olivecrona et al. (2017), Putin (2018), Popova, Isayev, and Tropsha (2018) applied RL techniques on top of a string generator to generate SMILES strings. They struggled with validity of the generated chemical structures. Recently, You et al. (2018), Zhou et al. (2019) leverage deep reinforcement learning to generate molecular graph, achieving perfect validity. However, all these methods require pre-training on a specific dataset, which makes their exploration ability limited by the biases present in the training data. More recently, Jin, Barzilay, and Jaakkola (2020b) focused on molecule generation method for creating molecules with multiple properties. However, this approach can lead to arbitrary diverse structures (not optimized for a specific input molecule) and assumes each property is associated with specific molecular substructures which are not applicable to all properties.

In this paper, we proposed a new molecule optimization method that casts molecule optimization as a sampling problem, which provides an efficient and flexible framework for optimizing multiple constraints (e.g., similarity constraint, multiple property constraints) simultaneously.

3 The MIMOSA Method

3.1 Molecule Optimization via Sampling

Slightly different from general molecule generation that focuses on generating valid and diverse molecules, the *molecule optimization* task takes a molecule X as input, and aims to obtain a new molecule Y that is not only similar to X but also have more desirable drug properties than X .

We formulate a Markov Chain Monte Carlo (MCMC) based sampling strategy. The MCMC methods are popular Bayesian sampling approaches of estimating posterior distributions. They allow drawing samples from complex distributions with desirable sampling efficiency (Welling and Teh 2011) as long as unnormalized probability density for samples can be calculated.

Here to formulate molecule optimization that aim to optimize on similarity between the input molecule X and the target molecules Y as well as M molecular properties of Y , $\mathcal{P}_1, \dots, \mathcal{P}_M$ (the higher score the better). We propose to draw Y from the *unnormalized target distribution* in Eq. (1).

$$p_X(Y) \propto \mathbf{1}(Y) \exp \left(\eta_0 \text{sim}(X, Y) + \eta_1 \left(\mathcal{P}_1(Y) - \mathcal{P}_1(X) \right) + \dots + \eta_M \left(\mathcal{P}_M(Y) - \mathcal{P}_M(X) \right) \right) \quad (1)$$

where $\eta_0, \eta_1, \dots, \eta_M \in \mathbb{R}_+$ are the hyperparameters that control the strength of various terms, $\mathbf{1}(Y)$ is an indicator function measuring whether the molecule Y is a valid molecule. It is added to ensure the validity of the generated molecule Y . The target distribution is designed to encode any number of type of constraints, including similarity constraint and multiple drug property constraints. Here the use of \exp is to guarantee $p_X(Y)$ is valid probability distribution. Usually we define the similarity $\text{sim}(X, Y)$ as in Def. 1 and measured using Eq. (2).

Definition 1 (Tanimoto Similarity of Molecules). Denote \mathcal{S}_X and \mathcal{S}_Y as fragment descriptor¹ sets of molecule X and Y , respectively. The Tanimoto similarity between X and Y is given by

$$\text{sim}(X, Y) = \frac{|\mathcal{S}_X \cap \mathcal{S}_Y|}{|\mathcal{S}_X \cup \mathcal{S}_Y|} \in [0, 1], \quad (2)$$

where \cap, \cup represent the intersection and union of two binary vectors respectively; $|\cdot|$ denotes the cardinality of a set. Higher value means more similar.

3.2 The MIMOSA Method for Molecule Sampling

Fig. 1 illustrates the overall procedure of MIMOSA, which can be decomposed into the following steps: (1) Pre-train GNN. MIMOSA pre-trains two graph neural networks (GNNs) using a large number of unlabeled molecules, which will be used in the sampling process. Then MIMOSA iterates

¹Fragment descriptors, represent selected substructures (fragments) of 2D molecular graphs and their occurrences in molecules; they constitute one of the most important types of molecular descriptors (Baskin and Varnek 2009).

MIMOSA: Multi Constraint Molecule Sampling for Molecule Optimization

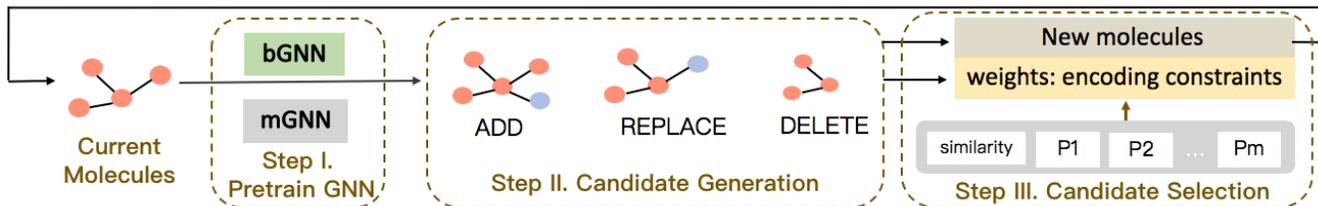


Figure 1: The Multi-constraint Molecule Sampling for Molecule Optimization(MIMOSA) framework illustrated using a single molecule. In Step I (Pretrain GNN), MIMOSA pretrains two property-agnostic GNNs for molecule topology and substructure-type prediction. Then, in Step II (Candidate Generation), MIMOSA uses the prediction and employs three basic substructure operations (ADD, REPLACE and DELETE) to generate new molecule candidates. In Step III (Candidate Selection), MIMOSA assigns weights for new molecule. The weights can encode multiple constraints including similarity and drug property constraints, upon which we accept promising molecules for next iteration. MIMOSA iteratively edits the molecule and can efficiently draw molecule samples.

Notations	short explanation
X, Y	Input molecule, target molecule.
$\text{sim}(X, Y) \in [0, 1]$	Similarity of molecules X and Y .
$p_X(Y)$	Target dist. when optimizing X , Eq. 1.
M	# of properties to optimize.
$\gamma_0, \gamma_1, \dots, \gamma_M \in \mathbb{R}_+$	Hyperparameter in Target dist. $p_X(Y)$.
$\mathcal{P}_1, \dots, \mathcal{P}_M$	Molecular properties to optimize.
$\mathbf{1}(Y)$	Validity Indicator func. of molecule Y .
K	Depth of GNN.
$\mathbf{h}_v^{(k)} \in \mathbb{R}^{300}$	Node embedding v in the k -th layer.
C_1/C_2	# of all possible substructures/bonds.
$v; s_v/s'_v$	node v ; substructures of v .
$\mathbf{f}_v/\mathbf{g}_e$	one-hot node/edge feature.
$\hat{y}_v/\text{mGNN}(Y, v)$	substructure distribution. Eq. (6).
$\hat{z}_v/\text{bGNN}(Y, v)$	probability of v will expand. Eq. (9).
y_v/z_v	ground truth label of node v
Y/Y'	current/next Sample.
$S_{\text{add}}, S_{\text{replace}}, S_{\text{delete}}$	sampling operation from Y to Y' .

Table 1: Notations used in the paper.

over the following two steps. (2) Candidate Generation. We generate and score molecule candidates via modification operations (*add*, *delete*, *replace*) to the current molecule. (3) Candidate Selection. We perform MCMC sampling to select promising molecule candidates for the next sampling iteration by repeating Step 2 and 3. Note that all modification operations are on the *substructure* level, where a substructure can be either an atom or a single ring. The substructure set includes all 118 atoms and 31 single rings.

(I) Pretrain GNNs for Substructure-type and Molecule Topology Prediction

To provide accurate molecule representation, we propose to pretrain molecule embeddings on large molecule datasets. Since we consider molecules in graph representations where each substructure is a node, we develop two GNN based pretraining tasks to assist molecule modification. These two GNNs will assess the probability of each substructure conditioned on all the other substructures in the molecule graph.

Mathematically, in molecular graph $Y = (V, E)$, we have one-hot node feature $\mathbf{f}_v \in \{0, 1\}^{C_1}$ for every node $v \in V$ and one-hot edge feature $\mathbf{g}_e \in \{0, 1\}^{C_2}$ for every edge $e = (u, v) \in E$. C_1 and C_2 are the number of substructures and the number of bond types, respectively. In our experiment, $C_1 = 149$, including 118 atoms, 31 single rings, and $C_2 = 4$ correspond to the four bond types. The node and edge features can be found in Fu et al. (2020b).

The two Graph Neural Networks (GNN) (Hu et al. 2019) are learned with these node and edge features and the same molecule graph to learn an embedding vector \mathbf{h}_v for every node $v \in V$.

$$\mathbf{h}_v^{(k)} = \text{ReLU} \left(\text{MLP} \left(\text{CONCAT} \left(\left(\sum_{u \in \mathcal{N}(v) \cup \{v\}} \mathbf{h}_u^{(k-1)}, \sum_{e=(u,v):u \in \mathcal{N}(v)} \mathbf{g}_e^{(k-1)} \right) \right) \right) \right), \quad (3)$$

where the layer $k = 1, \dots, K$; $\text{CONCAT}(\cdot, \cdot)$ is the concatenation of two vectors; $\mathcal{N}(v)$ is the set of all neighbors of v ; $\mathbf{h}_v^{(0)}$ is the initial node embedding \mathbf{f}_v . After K layers of GNN, we have the final node embedding $\mathbf{h}_v^{(K)}$ for node v . In our experiment, $K = 5$.

Using the same GNN architecture, we trained two GNN models: one for substructure-type prediction called *mGNN* and one for molecule topology prediction called *bGNN*: We choose to train two separate GNNs instead of sharing a single GNN because sufficient unlabeled molecule samples exist and the two tasks are very different in nature.

The mGNN model aims at multi-class classification for predicting the substructure type of a masked node. The mGNN model outputs the type of an individual substructure conditioned on all other substructures and their connections. We mask the individual substructure, replace it with a special masked indicator following (Hu et al. 2019). Suppose we only mask one substructure for each molecule during training and v is the masked substructure (i.e., node), y_v is the node label corresponding to masked substructure type, we add fully-connected (FC) layers with softmax activation

(Eq. (4)) to predict the type of the node v .

$$\hat{\mathbf{y}}_v = \text{Softmax}(\text{FC}(\mathbf{h}_v^{(K)})). \quad (4)$$

where $\hat{\mathbf{y}}_v$ is a C_1 dimension vector, indicating the predicted probability of all possible substructures. Multi-class cross entropy loss (Eq. (5)) is used to guide the training of GNN:

$$\mathcal{L}(\mathbf{y}_v, \hat{\mathbf{y}}_v) = - \sum_{i=1}^{C_1} ((\mathbf{y}_v)_i \log(\hat{\mathbf{y}}_v)_i), \quad (5)$$

where \mathbf{y}_v is the groundtruth, one-hot vector. C_1 is number of all substructures (atoms and single rings), $(\mathbf{y}_v)_i$ is i -th element of vector \mathbf{y}_v .

To summarize, the prediction of mGNN is defined as

$$\hat{\mathbf{y}}_v \triangleq \text{mGNN}(Y, \text{mask} = v) = \text{mGNN}(Y, v), \quad (6)$$

where in a given molecule Y the node v is masked, mGNN predicts the substructure distribution on masked node v , which is denoted $\hat{\mathbf{y}}_v$.

The bGNN model aims at binary classification for predicting the molecule topology. The goal of bGNN is to predict whether a node will expand. To provide training labels for bGNN, we set the leaf nodes (nodes with degree 1) with label $z_v = 0$ as we assume they are no longer expanding. And we set label $z_v = 1$ on the non-leaf nodes that are adjacent to leaf nodes as those nodes expanded (to the leaf nodes). The prediction is done via

$$\hat{z}_v = \text{Sigmoid}(\text{FC}(\mathbf{h}_v^{(K)})), \quad (7)$$

where FC is two-layer fully-connected layers (of 50 neurons followed by 1 neuron). $\mathbf{h}_v^{(K)}$ is defined in Eq. (3), the node embedding of v produced by GNN. Binary cross-entropy loss is used to guide the training:

$$\mathcal{L}(z_v, \hat{z}_v) = -z_v \log(\hat{z}_v) - (1 - z_v) \log(1 - \hat{z}_v). \quad (8)$$

Since the total number of unlabeled molecules is large, when training bGNN we randomly select one substructure v for each molecule to speed up the pretraining.

In sum, prediction of bGNN is defined as

$$\hat{z}_v \triangleq \text{bGNN}(Y, v), \quad (9)$$

where v is a node in molecule Y , \hat{z}_v is the probability that v will expand.

(II) Candidate Generation via Substructure Modification Operation

With the help of mGNN and bGNN, we define substructure modification operations namely *replace*, *add* or *delete* on input molecule Y :

- **Replace a substructure.** At node v , the original substructure category is s_v .
 1. We mask v in Y , evaluate the substructure distribution in v via mGNN, i.e., $\hat{\mathbf{y}}_v = \text{mGNN}(Y, v)$, as Eq. (6).
 2. Then we sample a new substructure s'_v from the multinomial distribution $\hat{\mathbf{y}}_v$, denoted by $s'_v \sim \text{Multinomial}(\hat{\mathbf{y}}_v)$.

3. At node v , we replace the original substructure s_v with new substructure s'_v to produce the new molecule Y' .

The whole operation is denoted as

$$Y' \sim S_{\text{replace}}(Y'|Y). \quad (10)$$

- **Add a substructure.** Suppose we want to add a substructure as leaf node (denoted as v) connecting to an existing node u in current molecule Y . The substructure category of v is denoted s_v , which we want to predict.
 1. We evaluate the probability that node u has a leaf node v with help of bGNN in Eq. (9), i.e.,

$$\hat{z}_u = \text{bGNN}(Y, u) \in [0, 1].$$
 2. Suppose the above prediction is to add a leaf node v . We then generate a new molecule Y' via adding v to Y via a new edge (u, v) .
 3. In Y' , s_v , the substructure of v is unknown. We will predict its substructure using mGNN, i.e., $\hat{\mathbf{y}}_v = \text{mGNN}(Y', v)$, following Eq. (6).
 4. We sample a new substructure s'_v from the multinomial distribution $\hat{\mathbf{y}}_v$ and complete the new molecule Y' .

The whole operation is denoted as

$$Y' \sim S_{\text{add}}(Y'|Y). \quad (11)$$

- **Delete a substructure.** We delete a leaf node v in current molecule Y . It is denoted

$$Y' \sim S_{\text{delete}}(Y'|Y). \quad (12)$$

In the MCMC process, $S_*(Y'|Y)$ indicates the sequential sampling process from previous sample Y to next sample Y' . And the very first sample is the input X .

Handling Bond Types and Rings. Since the number of possible bonds are small (single, double, triple, aromatic), we enumerate all and choose the one with largest $p_X(Y)$. In some case, basic operation would generate invalid molecules. Based on the indicator function in target distribution in Eq. (1), the density is equal to 0. Thus, we perform validity check using RDKit (<https://www.rdkit.org/>) to filter out the new molecule graphs that are not valid. When adding/replacing a ring, there might be multiple choices to connect to its neighbor. We enumerate all possible choices and retain all valid molecules.

(III) Candidate Selection via MCMC Sampling

The set of generated candidate molecules can be grouped as three sets based on the type of substructure modification they received, namely, *replace* set S_{replace} , *add* set S_{add} , and *delete* set S_{delete} . MIMOSA uses the Gibbs sampling (Geman and Geman 1984), a particular type of MCMC, for molecule candidate selection. Gibbs sampling algorithm generates an instance from the distribution of each variable in sequential or random order (Levine and Casella 2006), conditional on the current values of the other variables. Here molecules from the three sets will be sampled with different sampling weights. Their weights are designed to satisfy the detailed balance condition (Brooks et al. 2011).

Sampling S_{replace} . For molecules produced by the ‘‘replace’’ operation, the weight in sampling w_r is given by Eq. (13).

$$w_r = \frac{p_X(Y') \cdot [\text{mGNN}(Y, v)]_{s'_v}}{p_X(Y) \cdot [\text{mGNN}(Y, v)]_{s_v}}, \quad (13)$$

where $P_X(\cdot)$ is the unnormalized target distribution for optimizing X , defined in Eq. (1), $[\text{mGNN}(Y, v)]_{s_v}$ is the predicted probability of the substructure s_v in the prediction distribution $\text{mGNN}(Y, v)$. The acceptance rate in the proposal is $\min\{1, w_r\}$. If the proposal is accepted, we use the new prediction s'_v to replace origin substructure s_v in current molecule Y and produce the new molecule Y' .

Sampling S_{add} . For molecules produced by the ‘‘add’’ operation, the weight in sampling is given by Eq. (14).

$$w_a = \frac{p_X(Y') \cdot \text{bGNN}(Y, u) \cdot [\text{mGNN}(Y', v)]_{s_v}}{p_X(Y) \cdot (1 - \text{bGNN}(Y, u))}, \quad (14)$$

where The acceptance rate in the proposal is $\min\{1, w_a\}$.

Sampling S_{delete} . For these molecules produced by ‘‘delete’’ operation, the weight in sampling is given by Eq. (15).

$$w_d = \frac{p_X(Y') \cdot (1 - \text{bGNN}(Y', u))}{p_X(Y) \cdot \text{bGNN}(Y', u) \cdot [\text{mGNN}(Y, v)]_{s_v}}, \quad (15)$$

where v is the deleted node, leaf node (with degree 1) in molecular graph of Y . u and v are connected in Y . The acceptance rate in the proposal is $\min\{1, w_d\}$.

Soft-constraint Encoding. For these operations, any number or type of constraints (e.g., here the similarity and drug property constraints) can be encoded in $p_X(Y)$ and $p_X(Y')$ and thus reflected in the weights w_r, w_a, w_d .

For a single-chain MCMC, we construct the transition kernel as given by Eq. (16).

$$Y' \sim \begin{cases} S_{\text{replace}}(Y'|Y), & \text{prob } \gamma_1, \text{ accept w. } \min\{1, w_r\}, \\ S_{\text{add}}(Y'|Y), & \text{prob. } \gamma_2, \text{ accept w. } \min\{1, w_a\}, \\ S_{\text{delete}}(Y'|Y), & \text{prob. } \gamma_3, \text{ accept w. } \min\{1, w_d\}, \end{cases} \quad (16)$$

where $\gamma_1, \gamma_2, \gamma_3 \in \mathbb{R}_+$ are hyperparameters that determine the sampling probabilities from the three molecule sets. In Section 3.3, we show the transition kernel will leave the target distribution $p_X(Y)$ invariant for arbitrary $\gamma_1, \gamma_2, \gamma_3$ satisfying $\gamma_1 + \gamma_2 + \gamma_3 = 1$ and $\gamma_2 = \gamma_3$. After molecules are sampled, they will be accepted with their corresponding acceptance rates related to w_r, w_a, w_d .

The MIMOSA method is summarized in Algorithm 1. To accelerate the sampling procedure, we also deploy a multi-chain strategy (Liu, Liang, and Wong 2000): during each step, we use N samples for each state, with each sample generating multiple proposals. Also, during burn-in period (Step 12 in Algorithm 1), we pick the molecules with highest density for efficiency (Brooks et al. 2011). We retain N proposals in iterative sampling.

3.3 Analysis of the MCMC Algorithm

Our MCMC method draws unbiased samples from the target distribution, i.e., exhibiting ergodicity and convergence. The proofs of Lemma 1 and 2 can be found in Fu et al. (2020b).

Algorithm 1 MIMOSA for Molecule Optimization

- 1: Input: molecule X , # of Particle N , max # of sampling iter. T_{max} , # of burn-in iter. T_{burnin}
 - 2: Output: Generated molecules Φ .
 - 3: # Step (I) Pretrain GNN
 - 4: Train mGNN (Eq.6), bGNN (Eq.9).
 - 5: Candidate set $\Theta = \{X\}$, Output set $\Phi = \{\}$.
 - 6: **for** iter = 1, \dots , T_{max} **do**
 - 7: # Step (II) Candidate Generation.
 - 8: Candidate Pool $\Psi = \{\}$.
 - 9: **for** molecule Z in Θ **do**
 - 10: Generate candidates Z' via editing Z using sub-structure operations; validity check; add Z' in Ψ .
 - 11: **end for**
 - 12: $\Theta = \{\}$.
 - 13: # Step (III) Candidate Selection.
 - 14: **if** iter < T_{burnin} **then**
 - 15: Select N molecules with highest density value (Eq. 1) from Ψ and add them into Θ .
 - 16: **else**
 - 17: Draw N molecules from Ψ using importance sampling (\propto weight w_r in Eq. (13), w_a in Eq. (14) or w_d in Eq. (15)) and add to Θ .
 - 18: **end if**
 - 19: $\Phi = \Phi \cup \Theta$.
 - 20: **end for**
-

Theorem 1. Suppose $\{Y_1, Y_2, \dots, Y_n\}$ is the chain of molecules sampled via MCMC based on transition kernel defined in Eq. (16), with initial state X , then the Markov chain is ergodic with stationary distribution $p_X(Y)$ in Eq. (1). That is, empirical estimate (time average over Y_1, Y_2, \dots, Y_n) is equal to target value (space average over $p_X(Y)$), i.e., $\lim_{n \rightarrow +\infty} \frac{1}{n} \sum_{i=1}^n f(Y_i) = \int f(Y) p_X(Y) dY$ holds for any integratable function f .

Proof Sketch. We split the proof into Lemma 1 and 2. First, regarding the ergodicity, it is sufficient to prove the irreducibility, aperiodicity of the Markov chain (Lemma 1). Then, to show that $p_X(Y)$ is maintained invariant for the whole chain, in Lemma 2, we show that detailed balance condition holds for any neighboring samples (Y_i and Y_{i+1}). Then we strengthen this results on the whole chain.

Lemma 1. The Markov chain of the sampled molecules ($\{Y_1, \dots, Y_n\}$, starting at X , based on transition kernel in Eq. (16)) is ergodic over the target distribution $p_X(Y)$.

Lemma 2. $p_X(Y)$ is maintained as the invariant distribution for the whole Markov chain produced by MCMC transition kernel defined in Eq. (16).

4 Experiment

4.1 Experimental Setup

Dataset and Molecular Properties. We use 2 million molecules from ZINC database (Sterling and Irwin 2015; Hu et al. 2019) to train both mGNN and bGNN. Following (You et al. 2018; Jin et al. 2019; Zhou et al. 2019; Jin,

Barzilay, and Jaakkola 2020b; Fu, Xiao, and Sun 2020; Fu et al. 2020a), we focus on the molecular properties below. For all scores, the higher the better.

- QED (Quantitative Estimate of Drug likeness) is an indicator of drug-likeness (Bickerton et al. 2012).
- DRD (Dopamine Receptor) measures a molecule’s biological activity against a biological target dopamine type 2 receptor (Comings, Muhleman, and Gysin 1996).
- PLogP (Penalized LogP) is the log of the partition ratio of the solute between octanol and water minus the synthetic accessibility score and number of long cycles (Ertl and Schuffenhauer 2009).

Note that PLogP is more sensitive to the change of local molecule structures, while DRD and QED are related to both local and global molecule structures. For chemically valid molecules, their QED and LogP scores can be evaluated using the RDKit package (<https://www.rdkit.org/>). DRD2 can be evaluated using well-trained model (Jin et al. 2019; Fu, Xiao, and Sun 2020; Fu et al. 2020a).

Baseline Methods. We compare MIMOSA with the following molecule optimization baselines. The parameter setting of these methods are provided in Fu et al. (2020b).

- JTVAE (Junction Tree Variational Auto-Encoder) (Jin, Barzilay, and Jaakkola 2018) is a generative model that learns latent space to generate desired molecule. It also uses an encoder-decoder architecture and leverage a junction tree to simplify molecule generation procedure.
- VJTNN (Variational Junction Tree Encoder-Decoder) (Jin et al. 2019) improves over JTVAE by leveraging adversarial learning and attention.
- GCPN (Graph Convolutional Policy Network) (You et al. 2018). GCPN is state-of-the-art reinforcement learning based approach on molecule optimization.
- GA (Genetic Algorithm) (Nigam et al. 2020) is a genetic algorithm that explores chemical space efficiently.

Details on Implementation, Features, Dataset Construction, Evaluation Strategies are in Fu et al. (2020b).

Metrics We consider the following metrics for evaluation.

- Similarity between the input and generated molecule, measured by Tanimoto similarity over Morgan fingerprints (Rogers and Hahn 2010), defined in Eq. (2).
- Property Improvement of generated molecule in QED, DRD, and PLogP. It is defined as the difference of the property score between generated molecules Y and input molecule X , i.e., $\text{property}(Y) - \text{property}(X)$.
- Success Rate (SR) based on similarity and property improvement between input molecule X and generated molecule Y . We follow the same definitions of SR as in (Jin et al. 2019) (See details in Fu et al. (2020b)).

4.2 Results

Exp 1. Optimize Multiple Properties

To evaluate model performance on optimizing multiple drug properties, we consider the following combinations of property constraints:

- (1) optimize QED (drug likeness) and PLogP (solubility);
- (2) optimize DRD (biological activity against dopamine type 2 receptor) and PLogP (solubility).

Optimizing PLogP and QED				
Method	Similarity	PLogP-Imp.	QED-Imp.	Success
JTVAE	0.16±0.08	0.14±0.27	0.01±0.10	0.4%
VJTNN	0.17±0.06	0.46±0.35	0.02±0.09	1.0%
GCPN	0.25±0.15	0.56±0.25	0.06±0.08	11.3%
GA	0.35±0.16	0.93±0.67	0.09±0.07	24.9%
MIMOSA	0.42±0.17	0.93±0.48	0.10±0.09	32.0%
Optimizing PLogP and DRD				
Method	Similarity	PLogP-Imp.	DRD-Imp.	Success
JTVAE	0.18±0.08	0.20±0.18	0.18±0.09	0.8%
VJTNN	0.18±0.08	0.55±0.16	0.27±0.05	3.4%
GCPN	0.23±0.12	0.38±0.25	0.25±0.11	20.4%
GA	0.38±0.16	0.68±0.49	0.20±0.16	29.3%
MIMOSA	0.54±0.16	0.75±0.48	0.35±0.20	43.7%

Table 2: Exp 1. Optimizing Multiple Properties.

From Table 2, MIMOSA has significantly better and stable performance on all metrics, with 28.5% relative higher success rate in optimizing both QED and PLogP, and 49.1% relative higher success rate in optimizing both DRD and PLogP compared with the second best algorithm GA. The GA algorithm uses genetic algorithm for local structure editing, hence is expected to work well on optimizing properties that are sensitive to local structural changes, such as joint optimizing both QED and PLogP where PLogP is related to the polarity of a molecule and is sensitive to the change of local structure. Because of the local editing of GA, GA does not perform well on optimizing both DRD and PLogP since DRD is less sensitive to the change of local structures.

Optimizing QED			
Method	Similarity	QED-Improve	Success
JTVAE	0.30±0.09	0.17±0.12	17.4%
VJTNN	0.37±0.11	0.20±0.05	37.6%
GCPN	0.32±0.14	0.20±0.09	26.5%
GA	0.43±0.17	0.17±0.11	42.5%
MIMOSA	0.50±0.30	0.20±0.14	47.8%
Optimizing DRD			
Method	Similarity	DRD-Improve	Success
JTVAE	0.31±0.07	0.34±0.17	25.6%
VJTNN	0.36±0.09	0.40±0.20	40.5%
GCPN	0.30±0.07	0.35±0.20	27.8%
GA	0.46±0.14	0.25±0.10	37.5%
MIMOSA	0.57±0.29	0.43±0.29	48.3%
Optimizing PLogP			
Method	Similarity	PLogP-Improve	Success
JTVAE	0.30±0.09	0.28±0.17	2.9%
VJTNN	0.38±0.08	0.47±0.24	14.3%
GCPN	0.32±0.07	0.33±0.19	7.8%
GA	0.53±0.15	0.99±0.54	92.8%
MIMOSA	0.56±0.17	0.94±0.47	94.0%

Table 3: Exp 2. Optimizing Single Property.

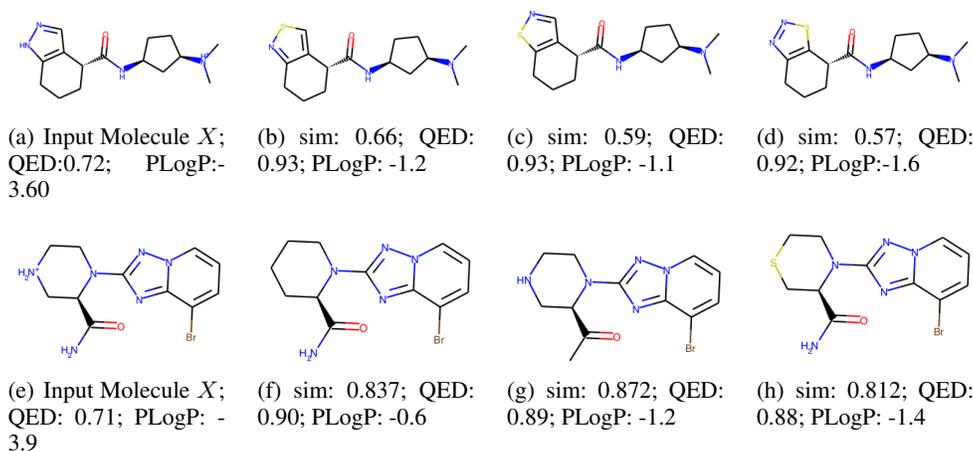


Figure 2: Exp 3. Examples of “QED & PLogP” optimization. (Upper), the imidazole ring in the input molecule (a) is replaced by less polar rings thiazole (b and c) and thiadiazol (d). Since more polar indicates lower PLogP, the output molecules increase PLogP while maintaining the molecular scaffold. (Lower), the PLogP of input molecule (e) is increased by neutralizing the ionized amine (g) or replacing with substructures with less electronegativity (f and h). These changes improve the QED.

Exp 2. Optimize Single Property

Since most baseline models were designed to optimize single drug properties, we also conduct experiments to compare MIMOSA with them on optimizing the following single properties: (1) DRD; (2) QED and (3) PLogP.

From the results shown in Table 3, we can see that when optimizing a single drug property, MIMOSA still achieved the best performance overall, with 12.5% relative higher success rate in optimizing QED compared with the second best model GA, and 28.8% relative higher success rate in optimizing both DRD compared with the second best algorithm VJTNN. Among the baseline models, algorithms such as JTVAE, VJTNN, and GCPN that were designed to optimize single property have good performance in property improvement as expected, however they generate molecules that have lower similarity hence the final success rates. Also, GA has the lowest QED and DRD improvement maybe due to its limitation in capturing global properties. High similarity between the output and input molecules is a unique requirement for the molecule optimization task, on which MIMOSA significantly outperformed the other baselines.

Exp 3. Case Study

To further examine how MIMOSA can also effectively improve properties that are sensitive to local structural change, e.g., PLogP, we show two examples in Fig. 2. For the first row, the imidazole ring in the input molecule (a) is replaced by less polar five-member rings thiazole (b and c) and thiadiazol (d). Since PLogP is related to the polarity of a molecule: more polar indicates lower PLogP. The generation results in the increase of PLogP while maintaining the molecular scaffold. For the second row, the PLogP of input molecule (e) is increased by neutralizing the ionized amine (g) or replacing with substructures with less electronegativity (f and h). These changes would also help improve the drug likeness,

i.e., QED value.

Sampling Efficiency. The sampling complexity is $O(NN_2)$ where N the size of candidate set (e.g., 20) and N_2 is the size of the possible proposal set (< 200). Empirically, this entire sampling process takes about 10-20 minutes for optimizing one source molecule, which is very acceptable for molecule optimization. And MCMC can directly operate with an un-normalized distribution which is more efficient. Note that all the existing methods for molecule optimization also utilize RDKit in their learning process, either in preprocessing steps for creating training data (Jin, Barzilay, and Jaakkola 2018; Jin et al. 2019), or inside their training procedure such as using RDKit to evaluate reward for reinforcement learning (You et al. 2018; Popova, Isayev, and Tropsha 2018; Zhou et al. 2019).

5 Conclusion

In this paper, we proposed MIMOSA, a new MCMC sampling based GNNs method for molecule optimization. MIMOSA pretrains GNNs and employs three basic substructure operations to generate new molecules and associated weights that can encode multiple drug property constraints, upon which we accept promising molecules for next iteration. MIMOSA iteratively produces new molecule candidates and can efficiently draw molecules that satisfy all constraints. MIMOSA significantly outperformed several state of the arts baselines for molecule optimization with 28.5% to 49.1% improvement when optimizing PLogP+QED, and PLogP+DRD, respectively.

Acknowledgments

This work was in part supported by the National Science Foundation award SCH-2014438, IIS-1418511, CCF-1533768, IIS-1838042, the National Institute of Health award NIH R01 1R01NS107291-01 and R56HL138415.

References

- Baskin, I.; and Varnek, A. 2009. Fragment descriptors in SAR/QSAR/QSPR studies, molecular similarity analysis and in virtual screening. *ChemInform* 40(20): i.
- Bickerton, G. R.; Paolini, G. V.; Besnard, J.; Muresan, S.; and Hopkins, A. L. 2012. Quantifying the chemical beauty of drugs. *Nature chemistry* 4(2): 90.
- Blaschke, T.; Olivecrona, M.; Engkvist, O.; Bajorath, J.; and Chen, H. 2018. Application of generative autoencoder in de novo molecular design. *Mol. Inform.* .
- Brooks, S.; Gelman, A.; Jones, G.; and Meng, X.-L. 2011. *Handbook of markov chain monte carlo*. CRC press.
- Comings, D.; Muhleman, D.; and Gysin, R. 1996. Dopamine D2 receptor (DRD2) gene and susceptibility to posttraumatic stress disorder: a study and replication. *Biological psychiatry* 40(5): 368–372.
- Dai, H.; Tian, Y.; Dai, B.; Skiena, S.; and Song, L. 2018. Syntax-directed variational autoencoder for structured data. *International Conference on Learning Representations (ICLR)* .
- Ertl, P.; and Schuffenhauer, A. 2009. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of cheminformatics* 1(1): 8.
- Fu, T.; Xiao, C.; Glass, L.; and Sun, J. 2021. MOLER: Incorporate Molecule-Level Reward to Enhance Deep Generative Model for Molecule Optimization. *IEEE Transactions on Knowledge and Data Engineering* .
- Fu, T.; Xiao, C.; Glass, L. M.; and Sun, J. 2020a. α -MOP: Molecule optimization with α -divergence. In *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 240–244. IEEE.
- Fu, T.; Xiao, C.; Li, X.; Glass, L. M.; and Sun, J. 2020b. MI-MOSA: Multi-constraint Molecule Sampling for Molecule Optimization. *arXiv preprint arXiv:2010.02318* .
- Fu, T.; Xiao, C.; and Sun, J. 2020. CORE: Automatic molecule optimization using copy & refine strategy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 638–645.
- Geman, S.; and Geman, D. 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence* 6: 721–741.
- Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; and Aspuru-Guzik, A. 2018. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science* 4(2): 268–276.
- Hu, W.; Liu, B.; Gomes, J.; Zitnik, M.; Liang, P.; Pande, V.; and Leskovec, J. 2019. Strategies for Pre-training Graph Neural Networks. In *International Conference on Learning Representations (ICLR)*.
- Huang, K.; Fu, T.; Glass, L.; Zitnik, M.; Xiao, C.; and Sun, J. 2020a. DeepPurpose: a deep learning library for drug-target interaction prediction and applications to repurposing and screening. *Bioinformatics* .
- Huang, K.; Fu, T.; Khan, D.; Abid, A.; Abdalla, A.; Abid, A.; Glass, L. M.; Zitnik, M.; Xiao, C.; and Sun, J. 2020b. MolDesigner: Interactive Design of Efficacious Drugs with Deep Learning. *NEURIPS Demo* .
- Jin, W.; Barzilay, R.; and Jaakkola, T. 2018. Junction tree variational autoencoder for molecular graph generation. *ICML* .
- Jin, W.; Barzilay, R.; and Jaakkola, T. 2020a. Hierarchical generation of molecular graphs using structural motifs. In *International Conference on Machine Learning*, 4839–4848. PMLR.
- Jin, W.; Barzilay, R.; and Jaakkola, T. 2020b. Multi-objective molecule generation using interpretable substructures. In *International Conference on Machine Learning*, 4849–4859. PMLR.
- Jin, W.; Yang, K.; Barzilay, R.; and Jaakkola, T. 2019. Learning Multimodal Graph-to-Graph Translation for Molecular Optimization. *International Conference on Learning Representations (ICLR)* .
- Kusner, M. J.; Paige, B.; and Hernández-Lobato, J. M. 2017. Grammar variational autoencoder. In *ICML, 1945–1954*. JMLR. org.
- Levine, R. A.; and Casella, G. 2006. Optimizing random scan Gibbs samplers. *Journal of Multivariate Analysis* 97(10): 2071–2100.
- Liu, J. S.; Liang, F.; and Wong, W. H. 2000. The multiplicity method and local optimization in Metropolis sampling. *Journal of the American Statistical Association* 95(449): 121–134.
- Liu, Q.; Allamanis, M.; Brockschmidt, M.; and Gaunt, A. 2018. Constrained graph variational autoencoders for molecule design. In *Advances in Neural Information Processing Systems*, 7795–7804.
- Nigam, A.; Friederich, P.; Krenn, M.; and Aspuru-Guzik, A. 2020. Augmenting Genetic Algorithms with Deep Neural Networks for Exploring the Chemical Space. In *International Conference on Learning Representations (ICLR)*.
- Olivecrona, M.; Blaschke, T.; Engkvist, O.; and Chen, H. 2017. Molecular de-novo design through deep reinforcement learning. *J. Cheminform.* 9(1): 48.
- Polishchuk, P. G.; Madzhidov, T. I.; and Varnek, A. 2013. Estimation of the size of drug-like chemical space based on GDB-17 data. *Journal of computer-aided molecular design* 27(8): 675–679.
- Popova, M.; Isayev, O.; and Tropsha, A. 2018. Deep reinforcement learning for de novo drug design. *Science advances* 4(7): eaap7885.
- Putin, E. 2018. Reinforced adversarial neural computer for de novo molecular design. *J. Chem. Inf. Model.* .

Rogers, D.; and Hahn, M. 2010. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling* 50(5): 742–754.

Sterling, T.; and Irwin, J. J. 2015. ZINC 15–ligand discovery for everyone. *Journal of chemical information and modeling* 55(11): 2324–2337.

Welling, M.; and Teh, Y. W. 2011. Bayesian learning via stochastic gradient Langevin dynamics. In *ICML*, 681–688.

You, J.; Liu, B.; Ying, R.; Pande, V.; and Leskovec, J. 2018. Graph Convolutional Policy Network for Goal-directed Molecular Graph Generation. In *Advances in Neural Information Processing Systems*, 6412–6422.

Zhou, Z.; Kearnes, S.; Li, L.; Zare, R. N.; and Riley, P. 2019. Optimization of molecules via deep reinforcement learning. *Scientific reports* 9(1): 10752.