GRADIENT GA: GRADIENT GENETIC ALGORITHM FOR DRUG MOLECULAR DESIGN

Chris Zhuang[1]*, **Debadyuti Mukherjee**[1]*, **Yingzhou Lu**[2], **Tianfan Fu**[3] & **Ruqi Zhang**[1]
[1]Purdue University
[2]Stanford University
[3]Rensselaer Polytechnic Institute

ABSTRACT

Molecular discovery has brought great benefits to the chemical industry. Various molecule design techniques are developed to identify molecules with desirable properties. Traditional optimization methods, such as genetic algorithms, continue to achieve state-of-the-art results across multiple molecular design benchmarks. However, these techniques rely solely on random walk exploration, which hinders both the quality of the final solution and the convergence speed. To address this limitation, we propose a novel approach called *Gradient Genetic Algorithm* (Gradient GA), which incorporates gradient information from the objective function into genetic algorithms. Instead of random exploration, each proposed sample iteratively progresses toward an optimal solution by following the gradient direction. We achieve this by designing a differentiable objective function parameterized by a neural network and utilizing the Discrete Langevin Proposal to enable gradient guidance in discrete molecular spaces. Experimental results demonstrate that our method significantly improves both convergence speed and solution quality, outperforming cutting-edge techniques. For example, it achieves up to a 25% improvement in the top-10 score over the vanilla genetic algorithm. The code is publicly available at `https://github.com/debadyuti23/GradientGA`.

1 INTRODUCTION

Designing molecules with desirable biological and chemical properties has become a demanding research topic since its outcome can benefit various domains, such as drug discovery (Huang et al., 2022), material design (Yang et al., 2017), etc. However, a limited number of molecules can be tested in real-life laboratories Altae-Tran et al. (2017) and clinical trials (Chen et al., 2024b;a). Therefore, numerous effective techniques for molecule discovery are proposed to discover favorable molecules throughout the vast sample space. Some evolutionary algorithms, such as the molecular graph-based genetic algorithm (Graph GA) (Jensen, 2019a), remain strong performance, often outperforming recently proposed machine learning-based algorithms (Huang et al., 2021; Gao et al., 2022b). Genetic algorithms are cheap, easy to implement, and are often regarded as simple baselines for molecular discovery. However, key GA operators, such as selection, crossover, and mutation, are random and do not leverage knowledge of the objective function. Given the vast molecular search space, this random walk approach is like searching for a needle in a haystack. As a result, GA tends to converge slowly and its final performance can be unstable. To address this issue, we introduce a novel molecule design method, *Gradient Genetic Algorithm* (Gradient GA), which leverages gradient information to navigate chemical space efficiently. First, we learn a differentiable objective function using a Graph Neural Network (GNN) (Scarselli et al., 2009), which maps the graph-structured information of molecules to vector embeddings. We then apply the Discrete Langevin Proposal (DLP) (Zhang et al., 2022) to incorporate gradient information from this objective, enabling more informed exploration in the discrete molecular space. Our main contributions are summarized as follows:

- We introduce Gradient GA, a gradient-based genetic algorithm for more informative and effective exploration in molecular spaces, mitigating the random-walk behavior in genetic algorithms. To

*Equal Contribution, Emails:{zhuang80, mukher83}@purdue.edu

the best of our knowledge, this is the first method to leverage gradient information for DLP within a genetic algorithm framework.

- The experimental results demonstrate that the proposed method achieves a significant and consistent improvement over a number of cutting-edge approaches (e.g., Graph-GA, SMILES-GA), for example, achieving an improvement of up to 25% over the traditional genetic algorithm when optimizing the mestranol similarity property.

2 Related Work

2.1 AI-aided Drug Molecular Design

Current AI-aided drug molecular design techniques can be primarily classified into two categories: deep generative models and combinatorial optimization methods. (I) Deep Generative Models (DGMs) model the distribution of general molecular structures using deep network models, enabling the generation of molecules by sampling from the learned distribution. Typical algorithms include Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), energy-based models, and flow-based models (Gómez-Bombarelli et al., 2018; Jin et al., 2018; De Cao & Kipf, 2018; Segler et al., 2018; Fu et al., 2020; Honda et al., 2019; Madhawa et al., 2019; Liu et al., 2021; Fu & Sun, 2022; Chen et al., 2024b; Bagal et al., 2021). However, these approaches often require a smooth and discriminative latent space, necessitating careful network architecture design and well-distributed datasets. This requirement can be restrictive in certain scenarios, such as multi-objective optimization. Furthermore, since DGMs learn the distribution of reference data, their ability to explore diverse chemical space is relatively limited, as demonstrated by recent molecular optimization benchmarks (Brown et al., 2019; Huang et al., 2021; Gao et al., 2022b). (II) On the other hand, combinatorial optimization methods directly search the discrete chemical space, mainly including deep reinforcement learning (DRL) (You et al., 2018; Zhou et al., 2019; Jin et al., 2020; Gottipati et al., 2020), evolutionary learning methods (Nigam et al., 2020; Jensen, 2019b; Fu et al., 2022a) and sampling methods (Xie et al., 2021; Fu et al., 2021). Specifically, Jensen (2019a) have proposed a molecular graph-based Genetic Algorithm (Graph GA). In the context of drug discovery, the algorithm samples two parent molecules and produces children molecules based on different combinations of the parent molecules with a chance of a random action or mutation to happen to the children. The entire population is then cut down based off highest scores of the entire population. Also, MARS have sampled potential molecules through the usage of Markov Chain Monte Carlo Method (MCMC) (Xie et al., 2021). Each sample molecule, forms a Markov Chain formulated as a chemical product of the previous sample. The formation of the product has been assumed between two choices- (i) the addition of fragments, and (ii) the removal of a chemical bond.

2.2 Discrete Langevin Proposal

Suppose that the target distribution is $\pi(v) \propto \exp(U(v))$, where $U(\cdot)$ is the energy function, v is an n-dimensional variable in the space $S \subseteq \mathbb{R}^n$. Langevin Dynamics samples from π by iteratively updating v as follows:

$$v' = v + \frac{\alpha}{2}\nabla U(v) + \sqrt{\alpha}\varepsilon, \;\; \varepsilon \sim \mathcal{N}(0, I_{n\times n}), \tag{1}$$

where α is the step size; $I_{n\times n}$ is n-dimensional identity matrix; $\mathcal{N}(\cdot,\cdot)$ denotes high-dimensional normal distribution. From Equation 1, it can inferred that v' follows a Gaussian distribution with mean $v + \frac{\alpha}{2}\nabla U(v)$ and covariance $\alpha I_{n\times n}$. Therefore, the probability of selecting v' from V, i.e., $p(v'|v)$, can be rewritten as

$$p(v'|v) = \frac{\exp\left(-\frac{1}{2\alpha}||v' - v - \frac{\alpha}{2}\nabla U(v)||_2^2\right)}{Z}, \tag{2}$$

where Z is the normalizing constant. The distribution $p(v'|v)$ has its mean shifted from v towards optimum due to the gradient factor $\frac{\alpha}{2}\nabla U(v)$ being present. Therefore, the high-probability samples from $p(v'|v)$ will be near the optimum compared to v as the Gaussian distribution has its density centered around the mean.

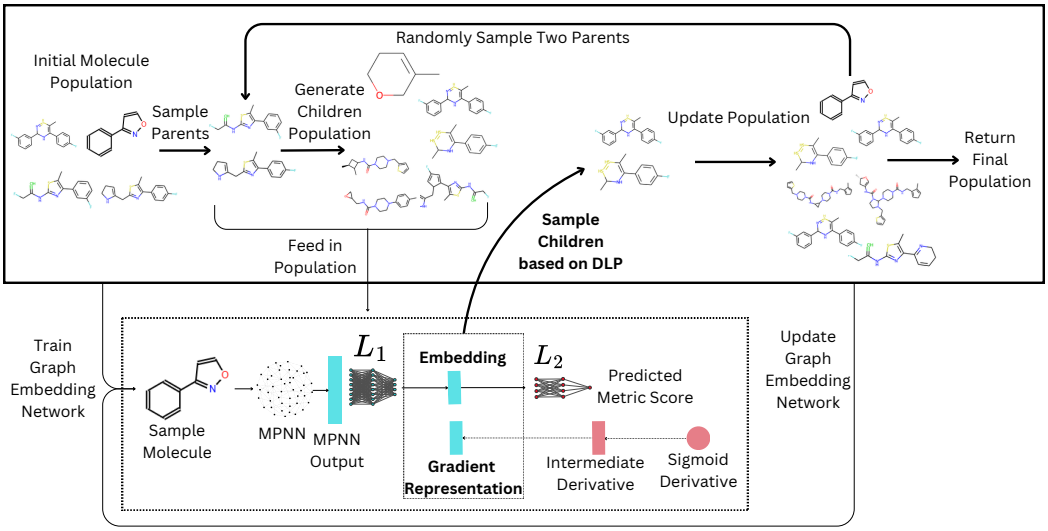

Figure 1: Gradient GA pipeline.

3 Methodology: Gradient Genetic Algorithm

Drug molecular design aims at identifying novel molecules with desirable pharmaceutical properties, which are evaluated by *oracle*. Oracles serve as objective functions in molecular optimization tasks, formally defined as follows.

Definition 3.1 (Oracle)**.** Oracle $\mathcal{O}(\cdot) : \mathcal{Q} \to \mathbb{R}$ is a black-box function that evaluates certain physical, chemical, or biological properties of a molecule X and yields an approximation of the ground-truth property $\mathcal{O}(X)$.

Mathematically, the drug molecular design problem (a.k.a. molecular optimization) can be formulated as

$$\arg\max_{X \in \mathcal{Q}} \mathcal{O}(X), \tag{3}$$

where X is a molecule, $\mathcal{Q}$ denotes the whole molecular space, i.e., the set of all the chemically valid molecules. The size of the whole molecular space is around 10^{60} Bohacek et al. (1996). Following MARS (Xie et al., 2021), we regard $\mathcal{O}(\cdot)$ as an unnormalized probability distribution and introduce a vector embedding v for each molecule X. The target distribution is then defined as $\pi(v) \propto \mathcal{O}(X)$.

Gradient Definition in Molecular Design. We now define the gradient information used to guide exploration in molecular spaces. To apply (discrete) Langevin dynamics, we need the gradient of the energy function. Since $\pi(v) \propto \exp(U(v))$ and given that $\pi(v) \propto \mathcal{O}(X)$, using chain rule, we have

$$\nabla U(v) = \frac{\nabla \mathcal{O}(X)}{\mathcal{O}(X)} = \frac{\nabla f(v)}{\mathcal{O}(X)}, \tag{4}$$

where f is a differentiable function that approximates the oracle $\mathcal{O}$, i.e., $f(v) = \mathcal{O}(X)$ for all $X \in \mathcal{Q}$. It is also possible to define $U(v) = \mathcal{O}(X)$ and $\nabla U(v) = \nabla f(v)$. Our approach includes $\mathcal{O}(X)$ in the denominator, effectively playing the role of adaptive step sizes. We found that this formulation leads to better performance. An empirical comparison is provided in the Appendix C.

Implementing gradient-based methods in molecular discovery is a challenging task due to two primary obstacles: (1) representing sample molecules in a vector format suitable for gradient-based methods, and (2) establishing a differentiable relationship between the probability distribution and the vector representation.

Finding the Embedding. We use the Message Passing Neural Network (MPNN) (Gilmer et al., 2017), which is a state-of-the-art approach for molecular activity prediction. We adopt the MPNN

architecture from Xie et al. (2021), which consists of (1) a simple neural network for message passing between neighboring nodes and (2) a Gated Recurrent Unit (GRU) (Cho, 2014) for updating node representations. For the readout function, we use Set2Set (Vinyals et al., 2015), which is particularly effective for isomorphic graphs due to its set-invariant property. To reduce the dimensionality of the Set2Set output, we introduce a two-layer multilayer perceptron (MLP) L_1, which compresses the output dimension from $2n$ to n. We define the graph representation function, which converts molecular data into graph data (nodes representing atomic features and edges representing bond features), as $G(\cdot)$. The output of L_1 becomes the embedding v for the molecular graph $G_X = G(X)$.

Finding the Gradient. We can learn the differentiable approximation $f(\cdot)$ in Equation 4 using an MLP-based architecture. After the first two-layer MLP L_1, we introduce another two-layer MLP L_2, which produces a scalar output $\widehat{y}$. The gradient $\nabla f(\cdot)$ can then be computed through backpropagation, as illustrated in Figure 1. Our objective is to make $\widehat{y}$ approximate the behavior of the oracle function $\mathcal{O}$ so that the entire architecture (from the MPNN to L_2), denoted by $\mathcal{M}$, effectively learns both the embedding and the differentiable function. To introduce nonlinearity, we apply LeakyReLU and sigmoid activation functions before and after L_2, respectively. With v and $\nabla U(\cdot)$ now properly defined for Equation 2, we proceed to describe the training procedure. We train the model $\mathcal{M}$ to fit the oracle function $\mathcal{O}$ using the initial molecule population D. The loss function is defined as the mean squared error. During the molecular optimization process, we continuously expand the training set by adding newly generated molecules that surpass a threshold criterion T. This provides additional information about the distribution and allows us to retrain the model $\mathcal{M}$.

Iterative Sampling. We propose a sampling technique inspired by both Graph GA (Jensen, 2019a) and DLP (Zhang et al., 2022). The workflow for each iteration is illustrated in Figure 11. Similar to Graph GA, we begin by selecting parent molecules based on their scores $\mathcal{O}(\cdot)$ and generating child molecules through crossover. We define the sample space S for DLP as the set of all possible crossovers between the selected parents d_1 and d_2 ($\in D$). Ideally, both parents should be considered as current samples. However, since DLP is designed to use a single sample, we aggregate the information from both parents into a single embedding v and gradient $\nabla U(v)$ using the following equation:

$$\{v, \nabla U(v)\} = \sum_{i=1,2} w_i \cdot \{v_i, \nabla U(v_i)\}, \tag{5}$$

where $\{v_i, \nabla U(v_i)\}$ represents the embedding and gradient information for parent d_i. Empirically, we found that a simple strategy of using only the best parent as the current sample works well. Specifically, the weights w_i are assigned as follows: $w_i = 1$ if $i = \arg\max(\mathcal{O}(d_i))$, and $w_i = 0$ otherwise. Applying DLP updates the embedding v' by moving it closer to the optimum, guided by the gradient information. Figure 11 illustrates this process. In the final step, DLP generates the next sample set D' of fixed size k from the sample space S. Following the Graph GA approach, each molecule in D' is mutated. Before the next iteration begins, we update both the population and the model. The population D is refreshed by selecting the top $|D|$ molecules based on their oracle scores from the combined set $\{D, D'\}$. To further enhance the graph embedding model $\mathcal{M}$'s understanding of $\pi(\cdot)$, we retrain $\mathcal{M}$ using a training set D'', which is updated in each iteration according to the following rule:

$$D'' = D'' \cup \{d \mid d \in D' \text{ and } T(d)\}, \tag{6}$$

where $T(d)$ is a threshold criterion for adding new samples to the training set. The complete Gradient GA workflow is detailed in Algorithm 1.

4 EXPERIMENTS

Baseline Methods. We use the practical molecule optimization (PMO) benchmark (Gao et al., 2022a) as our code base to compare results between the state-of-the-art methods. We select (1) genetic algorithm, including Graph GA (molecular graph-level genetic algorithm) method (Jensen, 2019b) and SMILES GA (SMILES string-level), (2) sampling-based methods, including MIMOSA (Multi-constraint Molecule Sampling) (Fu et al., 2021), MARS (Markov Molecular Sampling) (Xie et al., 2021) and (3) gradient-based method, DST (Differentiable Scaffolding Tree) (Fu et al., 2022b).

Dataset. For all the methods, we use ZINC 250K database Irwin et al. (2012) to select initial molecule population, extract chemical fragments, and pretraining.

Table 1: Comparison of Average Top 10, AUC Top 1, AUC Top 10, and AUC Top 100 with several GuacaMol objectives (mestranol similarity, amlodipine MPO, perindopril MPO, deco hop, median1, and isomers c9h10n2o2pf2cl) under 2500 oracle calls. For each metric, the best method is **bolded**. We conduct five independent runs using different random seeds for each method, and report the average scores and their standard deviation.

Method	mestranol similarity				amlodipine MPO			
	Average Top 10	AUC Top 1	AUC Top 10	AUC Top 100	Average Top 10	AUC Top 1	AUC Top 10	AUC Top 100
Gradient GA	**0.5130±0.0393**	**0.4433±0.0310**	**0.4082±0.0315**	**0.3534±0.0355**	**0.5667±0.0336**	**0.5614±0.0177**	**0.5176±0.0187**	0.4658±0.0199
Graph GA	0.4452±0.0241	0.3556±0.0268	0.3208±0.0199	0.2717±0.0147	0.5605±0.0364	0.5067±0.0270	0.4734±0.0215	0.4152±0.0142
SMILES GA	0.2582±0.0097	0.3777±0.0381	0.3634±0.0352	0.3347±0.0279	0.4480±0.0161	0.5016±0.0156	0.4956±0.0143	**0.4748±0.0158**
MIMOSA	0.4262±0.0246	0.4162±0.0115	0.3619±0.0181	0.2887±0.0252	0.5245±0.0143	0.5431±0.0261	0.4953±0.0109	0.4436±0.0075
MARS	0.3411±0.0160	0.3760±0.0003	0.3215±0.0096	0.2523±0.0081	0.4843±0.0210	0.4812±0.0144	0.4583±0.0098	0.3816±0.0157
DST	0.4131±0.0179	0.4148±0.0323	0.3507±0.0088	0.2780±0.0029	0.5192±0.0122	0.5411±0.0303	0.4908±0.0115	0.4257±0.0044

	perindopril MPO				deco hop			
	Average Top 10	AUC Top 1	AUC Top 10	AUC Top 100	Average Top 10	AUC Top 1	AUC Top 10	AUC Top 100
Gradient GA	0.4786±0.0257	**0.4542±0.0164**	**0.4361±0.0176**	0.3882±0.0193	0.6026±0.0053	**0.5883±0.0032**	0.5763±0.0050	0.5602±0.0053
Graph GA	**0.4788±0.0067**	0.4519±0.0055	0.4317±0.0045	0.3770±0.0049	**0.6039±0.0043**	0.5186±0.0037	0.5028±0.0032	0.4708±0.0033
SMILES GA	0.3698±0.0117	0.4346±0.0124	0.4271±0.0115	**0.4065±0.0102**	0.5548±0.0059	0.5862±0.0047	**0.5817±0.0042**	**0.5733±0.0036**
MIMOSA	0.4629±0.0176	0.4500±0.0144	0.4289±0.0116	0.3783±0.0085	0.6008±0.0053	0.5882±0.0061	0.5773±0.0035	0.5600±0.0021
MARS	0.4564±0.0167	0.4538±0.0087	0.4278±0.0065	0.3648±0.0042	0.5944±0.0070	0.5830±0.0227	0.5711±0.0301	0.5493±0.0421
DST	0.4615±0.0100	0.4530±0.0041	0.4210±0.0041	0.3564±0.0028	0.6034±0.0083	0.5860±0.0071	0.5721±0.0025	0.5518±0.0009

	median1				isomers c9h10n2o2pf2cl			
	Average Top 10	AUC Top 1	AUC Top 10	AUC Top 100	Average Top 10	AUC Top 1	AUC Top 10	AUC Top 100
Gradient GA	**0.3033±0.0074**	**0.2581±0.0115**	**0.2298±0.0151**	**0.1906±0.0183**	0.7783±0.0959	0.6628±0.0731	0.5444±0.0693	0.4033±0.0614
Graph GA	0.2599±0.0182	0.2315±0.0206	0.1959±0.0148	0.1442±0.0070	0.7222±0.1119	0.6648±0.0957	0.5436±0.0770	0.3891±0.0425
SMILES GA	0.1310±0.0172	0.1832±0.0281	0.1795±0.0272	0.1697±0.0251	0.3180±0.3583	**0.8244±0.0848**	**0.7825±0.0752**	**0.7055±0.0668**
MIMOSA	0.2391±0.0080	0.2271±0.0103	0.1969±0.0044	0.1537±0.0030	**0.7866±0.0824**	0.6965±0.0562	0.5949±0.0440	0.3965±0.0265
MARS	0.2094±0.0181	0.2239±0.0140	0.2019±0.0116	0.1671±0.0158	0.6639±0.1606	0.6751±0.1032	0.5909±0.1057	0.4424±0.1499
DST	0.2179±0.0162	0.2097±0.0086	0.1765±0.0021	0.1331±0.0024	0.6748±0.0304	0.6305±0.0435	0.4932±0.0216	0.2293±0.0093

Evaluation Metrics. (1) **Average Top-K** ($K = 10$) is top-K average property value, which measures the algorithm's optimization ability. We limit the number of oracle calls to 2,500 to mimic the real experimental setup, though we expect methods to optimize well within hundreds of calls. (2) **AUC top-K** ($K = 1, 10, 100$). To consider optimization ability and sample efficiency simultaneously, we report the area under the curve (AUC) of top-K average property value of the top K versus the number of oracle calls (*AUC top-K*) All these metrics can be calculated via the evaluation function in Therapeutics data commons (TDC) (Huang et al., 2021; 2022)[1].

Experimental Results. To assess our method overall, we look at various metrics over multiple oracles. In Table 1, we look at an evaluation comparison between the Average Top 10 and AUC Top-1, 10, 100. From Table 1, Gradient GA performs better than Graph GA, MIMOSA, MARS, and SMILES GA. In the Average Top 10, Gradient GA was the best performing. We believe that for these oracles, the gradient information allows the method to better explore the local space for these oracles to converge to a better Average score. This is further supported by the AUC Top-K scores where the AUC scores methods that find a higher score within fewer oracle calls. We can see that for the AUC Top 1 and Top 10, Gradient GA dominates these metrics, while for the AUC Top 100 it is tied in top scores with SMILES GA. This leads us to believe that for these metrics Gradient GA is the best method overall.

5 Conclusion and Future Work

Genetic Algorithm (GA) is dominating drug molecular design thanks to its flexibility to manipulate molecular space. However, GA usually suffers from slow convergence due to its random walk nature. We address this problem by introducing a novel approach called Gradient GA, in which each proposed sample iteratively progresses toward the optimal solution. Our method leverages Discrete Langevin Proposal (DLP) as the foundational sampler, enabling gradient-based exploration in the discrete molecular space. Thorough experimental results validate that our proposed approach demonstrates faster and superior convergence. Future work will expand the current method in the following aspects: (1) Explore using gradient information for generated molecules, as shown in Appendix B; (2) Explore better ways to fit both parents into DLP; (3) Explore more DLP-oriented molecular optimizations with Metropolis-Hastings criterion involved.

[1] `https://tdcommons.ai/functions/data_evaluation/` and `https://tdcommons.ai/functions/oracles/`

REFERENCES

Han Altae-Tran, Bharath Ramsundar, Aneesh S Pappu, and Vijay Pande. Low data drug discovery with one-shot learning. *ACS central science*, 3(4):283–293, 2017.

Viraj Bagal, Rishal Aggarwal, PK Vinod, and U Deva Priyakumar. LigGPT: Molecular generation using a transformer-decoder model. 2021.

Regine S Bohacek, Colin McMartin, and Wayne C Guida. The art and practice of structure-based drug design: a molecular modeling perspective. *Medicinal research reviews*, 16(1):3–50, 1996.

Nathan Brown, Marco Fiscato, Marwin HS Segler, and Alain C Vaucher. Guacamol: benchmarking models for de novo molecular design. *Journal of chemical information and modeling*, 59(3): 1096–1108, 2019.

Jintai Chen, Yaojun Hu, Yue Wang, Yingzhou Lu, Xu Cao, Miao Lin, Hongxia Xu, Jian Wu, Cao Xiao, Jimeng Sun, et al. Trialbench: Multi-modal artificial intelligence-ready clinical trial datasets. *arXiv preprint arXiv:2407.00631*, 2024a.

Tianyi Chen, Nan Hao, Yingzhou Lu, and Capucine Van Rechem. Uncertainty quantification on clinical trial outcome prediction. *arXiv preprint arXiv:2401.03482*, 2024b.

Kyunghyun Cho. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.

Nicola De Cao and Thomas Kipf. MolGAN: An implicit generative model for small molecular graphs. *arXiv preprint arXiv:1805.11973*, 2018.

Tianfan Fu and Jimeng Sun. Antibody Complementarity Determining Regions (CDRs) design using constrained energy model. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 389–399, 2022.

Tianfan Fu, Cao Xiao, and Jimeng Sun. CORE: Automatic molecule optimization using copy and refine strategy. *AAAI*, 2020.

Tianfan Fu, Cao Xiao, Xinhao Li, Lucas M Glass, and Jimeng Sun. MIMOSA: Multi-constraint molecule sampling for molecule optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 125–133, 2021.

Tianfan Fu, Wenhao Gao, Connor W Coley, and Jimeng Sun. Reinforced genetic algorithm for structure-based drug design. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2022a.

Tianfan Fu, Wenhao Gao, Cao Xiao, Jacob Yasonik, Connor W Coley, and Jimeng Sun. Differentiable scaffolding tree for molecular optimization. *International Conference on Learning Representations*, 2022b.

Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor W. Coley. Sample efficiency matters: A benchmark for practical molecular optimization, 2022a. URL `https://arxiv.org/abs/2206.12411`.

Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor W Coley. Sample efficiency matters: benchmarking molecular optimization. *Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks*, 2022b.

Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, pp. 1263–1272. PMLR, 2017.

Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science*, 4(2):268–276, 2018.

Sai Krishna Gottipati, Boris Sattarov, Sufeng Niu, Yashaswi Pathak, Haoran Wei, Shengchao Liu, Simon Blackburn, Karam Thomas, Connor Coley, Jian Tang, et al. Learning to navigate the synthetically accessible chemical space using reinforcement learning. In *International Conference on Machine Learning*, pp. 3668–3679. PMLR, 2020.

Shion Honda, Hirotaka Akita, Katsuhiko Ishiguro, Toshiki Nakanishi, and Kenta Oono. Graph residual flow for molecular graph generation. *arXiv preprint arXiv:1909.13521*, 2019.

Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Therapeutics data commons: machine learning datasets and tasks for therapeutics. *NeurIPS Track Datasets and Benchmarks*, 2021.

Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Artificial intelligence foundation for therapeutic science. *Nature Chemical Biology*, pp. 1–4, 2022.

John J Irwin, Teague Sterling, Michael M Mysinger, Erin S Bolstad, and Ryan G Coleman. ZINC: a free tool to discover chemistry for biology. *Journal of chemical information and modeling*, 52(7): 1757–1768, 2012.

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

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

Wengong Jin, Regina Barzilay, and Tommi S. Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International Conference on Machine Learning (ICML)*, 2018.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Multi-objective molecule generation using interpretable substructures. In *International Conference on Machine Learning*, pp. 4849–4859. PMLR, 2020.

Meng Liu, Keqiang Yan, Bora Oztekin, and Shuiwang Ji. Graphebm: Molecular graph generation with energy-based models. *arXiv preprint arXiv:2102.00546*, 2021.

Kaushalya Madhawa et al. GraphNVP: An invertible flow model for generating molecular graphs. *arXiv*, 2019.

AkshatKumar Nigam, Pascal Friederich, Mario Krenn, and Alán Aspuru-Guzik. Augmenting genetic algorithms with deep neural networks for exploring the chemical space. In *The International Conference on Learning Representations (ICLR)*, 2020.

M. Olivecrona, T. Blaschke, O. Engkvist, and H. Chen. Molecular de-novo design through deep reinforcement learning. *Journal of Cheminformatics*, 2017.

Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009. doi: 10.1109/TNN.2008.2005605.

Marwin HS Segler, Thierry Kogej, Christian Tyrchan, and Mark P Waller. Generating focused molecule libraries for drug discovery with recurrent neural networks. *ACS central science*, 4(1): 120–131, 2018.

Oriol Vinyals, Samy Bengio, and Manjunath Kudlur. Order matters: Sequence to sequence for sets. *arXiv preprint arXiv:1511.06391*, 2015.

Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, and Lei Li. MARS: Markov molecular sampling for multi-objective drug discovery. In *ICLR*, 2021.

Xiufeng Yang, Jinzhe Zhang, Kazuki Yoshizoe, Kei Terayama, and Koji Tsuda. ChemTS: an efficient python library for de novo molecular generation. *Science and technology of advanced materials*, 18(1):972–976, 2017.

Jiaxuan You et al. Graph convolutional policy network for goal-directed molecular graph generation. In *Proceedings of the 32Nd International Conference on Neural Information Processing Systems*, pp. 6412–6422. Curran Associates Inc., 2018.

Ruqi Zhang, Xingchao Liu, and Qiang Liu. A langevin-like sampler for discrete distributions. In *International Conference on Machine Learning*, pp. 26375–26396. PMLR, 2022.

Zhenpeng Zhou, Steven Kearnes, Li Li, Richard N Zare, and Patrick Riley. Optimization of molecules via deep reinforcement learning. *Scientific reports*, 9(1):1–10, 2019.

A TABLE OF RESULTS BETWEEN GRADIENT GA AND OTHER BASELINES

The overall experimental setup includes 10,000 oracle calls, 5 runs, and early stopping enabled. The parameters for Gradient GAare set to 200 epochs, with D'' being cleared after each retraining. REINVENT is a reinforcement learning method Olivecrona et al. (2017). Additionally, we evaluate another metric, Diversity, which measures the Tanimoto Similarity between two molecules. The higher the score the better for all metrics except Average SA.

From the results, we can see that Gradient GA and Graph GA are near the top in terms of performance with REINVENT. These results follow the results gathered from PMO Gao et al. (2022b). We notice that for Gradient GA, it usually performs better than Graph GA whenever it is dealing with a molecular objective that is on the lower end, so the exploration near optimal molecules matter more than as random walk behavior may have a harder time finding the optimal molecules. Overall, the best performing models are REINVENT, Gradient GA, and Graph GA.

Table 2: Zaleplon_MPO Results

zaleplon_mpo	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.5156±0.0276	0.4857±0.0177	0.4489±0.0124	0.4728±0.0246	0.4435±0.0165	0.3917±0.0141	3.4679±0.3267	0.7208±0.0373
Graph GA	0.5190±0.0368	0.4993±0.0246	0.4509±±0.0169	0.4933±0.0295	0.4607±0.0163	0.3988±0.0115	3.0682±0.2047	0.7549±0.0399
SMILES GA	0.5000±0.0364	0.4915±0.0361	0.4735±0.0445	0.4795±0.0309	0.4652±0.0299	0.4397±0.0364	3.4405±0.2295	0.5248±0.1233
MARS	0.4544±0.0198	0.3972±0.0600	0.2850±0.1221	0.4429±0.0176	0.3772±0.0494	0.2510±0.0979	3.1777±0.3551	0.8546±0.0161
MIMOSA	0.4031±0.0537	0.3866±0.0543	0.3501±0.0547	0.2737±0.0314	0.2560±0.0303	0.2211±0.0251	3.1036±0.5144	0.7343±0.0349
REINVENT	0.5915±0.0331	0.5884±0.0334	0.5709±0.0303	0.3711±0.0729	0.3545±0.0697	0.3207±0.0636	2.2110±0.2345	0.5620±0.0280

Table 3: troglitazone_rediscovery Results

troglitazone_rediscovery	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.4623±0.0301	0.4482±0.0233	0.4195±0.0181	0.36757±0.0219	0.35207±0.0200	0.32767±0.0188	4.07857±0.3021	0.51887±0.0210
Graph GA	0.4157±0.0497	0.3978±0.0455	0.3711±0.0405	0.3662±0.0344	0.3408±0.0314	0.3073±0.0281	3.6433±0.2642	0.5947±0.0880
SMILES GA	0.3226±0.0000	0.2811±0.0086	0.2528±0.0259	0.3199±0.0000	0.2782±0.0072	0.2468±0.0223	4.5480±0.8299	0.6994±0.1207
MARS	0.2620±0.0180	0.2473±0.0167	0.2217±0.0180	0.2550±0.0121	0.2388±0.0109	0.2115±0.0132	3.4745±0.1514	0.8345±0.0156
MIMOSA	0.3010±0.0178	0.2863±0.0215	0.2688±0.0161	0.2573±0.0045	0.2194±0.0060	0.1661±0.0048	3.6586±0.2510	0.7209±0.0213
REINVENT	0.6159±0.0718	0.6091±0.0738	0.5864±0.0797	0.4623±0.0482	0.4250±0.0471	0.3712±0.0463	4.1275±0.1784	0.3759±0.1137

Table 4: thiothixene_rediscovery Results

thiothixene_rediscovery	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.5479±0.0624	0.51747±0.0422	0.47897±0.0256	0.4416±0.0265	0.41907±0.0248	0.38287±0.0247	2.91647±0.4819	0.53907±0.0296
Graph GA	0.5316±0.0899	0.5012±0.1001	0.4414±0.1111	0.4776±0.0492	0.4411±0.0565	0.3760±0.0684	2.6912±0.3586	0.5989±0.1548
SMILES GA	0.3115±0.0215	0.3066±0.0226	0.3036±0.0227	0.3083±0.0207	0.3010±0.0207	0.2945±0.0196	4.1085±0.4983	0.5647±0.0621
MARS	0.4067±0.0528	0.3731±0.0565	0.3145±0.0509	0.3747±0.0367	0.3380±0.0393	0.2822±0.0348	3.0312±0.3307	0.7815±0.0555
MIMOSA	0.3893±0.0384	0.3736±0.0298	0.3484±0.0233	0.2911±0.0191	0.2749±0.0151	0.2465±0.0111	3.2976±0.4421	0.6298±0.0597
REINVENT	0.6673±0.0055	0.6653±0.0077	0.6446±0.0138	0.3618±0.0354	0.3390±0.0325	0.3003±0.0287	2.4271±0.1250	0.3968±0.0099

Table 5: sitagliptin_mpo Results

sitagliptin_mpo	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.6750±0.0907	0.6391±0.1118	0.5511±0.1411	0.5633±0.0934	0.5013±0.1028	0.3836±0.09522	5.281±0.7692	0.7067±0.480
Graph GA	0.6562±0.0746	0.6144±0.0721	0.5459±0.0801	0.5494±0.0858	0.4911±0.0866	0.3879±0.0857	4.9441±0.2630	0.7078±0.0965
SMILES GA	0.5131±0.1061	0.5129±0.1061	0.5061±0.1012	0.4770±0.0829	0.4640±0.0773	0.4331±0.0659	6.0436±0.5582	0.4316±0.1368
MARS	0.3817±0.0325	0.3071±0.0416	0.1789±0.0570	0.3436±0.0315	0.2647±0.0282	0.1421±0.0375	3.7138±0.4830	0.8700±0.0088
MIMOSA	0.4645±0.0872	0.4032±0.0955	0.2874±0.0849	0.3072±0.0267	0.2533±0.0293	0.1571±0.0219	4.4924±1.1530	0.7621±0.0831
REINVENT	0.5323±0.0657	0.5050±0.0752	0.4593±0.1067	0.2080±0.0792	0.1871±0.0738	0.1477±0.0617	2.7170±0.2333	0.6421±0.1489

Table 6: scaffold_hop Results

scaffold_hop	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.6289±0.1265	0.6199±0.1236	0.6048±0.1178	0.5388±0.0307	0.5225±0.0178	0.5017±0.0089	3.6897±0.3603	0.4848±0.0788
Graph GA	0.5340±0.0251	0.5190±0.0251	0.5035±0.0225	0.5109±0.0159	0.4936±0.0145	0.4715±0.0118	3.3934±0.0797	0.6937±0.0436
SMILES GA	0.5117±0.0064	0.5117±0.0063	0.5113±0.0064	0.5032±0.0057	0.5011±0.0057	0.4968±0.0056	4.8192±0.6982	0.5346±0.0768
MARS	0.4774±0.0065	0.4689±0.0047	0.4488±0.0106	0.4736±0.0055	0.4642±0.0036	0.4414±0.0082	3.6729±0.1009	0.8538±0.0085
MIMOSA	0.4819±0.0146	0.4792±0.0161	0.4670±0.0152	0.3334±0.0093	0.3234±0.0093	0.3006±0.0072	3.4772±0.4528	0.6221±0.0685
REINVENT	0.6337±0.0396	0.6272±0.0474	0.6225±0.0489	0.5186±0.0267	0.5067±0.0256	0.4874±0.0239	2.9716±0.1661	0.4699±0.0739

Table 7: ranolazine_mpo Results

ranolazine_mpo	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.8006±0.0421	0.7932±0.0440	0.7784±0.0435	0.7474±0.0458	0.7244±0.0478	0.6790±0.0534	4.8047±0.5400	0.5629±0.0646
Graph GA	0.7982±0.0213	0.7833±0.0198	0.7546±0.0303	0.7426±0.0227	0.7119±0.0211	0.6527±0.0258	4.3215±0.4771	0.6879±0.1120
SMILES GA	0.7606±0.0263	0.7577±0.0282	0.7522±0.0307	0.7088±0.0233	0.6993±0.0234	0.6786±0.0244	6.1237±0.4029	0.5597±0.0584
MARS	0.7633±0.0217	0.7452±0.0128	0.6955±0.0144	0.7526±0.0207	0.7305±0.0116	0.6667±0.0108	4.7824±0.1203	0.8360±0.0086
MIMOSA	0.7308±0.0464	0.7197±0.0457	0.6841±0.0426	0.5630±0.0189	0.5265±0.0189	0.4631±0.0194	4.2304±0.4489	0.7087±0.0331
REINVENT	0.8555±0.0091	0.8528±0.0098	0.8482±0.0108	0.7428±0.0056	0.7239±0.0057	0.6882±0.0082	3.0679±0.3308	0.3919±0.0794

Table 8: qed Results

qed	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.9478±0.0003	0.9471±0.0004	0.9423±0.0037	0.9420±0.0008	0.9384±0.0011	0.9218±0.0022	3.4141±0.4995	0.7193±0.0153
Graph GA	0.9469±0.0012	0.9437±0.0020	0.9304±0.0046	0.9132±0.0011	0.9052±0.0016	0.8799±0.0036	3.2943±0.1225	0.8146±0.0139
SMILES GA	0.9468±0.0019	0.9465±0.0017	0.9458±0.0014	0.9150±0.0019	0.9113±0.0020	0.8982±0.0025	5.4677±0.8324	0.6377±0.0352
MARS	0.9401±0.0066	0.9287±0.0142	0.8771±0.0387	0.9311±0.0077	0.9172±0.0135	0.8565±0.0334	3.1873±0.0568	0.8776±0.0017
MIMOSA	0.9314±0.0014	0.9291±0.0029	0.9185±0.0086	0.7143±0.0020	0.6887±0.0028	0.6392±0.0042	3.7220±0.1559	0.8001±0.0241
REINVENT	0.9483±0.0001	0.9483±0.0001	0.9480±0.0001	0.7834±0.0847	0.7807±0.0848	0.7708±0.0845	2.5082±0.4728	0.6879±0.0783

Table 9: perindopril_mpo Results

perindopril_mpo	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.5741±0.0104	0.5613±0.0109	0.5443±0.0113	0.5228±0.0113	0.5102±0.0109	0.4874±0.0111	4.735366737±0.4388	0.5477±0.0594
Graph GA	0.5350±0.0612	0.5245±0.0585	0.5000±0.0528	0.5010±0.0393	0.4848±0.0371	0.4515±0.0339	3.9596±0.1366	0.6685±0.1190
SMILES GA	0.4495±0.0144	0.4495±0.0144	0.4478±0.0145	0.4455±0.0127	0.4433±0.0123	0.4355±0.0121	4.7639±0.4927	0.4968±0.0538
MARS	0.4793±0.0137	0.4647±0.0128	0.4357±0.0135	0.4751±0.0112	0.4570±0.0098	0.4173±0.0090	5.0202±0.2178	0.8231±0.0053
MIMOSA	0.4703±0.0177	0.4569±0.0107	0.4403±0.0076	0.3247±0.0046	0.3116±0.0042	0.2856±0.0031	3.9169±0.2835	0.6964±0.0230
REINVENT	0.6196±0.0460	0.6164±0.0485	0.6132±0.0511	0.4563±0.0764	0.4428±0.0755	0.4188±0.0746	4.1524±0.3994	0.3490±0.0603

Table 10: osimertinib_mpo Results

osimertinib_mpo	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.8536±0.0117	0.8489±0.0121	0.8393±0.0114	0.8263±0.0130	0.8144±0.0123	0.7872±0.0103	4.1230±0.4389	0.5518±0.0579
Graph GA	0.8481±0.0174	0.8375±0.0192	0.8204±0.0173	0.8052±0.0105	0.7891±0.0104	0.7575±0.0106	3.8223±0.4355	0.7062±0.0515
SMILES GA	0.8243±0.0293	0.8232±0.0289	0.8222±0.0287	0.7830±0.0241	0.7780±0.0227	0.7668±0.0206	6.7199±0.3103	0.5971±0.0263
MARS	0.8058±0.0129	0.7923±0.0112	0.7669±0.0230	0.7925±0.0111	0.7727±0.0070	0.7240±0.0156	5.2209±0.2640	0.8473±0.0133
MIMOSA	0.7883±0.0089	0.7843±0.0095	0.7732±0.0118	0.6335±0.0045	0.6046±0.0049	0.5516±0.0049	4.0396±0.3371	0.6819±0.0238
REINVENT	0.8905±0.0147	0.8869±0.0142	0.8772±0.0120	0.7931±0.0075	0.7804±0.0067	0.7539±0.0073	3.1860±0.4682	0.4975±0.0826

Table 11: mestranol_similarity Results

mestranol_similarity	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.7348±0.0321	0.7140±0.0282	0.6618±0.0298	0.6179±0.0401	0.5833±0.0349	0.5273±0.0318	4.6019±0.3707	0.5329±0.0394
Graph GA	0.7335±0.1325	0.6926±0.1161	0.6401±0.0928	0.6123±0.0516	0.5727±0.0429	0.5173±0.0307	4.2850±0.5503	0.5599±0.0471
SMILES GA	0.4488±0.0456	0.4477±0.0454	0.4441±0.0433	0.4196±0.0425	0.4129±0.0419	0.4007±0.0395	5.1324±0.6792	0.5278±0.1021
MARS	0.4142±0.0662	0.3782±0.0718	0.3202±0.0756	0.4059±0.0638	0.3671±0.0655	0.3047±0.0645	4.0433±0.3763	0.8546±0.0105
MIMOSA	0.5239±0.0145	0.4907±0.0147	0.4450±0.0186	0.5079±0.0010	0.4688±0.0021	0.4069±0.0044	3.8981±0.3227	0.8052±0.0398
REINVENT	0.7838±0.0823	0.7809±0.0835	0.7627±0.0834	0.3924±0.0836	0.3709±0.0838	0.3346±0.0824	3.6929±0.4919	0.2989±0.0471

Table 12: median2 Results

median2	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.2828±0.0310	0.2786±0.0305	0.2669±0.0292	0.2455±0.0223	0.2354±0.0213	0.2201±0.0184	3.7643±0.4034	0.4811±0.0656
Graph GA	0.2364±0.0254	0.2240±0.0278	0.2085±0.0249	0.2270±0.0189	0.2081±0.0217	0.1880±0.0189	3.1125±0.3191	0.6496±0.0693
SMILES GA	0.2070±0.0000	0.1869±0.0027	0.1788±0.0049	0.2051±0.0000	0.1833±0.0025	0.1709±0.0045	4.3820±0.3404	0.4846±0.0718
MARS	0.1916±0.0023	0.1732±0.0100	0.1327±0.0342	0.1906±0.0022	0.1716±0.0095	0.1305±0.0332	3.6812±0.4905	0.8670±0.0115
MIMOSA	0.2302±0.0104	0.2255±0.0092	0.2120±0.0069	0.2106±0.0083	0.2021±0.0075	0.1832±0.0054	3.0707±0.2480	0.5688±0.0862
REINVENT	0.4577±0.0000	0.3914±0.0001	0.3074±0.0264	0.2711±0.1986	0.2315±0.1698	0.1776±0.1327	3.0256±0.2779	0.7054±0.1610

Table 13: median1 Results

median1	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.3923±0.0173	0.3714±0.0183	0.3378±0.0095	0.3376±0.0215	0.3108±0.0153	0.2789±0.0125	4.0045±0.1275	0.6191±0.0278
Graph GA	0.3093±0.0345	0.2906±0.0257	0.2593±0.0180	0.2922±0.0311	0.2680±0.0228	0.2318±0.0158	4.0053±0.1491	0.7091±0.0363
SMILES GA	0.2004±0.0300	0.1989±0.0298	0.1980±0.0293	0.1977±0.0296	0.1936±0.0284	0.1889±0.0268	6.0024±1.0691	0.6396±0.0382
MARS	0.2322±0.0201	0.2094±0.0181	0.1777±0.0234	0.2239±0.0140	0.2019±0.0116	0.1671±0.0158	4.2508±0.3090	0.8458±0.0196
MIMOSA	0.3275±0.0130	0.3011±0.0036	0.2686±0.0003	0.2675±0.0043	0.2278±0.0013	0.1654±0.0013	3.9701±0.0735	0.7643±0.0091
REINVENT	0.4579±0.0004	0.4384±0.0193	0.4181±0.0344	0.2571±0.0514	0.2282±0.0447	0.1852±0.0373	4.7140±0.6151	0.3136±0.1278

Table 14: jnk3 Results

jnk3	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.5980±0.2128	0.5748±0.2080	0.5392±0.1999	0.4513±0.1305	0.4199±0.1210	0.3740±0.1085	4.6769±0.7752	0.4923±0.0770
Graph GA	0.8160±0.1935	0.8082±0.1958	0.7862±0.1965	0.6615±0.1340	0.6302±0.1267	0.5749±0.1152	2.8660±0.4549	0.4021±0.0378
SMILES GA	0.5530±0.0095	0.4450±0.0761	0.3974±0.0827	0.5465±0.0010	0.3813±0.0469	0.3137±0.0491	6.4713±0.6045	0.4436±0.1183
MARS	0.4950±0.0328	0.4484±0.0358	0.3436±0.0879	0.4598±0.0712	0.4085±0.0594	0.2841±0.0593	4.4597±0.4988	0.7913±0.0798
MIMOSA	0.4480±0.0476	0.4274±0.0527	0.3894±0.0514	0.3191±0.0358	0.2884±0.0391	0.2328±0.0470	4.6500±0.5810	0.6408±0.0239
REINVENT	0.8575±0.2237	0.8429±0.2526	0.8164±0.3004	0.5141±0.3014	0.4928±0.2945	0.4569±0.2842	3.5782±0.5887	0.3350±0.3209

Table 15: isomers_c9h10n2o2pf2cl Results

isomers_c9h10n2o2pf2cl	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.9394±0.0000	0.9237±0.0196	0.8733±0.0264	0.8286±0.0704	0.7878±0.0758	0.7018±0.0808	5.4939±0.2025	0.7481±0.0389
Graph GA	0.9137±0.0291	0.8869±0.0285	0.8316±0.0177	0.8441±0.0206	0.8024±0.0158	0.7171±0.0141	4.5273±0.3671	0.7995±0.0265
SMILES GA	0.9341±0.0366	0.9310±0.0363	0.8895±0.0457	0.8955±0.0351	0.8640±0.0352	0.8083±0.0450	5.8310±0.2092	0.7161±0.0519
MARS	0.7268±0.1260	0.6639±0.1606	0.5268±0.2389	0.6751±0.1032	0.5989±0.1057	0.4424±0.1499	3.1615±0.8526	0.8335±0.0970
MIMOSA	0.8352±0.0458	0.8081±0.0637	0.7564±0.0661	0.6092±0.0248	0.5745±0.0244	0.5007±0.0200	4.1175±0.8821	0.7669±0.0722
REINVENT	0.9166±0.0312	0.9030±0.0267	0.8718±0.0256	0.3557±0.0381	0.3331±0.0367	0.2874±0.0323	3.1677±0.7806	0.6673±0.0776

Table 16: Isomers_C7H8N2O2 Results

Isomers_C7H8N2O2	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	1.0000±0.0000	0.9946±0.0086	0.9372±0.0255	0.9554±0.0152	0.9119±0.0299	0.8208±0.0294	4.9965±0.1963	0.8306±0.0386
Graph GA	1.0000±0.0000	0.9965±0.0077	0.9349±0.0309	0.9649±0.0051	0.9413±0.0156	0.8514±0.0236	3.8005±0.2170	0.8363±0.0196
SMILES GA	0.7972±0.1597	0.7731±0.1713	0.7577±0.1796	0.7816±0.1527	0.7373±0.1660	0.7080±0.1727	6.2068±0.8572	0.7471±0.0912
MARS	0.9384±0.0569	0.8843±0.0741	0.7124±0.1851	0.8646±0.0975	0.7936±0.0734	0.5976±0.0968	2.9104±0.4353	0.8719±0.0181
MIMOSA	1.0000±0.0000	1.0000±0.0000	0.9586±0.0238	0.7046±0.0024	0.6735±0.0032	0.6012±0.0103	2.9549±0.0676	0.8403±0.0119
REINVENT	1.0000±0.0000	0.9985±0.0033	0.9314±0.0135	0.2286±0.0390	0.2047±0.0287	0.1603±0.0187	2.6110±0.2461	0.7901±0.0476

Table 17: gsk3b Results

gsk3b	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.9160±0.0991	0.9002±0.1005	0.8792±0.1086	0.8059±0.0991	0.7656±0.1023	0.7114±0.1078	3.3257±1.0183	0.4423±0.0746
Graph GA	0.8900±0.0869	0.8728±0.0890	0.8455±0.0928	0.7653±0.0460	0.7379±0.0511	0.6882±0.0539	3.1720±0.2601	0.4990±0.0688
SMILES GA	0.7740±0.0695	0.7634±0.0659	0.7430±0.0581	0.7225±0.0291	0.7014±±0.0276	0.6597±0.0283	6.4360±0.7540	0.3851±0.1339
MARS	0.6433±0.0494	0.6068±0.0623	0.5281±0.0553	0.5770±0.0808	0.5357±0.0871	0.4427±0.0631	4.2383±0.9253	0.8121±0.0206
MIMOSA	0.7320±0.1310	0.7018±0.1334	0.6685±0.1443	0.6470±0.1367	0.5991±0.1397	0.5326±0.1428	4.9866±0.7911	0.6864±0.0976
REINVENT	0.9820±0.0110	0.9760±0.0191	0.9682±0.0250	0.5906±0.1020	0.5629±0.0996	0.5199±0.0959	3.2361±0.3686	0.3586±0.1227

Table 18: fexofenadine_mpo Results

fexofenadine_mpo	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.8282±0.0355	0.8191±0.0335	0.8057±0.0335	0.7607±0.0257	0.7477±0.0244	0.7201±0.0241	5.1820±0.4382	0.5920±0.0694
Graph GA	0.8292±0.0337	0.8155±0.0306	0.7971±0.0290	0.7814±0.0178	0.7650±0.0164	0.7327±0.0145	4.8748±0.2502	0.6605±0.0571
SMILES GA	0.7431±0.0378	0.7354±0.0438	0.7316±0.0455	0.7223±0.0246	0.7115±0.0327	0.6970±0.0345	4.7677±0.4361	0.5022±0.0945
MARS	0.7311±0.0329	0.7145±0.0289	0.6898±0.0256	0.6932±0.0638	0.6699±0.0671	0.6262±0.0701	5.3380±0.5715	0.8163±0.0363
MIMOSA	0.7199±0.0297	0.7084±0.0306	0.6875±0.0331	0.5955±0.0166	0.5648±0.0171	0.5115±0.0153	4.3367±0.5893	0.7008±0.0580
REINVENT	0.9011±0.0237	0.8929±0.0247	0.8786±0.0277	0.7309±0.0339	0.7145±0.0345	0.6849±0.0359	4.2292±0.2216	0.4383±0.0780

Table 19: drd2 Results

drd2	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.9988±0.0026	0.9982±0.0029	0.9967±0.0049	0.9526±0.0164	0.9329±0.0248	0.8850±0.0398	3.2903±0.7108	0.5366±0.0575
Graph GA	0.9128±0.1950	0.9010±0.2213	0.8822±0.2511	0.8824±0.1798	0.8549±0.1977	0.8037±0.2108	2.7404±0.2834	0.7149±0.0882
SMILES GA	0.9949±0.0054	0.9949±0.0054	0.9948±0.0055	0.9398±0.0161	0.9301±0.0171	0.9099±0.0185	6.5066±0.4736	0.6116±0.0576
MARS	0.8646±0.1256	0.7525±0.1770	0.5285±0.2969	0.7698±0.1365	0.6517±0.1348	0.4190±0.1945	3.8385±0.1248	0.8335±0.0623
MIMOSA	0.8396±0.1932	0.7931±0.2114	0.7139±0.2337	0.5757±0.1220	0.5312±0.1180	0.4549±0.1063	3.8483±0.3802	0.7396±0.0731
REINVENT	1.0000±0.0000	1.0000±0.0000	1.0000±0.0000	0.5679±0.0701	0.5485±0.0851	0.5153±0.0936	2.6691±0.5179	0.4427±0.0443

Table 20: deco_hop Results

deco_hop	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.6475±0.0092	0.6413±0.0091	0.6353±0.0092	0.6200±0.0060	0.6124±0.0059	0.6024±0.0060	3.9402±0.4094	0.5562±0.0621
Graph GA	0.6794±0.1400	0.6738±0.1420	0.6632±0.1411	0.6452±0.0840	0.6326±0.0799	0.6120±0.0729	3.2752±0.2231	0.7074±0.1102
SMILES GA	0.6147±0.0061	0.6147±0.0061	0.6144±0.0059	0.5927±0.0060	0.5903±0.0060	0.5844±0.0057	4.9016±0.5634	0.5258±0.0800
MARS	0.6014±0.0069	0.5944±0.0070	0.5830±0.0095	0.5830±0.0227	0.5711±0.0301	0.5493±0.0421	3.7003±0.1504	0.8182±0.0646
MIMOSA	0.6051±0.0122	0.6032±0.0130	0.5896±0.0100	0.5428±0.0053	0.5173±0.0049	0.4726±0.0029	4.1345±0.2664	0.6712±0.0811
REINVENT	0.8014±0.1476	0.7915±0.1430	0.7853±0.1413	0.6577±0.0628	0.6415±0.0543	0.6231±0.0478	3.0143±0.2320	0.4356±0.0308

Table 21: celecoxib_rediscovery Results

celecoxib_rediscovery	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.6423±0.2072	0.6043±0.1705	0.5574±0.1350	0.5356±0.1032	0.4984±0.0834	0.4462±0.0593	2.6473±0.1911	0.5293±0.1011
Graph GA	0.6611±0.1368	0.6222±0.1323	0.5639±0.1095	0.5845±0.0683	0.5443±0.0664	0.4841±0.0530	2.5538±0.1068	0.5326±0.1117
SMILES GA	0.3681±0.0545	0.3651±0.0576	0.3630±0.0596	0.3595±0.0502	0.3552±0.0520	0.3483±0.0518	4.1684±0.4408	0.6008±0.0567
MARS	0.4828±0.0233	0.4215±0.0344	0.2677±0.1257	0.4763±0.0140	0.4132±0.0213	0.2476±0.1038	2.7526±0.3248	0.8393±0.0606
MIMOSA	0.4090±0.0265	0.3915±0.0329	0.3726±0.0303	0.3640±0.0507	0.3348±0.0422	0.2961±0.0335	3.4475±0.2161	0.6614±0.0863
REINVENT	0.9550±0.1006	0.8545±0.0639	0.8050±0.0617	0.4196±0.1028	0.3739±0.0877	0.3237±0.0794	2.6667±0.1317	0.3329±0.0578

Table 22: amlodipine_mpo Results

amlodipine_mpo	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.6942±0.0499	0.6848±0.0510	0.6698±0.0456	0.6309±0.0361	0.6120±0.0348	0.5845±0.0309	4.0448±0.0917	0.4402±0.0859
Graph GA	0.7150±0.0407	0.6958±0.0290	0.6730±0.0258	0.6569±0.0202	0.6350±0.0174	0.6004±0.0131	3.8835±0.1383	0.5982±0.0422
SMILES GA	0.5254±0.0335	0.5235±0.0375	0.5216±0.0379	0.5130±0.0295	0.5079±0.0324	0.4995±0.0317	4.6961±0.2230	0.5945±0.0615
MARS	0.5081±0.0303	0.4902±0.0279	0.4488±0.0392	0.5014±0.0259	0.4818±0.0223	0.4311±0.0315	3.6973±0.3809	0.8430±0.0302
MIMOSA	0.6045±0.0118	0.5789±0.0166	0.5540±0.0134	0.5908±0.0211	0.5429±0.0174	0.4979±0.0084	4.0689±0.5939	0.6628±0.0650
REINVENT	0.7382±0.0453	0.7334±0.0431	0.7259±0.0400	0.5576±0.0574	0.5409±0.0553	0.5147±0.0527	3.2323±0.2277	0.3946±0.0670

Table 23: albuterol_similarity Results

albuterol_similarity	Average Top 1	Top 10	Top 100	AUC Top 1	Top 10	Top 100	Average SA	Diversity
Gradient GA	0.9936±0.0090	0.9806±0.0306	0.9322±0.0484	0.8614±0.0466	0.8198±0.0531	0.7413±0.0546	3.3887±0.2129	0.5094±0.0336
Graph GA	0.9973±0.0060	0.9840±0.0255	0.9377±0.0441	0.8905± 0.0129	0.8528±0.0263	0.7732±0.0290	3.2491±0.2386	0.5262±0.0371
SMILES GA	0.6946 ± 0.1058	0.6942 ± 0.1055	0.6876 ± 0.1030	0.6694 ± 0.0960	0.6628 ± 0.0949	0.6517 ± 0.0922	6.4222 ± 0.4881	0.6826 ± 0.0169
MARS	0.9139 ± 0.0598	0.8569 ± 0.0436	0.7636 ± 0.0383	0.7601 ± 0.0289	0.6977 ± 0.0321	0.5974 ± 0.0275	3.4333 ± 0.2278	0.7377 ± 0.0293
MIMOSA	0.8220 ± 0.0700	0.7882 ± 0.0710	0.7382 ± 0.0583	0.6563 ± 0.0326	0.6161 ± 0.0360	0.5541 ± 0.0373	3.8936 ± 0.4931	0.7061 ± 0.0181
REINVENT	1.0000 ± 0.0000	1.0000 ± 0.0000	0.9961 ± 0.0073	0.3614 ± 0.0420	0.3393 ± 0.0399	0.3019 ± 0.0373	3.2091 ± 0.0681	0.4777 ± 0.0562

B MOLECULES GENERATED BY VARIOUS METHODS

Here we have a list of Top 10 molecules generated by each method that we tested from the experiments in table 1. Each molecule has its respective scores underneath them, and all molecules are for the bio-activity object mestranol similarity.

We notice that the molecules generated by Gradient GA, SMILES GA, and Graph GA are very similar, with Gradient GA and Graph GA having the most similarity. MARS and MIMOSA both have unique molecules generated, but the performance of those molecules is low. We notice that Gradient GA just needs to take a molecule from a good run, and that will lead to having an entire set of molecules. Most of the Gradient GA top 10 molecules are all from the same run, and we can see that molecular structure was found in the bottom right corner of figure 4. This further supports the idea that Gradient GA is exploring similar molecules to Graph GA, and is further exploring an area around its top-performing molecules.

Figure 2: Top 10 molecules generated by Gradient GA for the bio-activity objective mestranol similarity with their associated score underneath each molecule.

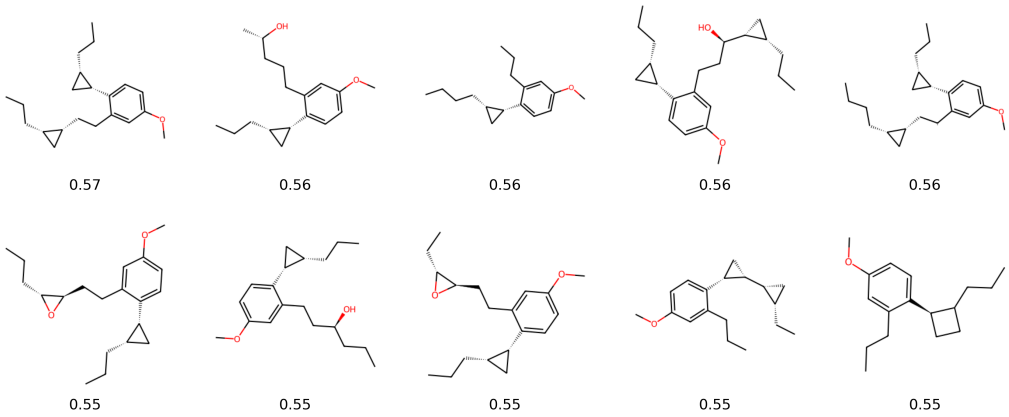

Figure 3: Top 10 molecules generated by SMILES GA for the bio-activity objective mestranol similarity with their associated score underneath each molecule.

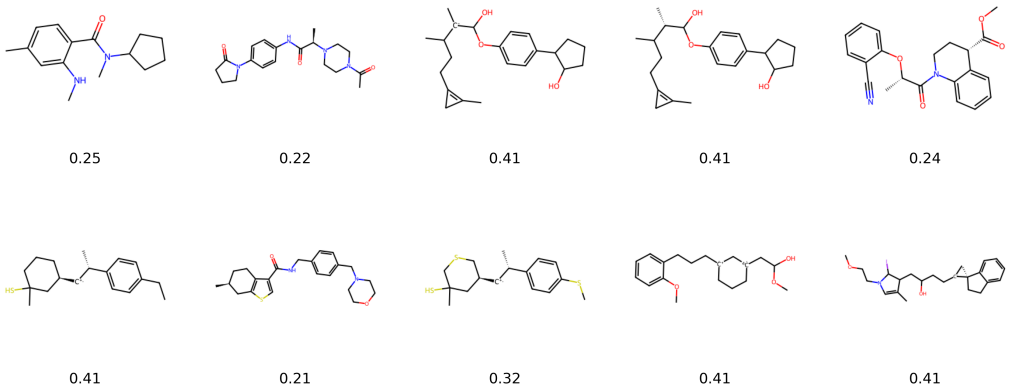

Figure 4: Top 10 molecules generated by Graph GA for the bio-activity objective mestranol similarity with their associated score underneath each molecule.

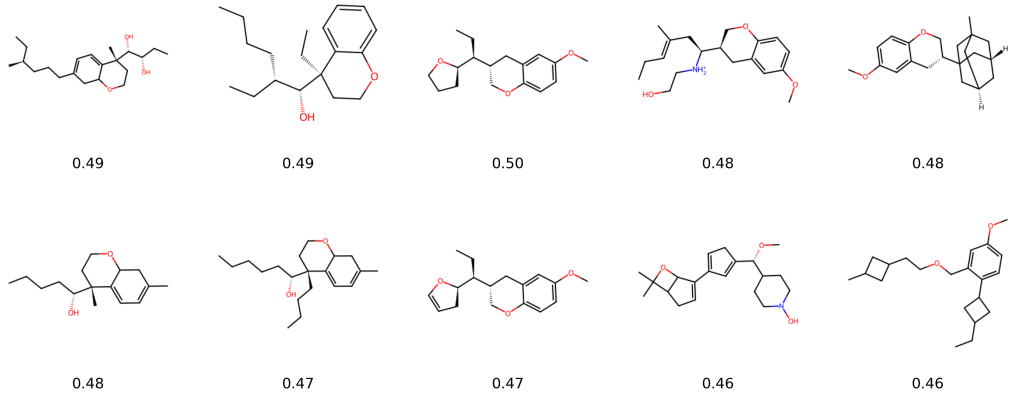

Figure 5: Top 10 molecules generated by MARS for the bio-activity objective mestranol similarity with their associated score underneath each molecule.

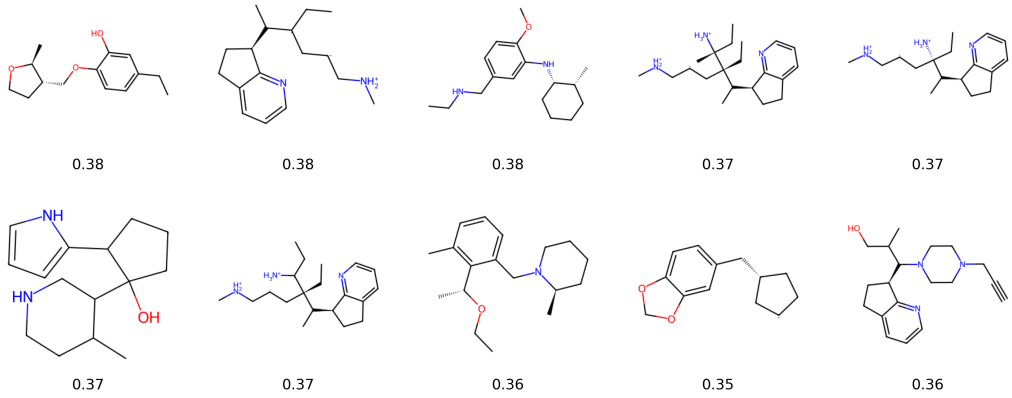

Figure 6: Top 10 molecules generated by MIMOSA for the bio-activity objective mestranol similarity with their associated score underneath each molecule.

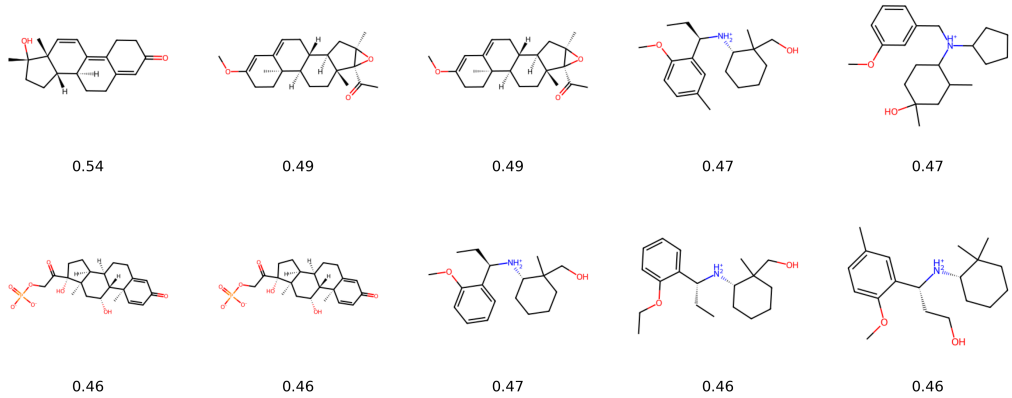

C COMPARISON OF USING $\mathcal{O}(x)$

In table below, we show the results of using $\nabla U(v) = \nabla f(v)$ vs $\nabla U(v) = \frac{\nabla f(v)}{\mathcal{O}(X)}$. We notice from the results, using $\mathcal{O}(x)$ leads to a better score throughout all metrics.

Table 24: Comparison of Average Top 10, AUC Top 1, AUC Top 10, and AUC Top 100 with GuacaMol objective, mestranol similarity, under 2500 oracle calls. The best Gradient GA setup is **bolded**. We conduct five independent runs using different random seeds for both versions of Gradient GA, and report the average scores and their standard deviation.

Gradient GA	mestranol similarity			
	Average Top 10	AUC Top 1	AUC Top 10	AUC Top 100
With $\mathcal{O}(x)$	**0.5130±0.0393**	**0.4433±0.0310**	**0.4082±0.0315**	**0.3534±0.0355**
Without $\mathcal{O}(x)$	0.5064±0.0312	0.4433±0.0319	0.4072±0.0367	0.3501±0.0419

D ORACLE CALL EFFICIENCY

To demonstrate that using gradient information accelerates convergence, we conduct experiments measuring AUC Top 10 and AUC Top 100 scores as the number of oracle calls increases. All methods are evaluated with 2,500 oracle calls over 5 runs. Our primary focus is on the bio-objective Mestranol Similarity. Figures 7 and 8 show that after the initialization phase, where each method achieves a baseline score based on the initial oracle calls, Gradient GA consistently outperforms almost all other methods at each step. This indicates that Gradient GA is not only more effective at finding optimal values but also more efficient, due to its use of gradient guidance rather than random walk exploration.

Figure 7: Mestranol similarity AUC Top 10 score comparison as more oracle calls occur.

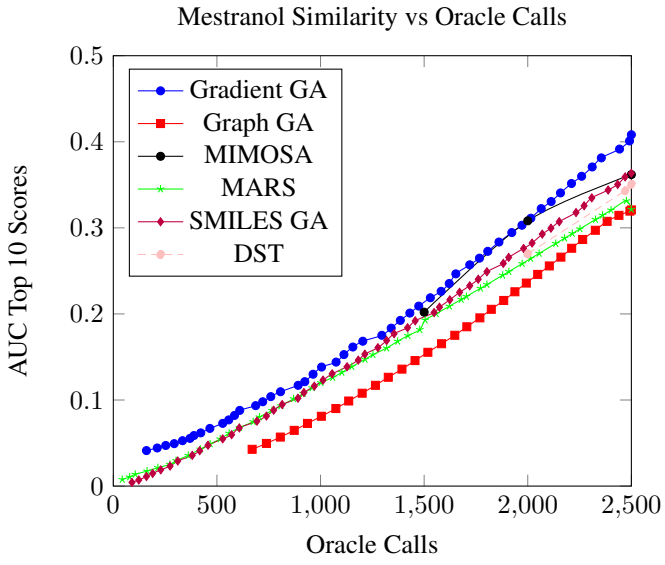

Figure 8: Mestranol similarity AUC Top 100 scores comparison as more oracle calls occur.

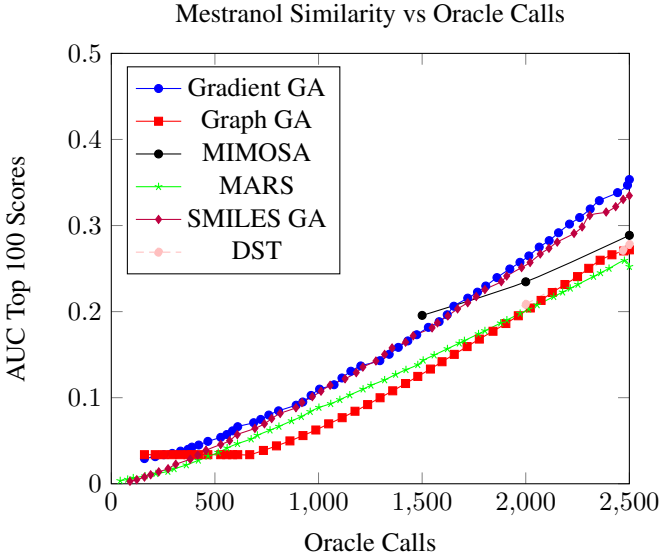

E SYNTHETIC ACCESSIBILITY AND DIVERSITY HEATMAPS

We further analyze additional metrics, synthetic accessibility (SA) in Figure 9 and diversity in Figure 10. It is important to note that these metrics are not explicitly optimized in our objective function. Therefore, their performance is a byproduct of the discovered molecules rather than a direct outcome of our method. From Figure 9, we observe that the SA scores of Gradient GA are comparable to those of Graph GA, indicating that there is no significant trade-off between improved performance and SA score. Additionally, in terms of overall SA performance, Gradient GA is also close to DST, another gradient-based method. In Figure 10, we observe that the diversity score for Gradient GA is lower than that of other methods. This outcome is expected, as our approach samples molecules near high-performing parent molecules. While this may reduce diversity, it can be advantageous when the goal is to perform a fine-grained local search over good regions.

Figure 9: Heatmap of synthetic accessibility (SA, ↓, lower is better) score of all methods and oracles.

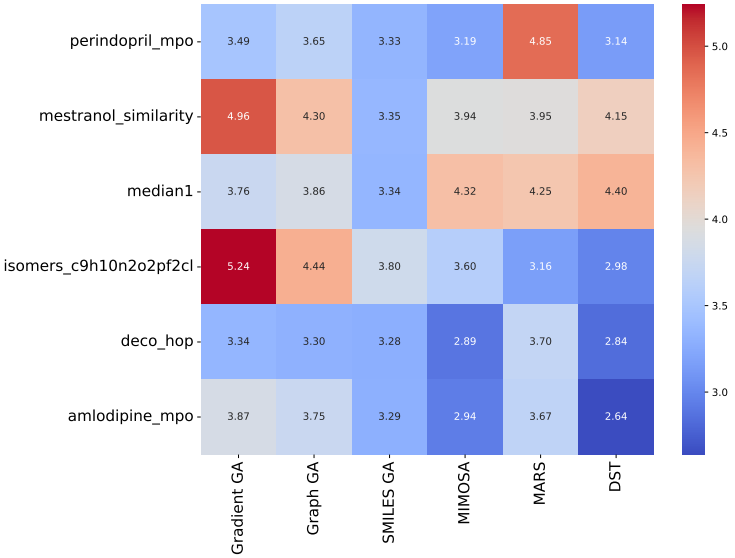

Figure 10: Heatmap of Diversity (↑, higher is better) score of all methods and oracles.

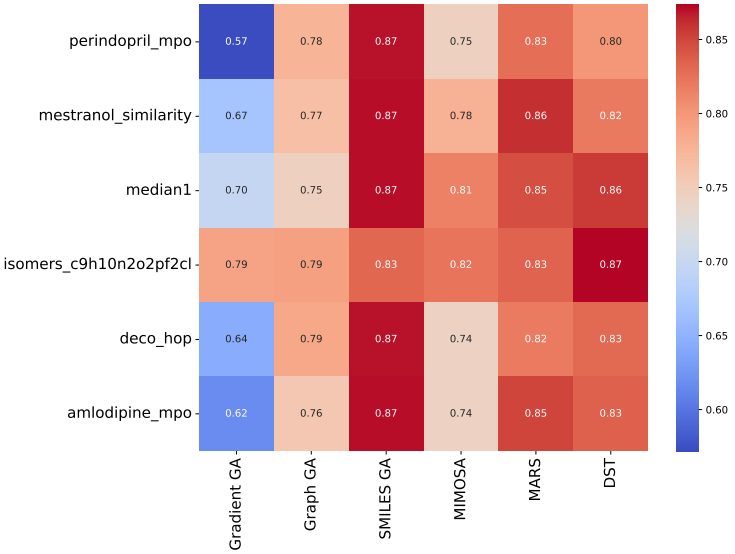

F NOTATION TABLE AND ALGORITHM

Here, we have put a notation table for all the variables introduced in the body of the paper. We also present the algorithm directly.

Table 25: Mathematical notations and their explanations.

Notations	Explanations
$\mathcal{O}$	oracle function
$\mathcal{Q}$	molecular space
X	molecule
$\mathcal{M}$	complete MPNN-based model for oracle prediction
$G(\cdot)$	graph representation
G_X	molecular graph of X
L_1, L_2	linear layer
v	embedding
$\pi(\cdot)$	target distribution (normalized score)
$U(\cdot)$	energy function
D	population, a set of molecules.
D'	new molecule set
D''	Re-training molecule set
n	embedding size
k	number of generated samples at each step
T	threshold criterion to add to training set
τ	retrainable threshold
α	step size for sampling

Algorithm 1 Gradient Genetic Algorithm

Input: oracle function $\mathcal{O}$, step size α, retrainable threshold τ
Initialize $D \leftarrow$ original population
Initialize new molecule set $D' \leftarrow \{\}$
Train predictive model $\mathcal{M}$ using
$\{(G(d), \mathcal{O}(d)) \ \ \forall d \in D\}$
Initialize retrained molecule set $D'' \leftarrow \{\}$
for $t = 1, 2, \ldots$ **do**
  $p(d) \propto \mathcal{O}(d) \quad \forall d \in D$
  Parent molecules $d_1, d_2 \sim p(d) \quad [\, d \in D]$
  Get parents' embedding v_i for each $G(d_i)$ using Eq. **??**
  Get parents' gradient $\nabla U(v_i)$ for each d_i using Eq. 4
  Evaluate $v, \nabla U(v)$ using Eq. 5
  Get crossover set: $S \leftarrow \texttt{CROSSOVER}(d_1, d_2)$
  Get sampling probability $probs$ of S using Eq. 2
  Evaluate sample set $D' \leftarrow \texttt{Sample}(S, probs, k)$
  Mutate each molecule in $D' \leftarrow \texttt{MUTATE}(D')$
  Update population $D \leftarrow \texttt{toporacle}(\{D, D_1\}, |D|)$
  Update training set D'' with T using Eq. 6
  if $|D''| \geq \tau$ **then**
    Retrain model $\mathcal{M}$ with $\{(d, \mathcal{O}(d)) \ \ \forall d \in D''\}$
  end if
  $D'' \leftarrow \{\}$ or $D'' \leftarrow D''$
end for

G ILLUSTRATION OF DLP SAMPLING PROCEDURE

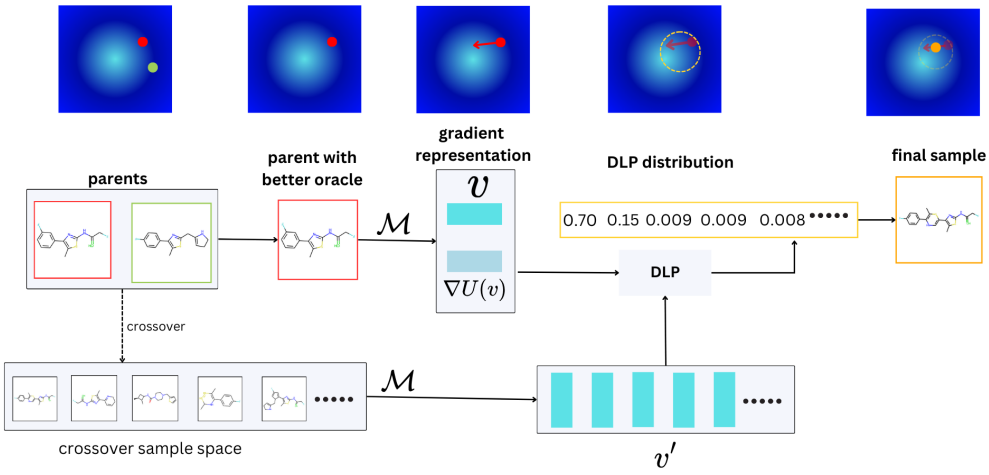

Figure 11: Overview of DLP-based sampling procedure in Gradient GA, illustrating how the sampled molecule moves toward the optimum.