

Dynamic Many-Objective Molecular Optimization: Unfolding Complexity with Objective Decomposition and Progressive Optimization

Dong-Hee Shin, Young-Han Son, Deok-Joong Lee, Ji-Wung Han and Tae-Eui Kam

Department of Artificial Intelligence, Korea University

{dongheeshin, yhson135, deokjoong, danielhan, kamte}@korea.ac.kr

Abstract

Molecular discovery has received significant attention across various scientific fields by enabling the creation of novel chemical compounds. In recent years, the majority of studies have approached this process as a multi-objective optimization problem. Despite notable advancements, most methods optimize only up to four molecular objectives and are mainly designed for scenarios with a predetermined number of objectives. However, in real-world applications, the number of molecular objectives can be more than four (many-objective) and additional objectives may be introduced over time (dynamic-objective). To fill this gap, we propose DyMol, the first method designed to tackle the dynamic many-objective molecular optimization problem by utilizing a novel divide-and-conquer approach combined with a decomposition strategy. Additionally, we comprehensively integrate convergence, Pareto diversity, and structural diversity into the optimization process to provide efficient exploration of the search space. We validate the superior performance of our method using the practical molecular optimization (PMO) benchmark. The source code and supplementary material are available online.

1 Introduction

Molecular discovery is foundational to progress in a variety of scientific fields, ranging from the development of new pharmaceuticals to the creation of innovative materials [Bilodeau *et al.*, 2022]. Basically, molecular discovery is a complex process that seeks to identify molecules with desirable properties [Son *et al.*, 2024]. In essence, this process is fundamentally a constrained multi-objective optimization problem, where the objectives are to simultaneously maximize or minimize certain attributes of molecules [Fromer and Coley, 2023].

Unlike single-objective optimization, the multi-objective optimization problem introduces distinct challenges that arise from the necessity to balance multiple and often conflicting objectives [Marler and Arora, 2004]. Therefore, it becomes infeasible to identify a single optimal solution that satisfies all

objectives. Instead, the focus shifts to finding Pareto optimal solution sets that represent various trade-offs among these objectives [Gunantara, 2018].

In the context of molecular discovery, the application of multi-objective optimization may exhibit unique characteristics compared to its use in other general domains. First, oracle calls in real-world molecular discovery are expensive, requiring resource-intensive wet-lab experiments or computer simulations to accurately evaluate molecular properties [Huang *et al.*, 2021]. Second, the discrete nature of molecules results in a complex and challenging optimization landscape [Wang *et al.*, 2006]. Lastly, the non-gradual transitions in molecular structures introduce additional complexity by creating an optimization landscape with sharp cliffs [Aldeghi *et al.*, 2022].

To tackle the multi-objective molecular optimization (MOMO) problem, much prior work has employed a range of generative models, including sampling-based methods [Fu *et al.*, 2021; Xie *et al.*, 2021], genetic algorithms [Jensen, 2019; Tripp *et al.*, 2021], probabilistic models [Bengio *et al.*, 2021], and reinforcement learning [Olivecrona *et al.*, 2017; Jin *et al.*, 2020]. However, given the necessity of simultaneously optimizing multiple objectives, they have commonly adopted two multi-objective optimization techniques: scalarization [Eichfelder, 2009] and Bayesian optimization [Lauermann and Ocenasek, 2002]. The scalarization method transforms multiple objectives into a single objective function by aggregating them using weighted sums or other combining strategies [Gunantara, 2018]. On the other hand, the Bayesian optimization method can address multiple objectives concurrently by leveraging acquisition functions to navigate the optimization landscape without needing to quantify the relative weights of each objective [Fromer and Coley, 2023]. While prior frameworks have shown effectiveness in molecular optimization, their applications exhibit distinct constraints. Specifically, they are typically limited to optimizing up to four objectives and they are designed to work with a fixed number of objectives, thereby lacking the capability to adapt to scenarios with varying numbers of objectives.

In real-world applications such as drug discovery, the importance of a dynamic many-objective molecular optimization setting becomes particularly evident [Luukkonen *et al.*, 2023]. From a many-objective perspective, the drug development process is inherently complex and multifaceted, typically requiring optimization of more than four objectives. In

<https://github.com/MolecularTeam/DyMol>

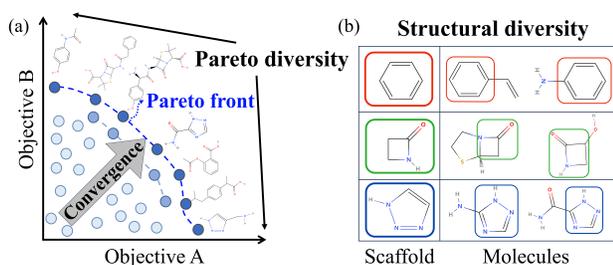


Figure 1: (a) Pareto front is a set of optimal solutions representing the best possible trade-offs among objectives. Convergence pushes the Pareto front’s boundary, while Pareto diversity pulls them apart for a broader spread. (b) Structural diversity, represented by colors, illustrates the variety in molecular scaffolds of chemical compounds.

particular, a new drug must meet various criteria, including potency, bioavailability, safety, solubility, stability, and synthesizability [Luukkonen *et al.*, 2023]. Concurrently, from a dynamic-objective standpoint, regulatory agencies such as the FDA consistently update their drug approval standards and requirements in response to new scientific insights, public health needs, and safety concerns [Darrow *et al.*, 2020]. These updates can introduce new optimization objectives within the drug development process. However, to our knowledge, no prior studies have tackled these issues.

The dynamic many-objective molecular optimization problem presents distinct challenges that set it apart from the typical MOMO problem. The many-objective aspect introduces an enormously large search space, which hinders efficient exploration and convergence to optimal Pareto solutions [Yuan *et al.*, 2015]. The dynamic-objective nature of this problem further exacerbates difficulties by continually altering the optimization landscape through the introduction of new objectives. To tackle these challenges, our approach diverges from most previous works, which typically aim to find Pareto optimal solutions for all objectives in a high-dimensional joint search space. Instead of tackling all objectives jointly, we propose a divide-and-conquer approach that decomposes the complex many-objective task into a series of manageable sub-problems, allowing us to address each of them sequentially. This progressive optimization scheme systematically unfolds the inherent complexities associated with the many-objective settings. Moreover, we incrementally introduce objectives over time in our method, allowing it to seamlessly adapt to the dynamic-objective settings with the incorporation of new objectives. To enhance our method’s adaptability for new objectives, we have also developed an objective adaptation technique that detects changes in the optimization landscape and helps the model to identify effective Pareto solutions.

To further promote the achievement of effective Pareto solutions, we propose the Pareto sampling to strategically emphasize both convergence and Pareto diversity during the optimization process. As shown in Figure 1, convergence refers to the ability of the optimization process to approach the true Pareto optimal front, indicating the proximity of the solutions to the best trade-offs among objectives [Yuan *et al.*, 2015].

On the other hand, Pareto diversity indicates the spread of solutions across the Pareto front, ensuring a wide range of solutions with varying trade-offs among objectives [Yuan *et al.*, 2015]. Pareto diversity is particularly important in molecular optimization, where the optimal trade-off is often unknown in advance [Fromer and Coley, 2023]. In this paper, we comprehensively address both convergence and Pareto diversity through progressive optimization using the divide-and-conquer approach alongside convergence and Pareto sampling. In addition to Pareto diversity, we also take into account for molecular structural diversity, which refers to the variety in chemical structures within the generated molecules. Promoting structural diversity is vital in molecular discovery fields such as drug design, as it offers multiple candidate molecules that meet desired criteria while exhibiting distinct chemical structures [Mathur and Hoskins, 2017].

The main contributions of our method can be summarized:

- To the best of our knowledge, DyMol, our proposed method, is the first to tackle the dynamic many-objective molecular optimization problem in molecular discovery.
- We propose a novel divide-and-conquer approach to decompose many-objective problems into sub-problems by starting with a single objective, then systematically adding over time to enable progressive optimization.
- Due to the incremental nature of adding objectives, our approach can handle dynamic-objective settings, where new objectives are introduced throughout the optimization process. In addition, we develop an objective adaptation technique that aids our model in adjusting more easily and effectively to newly introduced objectives.
- As far as we are aware, our method is the first to comprehensively integrate convergence, Pareto diversity, and structural diversity into molecular optimization by utilizing both convergence sampling and Pareto sampling.

2 Related Work

2.1 Generative Models for Molecular Discovery

In recent years, there has been a growing interest in the use of various generative models for molecular discovery. Generative models employed in molecular discovery can be broadly classified into four categories: 1) sampling-based methods, 2) genetic algorithms, 3) reinforcement learning (RL), and 4) probabilistic models. The sampling-based methods [Xie *et al.*, 2021; Fu *et al.*, 2021] focus on sampling from a distribution of possible molecules with desirable properties. The genetic algorithms [Jensen, 2019; Tripp *et al.*, 2021] employ a population-based approach that evolves molecules through iterative selection, crossover, and mutation guided by a fitness function. The RL-based methods [Olivecrona *et al.*, 2017; Jin *et al.*, 2020] involve an agent that interacts with an environment to generate molecular structures. In this paradigm, the agent receives rewards for taking actions that lead to desirable outcomes, thereby gradually refining its strategy through trial and error to learn an optimal policy [Shin *et al.*, 2024]. The probabilistic models, GFlowNets [Bengio *et al.*, 2021], generate molecular structures by identifying high-potential regions using probability distributions learned from data.

2.2 Multi-Objective Molecular Optimization

In the context of the MOMO problem, the challenge lies in simultaneously optimizing multiple molecular objectives, which often conflict with each other [Luukkonen *et al.*, 2023]. To address this, two prominent multi-objective optimization techniques have been widely adopted: scalarization and Bayesian optimization (BO). For instance, in the case of scalarization, MIMOSA [Fu *et al.*, 2021] has employed straightforward linear scalarization techniques to handle the MOMO problem. These techniques aim to aggregate multiple objectives into a single objective function by using weighted sums or Tchebycheff methods [Lin *et al.*, 2022]. On the other hand, BO offers a black box optimization approach that has been integrated into various molecular generative models to enhance sample efficiency [Laumanns and Ocenasek, 2002]. In particular, GPBO [Tripp *et al.*, 2021] exemplifies the integration of BO within the framework of GraphGA [Jensen, 2019] as the backbone model. Similarly, LaMBO [Stanton *et al.*, 2022] leverages BO on top of denoising autoencoders to address multi-objective sequence design problems. Recently, HN-GFN [Zhu *et al.*, 2023] proposes a multi-objective BO algorithm that leverages hypernetwork-based GFlowNets.

2.3 Dynamic Many-Objective Optimization

Dynamic many-objective optimization, while yet to be widely explored in molecular optimization, has found application in diverse fields such as manufacturing [Quan *et al.*, 2022], environmental management [Liu *et al.*, 2021], and mineral processing [Ding *et al.*, 2018]. Existing approaches in these domains have predominantly employed decomposition-based MOEA/D [Zhang and Li, 2007] and non-dominated sorting NSGA-III [Deb and Jain, 2013] frameworks due to their effectiveness in navigating high-dimensional search space.

3 Preliminary

3.1 Problem Formulation

In this work, we tackle the dynamic many-objective molecular optimization problem, characterized by the introduction of new objectives during the optimization process. In this context, ‘many-objective’ implies optimizing over four or more objectives, presenting a significant increase in complexity and dimensionality compared to a typical multi-objective problem [Hughes, 2005]. The dynamic many-objective molecular optimization problem can be formally defined as:

$$\begin{aligned} \underset{x \in \mathcal{X}}{\text{Maximize}} \quad & \mathbf{F}(x, t) = \{f_1(x), f_2(x), \dots, f_l(x)\}, \\ \text{subject to} \quad & g_j(x, t) \leq 0, \quad j = 1, 2, \dots, k; \\ & h_j(x, t) = 0, \quad j = 1, 2, \dots, l; \end{aligned} \quad (1)$$

where x denotes the molecule vector, with \mathcal{X} representing the feasible set in n -dimensional solution space. The $\mathbf{F}(x, t)$ represents the set of molecular objective functions at time stage $t \in T$. The objective function $f_i : \mathcal{X} \times T \rightarrow \mathbb{R}$ maps the molecular space to the real numbers. Constraints are twofold: inequality constraints $g_j(x, t)$ set the boundaries for feasible solutions, while equality constraints $h_j(x, t)$ specify exact conditions that feasible solutions must satisfy. The goal of this problem is to identify and track the evolving Pareto front.

3.2 Key Pareto Principles in Optimization

As mentioned earlier, the Pareto front is a set of solutions that are each Pareto optimal, which indicates that none of these solutions can be dominated by any other solutions. Therefore, we have the following definitions in our problem formulation:

Definition 1 (Pareto Dominance). Let $x^a, x^b \in \mathcal{X}$, x^a is said to Pareto dominate x^b (denoted as $x^a \succ x^b$) if and only if:

$$\begin{aligned} \forall i \in \{1, \dots, t\}, f_i(x^a) &\geq f_i(x^b) \\ \text{and } \exists j \in \{1, \dots, t\}, f_j(x^a) &> f_j(x^b). \end{aligned} \quad (2)$$

Here, $f_i(x^a)$ and $f_i(x^b)$ represent the values of the i^{th} objective function for solutions x^a and x^b , respectively.

Definition 2 (Pareto Optimality). A solution $x^* \in \mathcal{X}$ is defined as Pareto optimal if no other solution in the feasible set dominates it such as follows:

$$\nexists x \in \mathcal{X} : x \succ x^*. \quad (3)$$

Definition 3 (Pareto Front). The Pareto front, denoted as \mathbf{PF} , indicates the boundary of best possible trade-offs within the objective space. It is formally defined as:

$$\mathbf{PF} = \{\mathbf{f}(x^*) \mid x^* \in \mathcal{X}, \nexists x \in \mathcal{X} : x \succ x^*\}, \quad (4)$$

where $\mathbf{f}(x^*)$ denotes the vector of objective function values corresponding to a Pareto optimal solution.

4 Methods

4.1 Objective Decomposition

As shown in Figure 2, at time stage t_0 , our method begins by decomposing complex many-objective sets into more manageable sub-problems. This decomposition process is facilitated by our decomposition module, which analyzes the complexities of each objective and automatically determines the order in which they should be prioritized during optimization. Within this decomposition module, the generative model initially optimizes all objectives jointly for a limited number of iterations. The model uses an initial molecule as a starting point and produces an optimized molecule. Once complete, both the optimized and initial molecules are evaluated by objective functions $\mathbf{F}(x, t_0)$, also referred to as oracle functions.

From this evaluation, we calculate ordering scores by subtracting the objective scores of the initial molecules from the optimized molecules. These ordering scores offer insights into the extent of improvement in objective scores accomplished by the model. A lower ordering score indicates that optimizing a specific objective is more challenging, implying the need for early prioritization. It is worth noting that if a preferred order for the optimization process is available, this can be employed as the decomposition order as well.

4.2 Progressive Optimization

The main idea behind our method is to employ a divide-and-conquer approach to provide adaptability and efficiency when dealing with dynamic many-objective optimization. Consider a scenario where we have a total of five molecular objectives: objectives A through E, as shown in Figure 2. Following the

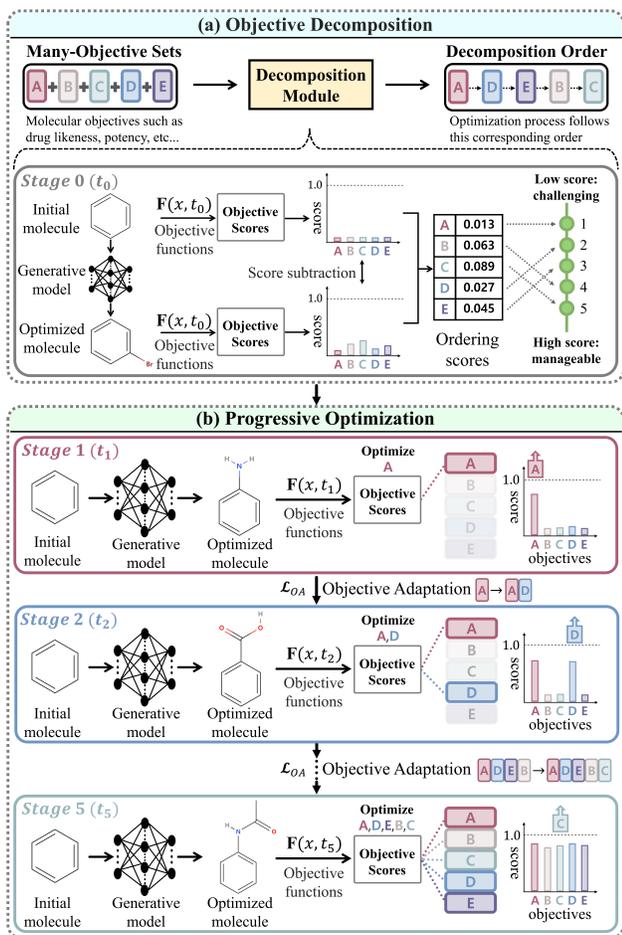


Figure 2: The overview of our DyMol method for dynamic many-objective molecular optimization problem, specifically exemplifies the scenario involving five molecular objectives. (a) Our method begins by decomposing complex many-objective sets into more manageable sub-problems facilitated by our decomposition module. (b) The optimization then progresses from a single objective, systematically incorporating additional objectives over time according to the decomposition order, thereby enabling progressive optimization.

decomposition order of $A \rightarrow D \rightarrow E \rightarrow B \rightarrow C$, our proposed model begins by solely optimizing objective A as follows:

$$\begin{aligned} & \underset{x \in \mathcal{X}_1}{\text{Maximize}} && \mathbf{F}(x, t_1) = \{f_A(x)\}, \\ & \text{subject to} && g_j(x, t_1) \leq 0, \quad j = 1, 2, \dots, k; \\ & && h_j(x, t_1) = 0, \quad j = 1, 2, \dots, l; \end{aligned} \quad (5)$$

where $f_A(x)$ denotes the value of objective A for a molecule x at time stage t_1 , and \mathcal{X}_1 represents the feasible set in the decision space specific to t_1 , potentially different from the general decision space \mathcal{X} due to the dynamic nature of the problem. The $g_j(x, t_1)$ and $h_j(x, t_1)$ represent inequality and equality constraints, respectively, derived from the physical, biological, or chemical requirements that a molecule must meet to be viable in a real-world environment.

When the model satisfies a certain score threshold related to objective A or reaches a predetermined number of itera-

tions, it progresses to the next time stage t_2 and incrementally incorporates additional objective D. Consequently, the optimization problem is expanded to maximize:

$$\mathbf{F}(x, t_2) = \{f_A(x), f_D(x)\}. \quad (6)$$

However, the introduction of a new objective function $f_D(\cdot)$ necessarily alters the optimization landscape by expanding the dimensions of the objective space. In this context, our objective adaptation technique plays a crucial role by enabling the model to adapt to this evolving optimization landscape. Specifically, it detects changes in the composition of objective scores, which provide learning feedback for model training and updates. For instance, at stage t_1 , the objective scores are solely based on the value of objective A. However, at t_2 , they evolve to encompass a composite value of both objectives A and D. The major role of the objective adaptation technique is to retrain the model using these updated objective scores, enabling the model to adjust to the evolving Pareto front, which is defined as:

$$\mathbf{PF}(t) = \{f(x^*) \mid x^* \in \mathcal{X}, \nexists x \in \mathcal{X} : x \succ x^* \text{ w.r.t. } \mathbf{F}(x, t)\}. \quad (7)$$

As time stages progress, the model systematically incorporates each new subsequent objective in line with the decomposition order and sequentially adjusts to the evolving Pareto front. Eventually, at the end of the time stage, the model can address the complete set of objectives. Thus, our method can be considered as a divide-and-conquer approach, as it strategically divides the complex optimization task into a series of simpler sub-problems, each focusing on a specific subset of the objectives. However, distinct from conventional divide-and-conquer methods that solve sub-problems independently and then combine their solutions, our approach is characterized by its sequential adaptation and refinement of solutions. As new objectives are introduced, the model dynamically adjusts its search process and integrates the incremental sub-problem solutions into a comprehensive solution that addresses all objectives. This adaptive nature of our method can make it effective in the dynamic-objective settings, where the optimization landscape progressively evolves over time.

Theoretical Analysis on Divide-and-Conquer

Here, we present a theoretical analysis demonstrating that our divide-and-conquer method, despite being incremental in nature, progressively converges toward solutions to those of the original complex problem that addresses all objectives jointly.

Theorem 1 (Convergence of the Divide-and-Conquer Approach to Global Near-Optimal Solutions). *Let \mathcal{P} be an original problem with objectives $\{f_1, \dots, f_i, \dots, f_n\}$, and \mathcal{P}_i be a sub-problem of \mathcal{P} focusing on $\{f_1, \dots, f_i\}$. The solution $\{x_1^*, x_2^*, \dots, x_i^*\}$ obtained at each stage i can be served as effective initial points for the next stage $i + 1$ to progressively converge towards a globally near-optimal solution for \mathcal{P} .*

Lemma 1 (Initial Convergence). For the initial base case of the divide-and-conquer approach where $i = 1$, the solution x_1^* for the sub-problem \mathcal{P}_1 can converge to a near-optimal solution for the objective f_1 even in cases of non-convexity. Let $x_1^* = \operatorname{argmax}_{x \in \mathcal{X}_1} f_1(x)$. Then, by the Bolzano-Weierstrass theorem [Brattka *et al.*, 2012], there exists a sequence $\{x_k\} \subset \mathcal{X}_1$ such that $x_k \rightarrow x_1^*$, and $\liminf_{k \rightarrow \infty} \frac{f_1(x_1^*) - f_1(x_k)}{\|x_1^* - x_k\|} \geq 0$.

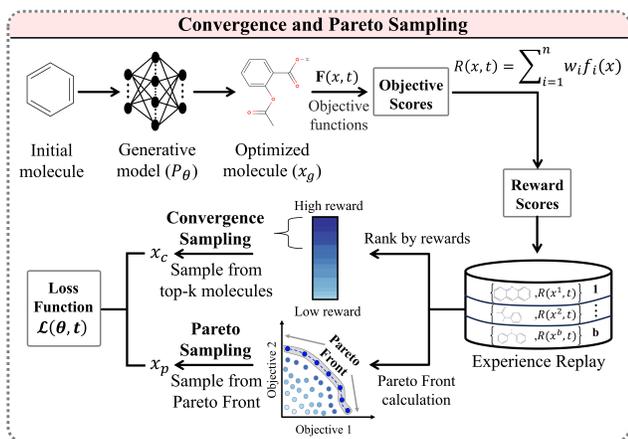


Figure 3: Reward scores are derived as a weighted sum of objective scores from optimized molecules. Molecules and their rewards are stored in Experience replay, which serves two major roles: ranking molecules by rewards for convergence sampling, and extracting the Pareto front for Pareto sampling. These sampled molecules are then integrated into the loss function for training the generative model.

Lemma 2 (Effective Initial Points). Let \tilde{x}_1 be identified as a near-optimal solution for f_1 in \mathcal{P}_1 . Extend the optimization problem to \mathcal{P}_2 by incorporating an additional objective function f_2 . Utilize \tilde{x}_1 as an initial point to define the restricted Pareto set \mathcal{S} by:

$$\mathcal{S} = \left\{ x \in \mathcal{X}_2 \mid \begin{array}{l} f_1(x) = f_1(\tilde{x}_1), \\ f_2(x) = \max\{f_2(z; \tilde{x}_1) \mid z \in \mathcal{X}_2\} \end{array} \right\}, \quad (8)$$

where f_2 is conditioned on \tilde{x}_1 and \mathcal{S} is initially formed by maximizing f_2 while fixing f_1 at the near-optimal level achieved by \tilde{x}_1 . Next, apply the notion of coordinate ascent [Tseng and others, 1988] by re-optimizing f_1 under the constraints imposed by the updated f_2 values. This process can be iterated until convergence. Let $x' \in \mathbf{PF}_2^*$ be a point on the true Pareto front for \mathcal{P}_2 . Assume f_1 and f_2 satisfy a Lipschitz condition [Hager, 1979] on \mathcal{X}_2 that there exists at least local Lipschitz constants $C_i \geq 0$:

$$\|f_i(x) - f_i(x')\| \leq C_i \|x - x'\|, \quad \forall i = 1, 2. \quad (9)$$

Under the conjecture that if $\|x - x'\|$ is sufficiently small (implying x and x' are close to each other), then $\|f_i(x) - f_i(x')\|$ is also small. Thus, with the notion of coordinate ascent, \mathcal{S} can progressively approach closer to \mathbf{PF}_2^* in the input space. Due to page constraints, additional proofs and their proof sketches are provided in the supplementary material 7.1.

4.3 Pareto Sampling and Objective Adaptation

Our proposed model is designed to generate molecules that satisfy given objectives, while promoting score-convergence, Pareto diversity, and structural diversity. Our model employs likelihood P_θ to autoregressively generate molecule x_g , from initial molecule sequence x_0 up to a maximum length L as:

$$P_\theta(x_g = x_L) = \prod_{j=0}^{L-1} P_\theta(x_j | x_{j-1}, x_{j-2}, \dots, x_0). \quad (10)$$

These generated molecules are evaluated by $\mathbf{F}(x, t)$ to obtain objective scores. To provide learning feedback, we define the reward scores $R(x, t)$ at each time stage t by computing the weighted sum of objective scores from $\mathbf{F}(x, t)$ such as:

$$R(x, t) = \sum_{i=1}^n w_i(t) f_i(x), \quad (11)$$

where n is the total number of objective functions and $w_i(t)$ denotes the relative weight for $f_i(x)$ at time stage t .

To further enhance the training efficiency, we utilize experience replay \mathcal{B} that stores previously optimized molecules with high reward scores. In contrast to traditional approaches that primarily emphasize score-convergence, we develop the Pareto sampling technique to also consider Pareto diversity.

Specifically, we perform two types of sampling: convergence sampling, where we sample molecules x_c with high reward scores from \mathcal{B} to promote score-convergence, and Pareto sampling, where we sample molecules x_p from the Pareto front to encourage Pareto diversity. Finally, the generative model parameters θ are optimized by the following loss function:

$$\mathcal{L}(\theta, t) = [-\log P_\theta(\mathbf{x}) + \log P_{\text{prior}}(\mathbf{x}) + R(\mathbf{x}, t)]^2, \quad (12)$$

where P_{prior} is the likelihood of a pre-trained model that imposes additional constraints based on the chemical grammar. It should be noted that \mathbf{x} encompasses a set of molecules x_g , x_c , and x_p , represented as $\mathbf{x} = \{x_g, x_c, x_p\}$.

As time stages advance $t \rightarrow t + 1$, the introduction of new objective changes the composition of objective scores and the reward scores. Although the generative model is initially unaware of these changes, we introduce the objective adaptation technique to update θ . This involves retraining the model using updated reward scores to account for the impact of new objectives. The objective adaptation loss can be expressed as:

$$\mathcal{L}_{\text{OA}}(\theta, t) = [-\log P_\theta(x_b) + \log P_{\text{prior}}(x_b) + R(x_b, t + 1)]^2, \quad (13)$$

where x_b denotes all molecules from \mathcal{B} . Note that we employ REINVENT [Olivecrona *et al.*, 2017] as our backbone generative model due to its superior performance. The pseudo-code for the entire process is in the supplementary material 7.2.

5 Experiments

5.1 Experimental Setup

We evaluated the performance of our proposed method using the practical molecular optimization (PMO) benchmark [Gao *et al.*, 2022]. In this setup, oracle call budgets are strictly limited to 10,000 evaluations to reflect the real-world constraints of molecular discovery. For the oracle functions in our experiments, we adopted the most commonly used molecular objective functions in previous MOMO studies [Jin *et al.*, 2020; Xie *et al.*, 2021]. These include biological objectives GSK3 β and JNK3, which represent inhibition scores against two target proteins related to Alzheimer’s disease, as well as non-biological objectives like QED and SA that quantify drug-likeness and synthesizability, respectively. To extend our approach to many-objective settings, we included further objectives such as DRD2, associated with dopamine receptor binding affinity, as well as Osimertinib MPO and Fexofenadine MPO objectives for discovering new therapeutics that optimize existing drugs with multiple desirable attributes.

Method	Four objectives		Five objectives		Six objectives	
	HV(\uparrow)	R2(\downarrow)	HV(\uparrow)	R2(\downarrow)	HV(\uparrow)	R2(\downarrow)
Random ZINC	0.065 \pm 0.013	4.809 \pm 0.243	0.005 \pm 0.003	8.867 \pm 0.202	0.001 \pm 0.000	14.456 \pm 0.344
SMILES-VAE	0.073 \pm 0.023	4.839 \pm 0.435	0.004 \pm 0.001	9.226 \pm 0.512	0.001 \pm 0.000	15.109 \pm 0.548
GFlowNet	0.063 \pm 0.011	4.952 \pm 0.196	0.011 \pm 0.004	9.011 \pm 0.314	0.004 \pm 0.001	14.462 \pm 0.422
MIMOSA	0.082 \pm 0.028	4.905 \pm 0.410	0.016 \pm 0.012	9.195 \pm 0.761	0.005 \pm 0.005	15.454 \pm 1.701
LaMBO	0.123 \pm 0.006	4.496 \pm 0.148	0.009 \pm 0.001	9.302 \pm 0.159	Out-of-memory	Out-of-memory
HN-GFN	0.120 \pm 0.000	4.013 \pm 0.060	0.004 \pm 9.861	9.861 \pm 0.096	0.002 \pm 0.000	15.349 \pm 0.102
MOEA/D	0.176 \pm 0.123	4.615 \pm 1.193	0.094 \pm 0.052	8.105 \pm 1.561	0.025 \pm 0.018	15.054 \pm 2.028
NSGA-III	0.234 \pm 0.107	3.477 \pm 0.837	0.071 \pm 0.047	7.130 \pm 1.349	0.016 \pm 0.020	14.366 \pm 2.393
GraphGA	0.254 \pm 0.069	3.379 \pm 0.666	0.100 \pm 0.057	7.676 \pm 1.312	0.051 \pm 0.030	11.814 \pm 1.787
GPBO	0.275 \pm 0.091	3.311 \pm 0.757	0.091 \pm 0.031	7.670 \pm 0.761	0.026 \pm 0.024	12.840 \pm 1.811
REINVENT BO	0.309 \pm 0.021	2.795 \pm 0.103	0.071 \pm 0.014	7.537 \pm 0.620	0.033 \pm 0.019	10.929 \pm 1.019
REINVENT	0.338 \pm 0.030	2.770 \pm 0.116	0.099 \pm 0.054	7.578 \pm 1.187	0.062 \pm 0.028	10.032 \pm 0.922
AugMem	0.395 \pm 0.038	2.496 \pm 0.192	0.090 \pm 0.043	8.011 \pm 0.955	0.071 \pm 0.072	12.472 \pm 3.758
DyMol (Ours)	0.422\pm0.023	2.297\pm0.095	0.247\pm0.087	4.943\pm0.990	0.143\pm0.056	8.842\pm1.632

Table 1: Performance comparison in many-objective optimization scenarios with Four objectives (GSK3 β +JNK3+QED+SA), Five objectives (GSK3 β +JNK3+QED+SA+DRD2), and Six objectives (GSK3 β +JNK3+QED+SA+DRD2+Osimertinib MPO) using 10 different seeds.

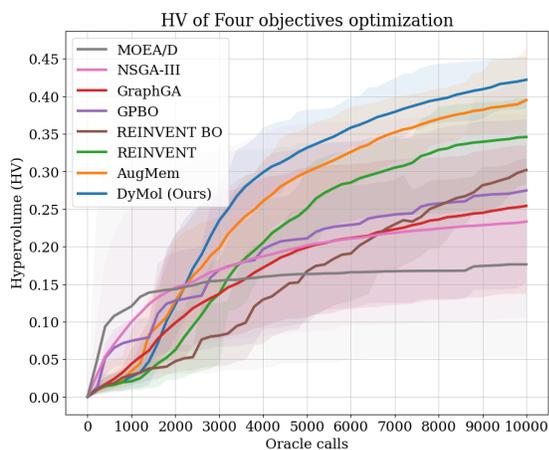


Figure 4: Average HV improvement curves for the top 8 methods.

5.2 Competing Methods

We compared the performance of our method against a range of competing methods, including Random ZINC [Sterling and Irwin, 2015], SMILES-VAE [Gómez-Bombarelli *et al.*, 2018], MIMOSA [Fu *et al.*, 2021], GFlowNet [Bengio *et al.*, 2021], and GraphGA [Jensen, 2019]. Additionally, we evaluated against BO methods such as GPBO [Tripp *et al.*, 2021], LaMBO [Stanton *et al.*, 2022], and HN-GFN [Zhu *et al.*, 2023]; well-known many-objective optimization algorithms like MOEA/D [Zhang and Li, 2007] and NSGA-III [Verhellen, 2022]; and RL-based methods, including REINVENT [Olivecrona *et al.*, 2017], REINVENT BO [Tripp *et al.*, 2021], and AugMem [Guo and Schwaller, 2024]. Note that REINVENT was acknowledged as the best-performing algorithm for molecular optimization, as evidenced by the PMO benchmark results [Gao *et al.*, 2022]. Hence, as mentioned earlier, we employed REINVENT as our backbone generative model. However, we distinctively developed it for a dynamic many-objective setting. More information on competing methods, experimental settings, and hyperparameter configurations is available in the supplementary material 7.3.

5.3 Experimental Results

The performances of our proposed method and the competing methods were assessed by two evaluation metrics: the hypervolume indicator (HV) [Zitzler *et al.*, 2003] and the R2 indicator (R2) [Brockhoff *et al.*, 2012]. The HV measures the volume of the objective space dominated by the Pareto front, while the R2 evaluates the quality of a solution set based on user-defined reference points. A higher HV value indicates a better solution set, while a lower R2 value is more desirable. Refer to the supplementary material 7.4 for detailed explanations of evaluation metrics. Note that each experiment was conducted with 10 different seeds to ensure result reliability.

Table 1 presents the HV and R2 performance with standard deviations for each method across many-objective optimization scenarios with different numbers of objectives as follows:

- Four objectives: GSK3 β + JNK3 + QED + SA;
- Five objectives: GSK3 β + JNK3 + QED + SA + DRD2;
- Six objectives: GSK3 β + JNK3 + QED + SA + DRD2 + Osimertinib MPO.

As shown in Table 1, our method outperforms all competing methods across all scenarios. Notably, in scenarios with Five and Six objectives, our method demonstrates a substantial performance improvement. This highlights the effectiveness of our divide-and-conquer approach, which successfully handles the inherent complexity of many-objective problems by decomposing them into manageable sub-problems. However, other competing methods struggle with exponential increases in complexity. Additional experiments for many-objective scenarios are provided in the supplementary material 7.5.

Figure 4 displays the average HV improvement curves for the top 8 methods in Four objective scenarios. As depicted, our method consistently outperforms others after reaching 2500 oracle calls. Genetic-based algorithms like MOEA/D, NSGA-III, and GPBO demonstrate rapid initial performance improvements but typically reach a plateau beyond 3000 oracle calls. In contrast, RL-based algorithms such as REINVENT and AugMem exhibit more consistent and continual improvement over time. Additional HV improvement curves and analyses can be found in supplementary material 7.6.

Ablation			Four objectives		Five objectives		Six objectives	
PS	DC	OA	HV(\uparrow)	R2(\downarrow)	HV(\uparrow)	R2(\downarrow)	HV(\uparrow)	R2(\downarrow)
-	-	-	0.338	2.770	0.099	7.578	0.062	10.032
✓	-	-	0.379	2.501	0.103	7.018	0.073	10.033
-	✓	-	0.363	2.692	0.150	6.276	0.101	10.408
-	-	✓	0.373	2.621	0.209	5.480	0.105	10.269
✓	✓	-	0.412	2.321	0.182	5.488	0.122	9.439
✓	✓	✓	0.422	2.297	0.247	4.943	0.143	8.842

Table 2: Ablation study for each technique: Pareto Sampling (PS), Divide-and-Conquer (DC), and Objective Adaptation (OA).

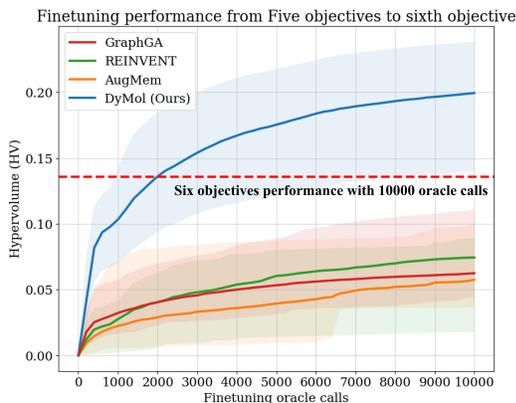


Figure 5: Finetuning performance of the top 4 methods in dynamic-objective scenarios, which introduce a new molecular objective.

5.4 Ablation Study

As shown in Table 2, we have conducted an ablation study to investigate the impact of key techniques on the performance of our method: Pareto Sampling (PS), Divide-and-Conquer (DC), and Objective Adaptation (OA). We observed that each of these techniques significantly contributes to improved performance. DC primarily focuses on improving convergence, bringing solutions closer to the optimal Pareto front values. PS enhances performance through emphasis on Pareto diversity. Remarkably, OA leads to substantial performance gains, especially in Five-objective scenarios, highlighting its capability to adapt effectively to newly introduced objectives.

5.5 Dynamic-Objective Scenarios

To assess our method’s adaptability in dynamic-objective scenarios, we propose a novel experimental setup where a model has initially been fully optimized for a set of Five objectives. Subsequently, a new, sixth objective (Osimertinib MPO) is introduced, requiring additional optimization. Instead of re-optimizing all objectives from scratch, we implement a finetuning approach that leverages the model already optimized for the initial Five objectives, and further optimizing the new objective. As depicted in Figure 5, our method effectively reaches the baseline performance of the Six objectives within fine-tuning 2000 oracle calls and continues to improve beyond that. This achievement can be attributed to our OA technique and the incremental nature of adding objectives within our method. Details on dynamic-objective scenarios and additional experiments are in the supplementary material 7.7.

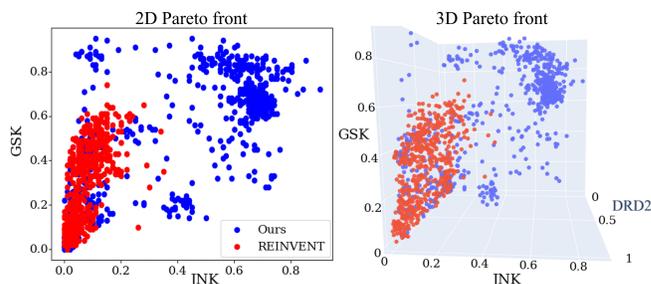


Figure 6: Visualization of the Pareto front for biological objectives.

	Four objectives		Five objectives		Six objectives	
Method	BM(\uparrow)	CS(\uparrow)	BM(\uparrow)	CS(\uparrow)	BM(\uparrow)	CS(\uparrow)
REINVENT	5026	2004	6298	3134	7490	3400
AugMem	6796	2557	7131	3566	6918	4161
DyMol (Ours)	7210	3952	7134	3970	7664	4251

Table 3: Analysis of structural diversity indicated by the number of unique Bemis-Murcko scaffolds (BM) and carbon skeletons (CS).

5.6 Visualization of the Pareto Front

To evaluate solution quality in our method, we have attempted to visualize the Pareto front. However, visualizing the Pareto front in settings with more than three objectives is challenging due to the human limitations in interpreting high-dimensional data [Tušar and Filipič, 2014]. Thus, we only focus on visualizing biological objectives in Five objective scenarios, as they are considered to be more important in drug discovery [Sun *et al.*, 2022]. Figure 6 illustrates the 2D and 3D Pareto fronts, comparing our method with baseline REINVENT method. In both cases, the solutions of our method dominate the baseline method by approaching the optimal Pareto front more closely. Moreover, our method exhibits a broader distribution of solutions along the Pareto front, suggesting a better exploration.

5.7 Structural Diversity Analysis

In the realm of drug discovery, structural diversity plays a pivotal role as it substantially enhances the chances of discovering compounds with distinctive and potent biological activities [Walters and Namchuk, 2003]. To quantitatively assess the structural diversity among the molecules generated by our method, we adopted a diversity metric based on the number of unique Bemis-Murcko scaffolds (BM) and corresponding carbon skeletons (CS) [Bemis and Murcko, 1996]. Refer to the supplementary material 7.8 for further details. As shown in Table 3, our method outperforms other methods in terms of BM and CS, indicating a higher average structural diversity.

6 Conclusion

In this work, we propose DyMol as a novel and first method to address the dynamic many-objective molecular optimization problem by leveraging the divide-and-conquer approach. DyMol decomposes complex many-objective sets into manageable sub-problems for progressive optimization. Our results demonstrate that DyMol outperforms competing methods in both many-objective and dynamic-objective scenarios. Future work can include extending research to material science.

Acknowledgments

This work was supported by grants from the Institute of Information & Communications Technology Planning & Evaluation (IITP)—specifically, the Artificial Intelligence Graduate School Program at Korea University (No. 2019-0-00079) and the Development of AI Autonomy and Knowledge Enhancement for AI Agent Collaboration (No. 2022-0-00871). Additionally, this research received funding from the National Research Foundation of Korea (NRF), supported by the Korea government (MSIT) under grant (No. RS202300212498).

Contribution Statement

Tae-Eui Kam is the corresponding author. Dong-Hee Shin and Young-Han Son contributed equally to this work as co-first authors. Deok-Joong Lee assisted with figure preparation and data analysis. Ji-Wung Han proofread the manuscript. All authors have read and approved the final manuscript.

References

- [Aldeghi *et al.*, 2022] Matteo Aldeghi, David E Graff, Nathan Frey, Joseph A Morrone, Edward O Pyzer-Knapp, Kirk E Jordan, and Connor W Coley. Roughness of molecular property landscapes and its impact on modellability. *Journal of Chemical Information and Modeling*, 62(19):4660–4671, 2022.
- [Bemis and Murcko, 1996] Guy W Bemis and Mark A Murcko. The properties of known drugs. 1. molecular frameworks. *Journal of medicinal chemistry*, 39(15):2887–2893, 1996.
- [Bengio *et al.*, 2021] Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow network based generative models for non-iterative diverse candidate generation. *Advances in Neural Information Processing Systems*, 34:27381–27394, 2021.
- [Bilodeau *et al.*, 2022] Camille Bilodeau, Wengong Jin, Tommi Jaakkola, Regina Barzilay, and Klavs F Jensen. Generative models for molecular discovery: Recent advances and challenges. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 12(5):e1608, 2022.
- [Brattka *et al.*, 2012] Vasco Brattka, Guido Gherardi, and Alberto Marcone. The Bolzano–Weierstrass theorem is the jump of weak König’s lemma. *Annals of Pure and Applied Logic*, 163(6):623–655, 2012.
- [Brockhoff *et al.*, 2012] Dimo Brockhoff, Tobias Wagner, and Heike Trautmann. On the properties of the r_2 indicator. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, pages 465–472, 2012.
- [Darrow *et al.*, 2020] Jonathan J Darrow, Jerry Avorn, and Aaron S Kesselheim. Fda approval and regulation of pharmaceuticals, 1983–2018. *Jama*, 323(2):164–176, 2020.
- [Deb and Jain, 2013] Kalyanmoy Deb and Himanshu Jain. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: solving problems with box constraints. *IEEE transactions on evolutionary computation*, 18(4):577–601, 2013.
- [Ding *et al.*, 2018] Jinliang Ding, Cuie Yang, Qiong Xiao, Tianyou Chai, and Yaochu Jin. Dynamic evolutionary multiobjective optimization for raw ore allocation in mineral processing. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 3(1):36–48, 2018.
- [Eichfelder, 2009] Gabriele Eichfelder. Scalarizations for adaptively solving multi-objective optimization problems. *Computational Optimization and Applications*, 44:249–273, 2009.
- [Fromer and Coley, 2023] Jenna C Fromer and Connor W Coley. Computer-aided multi-objective optimization in small molecule discovery. *Patterns*, 4(2), 2023.
- [Fu *et al.*, 2021] 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, pages 125–133, 2021.
- [Gao *et al.*, 2022] Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor Coley. Sample efficiency matters: a benchmark for practical molecular optimization. *Advances in Neural Information Processing Systems*, 35:21342–21357, 2022.
- [Gómez-Bombarelli *et al.*, 2018] 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.
- [Gunantara, 2018] Nyoman Gunantara. A review of multi-objective optimization: Methods and its applications. *Cogent Engineering*, 5(1):1502242, 2018.
- [Guo and Schwaller, 2024] Jeff Guo and Philippe Schwaller. Augmented memory: Sample-efficient generative molecular design with reinforcement learning. *JACS Au*, 2024.
- [Hager, 1979] William W Hager. Lipschitz continuity for constrained processes. *SIAM Journal on Control and Optimization*, 17(3):321–338, 1979.
- [Huang *et al.*, 2021] Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Therapeutics data commons: Machine learning datasets and tasks for drug discovery and development. *arXiv preprint arXiv:2102.09548*, 2021.
- [Hughes, 2005] Evan J Hughes. Evolutionary many-objective optimisation: many once or one many? In *2005 IEEE congress on evolutionary computation*, volume 1, pages 222–227. IEEE, 2005.
- [Jensen, 2019] Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. *Chemical science*, 10(12):3567–3572, 2019.

- [Jin *et al.*, 2020] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Multi-objective molecule generation using interpretable substructures. In *International conference on machine learning*, pages 4849–4859. PMLR, 2020.
- [Laumanns and Ocenasek, 2002] Marco Laumanns and Jiri Ocenasek. Bayesian optimization algorithms for multi-objective optimization. In *International Conference on Parallel Problem Solving from Nature*, pages 298–307. Springer, 2002.
- [Lin *et al.*, 2022] Xi Lin, Zhiyuan Yang, Xiaoyuan Zhang, and Qingfu Zhang. Pareto set learning for expensive multi-objective optimization. *Advances in Neural Information Processing Systems*, 35:19231–19247, 2022.
- [Liu *et al.*, 2021] Ruochen Liu, Ping Yang, and Jiangdi Liu. A dynamic multi-objective optimization evolutionary algorithm for complex environmental changes. *Knowledge-Based Systems*, 216:106612, 2021.
- [Luukkonen *et al.*, 2023] Sohvi Luukkonen, Helle W van den Maagdenberg, Michael TM Emmerich, and Gerard JP van Westen. Artificial intelligence in multi-objective drug design. *Current Opinion in Structural Biology*, 79:102537, 2023.
- [Marler and Arora, 2004] R Timothy Marler and Jasbir S Arora. Survey of multi-objective optimization methods for engineering. *Structural and multidisciplinary optimization*, 26:369–395, 2004.
- [Mathur and Hoskins, 2017] Sunil Mathur and Clare Hoskins. Drug development: Lessons from nature. *Biomedical reports*, 6(6):612–614, 2017.
- [Olivecrona *et al.*, 2017] Marcus Olivecrona, Thomas Blaschke, Ola Engkvist, and Hongming Chen. Molecular de-novo design through deep reinforcement learning. *Journal of cheminformatics*, 9(1):1–14, 2017.
- [Quan *et al.*, 2022] Zhen Quan, Yan Wang, and Zhicheng Ji. Multi-objective optimization scheduling for manufacturing process based on virtual workflow models. *Applied Soft Computing*, 122:108786, 2022.
- [Shin *et al.*, 2024] Dong-Hee Shin, Young-Han Son, Jun-Mo Kim, Hee-Jun Ahn, Jun-Ho Seo, Chang-Hoon Ji, Ji-Wung Han, Byung-Jun Lee, Dong-Ok Won, and Tae-Eui Kam. MARS: Multiagent reinforcement learning for spatial-spectral and temporal feature selection in EEG-based BCI. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2024.
- [Son *et al.*, 2024] Young-Han Son, Dong-Hee Shin, and Tae-Eui Kam. FTMMR: Fusion transformer for integrating multiple molecular representations. *IEEE Journal of Biomedical and Health Informatics*, 2024.
- [Stanton *et al.*, 2022] Samuel Stanton, Wesley Maddox, Nate Gruver, Phillip Maffettone, Emily Delaney, Peyton Greenside, and Andrew Gordon Wilson. Accelerating bayesian optimization for biological sequence design with denoising autoencoders. In *International Conference on Machine Learning*, pages 20459–20478. PMLR, 2022.
- [Sterling and Irwin, 2015] Teague Sterling and John J Irwin. Zinc 15–ligand discovery for everyone. *Journal of chemical information and modeling*, 55(11):2324–2337, 2015.
- [Sun *et al.*, 2022] Mengying Sun, Jing Xing, Han Meng, Huijun Wang, Bin Chen, and Jiayu Zhou. Molsearch: search-based multi-objective molecular generation and property optimization. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pages 4724–4732, 2022.
- [Tripp *et al.*, 2021] Austin Tripp, Gregor NC Simm, and José Miguel Hernández-Lobato. A fresh look at de novo molecular design benchmarks. In *NeurIPS 2021 AI for Science Workshop*, 2021.
- [Tseng and others, 1988] Paul Tseng et al. Coordinate ascent for maximizing nondifferentiable concave functions. 1988.
- [Tušar and Filipič, 2014] Tea Tušar and Bogdan Filipič. Visualization of pareto front approximations in evolutionary multiobjective optimization: A critical review and the projection method. *IEEE Transactions on Evolutionary Computation*, 19(2):225–245, 2014.
- [Verhellen, 2022] Jonas Verhellen. Graph-based molecular pareto optimisation. *Chemical Science*, 13(25):7526–7535, 2022.
- [Walters and Namchuk, 2003] W Patrick Walters and Mark Namchuk. Designing screens: how to make your hits a hit. *Nature reviews Drug discovery*, 2(4):259–266, 2003.
- [Wang *et al.*, 2006] Mingliang Wang, Xiangqian Hu, David N Beratan, and Weitao Yang. Designing molecules by optimizing potentials. *Journal of the American Chemical Society*, 128(10):3228–3232, 2006.
- [Xie *et al.*, 2021] Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, and Lei Li. Mars: Markov molecular sampling for multi-objective drug discovery. *arXiv preprint arXiv:2103.10432*, 2021.
- [Yuan *et al.*, 2015] Yuan Yuan, Hua Xu, Bo Wang, Bo Zhang, and Xin Yao. Balancing convergence and diversity in decomposition-based many-objective optimizers. *IEEE Transactions on Evolutionary Computation*, 20(2):180–198, 2015.
- [Zhang and Li, 2007] Qingfu Zhang and Hui Li. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on evolutionary computation*, 11(6):712–731, 2007.
- [Zhu *et al.*, 2023] Yiheng Zhu, Jialu Wu, Chaowen Hu, Jiahuan Yan, Chang-Yu Hsieh, Tingjun Hou, and Jian Wu. Sample-efficient multi-objective molecular optimization with GFlownets. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [Zitzler *et al.*, 2003] Eckart Zitzler, Lothar Thiele, Marco Laumanns, Carlos M Fonseca, and Viviane Grunert Da Fonseca. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on evolutionary computation*, 7(2):117–132, 2003.