Expert-Inspired Multi-Agent Coordination for Multi-Objective Molecular Optimization

Daojian Zeng[1], Tianle Li[1], Jiahao Yang[2,3], Jiacai Yi[4], Xieping Gao[1], Lincheng Jiang[5,*], Tengfei Ma[2,*], Xiangxiang Zeng[2]

[1]Hunan Provincial Key Laboratory of Philosophy and Social Sciences of Artificial Intelligence and International Communication, Hunan Normal University, Changsha 410081, China
[2]State Key Laboratory of Chemo and Biosensing, College of Computer Science and Electronic Engineering, Hunan University, Changsha 410082, China
[3]Lingang Laboratory, Shanghai 201306, China
[4]School of Computer Science, National University of Defense Technology, Changsha 410073, China
[5]College of Advanced Interdisciplinary Studies, National University of Defense Technology, Changsha 410073, China
{zengdj, linjiadegou2}@hunnu.edu.cn, tfma@hnu.edu.cn, linchengjiang@nudt.edu.cn

Abstract

Multi-objective molecular optimization is a fundamental yet inherently challenging task in drug discovery, as it requires simultaneously optimizing multiple, often conflicting, molecular properties. Although recent deep learning methods have shown promise, they often lack objective-specific specialization and dynamic coordination, which makes them ineffective at handling competing objectives and limits their scalability in complex, high-dimensional molecular design tasks. Inspired by the division of labor among domain experts in medicinal chemistry, we propose **MAMO**, a multi-agent framework for molecular design that simulates expert collaboration. Each agent specializes in optimizing a single objective, and their interactions are orchestrated by a central scheduling module that dynamically reallocates tasks based on evaluation feedback. This coordination mechanism enables interpretable and goal-conditioned optimization while adaptively balancing conflicting objectives. Extensive experiments on benchmark datasets demonstrate that MAMO consistently achieves superior performance in both objective quality and Pareto diversity, particularly in scenarios with strong inter-objective conflict. Our results highlight the potential of multi-agent coordination strategies for scalable and conflict-aware molecular design.

Code — https://github.com/linjiadegou/MAMO

Introduction

Multi-objective molecular optimization plays a pivotal role in early-stage drug discovery (De Rycker et al. 2018), where the goal is to refine a given lead compound to simultaneously satisfy multiple desirable properties such as bioactivity, drug-likeness, and synthesizability (He et al. 2021). Traditional approaches, including high-throughput screening (HTS) (Graff, Shakhnovich, and Coley 2021) and simulation-based techniques (Hsu et al. 2017), are inherently time-consuming and resource-intensive.

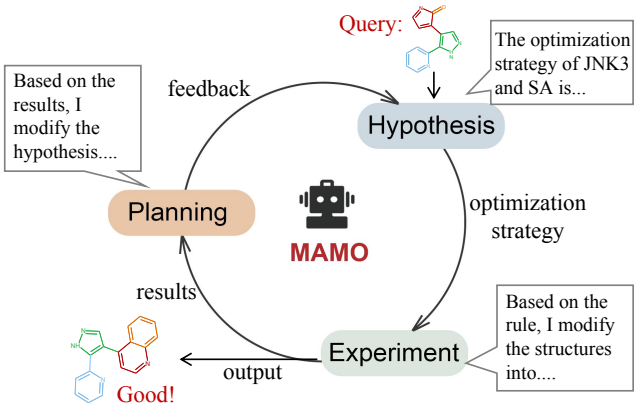

Figure 1: Conceptual illustration of MAMO. Inspired by expert-driven scientific workflows, MAMO follows a closed-loop process of hypothesis generation, structural modification, evaluation, and planning to iteratively optimize molecules across multiple conflicting objectives.

Influenced by the rapid development of artificial intelligence, machine learning has demonstrated strong potential in efficient candidate discovery (Hoffman et al. 2022). Previous methods for multi-objective molecular design typically reduce the problem to a single-objective formulation by assigning predefined weights to individual objectives (Ji et al. 2021; Xia et al. 2024). While effective in some settings, these methods heavily depend on expert-designed weightings, which are often difficult to calibrate and may lead to suboptimal solutions when objectives conflict (Xie et al. 2021; Fromer and Coley 2023). Alternatively, some approaches aim to identify the Pareto front through a two-stage process involving large-scale molecular sampling and non-dominated sorting (Verhellen 2022). However, such methods are computationally expensive and scale poorly with the number of objectives and candidate molecules. To address efficiency concerns, Bayesian optimization has been widely adopted due to its sample efficiency and principled

*Corresponding authors.

uncertainty modeling (Xie et al. 2021; Gao et al. 2022), yet it struggles with high-dimensional spaces and costly Gaussian process inference, limiting its practical use in multi-objective molecular design.

With the rapid advances in large language models (LLMs) (OpenAI et al. 2024; Bai et al. 2023; Dubey et al. 2024), recent research has begun to explore their potential in molecular generation (Brahmavar et al. 2024; Wang, Yang, and Shen 2025) and optimization (Yu et al. 2025; Ye et al. 2025). LLMs offer a flexible and scalable alternative, with the capacity to perform goal-conditioned generation and reasoning across diverse property objectives (Nguyen and Grover 2024). However, existing LLM-based methods often treat multi-objective optimization as a monolithic task, lacking explicit mechanisms for balancing trade-offs among conflicting goals (Liu et al. 2025).

To address the limitations of monolithic LLM-based approaches in managing trade-offs across multiple molecular objectives, we take inspiration from the collaborative workflows in medicinal chemistry (Pitt et al. 2025). In real-world lead optimization, domain experts such as biologists, pharmacokineticists, and synthetic chemists each contribute by focusing on distinct property dimensions, including activity, ADME, and synthetic accessibility. These experts iteratively refine molecular candidates through cycles of hypothesis generation, experimental testing, and adaptive planning (Kapustina et al. 2024; Mathur et al. 2025). Building on this paradigm, we present **MAMO**, a multi-agent framework that simulates expert collaboration through a closed-loop optimization process (Figure 1). To operationalize this process, MAMO is organized into four coordinated modules: **Formulation**, **Collaboration**, **Evaluation**, and **Scheduling**, each of which plays a distinct role in enabling iterative and goal-driven molecular optimization (see Section Method). At each iteration, LLM-based expert agents generate hypotheses and propose property-specific structural modifications based on their assigned optimization goals. These candidates are then evaluated by external tools, and results are fed back to a central scheduling agent that adaptively reallocates priorities for the next iteration. This iterative process mirrors the scientific discovery loop and allows MAMO to flexibly balance conflicting objectives, support modular reasoning, and converge on high-quality molecules. By integrating LLMs with structured agent coordination, MAMO offers an effective solution to multi-objective molecular optimization. In summary, our contributions include:

- We propose **MAMO**, an expert-inspired multi-agent framework that simulates domain-specific collaboration for multi-objective molecular optimization, addressing the intrinsic conflict between properties.
- We introduce a **centralized scheduling agent** that dynamically allocates optimization tasks based on real-time evaluation feedback, enabling adaptive coordination and effective trade-off balancing among expert agents.
- Extensive experiments on benchmark tasks demonstrate that MAMO consistently outperforms strong baselines in both optimization quality and objective balancing, particularly under high-conflict, high-dimensional scenarios.

Related Work

Conventional molecular optimization. Molecular optimization is a fundamental task in drug discovery and has evolved from manual experiments to computational methods (Gao et al. 2022). Current approaches fall into two main categories: (1) *Generative models*, including variational autoencoders (Gómez-Bombarelli et al. 2018; Jin, Barzilay, and Jaakkola 2018), GANs (Guimaraes et al. 2017; Kadurin et al. 2017), and diffusion models (Ho, Jain, and Abbeel 2020; Xu et al. 2022), which generate candidate molecules in a data-driven manner; and (2) *Combinatorial optimization*, which searches chemical space using techniques like Monte Carlo tree search (Yang et al. 2017), genetic algorithms (Fu et al. 2022), and reinforcement learning (Olivecrona et al. 2017; Gu and Dao 2023).Despite their progress in single-objective tasks, these methods often lack the flexibility needed to balance conflicting objectives.

LLM for molecular optimization. Recent advances in large language models (LLMs) have opened up new opportunities for molecule-related tasks, including property prediction and molecular generation (Luo et al. 2022; Yunxiang et al. 2023; Han et al. 2023; Wu et al. 2024; Fang et al. 2023). Emerging efforts have also explored the application of LLMs in molecular optimization. For instance, MOLLEO (Wang et al. 2024) integrated LLMs into genetic algorithms by using them as operators for crossover and mutation, improving both solution quality and convergence speed. LICO (Nguyen and Grover 2024) employed context-aware prompting for molecule refinement, enabling in-context optimization without retraining. DrugAssist (Ye et al. 2023) further proposed an interactive human-in-the-loop optimization framework built on LLMs to leverage human knowledge and LLM reasoning collaboratively. Despite these initial efforts, most existing LLM-based methods treat multi-objective optimization as a unified generation task, lacking explicit modeling of inter-objective conflicts and coordination. The full potential of LLMs for adaptive and multi-agent optimization in complex design tasks remains largely untapped.

Multi-agent LLM systems. Large language models (LLMs) have become core components of autonomous agents for complex reasoning and decision-making tasks (Chen et al. 2023). Early systems such as AutoGPT (Yang, Yue, and He 2023) demonstrated tool-augmented task execution but were limited to single-agent settings, making them less effective for solving multi-dimensional or interdependent tasks. To overcome these limitations, recent work has explored multi-agent collaboration with LLMs. AutoGen (Wu et al. 2023) enables structured communication among role-specific agents, while MetaGPT (Hong et al. 2023) improves task planning and execution through coordinated agent roles. These approaches demonstrate the potential of multi-agent LLM frameworks for solving complex tasks. In particular, tasks such as multi-objective molecular optimization can benefit from expert-like role decomposition and collaboration, which remain underexplored in current systems.

Method

Problem Definition. We formulate multi-objective molecular optimization as a discrete optimization problem over the chemical space $\mathcal{X}$, where each molecule is represented by a valid SMILES string. Given an initial molecule $x_0 \in \mathcal{X}$, the goal is to generate a set of molecules $\{x_1, x_2, \ldots, x_n\} \subset \mathcal{X}$ that jointly optimize multiple properties. Formally, we aim to solve a K-objective maximization problem:

$$\max_{x \in \mathcal{X}} \quad \mathbf{f}(x) = [f_1(x), f_2(x), \ldots, f_K(x)], \tag{1}$$

where each $f_k(x) : \mathcal{X} \rightarrow \mathbb{R}$ denotes a non-differentiable scoring function for a molecular property, such as GSK3β/JNK3 inhibition, drug-likeness (QED), and synthetic accessibility (SA). The objective is to approximate the Pareto front $\mathcal{P} \subset \mathcal{X}$, i.e., the set of non-dominated molecules. We enforce chemical validity as a hard constraint to ensure all generated molecules are syntactically and chemically feasible.

Overview of MAMO. We introduce **MAMO**, a modular multi-agent framework for multi-objective molecular optimization. As shown in Figure 2a, MAMO comprises four key components: *Formulation*, *Collaboration*, *Evaluation*, and *Scheduling*. Given a molecule in SMILES format (Weininger 1988), each expert agent—powered by a large language model (LLM)—focuses on optimizing a specific molecular property (e.g., QED, SA). Agents collaboratively propose candidates based on property-aware strategies, which are evaluated by external tools. A scheduling module adaptively coordinates agent interactions based on feedback to guide the search process. This design enables goal-directed, scalable optimization across conflicting objectives. The functional roles of agents, grounded in domain knowledge, are illustrated in Figure 2b.

Formulation. *Agent Initialization*: As illustrated in Figure 2a, we instantiate a set of expert agents, each dedicated to optimizing a specific molecular property. Each agent is built upon an LLM, such as GPT-4o mini (OpenAI et al. 2024), and is assigned a domain-specific role, including QED Expert, SA Expert, GSK3β/JNK3 (targets related to Alzheimer's) Experts, according to the optimization target (Figure 2b). Recent studies demonstrate that the performance of LLMs on biochemical tasks can be significantly improved through domain-aware prompt engineering (Li et al. 2025; Luo et al. 2025). Following this insight, we design tailored prompts for each agent to explicitly define the task scope, optimization objective, input format, and expected output structure. An example prompt for the QED expert is shown in Figure 3. To further enhance agent specialization and actionability, we equip each expert with tool-calling capabilities, allowing access to domain tools such as RDKit and property oracles for informed decision-making. Details on tool usage are provided in Appendix A. *Optimization Strategy Formulation*: MAMO adopts a progressive optimization mechanism that decomposes molecular design into an iterative sequence of strategy-driven steps. At each stage, expert agents adapt their behavior based on feedback from prior rounds, forming a coherent and adaptive optimization trajectory. This design allows the framework to maintain logical consistency while flexibly adjusting optimization plans to match the molecular context and evolving property profiles. For each objective, the corresponding expert agent generates a strategy through prompt-based reasoning. These strategies are dynamically tailored to the characteristics of the input molecule, ensuring that modifications are both targeted and interpretable. By iteratively refining the molecule toward predefined property goals, the system effectively narrows the gap between current and desired performance. Additionally, strategy prompts are updated at each step based on multi-agent feedback, enabling agents to respond adaptively to emerging trade-offs (Appendix B).

Collaboration. To enable effective coordination among expert agents during molecular optimization, MAMO introduces a shared-memory mechanism where all agents access and update a common molecular context. This shared memory stores intermediate molecules, strategy updates, and property evaluations, allowing agents to (i) avoid redundant edits, (ii) recognize emerging trade-offs, and (iii) adapt their context-driven strategies based on global optimization dynamics. However, as the number of objectives increases, simple peer-to-peer interaction becomes inefficient. To address this, MAMO incorporates a centralized scheduling agent that monitors the shared memory and dynamically assigns optimization tasks to experts based on real-time progress and remaining objectives. This scheduler improves coordination by minimizing inter-agent conflicts, accelerating convergence, and enhancing the quality of the Pareto front. The scheduling mechanism is detailed in Section Central Scheduling.

Evaluation. Following each round of multi-agent collaboration, MAMO employs a dedicated *evaluation agent* to assess the quality of generated candidate molecules across all relevant property dimensions. This agent operates independently of the expert agents and is responsible for verifying whether the current candidates meet the predefined optimization goals. To perform this assessment, the evaluation agent dynamically selects appropriate property prediction tools—such as RDKit-based QED or synthetic accessibility calculators—based on the task configuration. It focuses only on the properties under active optimization and computes the corresponding scores to determine task satisfaction. If any objective remains unmet, the evaluation agent generates structured feedback and forwards it to the central scheduling agent. This feedback is then used to adjust agent priorities and revise optimization strategies in subsequent rounds. The evaluation process ensures closed-loop coordination between local optimization and global objective satisfaction. Detailed tool configurations and the prompt design of the evaluation agent are provided in Appendix C.

Central Scheduling. In the multi-agent collaboration framework, we define the set of candidate molecules as

$$\mathcal{M} = \{x_1, x_2, \ldots, x_N\}, \tag{2}$$

and the set of expert agents as

$$\mathcal{E} = \{e_1, e_2, \ldots, e_K\}, \tag{3}$$

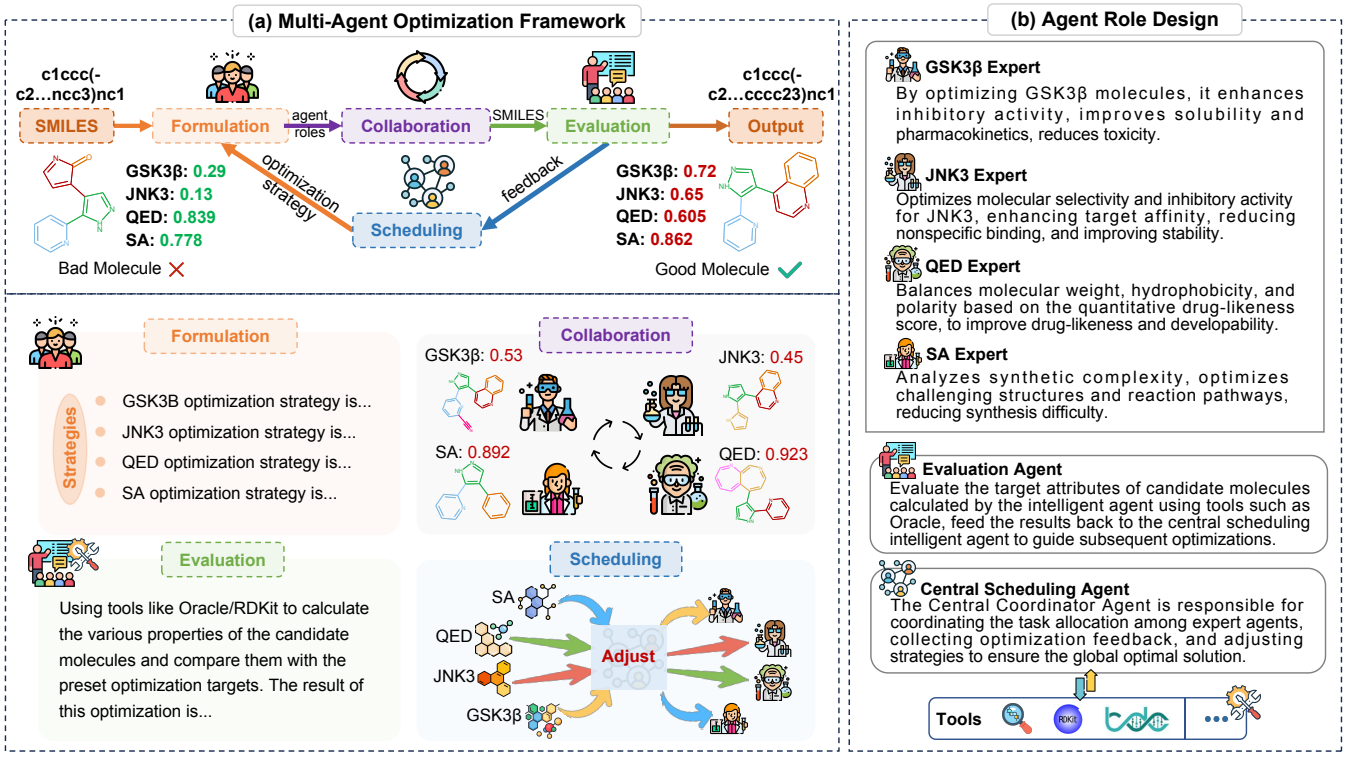

Figure 2: Overview of the **MAMO** framework. (a) A multi-agent optimization pipeline where each expert agent focuses on a specific molecular property. The system iteratively refines candidates through formulation, collaboration, evaluation, and dynamic scheduling. (b) Role definitions and tool configurations for expert, evaluation, and scheduling agents.

QED Expert Prompt

You are a molecular optimization expert specializing in improving the Drug–likeness (QED) of molecules...

Task Description:
Based on the provided molecular information and optimization strategies, use a Chain–of–Thought (COT) strategy generation mechanism...
Optimization Strategies:
{Generated through the optimization strategies}
Input Information:
SMILES: {}
Output Requirements:
Provide the optimized molecule SMILES as a list, returning only the SMILES strings in the following format:...

Figure 3: Prompt Template for the QED Expert Agent.

where N is the number of candidate molecules and K indicates the number of optimized objectives. The optimization proceeds in discrete rounds $t = 1, 2, \ldots, T$. After each round, each expert agent e_j independently optimizes the assigned molecules. A central scheduling agent (referred to as the scheduler) is responsible for task scheduling and allocation to enhance the efficiency of global optimization. After each round t, an evaluation agent provides feedback on molecular properties, represented as a property vector:

$$\mathbf{p}_i^t \in \mathbb{R}^K, \tag{4}$$

where $\mathbf{p}_i^t$ is the property values of i-th molecule from t-th round. Given a target property vector $\mathbf{p}^* \in \mathbb{R}^K$, we define the optimization gap as

$$g_i^t = d\big(\mathbf{p}_i^t, \mathbf{p}^*\big), \tag{5}$$

where $d(\cdot, \cdot)$ is defined as the Euclidean distance. The normalized priority score for each molecule is calculated as

$$\pi_i^t = \frac{g_i^t}{\sum_{j=1}^{N} g_j^t}. \tag{6}$$

Molecules with larger π_i^t are considered more critical for optimization in the next round. To allocate tasks for round $t+1$, the scheduler dynamically assigns molecules to agents based on the priority scores π_i^t. This ensures that each agent is assigned exactly one molecule, and each molecule is assigned to at most one agent, maximizing overall optimization priorities. The scheduler does not directly engage in molecular optimization but dynamically adjusts allocations based on updated evaluations $\{\mathbf{p}_i^t\}$ and molecule dependencies. Throughout the optimization process, the system iteratively executes the sequence:

$$\{A^1, A^2, \ldots, A^T\} \xrightarrow{\text{scheduling}} \text{global convergence} \tag{7}$$

By separating scheduling from optimization, the proposed multi-agent framework achieves higher flexibility and scalability in complex molecular design tasks. Detailed prompt designs for the scheduler are provided in Appendix D.

Methods	QED+SA	GSK3β+JNK3	JNK3+QED+SA	GSK3β+QED+SA	GSK3β+JNK3+QED	GSK3β+JNK3+QED+SA	Sum (↑)
REINVENT	0.894±0.000	0.145±0.000	0.476±0.005	0.335±0.000	0.057±0.006	0.092±0.000	1.999
Graph GA	**0.934±0.002**	0.302±0.067	0.240±0.073	0.526±0.045	0.185±0.076	0.342±0.096	2.529
SMILES GA	0.876±0.003	0.097±0.011	0.093±0.004	0.377±0.077	0.085±0.016	0.070±0.021	1.598
GP BO	0.925±0.001	0.241±0.043	0.201±0.039	0.466±0.030	0.320±0.110	0.304±0.101	2.457
REINVENT Trans	0.931±0.002	0.478±0.058	0.432±0.059	0.628±0.068	0.358±0.088	0.270±0.057	3.097
DrugAssist	0.812±0.002	0.583±0.012	0.565±0.036	0.597±0.045	0.387±0.023	0.444±0.063	3.388
EAG	0.912±0.001	0.669±0.021	0.665±0.043	0.631±0.061	0.585±0.069	0.517±0.075	3.979
GPT-4o mini	0.851±0.003	0.612±0.021	0.605±0.055	0.574±0.072	0.564±0.101	0.501±0.074	3.707
MAMO	0.928±0.001	**0.728±0.018**	**0.732±0.067**	**0.728±0.041**	**0.718±0.062**	**0.676±0.031**	**4.510**

Table 1: Performance (HV) comparison across methods. The **bold** and underline denote the best and second-best scores.

	QED	SA	GSK3β	JNK3
QED	1	0.051	-0.094	-0.220
SA	0.051	1	0.064	-0.055
GSK3β	-0.094	0.064	1	0.351
JNK3	-0.220	-0.055	0.351	1

Table 2: The correlation between multiple objectives.

EXPERIMENTS

Experimental Setup

Implementation Details. We consider four molecular property objectives and analyze their pairwise correlations on a subset of the ZINC20 dataset (Irwin et al. 2020). As shown in Table 2, most objective pairs exhibit low correlation, indicating substantial inter-objective independence. Based on this, we design five task settings with varying conflict intensities: 1) QED + SA (non-biological objectives): Drug-likeness (QED) and synthetic accessibility (SA) measure the developability and synthesizability of molecules, computed using RDKit. 2) GSK3β + JNK3 (biological objectives): The inhibition of GSK3β and JNK3, two kinase targets associated with Alzheimer's disease, predicted using random forest models. 3) QED + SA + GSK3β/JNK3: Optimization of either GSK3β or JNK3 inhibition under constraints of good QED and SA properties. 4) GSK3β + JNK3 + QED: Simultaneous optimization of GSK3β and JNK3 inhibition with QED constraints, without explicit control over synthetic accessibility. 5) QED + SA + GSK3β + JNK3: Joint optimization across all four objectives to balance activity, drug-likeness, and synthesizability.

Dataset and Metrics. We adopt the RationaleRL dataset (Jin, Barzilay, and Jaakkola 2020) to ensure a diverse initial molecular population for multi-objective optimization across GSK3β inhibition, JNK3 inhibition, QED, and SA. To assess optimization performance, we primarily use the hypervolume (HV) metric (Zitzler et al. 2003), which measures the volume of the dominated region in objective space relative to a reference point. HV effectively captures both convergence toward the Pareto front and the diversity of the solution set. In addition, we report standard evaluation metrics in molecular generation, including novelty and diversity. (details in Appendix F).

Baselines. We compare **MAMO** against both traditional baselines and LLM-based frameworks. The traditional methods include: 1) **REINVENT** (Olivecrona et al. 2017), which employs reinforcement learning to guide an RNN-based generative model; 2) **Graph GA** (Jensen 2019), a genetic algorithm operating on molecular graph structures; 3) **SMILES GA** (Brown et al. 2019), which performs genetic optimization over SMILES representations; 4) **GP BO** (Tripp, Simm, and Hernández-Lobato 2021), a Gaussian process-based Bayesian optimization method; and 5) **REINVENT Trans** (Pre-training and Fine-tuning 2024), a transformer-based variant of REINVENT. The LLM-based baselines include: 1) **DrugAssist** (Ye et al. 2025), which formulates molecule optimization as an interactive dialogue with an LLM; 2) **EAG** (Gu et al. 2025) decomposes complex tasks into pipeline stages to enhance multi-agent coordination; and 3) a direct **GPT-4o mini**-based optimization baseline, which applies a single-agent LLM to optimize multiple objectives simultaneously (details in Appendix G).

Main Results

Performance Comparison on Multi-Objective Tasks. Table 1 reports HV scores across six multi-objective molecular optimization tasks. MAMO consistently achieves the best performance in most settings, yielding the highest HV scores in each objective combination and the largest overall HV sum (4.51). This demonstrates its superior ability to explore diverse high-quality solutions and effectively balance conflicting objectives. On simpler dual-objective tasks such as QED+SA and GSK3β+JNK3, most baselines perform reasonably well, especially Graph GA and REINVENT Trans. However, their performance declines as objectives increase. Notably, SMILES GA and REINVENT perform poorly on tasks involving biological targets, highlighting their limited scalability in high-dimensional spaces. Among LLM-based approaches, single-agent methods like DrugAssist (HV sum: 3.39) and direct GPT-4o mini optimization (HV sum: 3.71) show constrained capability in handling objective conflicts. While the multi-agent framework EAG (HV sum: 3.98) demonstrates improved coordination, MAMO's specialized expert collaboration and dynamic scheduling significantly outperform all these LLM-based alternatives.

Performance on Top-ranked Candidate Molecules. To further understand the effectiveness of each method on individual molecular properties, we evaluate the top-10 and top-50 average scores for each objective, as shown in Tables 3 and 4. We also report the overall average score and average rank for each method. In the biological objective

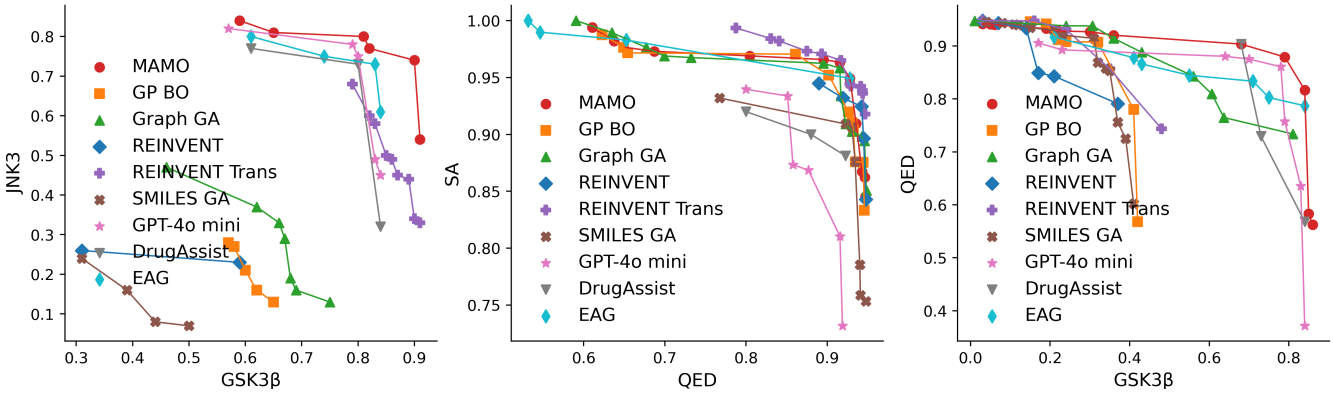

Figure 4: Non-dominated solutions of various methods on GSK3β+JNK3, QED+SA, and GSK3β+QED objectives.

Methods	Top-10 Avg GSK3β	Top-10 Avg JNK3	Top-50 Avg GSK3β	Top-50 Avg JNK3	Avg GSK3β	Avg JNK3	Avg Rank
REINVENT	0.495	0.234	0.395	0.167	0.077	0.036	9
Graph GA	0.688	0.343	0.636	0.244	0.296	0.061	7
SMILES GA	0.448	0.211	0.383	0.181	0.164	0.067	8
GP BO	0.619	0.229	0.586	0.179	0.331	0.059	6
REINVENT Trans	**0.880**	0.616	**0.865**	0.557	<u>0.429</u>	0.217	5
DrugAssist	0.791	0.698	0.729	0.593	0.382	0.255	4
EAG	0.812	<u>0.723</u>	0.755	<u>0.635</u>	0.412	<u>0.282</u>	2
GPT-4o mini	0.801	0.708	0.739	0.633	0.408	0.268	3
MAMO	<u>0.846</u>	**0.798**	<u>0.810</u>	**0.740**	**0.442**	**0.322**	1

Table 3: Performance of MAMO on top-ranked molecules under GSK3β+JNK3.

pair (GSK3β + JNK3), MAMO achieves the highest average rank (1.0), with top-10 and top-50 scores of 0.846 and 0.810 for GSK3β, and 0.798 and 0.740 for JNK3, respectively. While REINVENT Trans performs strongly on GSK3β (Top-50: 0.865), its performance on JNK3 is significantly lower (Avg: 0.217), suggesting limitations in handling dual bioactivity objectives. MAMO's superior performance on both properties highlights its ability to coordinate expert agents for biological optimization tasks, particularly when the targets are interdependent. In the non-biological task (QED + SA), multiple methods achieve competitive results. MAMO attains the best average QED score (0.806) and the second-best SA score (0.813), ranking second overall. REINVENT Trans slightly outperforms MAMO in average SA (0.854), but lags in QED. These results indicate that while traditional models may excel in optimizing individual non-biological objectives, MAMO maintains strong and balanced performance across both, demonstrating its capability for general-purpose multi-objective optimization. Taken together, these findings confirm that MAMO effectively balances trade-offs across diverse molecular properties, particularly in challenging biological settings where traditional or single-agent methods struggle to generalize.

Pareto Front Analysis in Objective Space. To visually assess the optimization performance, we plot the non-dominated solution sets obtained by different methods on three representative tasks: GSK3β + JNK3, QED + SA, and GSK3β + QED. As shown in Figure 4, MAMO consistently outperforms baselines in both convergence and diversity across the objective space. Compared to other methods, MAMO's solutions exhibit broader coverage and a more uniform distribution along the Pareto front, indicating stronger trade-off management between conflicting objectives. In particular, the improvement is most pronounced on biologically relevant tasks such as GSK3β + JNK3 and GSK3β + QED, where the optimization landscape is more challenging due to intricate interdependencies. In contrast, baseline methods often concentrate in limited regions or suffer from sparsity, failing to approximate the full Pareto front. These results further demonstrate the effectiveness of MAMO's multi-agent coordination mechanism, which adaptively adjusts optimization strategies in response to dynamic conflicts across objectives. Overall, the experimental results demonstrate the strong advantages of MAMO in tackling high-dimensional and highly conflicting optimization tasks. In particular, its dynamic multi-agent collaboration mechanism proves crucial in scenarios where target objectives exhibit strong interdependencies.

Ablation Study. To investigate the importance of the central scheduling agent in MAMO, we designed ablation experiments by replacing its dynamic coordination with predefined sequential optimization chains and by removing inter-agent collaboration. Specifically, we evaluated MAMO on two tasks of increasing complexity: the three-objective task (GSK3β+JNK3+QED) and the four-objective task (GSK3β+JNK3+QED+SA). In the sequential setting, agents optimize objectives one at a time in a fixed order without global context or real-time feedback.

Methods	Top-10 Avg QED	Top-10 Avg SA	Top-50 Avg QED	Top-50 Avg SA	Avg QED	Avg SA	Avg Rank
REINVENT	0.944	0.936	0.930	0.920	0.725	0.790	5
Graph GA	0.943	**0.996**	0.930	**0.968**	0.685	0.791	8
SMILES GA	0.939	0.908	0.926	0.877	0.761	0.735	6
GP BO	0.939	0.973	0.926	0.954	0.702	0.816	4
REINVENT Trans	0.944	0.979	**0.937**	0.962	0.790	**0.854**	1
DrugAssist	0.941	0.969	0.927	0.931	0.731	0.789	3
EAG	0.931	0.932	0.925	0.939	0.712	0.767	7
GPT-4o mini	0.877	0.916	0.833	0.894	0.576	0.793	9
MAMO	**0.945**	0.985	0.933	**0.968**	**0.798**	0.813	2

Table 4: Performance of MAMO on top-ranked molecules under QED+SA objectives.

Optimization Chain	GSK3β+JNK3+QED	GSK3β+JNK3+QED+SA
GSK3β-JNK3-QED	0.760	-
QED-JNK3-GSK3β	0.572	-
GSK3β-JNK3-QED-SA	-	0.501
QED-JNK3-GSK3β-SA	-	0.510
MAMO_NoCo	0.580	0.530
MAMO	**0.762**	**0.671**

Table 5: HV of MAMO under different optimization settings. '–' indicates configurations that are not applicable.

In the no-collaboration setting, each agent optimizes the SMILES independently without sharing optimization states across agents. As shown in Table 5, the complete MAMO framework consistently outperforms all sequential variants and the no-collaboration setting (MAMO_NOCO), achieving the highest HV scores on both tasks. The performance gap is particularly pronounced in the four-objective setting, where the absence of dynamic scheduling or agent coordination results in substantial declines. For example, the HV drops from 0.671 (MAMO) to 0.501/0.510 in sequential baselines and 0.530 in MAMO_NOCO, revealing the inefficiency of fixed update orders and isolated optimization in high-conflict scenarios. These results demonstrate that both sequential chains and independent optimization without collaboration cannot effectively coordinate multiple agents, often leading to conflicting or redundant updates. The lack of interaction across objectives and intermediate state sharing limits global optimization, especially when objectives are interdependent. In contrast, the MAMO central scheduling agent enables agents to exchange intermediate outputs and dynamically adjust their optimization focus based on task needs. This flexible orchestration promotes cooperative behavior and improves convergence toward the Pareto front.

Interpretable Optimization Path

We illustrate the iterative optimization process of the MAMO framework on a pyridine-based molecule (initial SMILES: Cc1ccn(Nc2cc(C)nc(-c3ccccc3)n2)c1). As shown in Figure 5, MAMO performs multi-round optimization to progressively improve molecular attributes. In the first iteration, the left-side pyridine ring is modified by replacing the methyl substituent with a hydrogen atom, and the central pyrazole loses its methyl group. Simultaneously, the right phenyl ring is substituted with a pyridine ring. These changes jointly lead to a strong boost in biological activity scores (GSK3β and JNK3) and drug-likeness (QED). In the second iteration, the right pyridine is replaced with a phenyl group, further enhancing JNK3 and synthetic accessibility (SA), with only a minor QED decrease. In the final iteration, the right ring is modified again, and the remaining methyl group on the central scaffold is removed. This results in a balanced gain across all properties, notably in GSK3β, JNK3, and SA. This trajectory showcases how MAMO adaptively coordinates expert agents to perform structure edits, achieving a consistent upward trend in overall molecular quality through goal-aware collaboration.

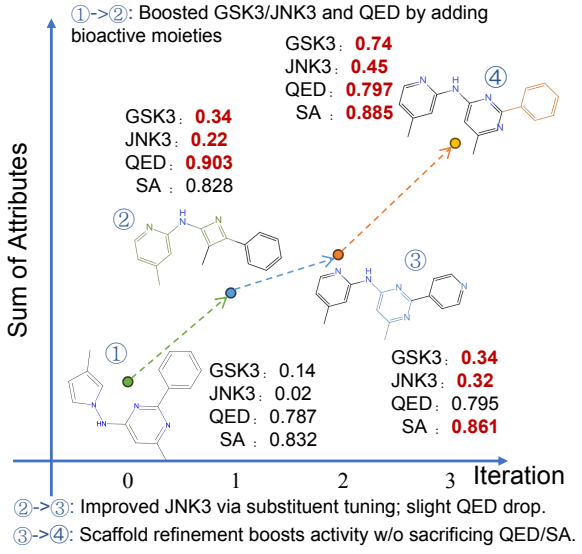

Figure 5: Molecular Optimization Trajectory with MAMO.

Conclusion

We propose MAMO, a novel multi-agent framework for multi-objective molecular optimization that resolves inter-objective conflicts through a centralized scheduling agent. By dynamically coordinating agent collaboration, MAMO achieves efficient optimization even in high-dimensional, conflicting objective spaces. Extensive experiments validate its effectiveness and robustness. Future work will focus on improving its scalability and extending its applicability to broader molecular design challenges.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62276095; 625B2067 to T.M.; U22A2037 to X.Z.; 62425204 to X.Z.; 62122025 to X.Z.; 62450002 to X.Z.; 62432011 to X.Z.; and 72204261). This work was supported by the Beijing Natural Science Foundation (L248013 to X.Z.)

References

Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Brahmavar, S. B.; Srinivasan, A.; Dash, T.; Krishnan, S. R.; Vig, L.; Roy, A.; and Aduri, R. 2024. Generating novel leads for drug discovery using LLMs with logical feedback. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 21–29.

Brown, N.; Fiscato, M.; Segler, M. H.; and Vaucher, A. C. 2019. GuacaMol: benchmarking models for de novo molecular design. *Journal of Chemical Information and Modeling*, 59(3): 1096–1108.

Chen, G.; Dong, S.; Shu, Y.; Zhang, G.; Sesay, J.; Karlsson, B. F.; Fu, J.; and Shi, Y. 2023. Autoagents: A framework for automatic agent generation. *arXiv preprint arXiv:2309.17288*.

De Rycker, M.; Baragaña, B.; Duce, S. L.; and Gilbert, I. H. 2018. Challenges and recent progress in drug discovery for tropical diseases. *Nature*, 559(7715): 498–506.

Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Fang, Y.; Liang, X.; Zhang, N.; Liu, K.; Huang, R.; Chen, Z.; Fan, X.; and Chen, H. 2023. Mol-instructions: A large-scale biomolecular instruction dataset for large language models. *arXiv preprint arXiv:2306.08018*.

Fromer, J. C.; and Coley, C. W. 2023. Computer-aided multi-objective optimization in small molecule discovery. *Patterns*, 4(2).

Fu, T.; Gao, W.; Coley, C.; and Sun, J. 2022. Reinforced genetic algorithm for structure-based drug design. *Advances in Neural Information Processing Systems*, 35: 12325–12338.

Gao, W.; Fu, T.; Sun, J.; and Coley, C. 2022. Sample efficiency matters: a benchmark for practical molecular optimization. *Advances in NeuralInformation Processing Systems*, 35: 21342–21357.

Gómez-Bombarelli, R.; Wei, J. N.; Duvenaud, D.; Hernández-Lobato, J. M.; Sánchez-Lengeling, B.; Sheberla, D.; Aguilera-Iparraguirre, J.; Hirzel, T. D.; Adams, R. P.; and Aspuru-Guzik, A. 2018. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2): 268–276.

Graff, D. E.; Shakhnovich, E. I.; and Coley, C. W. 2021. Accelerating high-throughput virtual screening through molecular pool-based active learning. *Chemical Science*, 12(22): 7866–7881.

Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.

Gu, W.; Han, J.; Wang, H.; Li, X.; and Cheng, B. 2025. Explain-analyze-generate: A sequential multi-agent collaboration method for complex reasoning. In *Proceedings of the International Conference on Computational Linguistics*, 7127–7140.

Guimaraes, G. L.; Sanchez-Lengeling, B.; Outeiral, C.; Farias, P. L. C.; and Aspuru-Guzik, A. 2017. Objective-reinforced generative adversarial networks (organ) for sequence generation models. *arXiv preprint arXiv:1705.10843*.

Han, T.; Adams, L. C.; Papaioannou, J.-M.; Grundmann, P.; Oberhauser, T.; Löser, A.; Truhn, D.; and Bressem, K. K. 2023. MedAlpaca–an open-source collection of medical conversational AI models and training data. *arXiv preprint arXiv:2304.08247*.

He, J.; You, H.; Sandström, E.; Nittinger, E.; Bjerrum, E. J.; Tyrchan, C.; Czechtizky, W.; and Engkvist, O. 2021. Molecular optimization by capturing chemist's intuition using deep neural networks. *Journal of Cheminformatics*, 13: 1–17.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.

Hoffman, S. C.; Chenthamarakshan, V.; Wadhawan, K.; Chen, P.-Y.; and Das, P. 2022. Optimizing molecules using efficient queries from property evaluations. *Nature Machine Intelligence*, 4(1): 21–31.

Hong, S.; Zheng, X.; Chen, J.; Cheng, Y.; Wang, J.; Zhang, C.; Wang, Z.; Yau, S. K. S.; Lin, Z.; Zhou, L.; et al. 2023. Metagpt: Meta programming for multi-agent collaborative framework. *arXiv preprint arXiv:2308.00352*.

Hsu, H.-H.; Hsu, Y.-C.; Chang, L.-J.; and Yang, J.-M. 2017. An integrated approach with new strategies for QSAR models and lead optimization. *BMC Genomics*, 18: 1–9.

Irwin, J. J.; Tang, K. G.; Young, J.; Dandarchuluun, C.; Wong, B. R.; Khurelbaatar, M.; Moroz, Y. S.; Mayfield, J.; and Sayle, R. A. 2020. ZINC20—a free ultralarge-scale chemical database for ligand discovery. *Journal of Chemical Information and Modeling*, 60(12): 6065–6073.

Jensen, J. H. 2019. A graph-based genetic algorithm and generative model/Monte Carlo tree search for the exploration of chemical space. *Chemical Ccience*, 10(12): 3567–3572.

Ji, C.; Zheng, Y.; Wang, R.; Cai, Y.; and Wu, H. 2021. Graph polish: A novel graph generation paradigm for molecular optimization. *IEEE Transactions on Neural Networks and Learning Systems*, 34(5): 2323–2337.

Jin, W.; Barzilay, R.; and Jaakkola, T. 2018. Junction tree variational autoencoder for molecular graph generation. In *International Conference on Machine Learning*, 2323–2332. PMLR.

Jin, W.; Barzilay, R.; and Jaakkola, T. 2020. Multi-objective molecule generation using interpretable substructures. In *International Conference on Machine Learning*, 4849–4859. PMLR.

Kadurin, A.; Nikolenko, S.; Khrabrov, K.; Aliper, A.; and Zhavoronkov, A. 2017. druGAN: an advanced generative adversarial autoencoder model for de novo generation of new molecules with desired molecular properties in silico. *Molecular Pharmaceutics*, 14(9): 3098–3104.

Kapustina, O.; Burmakina, P.; Gubina, N.; Serov, N.; and Vinogradov, V. 2024. User-friendly and industry-integrated AI for medicinal chemists and pharmaceuticals. *Artificial Intelligence Chemistry*, 2(2): 100072.

Li, J.; Liu, W.; Ding, Z.; Fan, W.; Li, Y.; and Li, Q. 2025. Large language models are in-context molecule learners. *IEEE Transactions on Knowledge and Data Engineering*.

Liu, Q.; Ruan, J.; Li, H.; Zhao, H.; Wang, D.; Chen, J.; Guanglu, W.; Cai, X.; Zheng, Z.; and Xu, T. 2025. AMoPO: Adaptive Multi-objective Preference Optimization without Reward Models and Reference Models. In *ACL 2025*.

Luo, F.; Zhang, J.; Wang, Q.; and Yang, C. 2025. Leveraging Prompt Engineering in Large Language Models for Accelerating Chemical Research. *ACS Central Science*.

Luo, R.; Sun, L.; Xia, Y.; Qin, T.; Zhang, S.; Poon, H.; and Liu, T.-Y. 2022. BioGPT: generative pre-trained transformer for biomedical text generation and mining. *Briefings in Bioinformatics*, 23(6): bbac409.

Mathur, Y.; Choudhury, A.; Prabha, S.; Saeed, M. U.; Sulaimani, M. N.; Mohammad, T.; and Hassan, M. I. 2025. Current advancement in AI-integrated drug discovery: Methods and applications. *Biotechnology Advances*, 108642.

Nguyen, T.; and Grover, A. 2024. Lico: Large language models for in-context molecular optimization. *arXiv preprint arXiv:2406.18851*.

Olivecrona, M.; Blaschke, T.; Engkvist, O.; and Chen, H. 2017. Molecular de-novo design through deep reinforcement learning. *Journal of Cheminformatics*, 9: 1–14.

OpenAI, J. A.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2024. Gpt-4 technical report, 2024. *URL https://arxiv. org/abs/2303.08774*, 2: 6.

Pitt, W. R.; Bentley, J.; Boldron, C.; Colliandre, L.; Esposito, C.; Frush, E. H.; Kopec, J.; Labouille, S.; Meneyrol, J.; Pardoe, D. A.; et al. 2025. Real-world applications and experiences of AI/ML deployment for drug discovery.

Pre-training, T.-b. M.; and Fine-tuning, M. A. 2024. REINVENT-Transformer: Molecular De Novo Design through Transformer-based Reinforcement Learning.

Tripp, A.; Simm, G. N.; and Hernández-Lobato, J. M. 2021. A fresh look at de novo molecular design benchmarks. In *NeurIPS 2021 AI for Science Workshop*.

Verhellen, J. 2022. Graph-based molecular Pareto optimisation. *Chemical Science*, 13(25): 7526–7535.

Wang, H.; Skreta, M.; Ser, C.-T.; Gao, W.; Kong, L.; Strieth-Kalthoff, F.; Duan, C.; Zhuang, Y.; Yu, Y.; Zhu, Y.; et al. 2024. Efficient evolutionary search over chemical space with large language models. *arXiv preprint arXiv:2406.16976*.

Wang, R.; Yang, M.; and Shen, Y. 2025. Bridging Molecular Graphs and Large Language Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 21234–21242.

Weininger, D. 1988. SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1): 31–36.

Wu, C.; Lin, W.; Zhang, X.; Zhang, Y.; Xie, W.; and Wang, Y. 2024. PMC-LLaMA: toward building open-source language models for medicine. *Journal of the American Medical Informatics Association*, ocae045.

Wu, Q.; Bansal, G.; Zhang, J.; Wu, Y.; Zhang, S.; Zhu, E.; Li, B.; Jiang, L.; Zhang, X.; and Wang, C. 2023. Autogen: Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*.

Xia, X.; Liu, Y.; Zheng, C.; Zhang, X.; Wu, Q.; Gao, X.; Zeng, X.; and Su, Y. 2024. Evolutionary Multiobjective Molecule Optimization in an Implicit Chemical Space. *Journal of Chemical Information and Modeling*, 64(13): 5161–5174.

Xie, Y.; Shi, C.; Zhou, H.; Yang, Y.; Zhang, W.; Yu, Y.; and Li, L. 2021. Mars: Markov molecular sampling for multi-objective drug discovery. *arXiv preprint arXiv:2103.10432*.

Xu, M.; Yu, L.; Song, Y.; Shi, C.; Ermon, S.; and Tang, J. 2022. Geodiff: A geometric diffusion model for molecular conformation generation. *arXiv preprint arXiv:2203.02923*.

Yang, H.; Yue, S.; and He, Y. 2023. Auto-gpt for online decision making: Benchmarks and additional opinions. *arXiv preprint arXiv:2306.02224*.

Yang, X.; Zhang, J.; Yoshizoe, K.; Terayama, K.; and Tsuda, K. 2017. ChemTS: an efficient python library for de novo molecular generation. *Science and Technology of Advanced Materials*, 18(1): 972–976.

Ye, G.; Cai, X.; Lai, H.; Wang, X.; Huang, J.; Wang, L.; Liu, W.; and Zeng, X. 2023. Drugassist: A large language model for molecule optimization. *arXiv preprint arXiv:2401.10334*.

Ye, G.; Cai, X.; Lai, H.; Wang, X.; Huang, J.; Wang, L.; Liu, W.; and Zeng, X. 2025. Drugassist: A large language model for molecule optimization. *Briefings in Bioinformatics*, 26(1): bbae693.

Yu, J.; Zheng, Y.; Koh, H. Y.; Pan, S.; Wang, T.; and Wang, H. 2025. Collaborative expert llms guided multi-objective molecular optimization. *arXiv preprint arXiv:2503.03503*.

Yunxiang, L.; Zihan, L.; Kai, Z.; Ruilong, D.; and You, Z. 2023. Chatdoctor: A medical chat model fine-tuned on llama model using medical domain knowledge. *arXiv preprint arXiv:2303.14070*, 2(5): 6.

Zitzler, E.; Thiele, L.; Laumanns, M.; Fonseca, C. M.; and Da Fonseca, V. G. 2003. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2): 117–132.