

# Multi-Objective Coverage via Constraint Active Search

Zakaria Shams Siam  
University at Albany, State University  
of New York  
Albany, United States  
zsiam@albany.edu

Xuefeng Liu  
University of Chicago  
Chicago, United States  
xuefeng@uchicago.edu

Chong Liu  
University at Albany, State University  
of New York  
Albany, United States  
cliu24@albany.edu

## ABSTRACT

In this paper, we formulate the new multi-objective coverage (MOC) problem where our goal is to identify a small set of representative samples whose predicted outcomes broadly cover the feasible multi-objective space. This problem is of great importance in many critical real-world applications, e.g., drug discovery and materials design, as this representative set can be evaluated much faster than the whole feasible set, thus significantly accelerating the scientific discovery process. Existing works cannot be directly applied as they either focus on sample space coverage or multi-objective optimization that targets the Pareto front. However, chemically diverse samples often yield identical objective profiles, and safety constraints are usually defined on the objectives. To solve this MOC problem, we propose a novel search algorithm, MOC-CAS, which employs an upper confidence bound-based acquisition function to select optimistic samples guided by Gaussian process posterior predictions. For enabling efficient optimization, we develop a smoothed relaxation of the hard feasibility test and derive an approximate optimizer. Compared to the competitive baselines, we show that our MOC-CAS empirically achieves superior performances across large-scale protein-target datasets for SARS-CoV-2 and cancer, each assessed on five objectives derived from SMILES-based features.

## KEYWORDS


Multi-objective coverage; Constraint active search; Upper confidence bound; Gaussian processes; Drug discovery

### ACM Reference Format:

Zakaria Shams Siam, Xuefeng Liu, and Chong Liu. 2026. Multi-Objective Coverage via Constraint Active Search. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/XTVI9400>

## 1 INTRODUCTION

Many critical real-world applications, such as drug discovery and materials design, increasingly rely on data-driven screening to identify a small set of promising designs for expensive downstream analysis. Often, the acceptance of candidate designs is governed by a feasible region defined by per-objective lower bounds or acceptance thresholds (e.g., for docking, solubility, or synthetic accessibility). Under limited experimental budgets, procedures are therefore sought that yield a compact yet diverse set of designs satisfying



**Figure 1: An example of the MOC problem where  $f_1, f_2$  are two objective functions and  $\tau_1, \tau_2$  are two thresholds defining the feasible region shown in light green. 15 samples are selected as representative feasible samples, each with a coverage ball whose radius is  $r$ .**

these acceptability criteria while requiring as few costly experiments as possible [15]. This motivates active and sample-efficient search algorithms that reason directly about the specific outcomes which will be most informative and cost-effective for the decision makers [6, 24].

Classical multi-objective optimization methods typically emphasize approximating the Pareto frontier and improving hypervolume [5, 12, 16, 30]. These objectives are well suited for managing trade-offs among competing objectives, but they do not directly target the discovery of threshold-compliant candidates. Contemporary works on *constraint active search* (CAS) contribute toward discovery under acceptance/rejection criteria, however, the state-of-the-art [19] frames coverage in the sample space and applies expected improvement (EI)-like utilities. However, sample-space coverage can be misaligned with many decision criteria (e.g., in molecular design), where chemically diverse inputs may map to similar objective profiles. In such settings, input-space novelty can waste evaluations on designs that look different yet yield nearly identical outcomes, while failing to cover the feasible region where decisions are made: the objective space.

In contrast, coverage in the output (objective) space aligns directly with downstream selection, as acceptability is inherently defined there. Optimizing output-space coverage therefore prioritizes the discovery of a few representative outcome patterns that satisfy all imposed constraints, enabling faster and more reliable down-selection than strategies that only spread samples in the input space or trace the Pareto front. Our goal is to select a compact set of samples whose outcomes are well distributed across the feasible region, rather than merely spreading inputs in sample space or tracing the Pareto frontier. This supports reliable down-selection



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems ([www.ifaamas.org](http://www.ifaamas.org)). <https://doi.org/10.65109/XTVI9400>

and risk-aware decision making under uncertainty. See Figure 1 for an example with two objective functions.

To solve this new and highly motivated problem, we formulate the new *multi-objective coverage* (MOC) problem, where per-objective thresholds define a feasible region in objective space. The goal is to identify a small set of diverse, representative samples whose outcomes satisfy the thresholds and collectively cover the feasible region under uncertainty—without attempting to enumerate all feasible points [8], which can be prohibitively expensive. This new problem yields a compact and diverse panel that supports decision-making under uncertainty while minimizing the number of costly evaluations. To solve the MOC problem, we model each objective with an independent Gaussian process (GP) [29] and design an optimistic acquisition function that trades off feasibility and incremental objective-space coverage gain, using upper confidence bounds (UCB) [25].

**Contributions.** Our contributions are summarized as follows.

- We formulate the new MOC problem, which explicitly targets fast coverage and diversity of the feasible objective region under per-objective thresholds. This is distinct from the CAS framed in the sample space [19] and also different from other frontier-centric objectives [5, 12, 30].
- To solve the problem, we propose the *multi-objective coverage via constraint active search* (MOC-CAS) algorithm, which is a novel constrained active search algorithm that evaluates each candidate based on the optimistic estimate of the new feasible volume it is expected to cover within the feasible region. It further employs a tie-breaking mechanism that promotes diversity by encouraging dispersion in the predicted objective values.
- To enable efficient optimization, we replace the hard feasibility indicators with smooth surrogates and thus derive an approximate optimizer for the proposed acquisition function.
- On large-scale protein-target datasets of 1 million compounds spanning cancer and SARS-CoV-2 [18], our MOC-CAS algorithm achieves better performance than competitive baselines under similar evaluation budgets.

**Technical Novelties.** Our technical novelties are:

- Our MOC-CAS selects samples by *maximizing newly uncovered neighborhood volume* in the objective space at a fixed coverage resolution, utilizing a UCB prediction to estimate feasibility before any coverage value is assigned. To maintain diversity in the objective space, we break ties by selecting the candidate with the greatest objective-space distance to prior feasible outcomes.
- We derive a mass-preserving relaxation of the geometric objective with hard indicators. We replace the ball indicator by a unit-mass Gaussian kernel. Additionally, we substitute a smooth per-objective probit gate for the orthant indicator, and further replace the union of covered balls by a soft kernel sum. The resulting local average therefore yields a closed-form novelty term and a fully differentiable acquisition function along with explicit gradients via the chain rule, thus enabling fast optimization in practice.

We released our Appendix on arXiv<sup>1</sup> and our code on GitHub<sup>2</sup>.

## 2 RELATED WORK

We briefly discuss related work from three different perspectives.

### 2.1 Constraint Active Search

Constraint active search (CAS) reexamines multiobjective design from Pareto-frontier tracing for finding a diverse set of threshold-satisfying samples. Past work in [19] presented CAS with the expected coverage improvement (ECI) acquisition function that scores samples by their expected increase in the covered feasible volume in the sample space, where where coverage is measured in the sample space, using fill distance [19] as a proxy for diversity. CAS has also been utilized for optimizing performance during machine learning model deployment by connecting offline and online experimentations [13]. A related but distinct line of work [21] learns active constraint sets in order to speed up repeated deterministic optimization by estimating which constraints bind at optimality where the objective is to find faster solution for parametric programs with theoretical guarantees, instead of applying probabilistic search for finding diverse satisfactory samples. In contrast, our proposed approach explicitly defines and directly optimizes coverage in the objective space under multi-threshold feasibility. It introduces an optimism-based feasibility gate linked to incremental objective-space coverage, thereby reshaping the acquisition geometry and the mechanism by which diversity is enforced, while preserving the core intuition of CAS methods - that a few well-chosen samples can effectively cover the feasible set.

### 2.2 Level Set Estimation

The level set estimation (LSE) problem was first approached in [2] through the straddle heuristic algorithm by utilizing a GP-based active learning technique. Subsequent work in [8] presented the LSE algorithm for active level-set estimation along with Gaussian-process confidence bounds, derived sample-complexity guarantees, and further extended the method to implicit thresholds and also batched sampling. Also, TRUVAR, a unified Bayesian framework for Bayesian optimization (BO) and LSE was proposed in [1] which greedily diminishes the truncated posterior variance over sample sets, controls pointwise costs and heteroscedastic noise, and also comes with theoretical guarantees. Zanette et al. [31] maximized the expected raise in the volume of the estimated level set. Wang et al. [28] developed a local-minimax framework for noisy zeroth-order optimization and proposed that instance-dependent rates could be substantially faster when the reference function’s level sets satisfy fast volume-growth conditions and regularity, explicitly showing that the problem’s difficulty is directly related to and also determined by the properties of the level-set geometry.

Iwazaki et al. [11] focused on the LSE problem with input uncertainty where the input parameters are vulnerable to perturbation from a known Normal distribution. This study is closely related to safe BO research [26, 27], where they sample only the configurations that surpass the threshold with high probability. Subsequent study in [20] proposed nearly-optimal (instance-dependent and

<sup>1</sup><https://arxiv.org/abs/2602.15595>

<sup>2</sup><https://github.com/ZakSiam/MOC-CAS>

non-asymptotic) sample-complexity bounds for explicit and implicit level-set estimation by connecting the problem to adaptive experimental design in reproducing kernel Hilbert space (RKHS), with methods which match known lower bounds in linear (kernel) settings.

### 2.3 Multi-Objective Optimization

Multi-objective optimization (MOO) aims to effectively approximate the Pareto front under several competing objectives. Early probabilistic techniques utilized estimation-of-distribution ideas to design Bayesian model-based evolutionary algorithms which maintain diverse approximations to the Pareto front [14]. The work in [14] combines a BO algorithm into a multi-objective evolutionary algorithm and investigate its behavior on combinatorial problems, highlighting the role of probabilistic modeling for convergence and diversity. Information-theoretic acquisition functions in Bayesian optimization, such as, PESMO select evaluations in order to maximally reduce the entropy of the posterior over the Pareto set and allow a decomposition across all objectives, thus enabling decoupled evaluation with computational cost which scales linearly in the number of objectives [10].

Another line of work utilizes random scalarizations where the authors proposed a flexible framework which samples scalarization weights to the target specified regions of the Pareto front and derive sublinear regret guarantees along with favorable computational cost [22]. Complementarily, it was also shown that dominated hypervolume could be written as an expectation over hypervolume scalarizations, producing algorithms with hypervolume regret bounds when paired with standard BO (UCB or Thompson sampling) methods and a simple predictor for hypervolume [32]. Recent scalable MOO-BO approaches develop parallel expected hypervolume improvement (EHVI) and trust-region techniques for higher dimensions and large batches (e.g., qEHVI/qNEHVI and MORBO) [3, 4, 16]. Park et al. [23] proposed BOTied which is based on a cumulative distribution function indicator. Unlike these Pareto front-centric approaches (e.g., hypervolume, entropy, or scalarization driven), our proposed work seeks coverage in objective space under explicit multi-threshold feasibility by optimizing a coverage-gain objective which rewards discovering diverse feasible outcomes instead of solely tracing the Pareto front.

## 3 PRELIMINARIES

### 3.1 Problem Statement

Let  $\mathcal{X} \subseteq \mathbb{R}^d$  be a compact sample space and  $f = (f_1, \dots, f_m) : \mathcal{X} \rightarrow \mathbb{R}^m$  a black-box, unknown, vector-valued function with  $m$  objectives. A noisy query at  $x \in \mathcal{X}$  gives  $y_i = f_i(x) + \varepsilon_i$ , where  $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$  are i.i.d. Gaussian noise terms, for  $i = 1, \dots, m$ . Given per-objective thresholds  $\tau = (\tau_1, \dots, \tau_m) \in \mathbb{R}^m$ , an outcome is called *feasible* provided that all objectives meet their thresholds. This defines the feasible region in the objective space:

$$\mathcal{S} = \{z \in \mathbb{R}^m : z_i \geq \tau_i \ \forall i\}.$$

In many critical real-world applications, e.g., drug discovery and materials design, the decision criterion is defined on the objective

values, not molecular inputs. Therefore, covering representative outcomes inside  $\mathcal{S}$  is more consistent and aligned with down-selection than finding the Pareto frontier.

We select designs  $x_1, x_2, \dots \in \mathcal{X}$  sequentially and observe  $y_t = f(x_t) + \varepsilon_t$ . At the end of  $t$  steps,  $\mathcal{D}_t = \{(x_s, y_s)\}_{s=1}^t$  and the observed outcomes are  $Z_t = \{y_s\}_{s=1}^t \subset \mathbb{R}^m$ . We fix a *resolution* parameter  $r > 0$ . In real-world problem settings, we can define  $r$  with respect to the robustness of the design to perturbations. Many experimental design problems typically have a sense of known resolution (e.g., manufacturing precision/tolerance or simulation accuracy). We interpret  $r$  as a coverage resolution: outcomes within distance  $r$  are treated as effectively redundant for representing  $\mathcal{S}$ . For  $r > 0$  and center  $z \in \mathbb{R}^m$ ,

$$B_r(z) = \{z' \in \mathbb{R}^m : \|z' - z\|_2 < r\}$$

represents the coverage neighborhood of  $z$  which is an  $r$ -ball in the output space. We thus define the covered volume for a set of points  $Z_t$  as

$$u_t = \text{Vol}(B_r(Z_t) \cap \mathcal{S}), \quad B_r(Z_t) = \bigcup_{z \in Z_t} B_r(z).$$

Our objective is to *maximize*  $u_t$  with as few evaluations as possible, thus generating a small and diverse set of threshold-compliant or feasible outcomes that broadly covers  $\mathcal{S}$  at resolution  $r$ .

**Evaluations.** No single metric fully captures performance in multi-objective coverage. We therefore report three complementary metrics.

**FILL DISTANCE.** Fill distance quantifies how evenly the sampled objective vectors cover the feasible region in objective space. The fill distance is formally defined as follows where  $d(\cdot, \cdot)$  denotes the Euclidean distance in the objective space.

$$\text{FILL}(Z_t, \mathcal{S}) = \sup_{z \in \mathcal{S}} \min_{z_j \in Z_t} d(z_j, z).$$

$\text{FILL}(Z_t, \mathcal{S})$  represents the radius of the largest empty ball that one can fit in  $\mathcal{S}$  in Euclidean space [19]. It measures the spacing of  $Z_t$  in  $\mathcal{S}$ . If a set's fill distance is smaller with respect to  $\mathcal{S}$ , then that set is better distributed within  $\mathcal{S}$ .

**POSITIVE SAMPLES.** This represents the number of feasible points in the objective space (inside the feasible set,  $\mathcal{S}$ ) for the samples selected by the algorithm.

**AREA UNDER THE POSITIVES CURVE (AUP).** We further summarize sample efficiency utilizing the area under the cumulative discoveries curve. For cumulative positives  $P(t)$  up to round  $t$ , we calculate

$$\text{AUP} = \sum_{t=1}^T P(t),$$

or its trapezoidal variant. AUP rewards those methods which find good samples both *early* and *consistently* across the budget. Higher AUP indicates better performance by the method (since earlier and steadier discoveries score higher in this case). This metric is analogous to the area under the learning curve metric which is widely utilized in active learning research [9].

**Additional Notations.** We denote the scalars as italics (e.g.,  $t, r$ ), vectors as bold (e.g.,  $z$ ), and sets as calligraphic (e.g.,  $\mathcal{X}, \mathcal{S}$ ). The Euclidean norm is denoted as  $\|\cdot\|$ .  $\text{Vol}(\cdot)$  represents the Lebesgue volume in  $\mathbb{R}^m$ . We use  $\mathbf{1}\{\cdot\}$  to represent an indicator where  $\mathbf{1}\{x\} = 1$

if the condition  $x$  holds and 0 otherwise. Similarly,  $\mathbf{1}(\cdot)$  represents a set-membership indicator where  $\mathbf{1}_A(z) = 1$  if  $z \in A$  and 0 if  $z \notin A$ . We use  $\mathcal{GP}$  for Gaussian processes.

### 3.2 Background

**Constraint Active Search.** Say  $\mathcal{X}$  is a discrete or finite sample set having unknown, black-box multi-objective response  $\tilde{y} : \mathcal{X} \rightarrow \mathbb{R}^m$ . For the given collections of per-objective minimum thresholds,  $\tau \in \mathbb{R}^m$ , constraint active search (CAS) targets the feasible subset of the samples

$$S_{\mathcal{X}} = \{x \in \mathcal{X} : \tilde{y}_i(x) \geq \tau_i \ \forall i \in [m]\},$$

and searches adaptively to identify a diverse set of feasible samples.

**Level Set Estimation.** The objective of level set estimation (LSE) is to determine the region  $\{x \in \mathcal{X} : f(x) \geq \tau\}$  by sequentially querying  $x$  and then observing  $y = f(x) + \epsilon$  given a (noisy) black-box objective  $f : \mathcal{X} \rightarrow \mathbb{R}$  and a threshold  $\tau \in \mathbb{R}$ . With a GP posterior at round  $t - 1$ , let  $\mu_{t-1}(x)$  and  $\sigma_{t-1}(x)$  represent the mean and standard deviation, a typical substitute for “being above the level” is

$$p_{t-1}(x) = \Pr(f(x) \geq \tau \mid \mathcal{D}_{t-1}) = \Phi\left(\frac{\mu_{t-1}(x) - \tau}{\sigma_{t-1}(x)}\right),$$

and the classic STRADDLE score selects points having high uncertainty near the level,

$$s_{\text{straddle}}(x) = \kappa \sigma_{t-1}(x) - |\mu_{t-1}(x) - \tau|,$$

mostly with  $\kappa \approx 1.96$ . LSE methods are inclined to concentrate queries close to the decision boundary that separates  $\{f(x) \geq \tau\}$  from its complement, instead of sampling the interior of the superlevel set.

**Multi-Objective Optimization.** Let  $f(x) = (f_1(x), \dots, f_m(x)) : \mathcal{X} \rightarrow \mathbb{R}^m$  be a function that needs to be maximized over a sample space  $\mathcal{X}$ . For  $x_1, x_2 \in \mathcal{X}$ , we say  $x_1$  Pareto-dominates  $x_2$  if  $f_i(x_1) \geq f_i(x_2)$  for all  $i$  and  $f_j(x_1) > f_j(x_2)$  for some  $j$ . The *Pareto set*  $X^* \subseteq \mathcal{X}$  accommodates all non-dominated samples and its image  $P = \{f(x) : x \in X^*\}$  is the *Pareto front*. A typical scalar summary is the (reference-point) hypervolume indicator

$$\begin{aligned} & \text{HV}_z(Y) \\ &= \text{vol}\left(\{y \in \mathbb{R}^m : y \geq z, \text{ and } \exists y' \in Y \text{ with } y' \text{ dominates } y\}\right), \end{aligned}$$

for a finite set  $Y \subset \mathbb{R}^m$  of achieved outcomes and a reference point  $z \in \mathbb{R}^m$ . Hypervolume enlarges as  $Y$  inflates the dominated portion of objective space beyond  $z$ . It is widely utilized to compare methods which trade off multiple objectives.

## 4 THE MOC-CAS ALGORITHM

In this section, we show full details of our MOC-CAS algorithm, summarized in Algorithm 1. Overall, MOC-CAS performs sequential experimental design in order to identify a diverse set of outcomes which satisfy all per-objective thresholds. At each round, we fit independent GPs for all the  $m$  objectives, then form optimistic predictions, score each sample by an objective space coverage gain acquisition function which rewards new feasible coverage, and finally select the maximizer, while breaking ties in order to preserve objective-space diversity.

---

### Algorithm 1 MOC-CAS Algorithm

---

**Require:** GP prior kernels  $\{k_i\}_{i=1}^m$ , thresholds  $\tau$ , radius  $r$ , confidence schedule  $(\beta_t)_{t \geq 1}$ , initial data  $\mathcal{D}_0$

- 1: **for**  $t = 1, 2, \dots, T$  **do**
- 2:     Update independent GP posteriors using  $\mathcal{D}_{t-1}$
- 3:      $x_t \leftarrow \arg \max_{x \in \mathcal{X}} \alpha_{t-1}^{\text{MOC-CAS}}(x)$
- 4:             **▷** break ties by farthest  $U_{t-1}(x)$  from  $Z_{t-1}$
- 5:     Query  $y_t = f(x_t) + \epsilon_t$
- 6:     Update  $\mathcal{D}_t = \mathcal{D}_{t-1} \cup \{(x_t, y_t)\}$
- 7: **end for**
- 8: **return**  $\{x_t\}_{t=1}^T$

---

**Gaussian Process Modeling.** We develop independent GP priors  $f_i \sim \mathcal{GP}(0, k_i)$ ,  $i = 1, \dots, m$ , to model all the objectives. Here, we refer to  $k_i(\cdot, \cdot)$  as the GP covariance function. Given  $\mathcal{D}_{t-1}$ , the posterior mean and variance at  $x$  are denoted as  $\mu_{t-1}^{(i)}(x)$  and  $\sigma_{t-1}^{(i)}(x)^2$ , shown as

$$\begin{aligned} \mu_{t-1}^{(i)}(x) &= (k_{t-1}^{(i)}(x))^\top (K_{t-1}^{(i)} + \sigma^2 I)^{-1} \mathbf{y}_{1:t-1}^{(i)}, \\ \sigma_{t-1}^{(i)}(x)^2 &= k_i(x, x) - (k_{t-1}^{(i)}(x))^\top (K_{t-1}^{(i)} + \sigma^2 I)^{-1} k_{t-1}^{(i)}(x), \end{aligned}$$

where  $X_{t-1} = \{x_1, \dots, x_{t-1}\}$  are the sample points queried so far and  $\mathbf{y}_{1:t-1}^{(i)} = [y_1^{(i)}, \dots, y_{t-1}^{(i)}]^\top$  are the observations for objective  $i$  with independent and identically distributed Gaussian noise variance  $\sigma^2$ , and  $K_{t-1}^{(i)} = [k_i(x_a, x_b)]_{a,b=1}^{t-1}$ ,  $k_{t-1}^{(i)}(x) = [k_i(x_1, x), \dots, k_i(x_{t-1}, x)]^\top$ .

Using a confidence schedule  $\beta_t > 0$ , we now define the optimistic UCB prediction

$$U_{t-1}^{(i)}(x) = \mu_{t-1}^{(i)}(x) + \sqrt{\beta_t} \sigma_{t-1}^{(i)}(x), \quad U_{t-1}(x) := (U_{t-1}^{(i)}(x))_{i=1}^m.$$

**Output-Coverage Acquisition.** We define the optimistic feasibility indicator in the objective space

$$\tilde{Z}_{t-1}(x) = \begin{cases} 1, & \text{if } U_{t-1}^{(i)}(x) \geq \tau_i \ \forall i, \\ 0, & \text{otherwise.} \end{cases}$$

Next, we define the incremental coverage gain at resolution  $r$  which is our acquisition function used in step 3 in Algorithm 1:

$$\alpha_{t-1}^{\text{MOC-CAS}}(x) := \tilde{Z}_{t-1}(x) \cdot \text{Vol}\left((B_r(U_{t-1}(x)) \cap \mathcal{S}) \setminus B_r(Z_{t-1})\right), \quad (1)$$

where  $B_r(Z_{t-1}) := \bigcup_{s=1}^{t-1} B_r(y_s)$ . Value is gained only if  $x$  is optimistically feasible, and then only by the new  $r$ -ball volume supplied within  $\mathcal{S}$ . We prefer the sample whose predicted outcome  $U_{t-1}(x)$  is farthest (in  $\ell_2$ ) from the existing set  $Z_{t-1}$  in order to diversify among ties.

**Smoothing for efficient optimization.** We substitute smooth surrogates (e.g., soft margins with respect to differentiable  $r$ -ball kernels and thresholds  $\tau$ ) for the hard feasibility test and set-difference to enable fast and gradient-based maximization of the acquisition function. This provides a smooth surrogate approximation of the discrete gain and could be optimized much more efficiently with the help of known gradient-based optimization methods like L-BFGS. The formal construction of smoothing is presented in Section 5.

## 5 EFFICIENT OPTIMIZATION FOR PRACTICE

In this section, we analyze and efficiently optimize Algorithm 1. Here, we delineate both the exact (hard) acquisition with an optimistic feasibility guard and a fully smooth surrogate function utilized for inner optimization. More analyses of acquisition approximation error and time complexity of MOC-CAS are deferred to the Appendix.

Beginning from the hard geometric definition (eq. (1)), Step A in Subsection 5.1 converts the set difference into a local Gaussian average. Then, Step B in Subsection 5.1 replaces the hard feasibility by a smooth multi-threshold gate evaluated at the local center. Finally, Step C in Subsection 5.1 replaces the union by a soft sum whose local average yields a closed-form expression. The resulting smooth acquisition function (eq. (7)) is cheap to evaluate and also easy to differentiate, thus yielding an efficient inner maximization routine for MOC-CAS.

### 5.1 From Hard to Soft Indicators: Deriving the MOC-CAS Acquisition

The hard new-coverage objective in output space at iteration  $t$  is given in eq. (1), that is, the  $S$ -constrained volume newly covered by the  $r$ -ball around the optimistic prediction  $U_{t-1}(x)$  while discarding the region already covered by the previous outputs  $\{y_s\}_{s=1}^{t-1}$ .

*Step A: Replace the Ball by a Mass-preserved Normalized Gaussian Kernel.* We write the set-volume in eq. (1) as an integral over three hard indicators:

$$\alpha_{t-1}^{\text{MOC-CAS}}(x) = \tilde{Z}_{t-1}(x) \int_{\mathbb{R}^m} \mathbf{1}_{B_r(U_{t-1}(x))}(z) \mathbf{1}_S(z) \cdot \left(1 - \mathbf{1}_{B_r(Z_{t-1})}(z)\right) dz. \quad (2)$$

The three factors inside the integral in (2) correspond to the new  $r$ -ball, the satisfactory set, and the not-yet-covered region, respectively.

We introduce two unit-mass radial kernels in output space. One is the uniform kernel on the open ball

$$u_r(z - U) := \frac{1}{V_m(r)} \mathbf{1}\{\|z - U\| < r\}, \quad V_m(r) := \frac{\pi^{m/2}}{\Gamma(\frac{m}{2} + 1)} r^m,$$

where  $\Gamma$  is the Gamma function, and another one is the Gaussian kernel

$$\kappa_r(z - U) := \mathcal{N}(z; U, r^2 I), \quad \int \kappa_r(z - U) dz = 1.$$

We then make the standard mass-preserving substitution

$$\mathbf{1}_{B_r(U)}(z) \approx V_m(r) \kappa_r(z - U),$$

which yields

$$\begin{aligned} \alpha_{t-1}^{\text{MOC-CAS}}(x) &\approx \tilde{Z}_{t-1}(x) V_m(r) \int_{\mathbb{R}^m} \kappa_r(z - U) \mathbf{1}_S(z) (1 - \mathbf{1}_{B_r(Z_{t-1})}(z)) dz \\ &= \tilde{Z}_{t-1}(x) V_m(r) \mathbb{E}_{Z \sim \mathcal{N}(U, r^2 I)} \left[ \mathbf{1}_S(Z) (1 - \mathbf{1}_{B_r(Z_{t-1})}(Z)) \right], \end{aligned} \quad (3)$$

where  $U := U_{t-1}(x)$ . Therefore, Step A converts the hard geometry into a local Gaussian average around  $U$ .

*Step B: Soften the Satisfactory Gate.* We replace the hard indicator of the thresholded orphant  $\mathbf{1}_S(z) = \prod_{i=1}^m \mathbf{1}\{z_i \geq \tau_i\}$  by a smooth *probit gate* with softness  $\lambda > 0$ :

$$p_{\text{sat}}(z) := \prod_{i=1}^m \Phi\left(\frac{z_i - \tau_i}{\lambda}\right), \quad (4)$$

where  $\Phi$  is the standard normal cumulative distribution function (CDF). Using the narrow and centered Gaussian kernel from eq. (3), we approximate the local average of the satisfactory gate by its value at the center:

$$\mathbb{E}_{Z \sim \mathcal{N}(U, r^2 I)} [p_{\text{sat}}(Z)] \approx p_{\text{sat}}(U). \quad (5)$$

Substituting eq. (4) and (5) into eq. (3) gives

$$\begin{aligned} \alpha_{t-1}^{\text{MOC-CAS}}(x) &\approx \tilde{Z}_{t-1}(x) V_m(r) p_{\text{sat}}(U) \mathbb{E}_{Z \sim \mathcal{N}(U, r^2 I)} [1 - \mathbf{1}_{B_r(Z_{t-1})}(Z)]. \end{aligned}$$

*Step C: Soften the Set Difference (Union) and Derive a Closed-Form Expectation.*  $\mathbf{1}_{B_r(Z_{t-1})}(z) = \mathbf{1}\{z \in \cup_{s=1}^{t-1} B_r(y_s)\}$  is the exact membership of the already-covered region. We utilize a bounded soft union by summing the normalized Gaussian kernels centered at the previous outputs, which yields the “soft-OR” of the covered neighborhoods:

$$\mathbf{1}_{B_r(Z_{t-1})}(z) \approx \sum_{s=1}^{t-1} \kappa_r(z - y_s).$$

Under the local Gaussian in eq. (3), the expectation of each term is reduced to a product-integral of two Gaussian distributions:

$$\begin{aligned} \mathbb{E}_{Z \sim \mathcal{N}(U, r^2 I)} [\kappa_r(Z - y_s)] &= \int \mathcal{N}(z; U, r^2 I) \mathcal{N}(z; y_s, r^2 I) dz \\ &= \mathcal{N}(U; y_s, 2r^2 I) \\ &= c_m(r) \exp\left(-\frac{\|U - y_s\|^2}{4r^2}\right). \end{aligned}$$

where  $c_m(r) := (4\pi r^2)^{-m/2}$ . Therefore,

$$\begin{aligned} \mathbb{E}_{Z \sim \mathcal{N}(U, r^2 I)} \left[ 1 - \mathbf{1}_{B_r(Z_{t-1})}(Z) \right] &\approx 1 - \sum_{s=1}^{t-1} c_m(r) \exp\left(-\frac{\|U - y_s\|^2}{4r^2}\right) \\ &=: n(U). \end{aligned} \quad (6)$$

### 5.2 Soft MOC-CAS Acquisition and Efficient Maximization

Once we combine eq. (3), (4), and (6), it produces the fully softened acquisition function:

$$\tilde{\alpha}_{t-1}^{\text{MOC-CAS}}(x) := V_m(r) \underbrace{p_{\text{sat}}(U_{t-1}(x))}_{\text{soft feasibility}} \underbrace{n(U_{t-1}(x))}_{\text{soft new-coverage}}. \quad (7)$$

This expression is not expensive to evaluate. This is because  $V_m(r)$  is a constant given  $r$  and  $m$ . Also,  $p_{\text{sat}}(U)$  is a product of 1-dimensional CDFs at  $U$ . Furthermore,  $n(U)$  is a short sum of exponentials in squared distances to the previous outputs. The hard gate  $\tilde{Z}_{t-1}(x)$  appears in the exact (hard) acquisition function defined in eq. (1) to enforce optimistic feasibility. For efficient and fully differentiable optimization, we drop this binary factor in eq. (7) and utilize the smooth feasibility term  $p_{\text{sat}}(U)$ . Empirically, both

variants behave identically. By default, we utilize (7) for inner optimization.

*Gradients for inner optimization.* Let  $U = U_{t-1}(x)$  and we write  $V_m(r)$  once as a constant factor.

$$\nabla_U p_{\text{sat}}(U) = p_{\text{sat}}(U) \left[ \frac{\phi\left(\frac{U_1 - \tau_1}{\lambda}\right)}{\lambda \Phi\left(\frac{U_1 - \tau_1}{\lambda}\right)}, \dots, \frac{\phi\left(\frac{U_m - \tau_m}{\lambda}\right)}{\lambda \Phi\left(\frac{U_m - \tau_m}{\lambda}\right)} \right]^\top,$$

Here, we use the standard normal CDF  $\Phi(u)$  and PDF  $\phi(u)$

$$\Phi(u) = \int_{-\infty}^u \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt, \quad \phi(u) = \frac{1}{\sqrt{2\pi}} e^{-u^2/2},$$

and in the gradient, the argument is  $u_i = \frac{U_i - \tau_i}{\lambda}$ .

$$\nabla_U n(U) = \sum_{s=1}^{t-1} c_m(r) \exp\left(-\frac{\|U - y_s\|^2}{4r^2}\right) \frac{U - y_s}{2r^2}.$$

By the chain rule, with  $U_i(x) = \mu_{t-1}^{(i)}(x) + \sqrt{\beta_t} \sigma_{t-1}^{(i)}(x)$ ,

$$\nabla_x \tilde{\alpha}_{t-1}^{\text{MOC-CAS}}(x) = V_m(r) \left( \nabla_U p_{\text{sat}}(U) n(U) + p_{\text{sat}}(U) \nabla_U n(U) \right)^\top \cdot \nabla_x U(x). \quad (8)$$

where, for each  $i$ ,

$$\nabla_x U_i(x) = \nabla_x \mu_{t-1}^{(i)}(x) + \sqrt{\beta_t} \nabla_x \sigma_{t-1}^{(i)}(x).$$

These derivatives are already available in standard GP toolkits. Eq. (7) together with eq. (8) yields a smooth objective and explicit gradients for multi-start L-BFGS (or similar methods) over  $x \in \mathcal{X}$ . The best  $x_t$  is then chosen by  $x_t = \arg \max_{x \in \mathcal{X}} \tilde{\alpha}_{t-1}^{\text{MOC-CAS}}(x)$ .

## 6 EXPERIMENTS

### 6.1 Experimental Settings

**Datasets.** We use the DrugImprover dataset [17, 18] that comprises curated, orderable subsets of the ZINC15 compound library, where 1 million in-stock or quickly shippable molecules were selected per target for structure-based drug discovery. Protein pockets were defined using crystallographic ligand data or FPocket, and a single binding site was used for docking per protein. The dataset includes docking scores (via OpenEye FRED) and trained surrogate models (with validation  $R^2$  of 0.842 for 3CLPro and 0.73 for RTCB) for 24 SARS-CoV-2 protein (e.g., 3CLPro, PLPro, RDRP, NSPs) and 5 human cancer-related targets (e.g., 6T2W, RTCB, WRN). Provided files include Simplified Molecular Input Line Entry System (SMILES)-docking score pairs, receptor structures, model weights, and inference code, supporting applications in high-throughput virtual screening and predictive modeling. From the available targets, we use three cancer proteins 6T2W, RTCB, and WRN, and one SARS-CoV-2 protease target 3CLPro.

**Baselines.** We benchmark MOC-CAS against the following four baseline methods on all the four drug-discovery tasks to delineate the comparison.

**RANDOM.** We utilize the uniform random selection method from the remaining pool each round.

**ONE-STEP.** We apply the one-step active search method that selects the sample maximizing the predicted feasibility probability

$p(Z=1 \mid \mathcal{D}_{t-1})$ , calculated as the product of per-objective Gaussian tail probabilities under independent GPs [7].

**STRADDLE.** We implement a multiobjective extension of STRADDLE [2] for level-set estimation that alternates the targeted objective  $i \in [m]$  each iteration and scores candidates by  $s_i(x) = \kappa \sigma_{t-1}^{(i)}(x) - |\mu_{t-1}^{(i)}(x) - \tau_i|$ ,  $\kappa = 1.96$ , then selects the maximizer using the shared shortlist.

**MOO+CLUSTER.** We implement the multiobjective BO with clustering that develops optimistic outcome predictions  $U_{t-1}(x)$  using UCB, keeps optimistically feasible candidates, clusters their estimated outcomes with  $k$ -means, scores each cluster by its uncovered feasible mass within radius  $r$  (via neighbor counts in objective space), then selects within the best cluster the point farthest (in objective space) from the previously observed outcomes.

All baseline methods share the same GP modeling pipeline, budget, initializations, and shortlist settings as MOC-CAS for ensuring controlled and fair comparisons.

**Evaluation.** As discussed in Section 3.1, we report the number of positives (samples with all objectives  $\geq \tau$ ), AUP, and fill distance inside the feasible outcome region, over 220 iterations. We finally aggregate the results as mean  $\pm$  adjusted standard error over four different trials for each experiment (i.e., mean  $\pm$  std/ $\sqrt{4}$ ).

### 6.2 Results on SARS-CoV-2 Dataset

We test all five methods on the SARS-CoV-2 3CLPro target dataset. MOC-CAS achieves the best results, as evidenced in Figure 2, where it consistently achieves the lowest fill distance and highest AUP and number of positives across multiple trials, outperforming all other algorithms. ONE-STEP active search method is competitive but underperforms compared to our MOC-CAS method. Overall, STRADDLE, RANDOM SEARCH, and MOO+CLUSTER did not perform well in terms of the three evaluation criteria where STRADDLE performs worst among all methods in terms of number of positives and AUP.

ONE-STEP selects points that maximize the estimated feasibility probability, therefore it tends to exploit the high-feasibility region (where feasibility probability is highest) early. This makes it competitive on number of positives, however, it can concentrate evaluations in a smaller portion of the feasible objective space, which is consistent with its slightly worse fill-distance trajectory than MOC-CAS. In contrast, MOC-CAS explicitly rewards coverage gain in the feasible objective space, so once it finds a feasible neighborhood, it is driven to expand into new feasible neighborhoods rather than repeatedly sampling near already-covered outcomes. This is exactly the behavior we want in molecular screening: beyond finding many feasible candidates, we want them to be diverse in objective space to support down-selection and follow-up synthesis. Finally, the weaker baselines align with their design intent: STRADDLE is primarily a level-set (boundary) learner, therefore it focuses on uncertain points near thresholds and consequently sacrifices feasible yield and AUP. MOO+CLUSTER relies on clustering of optimistic predictions and is less effective when posterior uncertainty is high, whereas RANDOM SEARCH provides a lower-bound reference.

### 6.3 Results on Cancer Dataset

We test all five methods on three cancer proteins 6T2W, RTCB, and WRN target datasets. Again, our MOC-CAS turns out to be the




Figure 2: Quantitative comparison on the SARS-CoV-2 3CLPro target dataset.

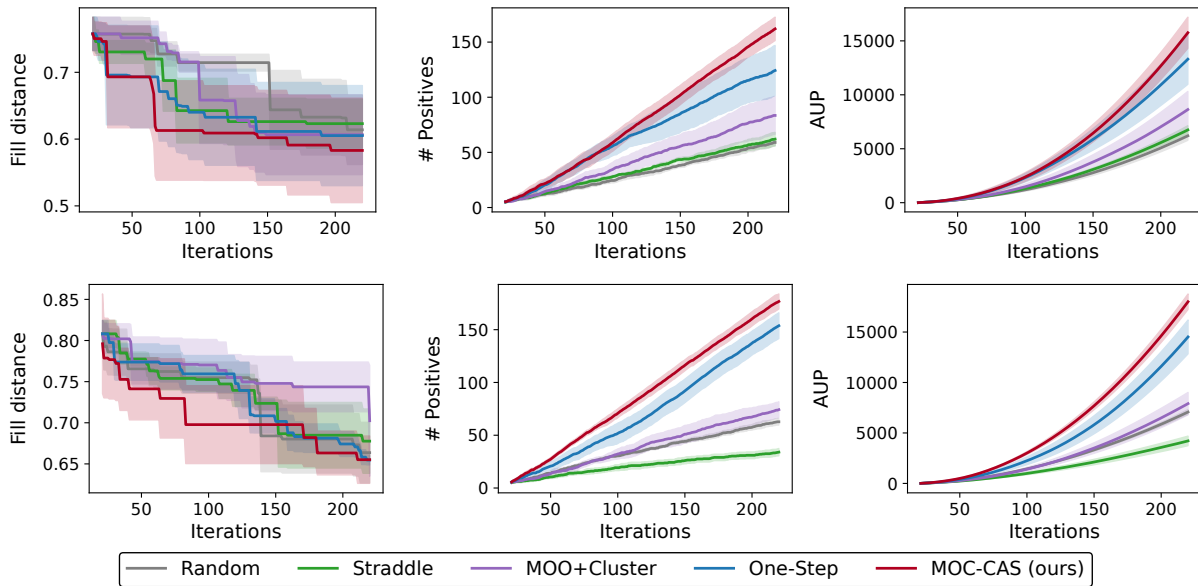


Figure 3: Quantitative comparison on the Cancer RTCB (top row) and Cancer 6T2W (bottom row) target datasets.

best method in terms of all the evaluation criteria, as evidenced in Figure 3. The performance gap between MOC-CAS and the baselines turns out to be statistically significant on the 6T2W dataset in Figure 3 in terms of AUP and number of positives - indicating that feasibility-only exploitation is not sufficient. The most plausible explanation is geometric: this target behaves like a harder coverage problem where feasible outcomes are either sparser or more fragmented in objective space, so repeatedly sampling the most confident feasible area leaves other feasible neighborhoods undiscovered. MOC-CAS is explicitly designed to exploit information from the posterior GPs and the history of discovered feasible outcomes, so its advantage grows as the posteriors sharpen and the soft feasibility/coverage terms become more accurate. In the initial iterations, all methods operate in a highly uncertain regime with very sparse data, so even sophisticated acquisitions behave close to exploratory baselines, which explains the small performance gaps.

Results on WRN dataset are deferred to the Appendix which show a similar pattern. Among the baseline methods, ONE-STEP

active search performed better than the other baselines across all the Cancer datasets.

### 6.4 Ablation Studies

MOC-CAS has two important hyperparameters carrying clear geometric meaning which govern its empirical efficiency and behavior. The coverage resolution,  $r$  assigns the granularity at which objective-space diversity is rewarded. Therefore, changes in the values of chosen  $r$  will have influence on the quality of dispersion, covered volume and cumulative positive samples found. Also, the optimism schedule,  $\beta_t$  governs how strongly the UCB gate considers feasibility, trading exploration close to the thresholds against conservative consolidation and so its time profile forms early vs. late discovery and therefore AUP. Since objective scales, threshold calibrations, and noise levels vary across the targets (e.g., cancer vs. viral proteins), a fixed setting is not always universally optimal. Therefore, we ablate  $r$  and  $\beta_t$  for characterizing sensitivity by reporting influences on all the evaluation criteria.




Figure 4: Ablation of coverage resolution  $r$  on the SARS-CoV-2 3CLpro dataset.




Figure 5: Ablation of optimism schedule  $\beta_t$  on the SARS-CoV-2 3CLpro dataset.

**Coverage resolution ( $r$ ).** We study the effect of the coverage resolution  $r$  by fixing all hyperparameters (including the UCB schedule) and reusing the same seeds and initial indices across different trials, then sweeping  $r \in \{0.035, 0.05, 0.10\}$  on the SARS-CoV-2 3CLpro dataset. We show the results in Figure 4. Across our runs, the highest resolution ( $r = 0.10$ ) delivered the most favorable trade-off yielding slightly lower fill distance and higher number of positives than others while maintaining competitive AUP, whereas the intermediate resolution ( $r = 0.05$ ) delivered slightly higher AUP. Figure 4 also shows that the smaller radius ( $r = 0.035$ ) comparatively under-covers the feasible region.

**Optimism schedule ( $\beta_t$ ).** In our implementation, the time-varying confidence is  $\beta_t = \beta_0 \cdot a_t$ , where the annealing multiplier  $a_t \in (0, 1]$  is held fixed across runs. The ablation changes only  $\beta_0$  while keeping  $a_t$  constant. To investigate exploration-exploitation, we fix the coverage radius and all other settings and then sweep the initial optimism level  $\beta_0 \in \{2.0, 3.0, 4.0\}$  with the same annealing schedule (with constant seeds and initial indices across four different trials). Figure 5 investigates the optimism schedule: a larger  $\beta_0$  early encourages broader exploration that can escape low-quality pockets, whereas smaller values promote later consolidation within the feasible region. Empirically, Figure 5 shows that increasing  $\beta_0$  modestly improves AUP and cumulative positives, while smaller  $\beta_0$  tends to slightly reduce fill distance.

## 7 CONCLUSION

We address a practically important setting common in drug discovery and materials design: when success depends on meeting

multiple performance thresholds, the objective is not to trace a Pareto front or spread samples in input space, but to rapidly identify a small, representative set whose outcomes broadly cover the feasible region. Such sets enable faster downstream evaluation and decision-making, accelerating scientific iteration under tight experimental budgets.

We formulate this as a multi-objective coverage problem and propose MOC-CAS, an acquisition strategy that integrates Gaussian process UCB predictions with an incremental coverage objective. To enable efficient inner optimization, we introduce a smoothed relaxation of the hard geometric objective, yielding a differentiable surrogate with closed-form components and simple gradients. Across large-scale protein-target benchmarks for cancer and SARS-CoV-2, MOC-CAS consistently achieves better coverage-quality tradeoffs than strong baselines, improving early discovery (higher cumulative positives and AUP) while maintaining diversity (lower fill distance). Our study employs independent Gaussian processes for each objective, though richer outcome models, such as deep kernels or correlated multi-output priors, represent promising directions for future work.

## ACKNOWLEDGMENTS

This work was partially supported by the IBM-UAlbany CEAIS Seed Grant 1201104-1-102522. The authors would like to thank anonymous reviewers and the area chair for helpful comments that improve this paper.

## REFERENCES

- [1] Ilija Bogunovic, Jonathan Scarlett, Andreas Krause, and Volkan Cevher. 2016. Truncated variance reduction: A unified approach to bayesian optimization and level-set estimation. *Advances in neural information processing systems* 29 (2016).
- [2] Brent Bryan, Robert C Nichol, Christopher R Genovese, Jeff Schneider, Christopher J Miller, and Larry Wasserman. 2005. Active learning for identifying function threshold boundaries. *Advances in neural information processing systems* 18 (2005).
- [3] Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. 2021. Parallel bayesian optimization of multiple noisy objectives with expected hypervolume improvement. *Advances in neural information processing systems* 34 (2021), 2187–2200.
- [4] Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. 2022. Multi-objective bayesian optimization over high-dimensional search spaces. In *Uncertainty in Artificial Intelligence*. PMLR, 507–517.
- [5] Michael TM Emmerich, André H Deutz, and Jan Willem Klinkenberg. 2011. Hypervolume-based expected improvement: Monotonicity properties and exact computation. In *2011 IEEE congress of evolutionary computation (CEC)*. IEEE, 2147–2154.
- [6] Peter I Frazier. 2018. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811* (2018).
- [7] Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann. 2012. Bayesian optimal active search and surveying. *arXiv preprint arXiv:1206.6406* (2012).
- [8] Alkis Gotovos, Nathalie Casati, Gregory Hitz, and Andreas Krause. 2013. Active learning for level set estimation. In *International Joint Conference on Artificial Intelligence*. 1344–1350.
- [9] Isabelle Guyon, Gavin C Cawley, Gideon Dror, and Vincent Lemaire. 2011. Results of the active learning challenge. In *Active Learning and Experimental Design workshop In conjunction with AISTATS 2010*. JMLR Workshop and Conference Proceedings, 19–45.
- [10] Daniel Hernández-Lobato, Jose Hernandez-Lobato, Amar Shah, and Ryan Adams. 2016. Predictive entropy search for multi-objective bayesian optimization. In *International conference on machine learning*. PMLR, 1492–1501.
- [11] Shogo Iwazaki, Yu Inatsu, and Ichiro Takeuchi. 2020. Bayesian experimental design for finding reliable level set under input uncertainty. *IEEE Access* 8 (2020), 203982–203993.
- [12] Joshua Knowles. 2006. ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE transactions on evolutionary computation* 10, 1 (2006), 50–66.
- [13] Junpei Komiya, Gustavo Malkomes, Bolong Cheng, and Michael McCourt. 2022. Bridging offline and online experimentation: Constraint active search for deployed performance optimization. *Transactions on Machine Learning Research* (2022).
- [14] Marco Laumanns and Jiri Ocenasek. 2002. Bayesian optimization algorithms for multi-objective optimization. In *International Conference on Parallel Problem Solving from Nature*. Springer, 298–307.
- [15] Diantong Li, Kyunghyun Cho, and Chong Liu. 2025. None To Optima in Few Shots: Bayesian Optimization with MDP Priors. *arXiv preprint arXiv:2511.01006* (2025).
- [16] Diantong Li, Fengxue Zhang, Chong Liu, and Yuxin Chen. 2025. Constrained Multi-objective Bayesian Optimization through Optimistic Constraints Estimation. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 370–378.
- [17] Xuefeng Liu, Songhao Jiang, Siyu Chen, Zhuoran Yang, Yuxin Chen, Ian Foster, and Rick Stevens. 2025. DrugImproverGPT: A Large Language Model for Drug Optimization with Fine-Tuning via Structured Policy Optimization. *arXiv preprint arXiv:2502.07237* (2025).
- [18] Xuefeng Liu, Songhao Jiang, Archit Vasan, Alexander Brace, Ozan Gokdemir, Thomas Brettin, Fangfang Xia, Ian Foster, and Rick Stevens. 2023. DRUGIMPROVER: Utilizing reinforcement learning for multi-objective alignment in drug optimization. In *NeurIPS 2023 Workshop on New Frontiers of AI for Drug Discovery and Development*, Vol. 1. 8.
- [19] Gustavo Malkomes, Bolong Cheng, Eric H Lee, and Mike Mccourt. 2021. Beyond the pareto efficient frontier: Constraint active search for multiobjective experimental design. In *International Conference on Machine Learning*. PMLR, 7423–7434.
- [20] Blake Mason, Lalit Jain, Subhojyoti Mukherjee, Romain Camilleri, Kevin Jamieson, and Robert Nowak. 2022. Nearly Optimal Algorithms for Level Set Estimation. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 7625–7658.
- [21] Sidhant Misra, Line Roald, and Yeesian Ng. 2022. Learning for constrained optimization: Identifying optimal active constraint sets. *INFORMS Journal on Computing* 34, 1 (2022), 463–480.
- [22] Biswajit Paria, Kirthevasan Kandasamy, and Barnabás Póczos. 2020. A flexible framework for multi-objective bayesian optimization using random scalarizations. In *Uncertainty in Artificial Intelligence*. PMLR, 766–776.
- [23] Ji Won Park, Nataša Tagasovska, Michael Maser, Stephen Ra, and Kyunghyun Cho. 2023. BOtied: Multi-objective Bayesian optimization with tied multivariate ranks. *arXiv preprint arXiv:2306.00344* (2023).
- [24] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems* 25 (2012).
- [25] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. 2010. Gaussian process optimization in the bandit setting: no regret and experimental design. In *International Conference on International Conference on Machine Learning*. 1015–1022.
- [26] Yanan Sui, Alkis Gotovos, Joel Burdick, and Andreas Krause. 2015. Safe exploration for optimization with Gaussian processes. In *International conference on machine learning*. PMLR, 997–1005.
- [27] Yanan Sui, Vincent Zhuang, Joel Burdick, and Yisong Yue. 2018. Stagewise safe Bayesian optimization with Gaussian processes. In *International conference on machine learning*. PMLR, 4781–4789.
- [28] Yining Wang, Sivaraman Balakrishnan, and Aarti Singh. 2019. Optimization of smooth functions with noisy observations: Local minimax rates. *IEEE Transactions on Information Theory* 65, 11 (2019), 7350–7366.
- [29] Christopher KI Williams and Carl Edward Rasmussen. 2006. *Gaussian processes for machine learning*. Vol. 2. MIT Press.
- [30] Kaifeng Yang, Michael Emmerich, André Deutz, and Thomas Bäck. 2019. Efficient computation of expected hypervolume improvement using box decomposition algorithms. *Journal of Global Optimization* 75, 1 (2019), 3–34.
- [31] Andrea Zanette, Junzi Zhang, and Mykel J Kochenderfer. 2018. Robust super-level set estimation using gaussian processes. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 276–291.
- [32] Richard Zhang and Daniel Golovin. 2020. Random hypervolume scalarizations for provable multi-objective black box optimization. In *International conference on machine learning*. PMLR, 11096–11105.