MULTI-OBJECTIVE OPTIMIZATION BY LEARNING SPACE PARTITIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

In contrast to single-objective optimization (SOO), multi-objective optimization (MOO) requires an optimizer to find the Pareto frontier, a subset of feasible solutions that are not dominated by other feasible solutions. In this paper, we propose LaMOO, a novel multi-objective optimizer that learns a model from observed samples to partition the search space and then focus on promising regions that are likely to contain a subset of the Pareto frontier. The partitioning is based on the *dominance number*, which measures "how close" a data point is to the Pareto frontier among existing samples. To account for possible partition errors due to limited samples and model mismatch, we leverage Monte Carlo Tree Search (MCTS) to exploit promising regions while exploring suboptimal regions that may turn out to contain good solutions later. Theoretically, we prove the efficacy of learning space partitioning via LaMOO under certain assumptions. Empirically, on the HyperVolume (HV) benchmark, a popular MOO metric, LaMOO substantially outperforms strong baselines on multiple real-world MOO tasks, by up to 225% in sample efficiency for neural architecture search on Nasbench201, and up to 10% for molecular design.

1 INTRODUCTION

Multi-objective optimization (MOO) has been extensively used in many practical scenarios involving trade-offs between multiple objectives. For example, in automobile design (Chang, 2015), we must maximize the performance of the engine while simultaneously minimizing emissions and fuel consumption. In finance (Gunantara, 2018), one prefers a portfolio that maximizes the expected return while minimizing risk.

Mathematically, in MOO we optimize M objectives $\mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_M(\mathbf{x})] \in \mathbb{R}^M$:

$$\begin{aligned} \min \quad & f_1(\mathbf{x}), f_2(\mathbf{x}), ..., f_M(\mathbf{x}) \\ \text{s.t.} \quad & \mathbf{x} \in \Omega \end{aligned} \tag{1}$$

While we could set arbitrary weights for each objective to turn it into a single-objective optimization (SOO) problem, modern MOO methods aim to find the problem's entire *Pareto frontier*: the set of solutions that are not *dominated* by any other feasible solutions[1] (see Fig. 1 for illustration). The Pareto frontier yields a global picture of optimal solution structures rather than focusing on one specific weighted combination of objectives.

As a result, MOO is fundamentally different from SOO. Instead of focusing on a single optimal solution, a strong MOO optimizer should cover the search space broadly to explore the Pareto frontier. Popular quality indicators in MOO, such as hypervolume (HV), capture this aspect by computing the volume of the currently estimated frontier. Specifically, given a reference point $R \in \mathbb{R}^M$, as shown in Fig. 1(a), the *hypervolume* of a finite approximate Pareto set $\mathcal{P}$ is the M-dimensional Lebesgue measure λ_M of the space dominated by $\mathcal{P}$ and bounded from below by R. That is, $HV(\mathcal{P}, R) = \lambda_M(\cup_{i=1}^{|\mathcal{P}|}[R, y_i])$, where $[R, y_i]$ denotes the hyper-rectangle bounded by reference point R and y_i. Consequently, the optimizer must consider the diversity of solutions in addition to their optimality.

[1]Here we define *dominance* $\mathbf{y} \prec_{\mathbf{f}} \mathbf{x}$ as $f_i(\mathbf{x}) \leq f_i(\mathbf{y})$ for all functions f_i, and exists at least one i s.t. $f_i(\mathbf{x}) < f_i(\mathbf{y})$, $1 \leq i \leq M$. That is, solution $\mathbf{x}$ is always better than solution $\mathbf{y}$, regardless of how the M objectives are weighted.

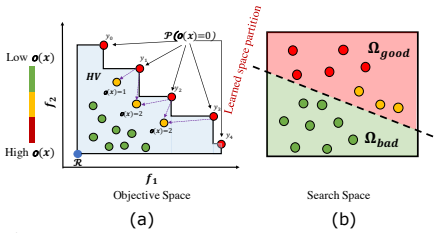

(a) (b)

MOO methods	Sampling Method	Objectives>3
MOEA/D (Zhang & Li, 2007)	Evolution	×
CMA-ES (Igel et al., 2007b)		×
NSGA-II (Deb et al., 2002a)		×
NAGA-III (Deb & Jain, 2014)		√
qPAREGO (Knowles, 2006)	Bayesian optimization	√
qEHVI (Daulton et al., 2020)		√
LaMOO (our approach)	Space partition	√

Figure 1: **Left**: A basic setting in Multi-objective Optimization (MOO), optimizing $M = 2$ objectives in Eqn. 1. (a) depicts the objective space (f_1, f_2) and (b) shows the search space $\mathbf{x} \in \Omega$. In (a), P denotes the Pareto frontier, R is the reference point, the hypervolume HV is the space of the shaded area, and $o(\mathbf{x})$ are the dominance numbers. In (b), once a few samples are collected within Ω, LaMOO *learns to partition* the search space Ω into sub-regions (i.e. Ω_{good} and Ω_{bad}) according to the dominance number in objective space, and then focuses future sampling on the good regions that are close to the Pareto Frontier. This procedure can be repeated to further partition Ω_{good} and Ω_{bad}. **Right**: A table shows the properties of MOO methods used in experiments.

While several previous works have proposed approaches to capture this diversity-optimality trade-off (Deb et al., 2002a; Deb & Jain, 2014; Igel et al., 2007a; Daulton et al., 2020; Knowles, 2006), in this paper, we take a fundamentally different route by *learning* promising candidate regions from past explored samples. Ideally, to find the Pareto frontier in as few function evaluations as possible, we want to sample heavily in the Pareto optimal set Ω_P, defined as the region of input vectors that corresponds to the Pareto frontier.

One way to focus samples on Ω_P is to gradually narrow the full search space down to the subregion containing Ω_P via partitioning. For example, in the case of quadratic objective functions, Ω_P can be separated from the non-optimal set $\Omega \backslash \Omega_P$ via simple linear classifiers (see Observation 1,2). Motivated by these observations, we thus design LaMOO, a novel MOO meta-optimizer that progressively partitions regions into sub-regions and then focuses on sub-regions that are likely to contain Pareto-optimal regions, where existing solvers can help. Therefore, LaMOO is a meta-algorithm.

Unlike cutting-plane methods (Loganathan & Sherali, 1987; Vieira & Lisboa, 2019; Hinder, 2018) that leverage the (sub)-gradient of convex objectives as the cutting plane, with global optimality guarantees, LaMOO is data-driven: it leverages previous samples to build classifiers to *learn* the partition and focuses future samples in these promising regions. No analytical formula of objectives or their sub-gradients is needed. LaMOO is a multi-objective extension of recent works (Wang et al., 2020; Yang et al., 2021a) that also learn space partitions but for a single black-box objective.

Empirically, LaMOO outperforms existing approaches on many benchmarks, including standard benchmarks in multi-objective black-box optimization, and real-world multi-objective problems like neural architecture search (NAS) (Cai et al., 2020; 2019) and molecule design. For example, as a meta-algorithm, LaMOO combined with CMA-ES as an inner routine requires only 62.5%, 8%, and 29% as many samples to reach the same hypervolume as the original CMA-ES (Igel et al., 2007b) in BraninCurrin (Belakaria et al., 2019), VehicleSafety (Liao et al., 2008) and Nasbench201 (Dong & Yang, 2020), respectively. On average, compared to qEHVI, LaMOO uses 50% samples to achieve the same performance in these problems. In addition, LaMOO with qEHVI (Daulton et al., 2020) and CMA-ES require 71% and 31% fewer samples on average, compared to naive qEHVI and CMA-ES, to achieve the same performance in molecule discovery.

2 Related Work

Bayesian Optimization (BO) (Knowles, 2006; Daulton et al., 2020; Zitzler et al., 2003; Yang et al., 2019; Couckuyt et al., 2014; Paria et al., 2018; Ponweiser et al., 2008) is a popular family of methods to optimize black-box single and multi-objectives. Using observed samples, BO learns a surrogate model $\hat{f}(\mathbf{x})$, search for new promising candidates based on *acquisition function* built on $\hat{f}(\mathbf{x})$, and query the quality of these candidates with the ground truth black-box objective(s). In multi-objective Bayesian optimization (MOBO), most approaches leverage Expected Hypervolume Improvement (EHVI) as their acquisition function (Zitzler et al., 2003; Yang et al., 2019; Couckuyt et al., 2014), since finding the Pareto frontier is equivalent to maximizing the hypervolume given a finite search space (Fleischer, 2003). There are methods (Knowles, 2006; Paria et al., 2018; Ponweiser et al., 2008) that use different acquisition functions like expected improvement (Jones et al., 1998) and Thompson

sampling (Thompson, 1933). EVHI is computationally expensive: its cost increases exponentially with the number of objectives. To address this problem, qEHVI (Daulton et al., 2020) accelerates optimization by computing EHVI in parallel, and has become the state-of-the-art MOBO algorithm. In this paper, we leverage qEHVI as a candidate inner solver in our proposed LaMOO algorithm.

Evolutionary algorithms (EAs) (Deb et al., 2002a; Fang et al., 2018; Zhang & Li, 2007; Beume et al., 2007; Igel et al., 2007b) are also popular methods for MOO tasks. MOO-EAs can be roughly categorized into three groups. One category (Srinivas & Deb, 1994; Deb et al., 2002a; Deb & Jain, 2014) leverages Pareto dominance to simultaneously optimize all objectives. A second category (e.g., (Zhang & Li, 2007)) decomposes a multi-objective optimization problem into a number of single-objective sub-problems, converting a difficult MOO into several SOOs. Another category is quality indicator-based methods, such as (Beume et al., 2007) and (Igel et al., 2007b). They scalarize the current Pareto frontier using quality indicators (e.g., HV) and transfer a MOO to a SOO. New samples are generated by crossover and mutation operations from existing ones. While EAs have shown success in two-objective problems, non-quality indicator-based methods (i.e., the first two categories) struggle to handle a larger number of objectives. Specifically, for MOO with many objectives, NSGA-II (Deb et al., 2002a) easily gets stuck in a dominance resistant solution (Pang et al., 2020) which is far from the true Pareto frontier. Determining the number of weight vectors, and efficiently generating and initializing them, are the main challenges of MOEA/D in many-objective (> 3) problems (Zhang & Li, 2007; Xu et al., 2020).

Quality Indicators. Besides hypervolume, there are several other quality indicators (Van Veldhuizen & Lamont, 1998; Bosman & Thierens, 2003; Zitzler et al., 2000; Bandyopadhyay et al., 2004; Jiang et al., 2014) for evaluating sample quality, which can be used to scalarize the MOO to SOO. The performance of a quality indicator can be evaluated by three metrics (Deng et al., 2007; Li et al., 2014), including convergence (closeness to the Pareto frontier), uniformity (the extent of the samples satisfy the uniform distribution), and spread (the extent of the obtained approximate Pareto frontier). Sec. B specifically illustrates the merits of each quality indicator. HyperVolume is the only metric we explored that can simultaneously satisfy the evaluation of convergence, uniformity, and spread without the knowledge of the true Pareto frontier. Therefore, throughout this work, we use HV to evaluate the optimization performance of different algorithms.

3 Learning Space Partitions: A Theoretical Understanding

Searching in high-dimensional space to find the optimal solution to a function is in general a challenging problem, especially when the function's properties are unknown to the search algorithm. The difficulty is mainly due to the curse of dimensionality: to adequately cover a d-dimensional space, in general, an exponential number of samples are needed.

For this, many works use a "coarse-to-fine" approach: partition the search space and then focusing on promising regions. Traditionally, manually defined criteria are used, e.g., axis-aligned partitions (Munos, 2011), Voronoi diagrams (Kim et al., 2020), etc. Recently, (Wang et al., 2019; 2020; Yang et al., 2021a) *learn* space partitions based on the data collected thus far, and show strong performance in NeurIPS black box optimization challenges (Sazanovich et al.; Kim et al.).

On the other hand, there is little quantitative understanding of space partition. In this paper, we first give a formal theoretical analysis on why learning plays an important role in space-partition approaches for SOO. Leveraging our understanding of how space partitioning works, we propose LaMOO which empirically outperforms existing SoTA methods on multiple MOO benchmarks.

3.1 Problem Setting

Intuitively, learning space partitions will yield strong performance if the classifier can determine which regions are promising given few data points. We formalize this intuition below and show why it is better than fixed and manually defined criteria for space partitioning.

Consider the following sequential decision task. We have N samples in a *discrete* subset S_0 and there exists one sample $\mathbf{x}^*$ that achieves a minimal value of a scalar function f. Note that f can be any property we want, e.g., in the Pareto optimal set. The goal is to construct a subset $S_T \subseteq S_0$ after T steps, so that (1) $\mathbf{x}^* \in S_T$ and (2) $|S_T|$ is as small as possible. More formally, we define the reward function r as the probability that we get $\mathbf{x}^*$ by randomly sampling from the resulting subset S_T:

$$r := \frac{1}{|S_T|} P(\mathbf{x}^* \in S_T) \tag{2}$$

It is clear that $0 \leq r \leq 1$. $r = 1$ means that we already found the optimal sample $\mathbf{x}^*$.

Here we use discrete case for simplicity and leave continuous case (i.e., partitioning a region Ω_0 instead of a discrete set S_0) to future work. Note N could be large, so here we consider it infeasible to enumerate S_0 to find $\mathbf{x}^*$. However, sampling from S_0, as well as comparing the quality of sampled solutions are allowed. An obvious baseline is to simply set $S_T := S_0$, then $r_\text{b} = N^{-1}$. Now the question is: can we do better? Here we seek help from the following *oracle*:

Definition 1 ((α, η)-Oracle)**.** *Given a subset S that contains $\mathbf{x}^*$, after taking k samples from S, the oracle can find a* good *subset S_good with $|S_\text{good}| \leq |S|/2$ and*

$$P\left(\mathbf{x}^* \in S_\text{good} | \mathbf{x}^* \in S\right) \geq 1 - \exp\left(-\frac{k}{\eta |S|^\alpha}\right) \tag{3}$$

Lemma 1. *The algorithm to uniformly draw k samples in S, pick the best and return is a $(1, 1)$-oracle.*

See Appendix for proof. Note that a $(1, 1)$-oracle is very weak, and is of little use in obtaining higher reward r. We typically hope for an oracle with smaller α and η (i.e., both smaller than 1). Intuitively, such oracles are more sample-efficient: with few samples, they can narrow down the region containing the optimal solution $\mathbf{x}^*$ with high probability.

Note that $\alpha < 1$ corresponds to *semi-parametric models*. In these cases, the oracle has *generalization property*: with substantially fewer samples than N (i.e., on the order of N^α), the oracle is able to put the optimal solution $\mathbf{x}^*$ on the right side. In its extreme case when $\alpha = 0$ (or *parametric models*), whether we classify the optimal solution $\mathbf{x}^*$ on the correct side only depends on the *absolute* number of samples collected in S, and is independent of its size. For example, if the function to be optimized is linear, then with $d + 1$ samples, we can completely characterize the property of all $|S|$ samples.

Relation with cutting plane. Our setting can be regarded as a data-driven extension of cutting plane methods (Loganathan & Sherali, 1987; Vieira & Lisboa, 2019; Hinder, 2018) in optimization, in which a cutting plane is found at the current solution to reduce the search space. For example, if f is convex and its gradient $\nabla f(\mathbf{x})$ is available, then we can set $S_\text{good} := \{\mathbf{x} : \nabla f(\mathbf{x}_0)^\top (\mathbf{x} - \mathbf{x}_0) \leq 0, \mathbf{x} \in S_0\}$, since for any $\mathbf{x} \in S_0 \setminus S_\text{good}$, convexity gives $f(\mathbf{x}) \geq f(\mathbf{x}_0) + \nabla f(\mathbf{x}_0)^\top (\mathbf{x} - \mathbf{x}_0) > f(\mathbf{x}_0)$ and thus $\mathbf{x}$ is not better than current $\mathbf{x}_0$.

However, the cutting plane method relies on certain function properties like convexity. In contrast, learning space partition can leverage knowledge about the function forms, combined with observed samples so far, to better partition the space.

3.2 Rewards under optimal action sequence

We now consider applying the (α, η)-oracle iteratively for T steps, by drawing k_t samples from S_{t-1} and setting $S_t := S_{\text{good},t-1}$. We assume a total sample budget K, so $\sum_{t=1}^T k_t = K$. Note that $T \leq \log_2 N$ since we halve the set size with each iteration.

Now the question is twofold. (1) How can we determine the action sequences $\{k_t\}$ in order to maximize the total reward r? (2) Following the optimal action sequences $\{k_t^*\}$, can r^* be better than the baseline $r_\text{b} = N^{-1}$? The answer is yes.

Theorem 1. *The algorithm yields a reward r^* lower bounded by the following:*

$$r^* \geq r_\text{b} \exp\left[\left(\log 2 - \frac{\eta N^\alpha \phi(\alpha, T)}{K}\right) T\right] \tag{4}$$

where $r_\text{b} := N^{-1}$ and $\phi(\alpha, T) := (1 - 2^{-\alpha T})/(1 - 2^{-\alpha})$.

Remarks. Following Theorem 1, a key condition to make $r^* > r_\text{b}$ is to ensure $\log 2 > \frac{\eta N^\alpha \phi(\alpha, T)}{K}$. This holds if when $\frac{\eta N^\alpha \phi(\alpha, T)}{K} \to 0$. Note that since $T \leq \log_2 N$, the final reward r^* is upper bounded by 1 (rather than goes to $+\infty$). We consider some common practical scenarios below.

Non-parametric models ($\alpha = 1$). In this case, $\phi(\alpha, T) \leq 2$ and the condition becomes $\frac{1}{2} \log 2 > \eta N/K$. This happens when the total sample budget $K = \Theta(N)$, i.e., on the same order of N, which means that the partitioning algorithm obtains little advantage over exhaustive search.

Semi-parametric models ($\alpha < 1$). In this case, $\phi(\alpha, T) \leq 1/(1 - 2^{-\alpha})$ and the condition becomes $(1 - 2^{-\alpha}) \log 2 > \eta N^\alpha / K$. This happens when the total sample budget $K = \Theta(N^\alpha)$. In this case,

we could use many fewer samples than exhaustive search to achieve better reward, thanks to the generalization property of the oracle.

Parametric models ($\alpha = 0$). Now $\phi(\alpha, T) = T$ and the condition becomes $\log 2 > \frac{\eta T}{K}$. Since $T \leq \log_2 N$, the total sample budget can be set to be $K = \Theta(\log N)$. Intuitively, the algorithm performs iterative halving (or binary search) to narrow down the search toward promising regions.

3.3 Extension to Multi-Objective Optimization

Given our understanding of space partitioning, we now extend this idea to MOO. Intuitively, we want "good" regions to be always picked by the space partition. For SOO, it is possible since the optimal solution is a single point. How about MOO?

Unlike SOO, in MOO we aim for a continuous region, the *Pareto optimal set* $\Omega_P := \{\mathbf{x} : \nexists \mathbf{x}' \neq \mathbf{x} : \mathbf{f}(\mathbf{x}') \prec \mathbf{f}(\mathbf{x})\}$. A key variable is the regularity of Ω_P: if it is highly non-regular and not captured by a simple partition boundary (ideally a parametric boundary), then learning a space partition would be difficult. Interestingly, the shape of Ω_P can be characterized for quadratic objectives:

Observation 1. *If all f_j are isotropic, $f_j(\mathbf{x}) = \|\mathbf{x} - \mathbf{c}_j\|_2^2$, then $\Omega_P = \mathrm{ConvexHull}(\mathbf{c}_1, \ldots, \mathbf{c}_q)$.*

Observation 2. *If $M = 2$ and $f_j(\mathbf{x}) = (\mathbf{x} - \mathbf{c}_j)^\top H_j(\mathbf{x} - \mathbf{c}_j)$ where H_j are positive definite symmetric matrices, then there exists $\mathbf{w}_1 := H_2(\mathbf{c}_2 - \mathbf{c}_1)$ and $\mathbf{w}_2 := H_1(\mathbf{c}_1 - \mathbf{c}_2)$, so that for any $\mathbf{x} \in \Omega_P$, $\mathbf{w}_1^\top(\mathbf{x} - \mathbf{c}_1) \geq 0$ and $\mathbf{w}_2^\top(\mathbf{x} - \mathbf{c}_2) \geq 0$.*

In both cases, Ω_P can be separated from non-Pareto regions $\Omega \backslash \Omega_P$ via a linear hyperplane. Empirically, Ω_P only occupies a small region of the entire search space (Sec. 4), and quickly focusing samples on the promising regions is critical for high sample efficiency.

In the general case, characterizing Ω_P is analytically hard and requires domain knowledge about the objectives (Li et al., 2014). However, for MOO algorithms in practice, knowing that Ω_P can be separated from $\Omega \backslash \Omega_P$ via simple decision planes is already useful: we could learn such decision planes given previous data that are already collected, and sample further in promising regions.

4 LaMOO: Latent Action Multi-objective Optimization

In Sec. 3, for convenience, we only analyze a greedy approach, which makes decisions on space partitions and never revises them afterwards. While this greedy approach indeed works (as shown in Sec. 5.3), an early incorrect partition could easily rule out regions that turn out to be good but weren't identified with few samples. In practice, we want to keep the decision softer: while *exploiting* the promising region, we also *explore* regions that are currently believed to be sub-optimal given limited samples. It is possible that these regions turn out to contain good solutions when more samples are available, and the oracle can then make a different partition.

To balance the trade-off between exploration and exploitation to cope with the generalization error of the learned classifier, we leverage Monte Carlo Tree Search (MCTS) (Kocsis & Szepesvári, 2006) and propose our algorithm LaMOO. As shown in Alg. 1, LaMOO has four steps: (1) learn to partition the search space given previous observed data points D_t, which are collected $\{\mathbf{x}_i, \mathrm{f}(\mathbf{x}_i)\}$ from iterations 0 to t. (2) With this information, we partition the region into promising and non-promising regions, and learn a classifier $h(\cdot)$ to separate them. (3) We select the region to sample from, based on the UCB value of each node. (4) We sample selected regions to obtain future data points D_{t+1}.

Learning Space Partitions. We construct the partition oracle using the *dominance number*. Let D_t be the collected samples up to iteration t and $D_{t,j} := D_t \cap \Omega_j$ be the samples within the region Ω_j we want to partition. For each sample $\mathbf{x} \in D_{t,j}$, its dominance number $o_{t,j}(\mathbf{x})$ at iteration t is defined as the number of samples in Ω_j that dominate $\mathbf{x}$ (here $\mathbb{I}[\cdot]$ is the indicator function):

$$o_{t,j}(\mathbf{x}) := \sum_{\mathbf{x}_i \in D_{t,j}} \mathbb{I}[\mathbf{x} \prec_{\mathbf{f}} \mathbf{x}_i,\ \mathbf{x} \neq \mathbf{x}_i] \tag{5}$$

While naive computation requires $O(|D_{t,j}|^2)$ operations, we use Maxima Set (Kung et al., 1975) which runs in $O(|D_{t,j}| \log |D_{t,j}|)$. For $\mathbf{x} \in \Omega_P$, $o(\mathbf{x}) = 0$.

For each $D_{t,j}$, we then get good (small $o(\mathbf{x})$) and bad (large $o(\mathbf{x})$) samples by ranking them according to $o(\mathbf{x})$. The smallest 50% are labeled to be positive while others are negative.

Algorithm 1 LaMOO Pseudocode.

Inputs: Initial D_0 from uniform sampling, sample budget T.
for $t = 0, \ldots, T$ **do**
  Set $\mathcal{L} \leftarrow \{\Omega_{\text{root}}\}$ (collections of regions to be split).
  while $\mathcal{L} \neq \emptyset$ **do**
    $\Omega_j \leftarrow \text{pop_first_element}(\mathcal{L}),\ \ D_{t,j} \leftarrow D_t \cap \Omega_j,\ \ n_{t,j} \leftarrow |D_{t,j}|$.
    Compute dominance number $o_{t,j}$ of $D_{t,j}$ using Eqn. 5 and train SVM model $h(\cdot)$.
    If $(D_{t,j}, o_{t,j})$ is splittable by SVM, **then** $\mathcal{L} \leftarrow \mathcal{L} \cup \text{Partition}(\Omega_j, h(\cdot))$.
  end while
  for $k = \text{root}$, k is not leaf node **do**
    $D_{t,k} \leftarrow D_t \cap \Omega_k,\ \ v_{t,k} \leftarrow \text{HyperVolume}(D_{t,k}),\ \ n_{t,k} \leftarrow |D_{t,k}|$.
    $k \leftarrow \arg\max_{c \in \text{children}(k)} \text{UCB}_{t,c}$, where $\text{UCB}_{t,c} := v_{t,c} + 2C_p\sqrt{\frac{2\log(n_{t,k}}{n_{t,c}}}$
  end for
  $D_{t+1} \leftarrow D_t \cup D_{\text{new}}$, where D_{new} is drawn from Ω_k based on qEHVI or CMA-ES.
end for

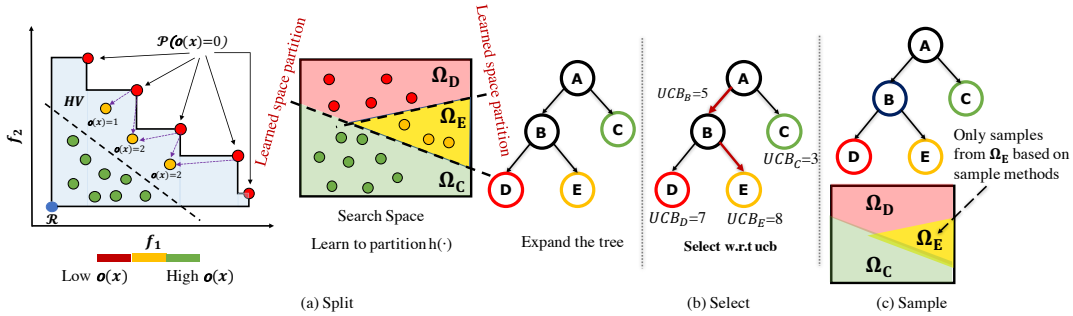

Figure 2: (a) The leaf nodes D and E that correspond to the non-splittable space Ω_D and Ω_E. (b). The node selection procedure based on the UCB value. (c). The new samples generation from the selected space Ω_E for bayesian optimization.

Based on the labeled samples, a classifier (e.g., Support Vector Machine (SVM)) is trained to learn a decision boundary as the latent action. We choose SVM since the classifier needs to be decent in regions with few samples, and has the flexibility of being parametric or non-parametric.

Exploration Using UCB. As shown in Fig. 2, LaMOO selects the final leaf node by always choosing the child node with larger UCB value. The *UCB value* for a node j is defined as $\text{UCB}_j := v_j + 2C_p\sqrt{2\log n_{\text{parent(j)}}/n_j}$, where n_j is the number of samples in node j, C_p is a tunable hyperparameter which controls the degree of exploration, and v_j is the hypervolume of the samples in node j. The selected leaf corresponds the partitioned region Ω_k as shown in Alg. 1.

Sampling in Search Region. We use existing algorithms as a sampling strategy in a leaf node, e.g., qEHVI (Daulton et al., 2020) and CMA-ES (Igel et al., 2007b). Therefore, LaMOO can be regarded as a *meta-algorithm*, applicable to any existing SOO/MOO solver to boost its performance.

LaMOO with qEHVI. As a multi-objective solver, qEHVI finds data points to maximize a parallel version of Expected Hypervolume Improvement (EHVI) via Bayesian Optimization (BO). To incorporate qEHVI into LaMOO's sampling step, we confine qEHVI's search space using the tree-structured partition to better search MOO solutions.

LaMOO with CMA-ES. CMA-ES is an evolutionary algorithm (EA) originally designed for single-objective optimization. As a leaf sampler, CMA-ES is used to pick a sample that maximizes the dominance number $o(\mathbf{x})$ within the leaf. Since $o(\mathbf{x})$ changes over iterations, at iteration t, we first update $o_{t'}(\mathbf{x})$ of all previous samples at $t' < t$ to $o_t(\mathbf{x})$, then use CMA-ES. Similar to the qEHVI case, we constrain our search to be within the leaf region.

Once a set of new samples D_{new} is obtained (as well as its multiple function values $\mathbf{f}(D_{\text{new}})$), we update all partitions along its path and the entire procedure is repeated.

5 EXPERIMENTS

We evaluate the performance of LaMOO in a diverse set of scenarios. This includes synthetic functions, and several real-world MOO problems like neural architecture search, automobile safety design, and molecule discovery. In such real problems, often a bunch of criteria needs to be optimized at the same time. For example, for molecule (drug) discovery, one wants the designed drug to be effective towards the target disease, able to be easily synthesized, and be non-toxic to human body.

5.1 SMALL-SCALE PROBLEMS

Synthetic Functions. Branin-Currin (Belakaria et al., 2019) is a function with 2-dimensional input and 2 objectives. DTLZ2 (Deb et al., 2002b) is a classical scalable multi-objective problem and is popularly used as a benchmark in the MOO community. We evaluate LaMOO as well as baselines in DTLZ2 with 18 dimensions and 2 objectives, and 12 dimensions and 10 objectives, respectively.

Structural Optimization in Automobile Safety Design (vehicle safety) is a real-world problem with 5-dimensional input and 3 objectives, including (1) the mass of the vehicle, (2) the collision acceleration in a full-frontal crash, and (3) the toe-board intrusion (Liao et al., 2008).

Nasbench201 is a public benchmark to evaluate NAS algorithms (Dong & Yang, 2020). There are 15625 architectures in Nasbench201, with groundtruth #FLOPs and accuracy in CIFAR10 (Krizhevsky, 2009). Our goal is to minimize #FLOPs and maximize accuracy in this search space. We normalized #FLOPs to range $[-1, 0]$ and accuracy to $[0, 1]$.

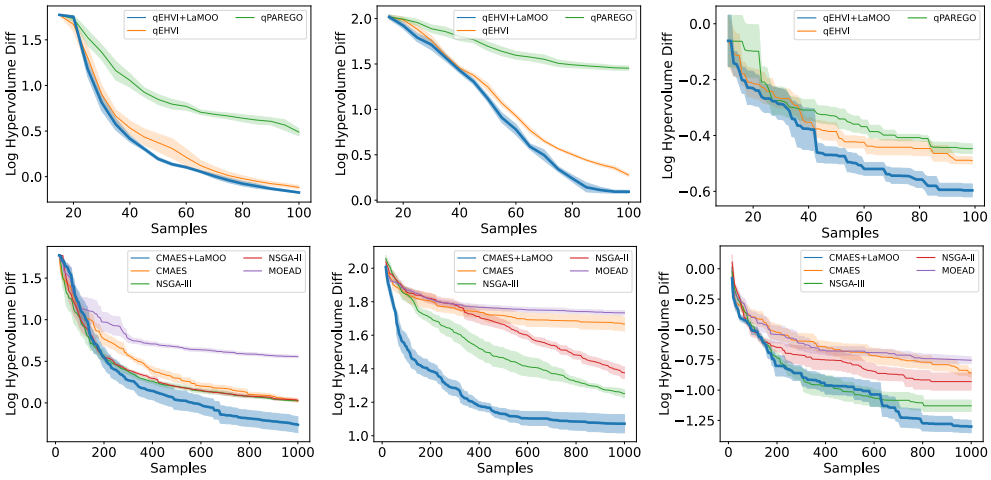

Figure 3: **Left**: Branin-Currin with 2 dimensions and 2 objectives. **Middle**: VehicleSafety with 5 dimensions and 3 objectives. **Right**: Nasbench201 with 6 dimensions and 2 objectives. We ran each algorithm 7 times (shaded area is $\pm$ std of the mean). **Top**: Bayesian Optimization w/o LaMOO. **Bottom**: evolutionary algorithms w/o LaMOO. Note the two algorithm families show very different sample efficiency in MOO tasks.

We compare LaMOO with 4 classical evolutionary algorithms (CMA-ES (Igel et al., 2007b), MOEA/D (Zhang & Li, 2007), NSGA-II (Deb et al., 2002a), and NSGA-III (Deb & Jain, 2014)) and 2 state-of-the-art BO methods (qEHVI (Daulton et al., 2020) and qParego (Knowles, 2006)).

Evaluation Criterion. we first obtain the maximal hypervolume (either by ground truth or from the estimation of massive sampling), then run each algorithm and compute the log hypervolume difference (Daulton et al., 2020):

$$HV_{\text{log_diff}} := \log(HV_{\max} - HV_{\text{cur}}) \tag{6}$$

where HV_{cur} is the hypervolume of current samples obtained by the algorithm with given budget.

Result. As shown in Fig. 3, LaMOO with qEHVI outperforms all our BO baselines and LaMOO with CMA-ES outperforms all our EA baselines, in terms of $HV_{\text{log_diff}}$.

Evolutionary algorithms rely on mutation and crossover of previous samples to generate new ones, and may be trapped into local optima. Thanks to MCTS, LaMOO also considers exploration and greatly

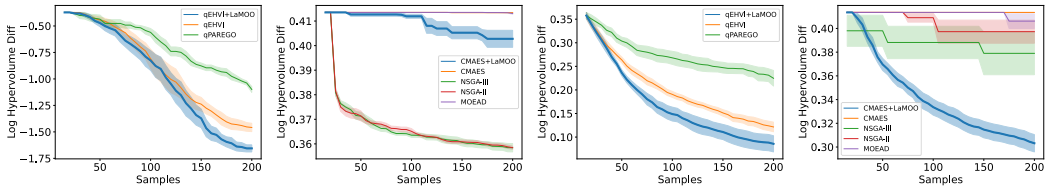

Figure 4: DTLZ2 with many objectives, We ran each algorithm 7 times (shaded area is $\pm$ std of the mean). From **left** to **right**: BO with 2 objectives; EA with 2 objectives; BO with 10 objectives; EA with 10 objectives.

improves upon vanilla CMA-ES over three different tasks are given 1000 samples. In addition, by plugging in BO, LaMOO+qEHVI achieve 225% sample efficiency compared to other BO algorithms on Nasbench201. This result indicates that for high-dimensional problems (6 in Nasbench201), space partitioning leads to faster optimization. We further analyze very high-dimensional problems on Sec. 5.2. For visualization of Pareto frontier by LaMOO+qEHVI, see Fig. 9 in Appendix.

Optimization of Many Objectives. While NSGA-II and NSGA-III perform well in the two-objective problems, all evolutionary-based baselines get stuck in the ten-objective problems. In contrast, LaMOO performs reasonably well. From Fig. 4, qEHVI+LaMOO shows strong performance in ten objectives. When combined with a CMA-ES, LaMOO help it escape the initial region to focus on a smaller promising region by space partitioning.

5.2 Multi-Objective Molecule Discovery

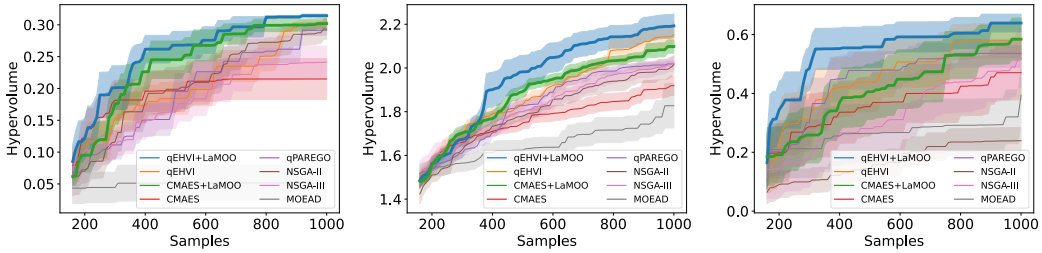

Figure 5: Molecule Discovery: **Left**: Molecule discovery with two objectives (GSK3β+JNK3). **Middle**: Molecule discovery with three objectives (QED+SA+SARS). **Right**: Molecule Discovery with four objectives (GSK3β+JNK3+QED+SA). We ran each algorithm 15 times (shaded area is $\pm$ std of the mean).

Next, we tackle the practical problem of multi-objective molecular generation, which is a high-dimensional problem (search space is 32-dimensional). Molecular generation models are a critical component of pharmaceutical drug discovery, wherein a cheap-to-run *in silico* model proposes promising molecules which can then be synthesized and tested in a lab (Vamathevan et al., 2019). However, one commonly requires the generated molecule to satisfy multiple constraints: for example, new drugs should generally be non-toxic and ideally easy-to-synthesize, in addition to their primary purpose. Therefore, in this work, we consider several multi-objective molecular generation setups from prior work on molecular generation (Yu et al., 2019; Jin et al., 2020b; Yang et al., 2021b): (1) activity against biological targets GSK3β and JNK3, (2) the same targets together with QED (a standard measure of "drug-likeness") and SA (a standard measure of synthetic accessibility), and (3) activity against SARS together with QED and SA. In each task, we propose samples from a pre-trained 32-dimensional latent space from (Jin et al., 2020a), which are then decoded into molecular strings and fed into the property evaluators from prior work.

Fig. 5 shows that LaMOO+qEHVI outperforms all baselines by up to 10% on various combinations of objectives. While EA struggles to optimize these high-dimensional problems due to the limitations mentioned in Sec. 2, LaMOO help them (e.g., CMA-ES) to performs much better.

5.3 Ablation Studies

Visualization of LaMOO. To understand how LaMOO works, we visualize its optimization procedure for Branin-Currin. First, the Pareto optimal set Ω_P is estimated from 10^6 random samples (marked as black stars), as shown in both search and objective space (Fig. 6(b) and bottom left of Fig. 6(c)). Over several iterations, LaMOO progressively prunes away unpromising regions so that the

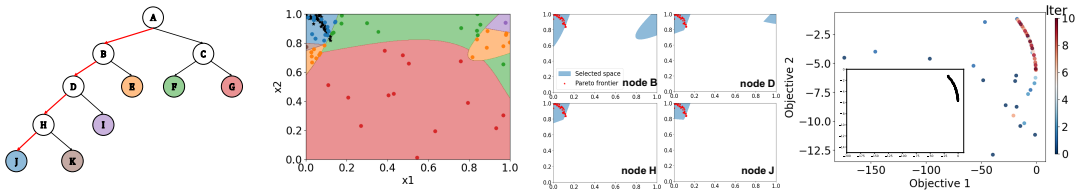

(a) The structure of Monte-Carlo Tree at the final search iteration.

(b) The samples and SVM split region of leaf node of the tree in the search space.

(c) The selected region in selected path of the tree.

(d) The samples in the objective space at different search iterations.

Figure 6: Visualization of selected region at different search iterations and nodes. (a) The Monte-Carlo tree with colored leaves. Selected path is marked in red. (b) Visualization of the regions($\Omega_J, \Omega_K, \Omega_I, \Omega_E, \Omega_F, \Omega_G$) that are consistent with leaves in (a) in the search space. (c) Visualization of selected path at final iteration. (d) Visualization of samples during search; bottom left is the Pareto frontier estimated from one million samples.

remaining regions approach Ω_P (Fig 6(c)). Fig 6(a) shows the final tree structure. The color of each leaf node corresponds to a region in the search space (Fig 6(b)). The selected region is recursively bounded by SVM classifiers corresponding to nodes on the selected path (red arrow in Fig 6(a)). The new samples are only generated from the most promising region Ω_J, improving sample efficiency.

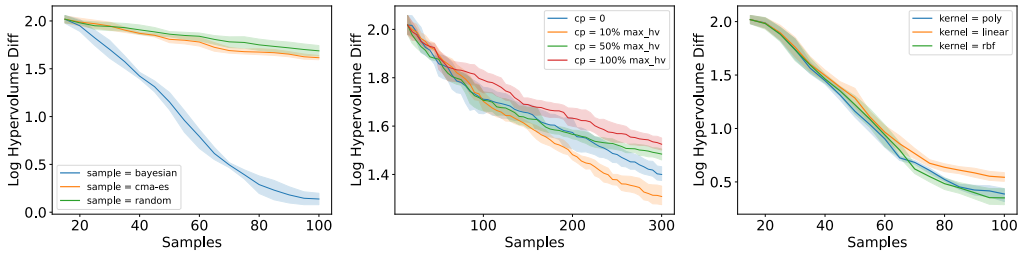

Figure 7: Ablation studies on hyperparameters and sampling methods in LaMOO. **Left**: Sampling without Bayesian/CMA-ES. **Middle**: Sampling with different C_p. **Right**: Partitioning with different svm kernels

Ablation of Design Choices. We show how different hyperparameters and sampling methods play a role in the performance. We perform the study in VehicleSafety below.

Sampling methods. LaMOO can be integrated with different sample methods, including Bayesian Optimization (e.g., qEHVI) and evolutionary algorithms (e.g., CMA-ES). Fig. 7(left) shows that compared to random sampling, qEHVI improves a lot while CMA-ES only improves slightly. This is consistent with our previous finding that for MOO, BO is much more efficient than EA.

The exploration factor C_p controls the balance of exploration and exploitation. A larger C_p guides LaMOO to visit the sub-optimal regions more often. Based on the results in Fig. 7(middle), greedy search ($C_p = 0$) leads to worse performance compared to a proper C_p value (i.e. 10% of maximum hypervolume), which justifies our usage of MCTS. On the other hand, over-exploration can also yield even worse results than greedy search. Therefore, a "rule of thumb" is to set the C_p to be roughly 10% of the maximum hypervolume $HV_{\max}$. When $HV_{\max}$ is unknown, C_p can be set empirically.

SVM kernels. As shown in Fig. 7(right), we find that the RBF kernel performs the best, in agreement with (Wang et al., 2020). Thanks to the non-linearity of the polynomial and RBF kernels, their region partitions perform better compared to a linear one.

6 Conclusion

We propose a search space partition optimizer called LaMOO as a meta-algorithm that extends prior single-objective works (Wang et al., 2020; Yang et al., 2021b) to multi-objective optimization. We demonstrated both theoretically and via experiments on multiple MOO tasks that LaMOO significantly improves the search performance compared to strong baselines like qEHVI and CMA-ES. In particular, LaMOO + qEHVI achieves state-of-the-art results on a variety of tasks including synthetic functions, neural architecture search, and molecule discovery.

7 REPRODUCIBILITY STATEMENT

Anonymous code is available at https://drive.google.com/drive/folders/1CMdg5iBdbKe3nkboIjiS998rnBEV09EB?usp=sharing. A full theoretical proof of LaMOO can be found in Appendix A.

REFERENCES

Sanghamitra Bandyopadhyay, Sankar K Pal, and B Aruna. Multiobjective gas, quantitative indices, and pattern classification. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 34(5):2088–2099, 2004.

Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Max-value entropy search for multi-objective bayesian optimization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper/2019/file/82edc5c9e21035674d481640448049f3-Paper.pdf.

Nicola Beume, Boris Naujoks, and Michael Emmerich. Sms-emoa: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181(3):1653–1669, 2007. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2006.08.008. URL https://www.sciencedirect.com/science/article/pii/S0377221706005443.

Peter AN Bosman and Dirk Thierens. The balance between proximity and diversity in multiobjective evolutionary algorithms. *IEEE transactions on evolutionary computation*, 7(2):174–188, 2003.

Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2019. URL https://arxiv.org/pdf/1812.00332.pdf.

Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once for all: Train one network and specialize it for efficient deployment. In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/pdf/1908.09791.pdf.

Kuang-Hua Chang. Chapter 5 - multiobjective optimization and advanced topics. In Kuang-Hua Chang (ed.), *Design Theory and Methods Using CAD/CAE*, pp. 325–406. Academic Press, Boston, 2015. ISBN 978-0-12-398512-5. doi: https://doi.org/10.1016/B978-0-12-398512-5.00005-0. URL https://www.sciencedirect.com/science/article/pii/B9780123985125000050.

Ivo Couckuyt, Dirk Deschrijver, and Tom Dhaene. Fast calculation of multiobjective probability of improvement and expected improvement criteria for pareto optimization. *Journal of Global Optimization*, 60(3):575–594, 2014.

Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable expected hypervolume improvement for parallel multi-objective bayesian optimization. *arXiv preprint arXiv:2006.05078*, 2020.

K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002a. doi: 10.1109/4235.996017.

K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable multi-objective optimization test problems. In *Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No.02TH8600)*, volume 1, pp. 825–830 vol.1, 2002b. doi: 10.1109/CEC.2002.1007032.

Kalyanmoy Deb and Himanshu Jain. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: Solving problems with box constraints. *IEEE Transactions on Evolutionary Computation*, 18(4):577–601, 2014. doi: 10.1109/TEVC.2013.2281535.

Guoqiang Deng, Zhangcan Huang, and Min Tang. Research in the performance assessment of multi-objective optimization evolutionary algorithms. In *2007 International Conference on Communications, Circuits and Systems*, pp. 915–918. IEEE, 2007.

Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. In *International Conference on Learning Representations (ICLR)*, 2020. URL https://openreview.net/forum?id=HJxyZkBKDr.

Xi Fang, Wenwen Wang, Lang He, Zhangcan Huang, Yang Liu, and Liang Zhang. Research on improved nsga-ii algorithm and its application in emergency management. *Research on Improved NSGA-II Algorithm and Its Application in Emergency Management*, 2018. doi: 10.1155/2018/1306341.

M. Fleischer. The measure of pareto optima applications to multi-objective metaheuristics. In Carlos M. Fonseca, Peter J. Fleming, Eckart Zitzler, Lothar Thiele, and Kalyanmoy Deb (eds.), *Evolutionary Multi-Criterion Optimization*, pp. 519–533, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.

Nyoman Gunantara. A review of multi-objective optimization: Methods and its applications. *Cogent Engineering*, 5(1):1502242, 2018. doi: 10.1080/23311916.2018.1502242.

Oliver Hinder. Cutting plane methods can be extended into nonconvex optimization. In *Conference On Learning Theory*, pp. 1451–1454. PMLR, 2018.

Christian Igel, Nikolaus Hansen, and Stefan Roth. Covariance matrix adaptation for multi-objective optimization. *Evolutionary Computation*, 15(1):1–28, 2007a. doi: 10.1162/evco.2007.15.1.1.

Christian Igel, Thorsten Suttorp, and Nikolaus Hansen. Steady-state selection and efficient covariance matrix update in the multi-objective cma-es. In Shigeru Obayashi, Kalyanmoy Deb, Carlo Poloni, Tomoyuki Hiroyasu, and Tadahiko Murata (eds.), *Evolutionary Multi-Criterion Optimization*, pp. 171–185, Berlin, Heidelberg, 2007b. Springer Berlin Heidelberg.

Siwei Jiang, Yew-Soon Ong, Jie Zhang, and Liang Feng. Consistencies and contradictions of performance metrics in multiobjective optimization. *IEEE transactions on cybernetics*, 44(12): 2391–2404, 2014.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Hierarchical generation of molecular graphs using structural motifs. In *International Conference on Machine Learning*, pp. 4839–4848. PMLR, 2020a.

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

Donald R Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13(4):455–492, 1998.

Beomjoon Kim, Kyungjae Lee, Sungbin Lim, Leslie Kaelbling, and Tomás Lozano-Pérez. Monte carlo tree search in continuous spaces using voronoi optimistic optimization with regret bounds. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 9916–9924, 2020.

Taehyeon Kim, Jaeyeon Ahn, Nakyil Kim, and Seyoung Yun. Adaptive local bayesian optimization over multiple discrete variables. URL https://valohaichirpprod.blob.core.windows.net/papers/kaist_osi.pdf.

J. Knowles. Parego: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1):50–66, 2006. doi: 10.1109/TEVC.2005.851274.

Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In Johannes Fürnkranz, Tobias Scheffer, and Myra Spiliopoulou (eds.), *Machine Learning: ECML 2006*, pp. 282–293, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. ISBN 978-3-540-46056-5.

Alex Krizhevsky. Learning multiple layers of features from tiny images. 2009.

H. T. Kung, F. Luccio, and F. P. Preparata. On finding the maxima of a set of vectors. *J. ACM*, 22(4):469–476, October 1975. ISSN 0004-5411. doi: 10.1145/321906.321910. URL https://doi.org/10.1145/321906.321910.

Miqing Li, Shengxiang Yang, and Xiaohui Liu. Diversity comparison of pareto front approximations in many-objective optimization. *IEEE Transactions on Cybernetics*, 44(12):2568–2584, 2014.

Xingtao Liao, Qing Li, Xujing Yang, Weigang Zhang, and Wei Li. Multiobjective optimization for crash safety design of vehicles using stepwise regression model. *Structural and multidisciplinary optimization*, 35(6):561–569, 2008.

GV Loganathan and Hanif D Sherali. A convergent interactive cutting-plane algorithm for multiobjective optimization. *Operations Research*, 35(3):365–377, 1987.

Rémi Munos. Optimistic optimization of a deterministic function without the knowledge of its smoothness. *Advances in neural information processing systems*, 24:783–791, 2011.

Lie Meng Pang, Hisao Ishibuchi, and Ke Shang. Nsga-ii with simple modification works well on a wide variety of many-objective problems. *IEEE Access*, 8:190240–190250, 2020. doi: 10.1109/ACCESS.2020.3032240.

Biswajit Paria, Kirthevasan Kandasamy, and Barnabás Póczos. A flexible multi-objective bayesian optimization approach using random scalarizations. *CoRR*, abs/1805.12168, 2018. URL http://arxiv.org/abs/1805.12168.

Wolfgang Ponweiser, Tobias Wagner, Dirk Biermann, and Markus Vincze. Multiobjective optimization on a limited budget of evaluations using model-assisted-metric selection. In Günter Rudolph, Thomas Jansen, Nicola Beume, Simon Lucas, and Carlo Poloni (eds.), *Parallel Problem Solving from Nature – PPSN X*, pp. 784–794, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

Mikita Sazanovich, Anastasiya Nikolskaya, Yury Belousov, and Aleksei Shpilman. Solving black-box optimization challenge via learning search space partition for local bayesian optimization. URL https://valohaichirpprod.blob.core.windows.net/papers/jetbrains.pdf.

Nidamarthi Srinivas and Kalyanmoy Deb. Muiltiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary computation*, 2(3):221–248, 1994.

William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.

Jessica Vamathevan, Dominic Clark, Paul Czodrowski, Ian Dunham, Edgardo Ferran, George Lee, Bin Li, Anant Madabhushi, Parantu Shah, Michaela Spitzer, et al. Applications of machine learning in drug discovery and development. *Nature Reviews Drug Discovery*, 18(6):463–477, 2019.

David A Van Veldhuizen and Gary B Lamont. Evolutionary computation and convergence to a pareto front. In *Late breaking papers at the genetic programming 1998 conference*, pp. 221–228. Citeseer, 1998.

Douglas AG Vieira and Adriano Chaves Lisboa. A cutting-plane method to nonsmooth multiobjective optimization problems. *European Journal of Operational Research*, 275(3):822–829, 2019.

Linnan Wang, Saining Xie, Teng Li, Rodrigo Fonseca, and Yuandong Tian. Sample-efficient neural architecture search by learning action space. *CoRR*, abs/1906.06832, 2019. URL http://arxiv.org/abs/1906.06832.

Linnan Wang, Rodrigo Fonseca, and Yuandong Tian. Learning search space partition for black-box optimization using monte carlo tree search. *arXiv preprint arXiv:2007.00708*, 2020.

Qian Xu, Zhanqi Xu, and Tao Ma. A survey of multiobjective evolutionary algorithms based on decomposition: Variants, challenges and future directions. *IEEE Access*, 8:41588–41614, 2020. doi: 10.1109/ACCESS.2020.2973670.

Kaifeng Yang, Michael Emmerich, André Deutz, and Thomas Bäck. Multi-objective bayesian global optimization using expected hypervolume improvement gradient. *Swarm and Evolutionary Computation*, 44:945–956, 2019. ISSN 2210-6502. doi: https://doi.org/10.1016/j.swevo.2018.10.007. URL https://www.sciencedirect.com/science/article/pii/S2210650217307861.

Kevin Yang, Tianjun Zhang, Chris Cummins, Brandon Cui, Benoit Steiner, Linnan Wang, Joseph E. Gonzalez, Dan Klein, and Yuandong Tian. Learning space partitions for path planning. *CoRR*, abs/2106.10544, 2021a. URL https://arxiv.org/abs/2106.10544.

Kevin Yang, Tianjun Zhang, Chris Cummins, Brandon Cui, Benoit Steiner, Linnan Wang, Joseph E Gonzalez, Dan Klein, and Yuandong Tian. Learning space partitions for path planning. *arXiv preprint arXiv:2106.10544*, 2021b.

Kaicheng Yu, Christian Sciuto, Martin Jaggi, Claudiu Musat, and Mathieu Salzmann. Evaluating the search phase of neural architecture search. *arXiv preprint arXiv:1902.08142*, 2019.

Qingfu Zhang and Hui Li. Moea/d: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731, 2007. doi: 10.1109/TEVC.2007.892759.

E. Zitzler, L. Thiele, M. Laumanns, C.M. Fonseca, and V.G. da Fonseca. Performance assessment of multiobjective optimizers: an analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, 2003. doi: 10.1109/TEVC.2003.810758.

Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary computation*, 8(2):173–195, 2000.