
Monte Carlo Tree Search based Space Transfer for Black-box Optimization

Shukuan Wang^{1,2*}, Ke Xue^{1,2*}, Lei Song^{1,2}, Xiaobin Huang^{1,2}, Chao Qian^{1,2†}

¹National Key Laboratory for Novel Software Technology, Nanjing University

²School of Artificial Intelligence, Nanjing University

{wangsk, xuek, huangxb, songl, qianc}@lamda.nju.edu.cn

Abstract

Bayesian optimization (BO) is a popular method for computationally expensive black-box optimization. However, traditional BO methods need to solve new problems from scratch, leading to slow convergence. Recent studies try to extend BO to a transfer learning setup to speed up the optimization, where search space transfer is one of the most promising approaches and has shown impressive performance on many tasks. However, existing search space transfer methods either lack an adaptive mechanism or are not flexible enough, making it difficult to efficiently identify promising search space during the optimization process. In this paper, we propose a search space transfer learning method based on Monte Carlo tree search (MCTS), called MCTS-transfer, to iteratively divide, select, and optimize in a learned subspace. MCTS-transfer can not only provide a well-performing search space for warm-start but also adaptively identify and leverage the information of similar source tasks to reconstruct the search space during the optimization process. Experiments on synthetic functions, real-world problems, Design-Bench and hyper-parameter optimization show that MCTS-transfer can demonstrate superior performance compared to other search space transfer methods under different settings. Our code is available at <https://github.com/lamda-bbo/mcts-transfer>.

1 Introduction

In many real-world tasks such as neural architecture search [56, 41, 42], hyper-parameter optimization [55, 4], and integrated circuit design [19, 30], we often need to solve black-box optimization (BBO) problems, where the objective function has no analytical form and can only be evaluated by different inputs, regarded as a “black-box” function. BBO problems are often accompanied by expensive computational costs of the evaluations, requiring a BBO algorithm to find a good solution with a small number of objective function evaluations.

Bayesian optimization (BO) [29, 8] is a widely used sample-efficient method for such problems. At each iteration, BO fits a surrogate model, typically Gaussian process (GP) [26], to approximate the objective function, and maximizes an acquisition function to determine the next query point. Under the limited evaluation budget, traditional BO methods only have a few observations, which are, however, insufficient for constructing a precise surrogate model, leading to slow convergence. Thus, traditional BO methods struggle to effectively solve expensive BBO problems, preventing their broader applications.

To tackle this issue, recent research tries to apply transfer learning methods for BO [2]. Transfer BO methods operate under the assumption that similar tasks may share common characteristics, and

*Equal Contribution

†Corresponding Author

the knowledge acquired from similar source tasks can be helpful for optimizing the current target task. They typically gather offline datasets from the source tasks and utilize them to expedite the target task. These methods can be categorized based on the learned components of BO, such as the acquisition function [51, 40], initialization point [45, 47], or search space [49, 23, 17].

Among these, learning the search space is a promising research area due to its effectiveness and orthogonal relationship with optimization methods. By learning and partitioning the search space, we can more effectively utilize potential subspaces and significantly accelerate the algorithm's search for optimal solutions. While the methods for partitioning the search space have demonstrated great potential [22, 14, 23, 17], they still have several limitations, particularly in terms of flexibly adjusting the search space for the current target task. In many scenarios, we are uncertain about which tasks are truly "similar" to the target task before optimizing it. Our comprehension about this similarity can only deepen gradually as we progress in optimizing the target task. Therefore, we hope that a search space transfer algorithm can automatically identify the most relevant source tasks during the optimization process, and give them more consideration when constructing the subspaces. However, existing methods have limited flexibility in adjusting in this manner.

In this paper, we propose a search space transfer learning method based on Monte Carlo tree search (MCTS), called MCTS-transfer, to iteratively divide, select, and optimize in a learned subspace. Each node of MCTS-transfer represents a subspace, whose potential is calculated as a weighted sum of the source and target sample values, assessing the node's utilization value. To better identify and leverage the information from source tasks, we assign different weights to different source tasks based on their similarity to the target task, which are adjusted dynamically during the optimization process. These weights are then used to calculate the potential of the node. Our proposed MCTS-transfer offers several notable advantages. First, it can provide a better initial search space for a warm-start of the optimization of the target task. Second, it can provide multiple promising compact subspaces by MCTS, to improve optimization efficiency. The upper confidence bound (UCB)-based node selection of MCTS can also automatically balance the trade-off between exploration and exploitation. Third, it can automatically identify source tasks similar to the current target task based on new observations, and re-construct the search space to make it more consistent with the current target task.

To evaluate the effectiveness of the proposed method, we compare MCTS-transfer with various search space transfer methods and conduct experiments on multiple BBO tasks, including synthetic functions, real-world problems, Design-Bench and hyper-parameter optimization problems. In different scenarios, such as varying similarity between the source tasks and target task, MCTS-transfer demonstrates superior performance. We also analyze the effectiveness of the adaptive weight adjustment, showing that MCTS-transfer can identify source tasks that are similar to the target task and assign them higher weights accordingly. Note that MCTS-transfer can be combined with any BBO algorithm. We only implemented it with the basic BO algorithm (i.e., using GP as the surrogate model and expected improvement (EI) as the acquisition function) in the experiments, and the performance can be further enhanced by advanced techniques.

2 Background

2.1 Bayesian Optimization

We consider the problem $\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$, where f is a black-box function and $\mathcal{X} \subseteq \mathbb{R}^D$ is the domain. The basic framework of BO contains two critical components: a surrogate model and an acquisition function. GP is the most popular surrogate model. Given the sampled data points $\{(\mathbf{x}_i, y_i)\}_{i=1}^{t-1}$, where $y_i = f(\mathbf{x}_i) + \epsilon_i$ and $\epsilon_i \sim \mathcal{N}(0, \eta^2)$ is the observation noise, GP at iteration t seeks to infer $f \sim \mathcal{GP}(\mu(\cdot), k(\cdot, \cdot) + \eta^2 \mathbf{I})$, specified by the mean $\mu(\cdot)$ and covariance kernel $k(\cdot, \cdot)$, where \mathbf{I} is the identity matrix of size D . After that, an acquisition function, e.g., probability of improvement (PI) [15], EI [12] or UCB [36], is optimized to determine the next query point \mathbf{x}_t , balancing exploration and exploitation.

2.2 Transfer Bayesian Optimization

When solving new BBO problems, traditional BO methods need to conduct optimization from scratch, leading to slow convergence. Transfer learning reuses knowledge from source tasks to boost the performance of current tasks, and thus can be naturally applied to address this issue, i.e., it can utilize

source task data to accelerate the convergence of current target task [2]. The main assumption of transfer learning for BBO is that many real-world problems exhibit certain similarities, and similar tasks often share similar characteristics that can be exploited.

Various transfer learning approaches have been explored concerning the surrogate model, acquisition function, initialization, and search space of BO. Several methods learn all available information from both source and target tasks in a single GP surrogate, and make the data comparable through multi-task GP [37, 16, 38], noisy GP model [31, 25, 13], and deep kernel learning model [47, 11, 10]. Additionally, certain approaches involve training multiple base surrogates and then combining them into a single surrogate [50, 28].

Transfer learning methods on surrogate models, however, may encounter difficulties when scaling to high-dimensional cases, because the influence of the source tasks is easy to diminish progressively with the increase of new observation points [2]. Alternatively, transferring knowledge via the acquisition function can circumvent these issues. Previous research has approached transfer learning for acquisition functions through multi-task BO [37, 20], ensemble GPs [51], and meta-learning strategies such as reinforcement learning-driven strategies [40].

Several approaches consider selecting multiple valuable points for better initial evaluations for the warm-start of optimization. Feurer et al. [7] simply selected the best point from the t most similar tasks as the initialization point. Other works [48, 45] achieved warm-start by constructing a meta-loss, organically combining the mean function of the GP model from source tasks, and using gradient descent to find a set of solutions that minimize the meta-loss. Wistuba and Grabocka [47] tried to find suitable initial points to warm-start BO by minimizing the loss on the source tasks using an evolutionary algorithm.

2.3 Search Space Transfer

Compared to the aforementioned methods, search space transfer has many advantages, especially in not limiting the transfer to a specific algorithm component (e.g., acquisition function in BO), but considering the search space that can be used by all BBO algorithms as the transfer object. That is, the learned search space is orthogonal to the optimization process and can be integrated with any advanced optimizer. A well-learned search space can greatly improve optimization efficiency and guide the optimization process to a good result.

The concept of search space transfer was first proposed by Wistuba et al. [49], which defined a region by a center point and a diameter, and pruned away regions deemed less promising. Instead of pruning space, Perrone and Shen [23] considered designing a promising search space for the target task. This approach extracts an optimal solution x_i^* from each task, and employs a simple geometric form (e.g., a box or ellipsoid) to define the smallest subspace encompassing all optimal solutions x_i^* . However, it ignores the correlation between tasks, and the space constructed with regular geometry may be too loose for the target task. To address this issue, Li et al. [17] selected the most similar source tasks to the target task in a certain proportion, and used a binary classification method to learn good spaces and bad spaces among these tasks. A voting mechanism is then employed to aggregate information from all selected tasks to determine a newly generated search space for the target task.

These current space transfer methods, however, lack mechanisms to adjust the search space when it is not well-suited for the target task. Furthermore, when source tasks differ significantly from the target task, they tend to devolve into full-space search. The binary nature of evaluating the search space as simply “good” or “bad” is another limitation. Our proposed method will overcome these drawbacks by adapting the search space dynamically according to the similarity between source tasks and the target task, applying MCTS to find a proper subspace, and using a more nuanced and precise numerical evaluation for the goodness of a search subspace.

2.4 Monte Carlo Tree Search

MCTS [5] is a search algorithm combining random sampling with a tree search structure, which has been widely applied in the field of game-playing and decision-making [32, 34]. A tree node represents a particular state in the search space, e.g., stone positions on the board in a GO game. Each tree node has an UCB [1] value during the search procedure, to balance exploration and exploitation.

The UCB of each tree node is calculated by its value v_m and its visit times n_m , defined as:

$$\text{ucb}_m = \frac{v_m}{n_m} + 2C_p \sqrt{2 \log(n_p) / n_m}, \quad (1)$$

where C_p is a hyper-parameter controlling the balance between exploration and exploitation, and n_p is the visit times of the parent node. In the iterative process of MCTS, a leaf node is systematically selected for expansion, involving four key steps: selection, expansion, simulation, and back-propagation. Initially, the algorithm recursively navigates from the root node to child nodes, prioritizing those with higher UCB values to identify the leaf node m . Subsequently, an action is executed based on the state of m , leading to the expansion of a new child node (state), k . In the simulation step, the node value v_k is determined through random sampling. Finally, through back-propagation, the value and visitation count of the ancestors of k are updated.

MCTS has been used to select important variables automatically for high dimensional BO [35]. LA-MCTS [43] is a scalable BBO algorithm based on learning space partition. Utilizing MCTS, it iteratively divides the search space into small subspaces for optimization. In this framework, the tree's root represents the entire search space denoted as Ω , and each tree node m represents a sub-region Ω_m . The value v_m is determined by the average objective value of the sampled points within the sub-region Ω_m . During each iteration, after selecting a leaf node m , LA-MCTS conducts optimization within Ω_m . The sampled points are then employed for clustering and classification, leading to the bifurcation of Ω_m into two distinct sub-regions: "good" and "bad". These sub-regions are expanded as two child nodes, with the left one denoted as "good" and the right as "bad".

3 MCTS-transfer

In this section, we propose a search space transfer learning method based on MCTS, called MCTS-transfer, which can be divided into two major stages: pre-learning stage and optimizing stage. The main idea is to apply MCTS to divide the entire space based on source task data in the pre-learning stage, and adaptively adjust the partition based on newly generated target task data during the optimization process. MCTS-transfer iteratively chooses one partitioned subspace for search and adjusts the partition after sampling each new data point. Assume that there are K source tasks $\{f_i\}_{i=1}^K$ and a target task T that we are going to optimize $\max_{\mathbf{x} \in \Omega} f_T(\mathbf{x})$, where Ω is the entire search space. For the i -th source task, we have offline dataset $D_i = \{(\mathbf{x}_{i,j}, y_{i,j})\}_{j=1}^{|D_i|}$, where $y_{i,j}$ is the observed objective value of $\mathbf{x}_{i,j}$, and $|D_i|$ is the number of data points. Let tree \mathcal{T} denotes the search space division process by MCTS.

In the pre-learning stage, \mathcal{T} initially has only one root node, representing the entire search space Ω . The expansion of a tree node m corresponds to the division of the search space Ω_m that the node m represents. The node expansion follows the rules below. With a set of source task samples $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ in the space Ω_m , we use k -means clustering to divide them into two groups. The cluster with a higher average objective value is regarded as the "good" cluster, while the other is the "bad" one. A binary classifier then establishes a decision boundary between the two clusters. In our approach, the space associated with the good cluster becomes the left node, while the space of the bad cluster becomes the right node. We recursively apply this partition process within each node, as shown in Figure 1. The depth of the resulting Monte Carlo tree \mathcal{T} is determined by a threshold hyperparameter θ : a node is divided if it contains more than θ samples and the contained samples can be clustered into two clusters. The tree \mathcal{T} generated at this stage can give a suitable space partition in advance based on source task data, which serves as a warm-start for the following optimization process. Details of pre-learning will be introduced in Section 3.1.

In the optimization process, we follow the four key steps of MCTS: selection, expansion, simulation, and back-propagation. At each iteration, we select a target node m by tracing the nodes' UCB values. That is, starting from the root, we recursively choose the child node with higher UCB value, until a leaf node m is reached, as Figure 1 displayed. The space Ω_m represented by m is then considered as a promising search space, where a BO optimizer is used to optimize. The BO optimizer can build a surrogate model using samples in either Ω_m or Ω . The newly sampled point will be evaluated and used for expanding the node m after updating the clustering in the node. It will further be utilized in back-propagation, where the number of visits and the potential value of each node will be updated. The potential value of a node is calculated by a weighted sum of objective values of the source and target task samples. Note that the tree structure will be reconstructed if there exists a node that the

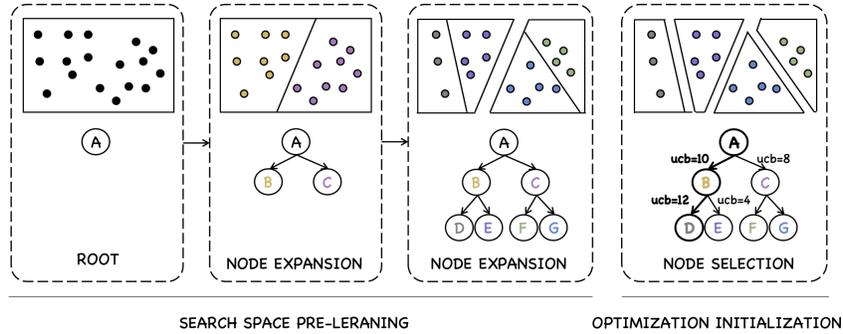


Figure 1: The workflow of MCTS-transfer. In pre-learning stage, MCTS-transfer builds the tree by clustering and classifying the samples apart recursively, until all nodes are not splittable. In optimization stage, the initial search space is based on the pre-learned tree. We will trace child node with greater UCB from ROOT and find the target leaf node to do optimization.

potential value of its right child exceeds that of the left one, violating the rule of our tree construction and implying that the current structure is not good. Please see Sections 3.2 and 3.3 for the details of node potential update and tree structure reconstruction.

Unlike LA-MCTS, MCTS-transfer not only considers the information of the target task when calculating node values but also takes into account the information from source tasks, enabling it to be used for search space transfer. Additionally, we propose adaptive weights for source tasks and validate their effectiveness through our experiments.

3.1 Search Space Pre-Learning

Compared to standard BBO algorithms, which sample randomly across the entire space at the beginning, our method leverages information from source tasks to concentrate sampling within a generally “good” space, providing a warm-start initialization.

In the pre-learning stage, tree \mathcal{T} initially has only one node, i.e., ROOT node. All source task samples are collected in this node, recursively clustered and classified, leading to node expansion. When none of the leaf nodes can be further bifurcated, i.e., contains more than θ samples and can be clustered into two clusters, \mathcal{T} is finally formed. In this process, we keep the rule that the left node is “better” (i.e., has a larger potential value) than the right node, so that we can easily identify that the leftmost leaf is the best and the rightmost leaf is the worst. The potential value of a node m in the pre-learning stage is estimated by the average $y_{i,j}$ of all source task points within it, i.e.,

$$p_m = \frac{\sum_{i \leq K} \sum_{(\mathbf{x}_{i,j}, y_{i,j}) \in D_i \cap \Omega_m} y_{i,j}}{\sum_{i \leq K} |D_i \cap \Omega_m|}. \quad (2)$$

This tree embeds historically good regions into specific tree nodes. Intuitively, the spaces represented by left leaves have higher potential of being good spaces than right ones.

The first iteration of MCTS-transfer will utilize the resulting tree \mathcal{T} generated in the pre-learning stage. Starting from the root node, it recursively selects nodes with higher UCB values until reaching a leaf node. The UCB value is calculated as Eq. (1), where v_m is replaced by the potential value p_m in Eq. (2). In MCTS-transfer, n_m and n_p represent the number of samples, including those from source and target tasks, in the subspace Ω_m and the parent subspace of Ω_m , respectively. Consequently, the algorithm preferentially samples from those regions that have shown to yield favorable outcomes in the source tasks.

3.2 Node Potential Update

In each iteration of MCTS-transfer, after evaluating a new sample, we use it to update the similarity between the source tasks and the target task, and then update each node potential. Note that the potential value p_m used in the pre-learning stage only contains information from source tasks. But in

the optimizing stage, the calculation of p_m includes both information from the source tasks $\{D_i\}_{i=1}^K$ and the target task $D_T = \{(\mathbf{x}_{T,j}, y_{T,j})\}_{j=1}^t$, and considers the similarity between them, i.e.,

$$p_m = \gamma^{t-1} \frac{\sum_{i \leq K} w_i \bar{y}_{i,m}}{\sum_{i \leq K} w_i} + \bar{y}_{T,m}, \quad (3)$$

where $\bar{y}_{i,m} = \sum_{(\mathbf{x}_{i,j}, y_{i,j}) \in D_i \cap \Omega_m} y_{i,j} / |D_i \cap \Omega_m|$, and $\bar{y}_{T,m} = \sum_{(\mathbf{x}_{T,j}, y_{T,j}) \in D_T \cap \Omega_m} y_{T,j} / |D_T \cap \Omega_m|$, which are the average objective values of the samples of D_i and D_T in Ω_m , respectively. There are two important parameters in Eq. (3).

- γ : A decay factor, which is used to adjust the overall influence of source tasks throughout the optimization process. A smaller γ accelerates the forgetting of the source task data.
- w_i : A weighting factor, which reflects the influence of the i -th source task, and is determined by the similarity between the i -th source task and the target task. A larger weight represents higher similarity and implies a greater influence of the i -th source task's samples on the potential calculation of a tree node.

To calculate w_i , we measure the distance $Distance(D_i, D_T)$ between the i -th source task and the target task by $Distance(\mathbf{x}_i^*, \mathbf{x}_T^*)$, where \mathbf{x}_i^* and \mathbf{x}_T^* denote the mean of the best N sampled points of these two tasks, respectively. The source tasks are ranked according to their distances to the target task. A smaller rank implies a smaller distance, i.e., a higher similarity. The weight w_i is then calculated as

$$w_i = \begin{cases} 1.0 - \frac{r_i}{\alpha N_m} & \text{if } r_i < \alpha N_m \\ 0.1 & \text{otherwise} \end{cases}, \quad (4)$$

where r_i is the rank of the i -th source task, N_m is the number of source tasks that have solutions located in Ω_m , and $\alpha \in [0, 1]$. After sampling a new data point in each iteration of MCTS-transfer, the ranks and weights are updated accordingly, which are then used to update the potential value of each node. We also consider other ways of calculating distances and weights, which are introduced and empirically compared in Appendix C.

3.3 Tree Structure Reconstruction

When constructing the search tree, the left child of a node is always better than the right child, i.e., the potential of the left child always exceeds that of the right child. However, after sampling a new data point and updating the potential of each node in each iteration of MCTS-transfer, this property may be violated. If this happens, it indicates that the current space partition is not ideal and needs adjustment. Specifically, we backtrack from the problematic leaf nodes to the highest-level ancestor node that upholds the desired property, and then proceed to reconstruct the subtree from that ancestor node. The subtree reconstruction process is consistent with the process of node expansion.

The detailed process is presented in Algorithm 2. We apply breadth-first search to traverse all tree nodes and use a queue \mathcal{Q} to manage the sequence of nodes. In addition, we set a queue \mathcal{N} to store the nodes that need to be reconstructed. If the potential of the right child of a node is better than that of the left child (line 6), the subtree of this node is deleted (line 7), and it should be reconstructed and is added into the queue \mathcal{N} (line 8). Otherwise, the two child nodes will enter into the queue \mathcal{Q} (lines 10–11) and will be examined later. After traversing the tree \mathcal{T} , we reconstruct the subtrees of nodes in \mathcal{N} (lines 15–19). If a node in \mathcal{N} is splittable, i.e., the contained samples in the node exceeds θ and can be clustered and divided by a binary classifier, it will be further expanded. Thus, Treeify can fine-tune the tree structure and reserve some history information; meanwhile, it can adaptively be more suitable to the target task.

3.4 Details of MCTS-transfer

The detailed procedure of MCTS-transfer is presented in Algorithm 1. It begins with a pre-learned MCTS model \mathcal{T} based on source task data (line 1), and initializes an empty set D_T to store samples of the target task (line 2). In each iteration, it selects a leaf node m based on the UCB value (line 4). If the target task already has samples in the space Ω_m represented by m , the algorithm conducts BO within Ω_m and gets a candidate point $\mathbf{x}_{T,t}$ (lines 8–9); otherwise, the candidate point is randomly selected in Ω_m (line 6). The sampled point $\mathbf{x}_{T,t}$ is then evaluated and added into D_T (line 11). To

Algorithm 1 MCTS-transfer

Input: Search space Ω , objective function $f_T(\cdot)$, data $\{D_i\}_{i=1}^K$ of K source tasks, and number N of objective evaluations

```
1: Initialize a MCTS model  $\mathcal{T}$  by search space pre-learning as introduced in Section 3.1;
2:  $D_T = \emptyset$ ;
3: for  $t = 1, 2, \dots, N$  do
4:   Select a leaf node  $m$  by UCB-based selection;
5:   if  $|D_T| = 0$  then
6:     Select a candidate point  $\mathbf{x}_{T,t}$  randomly in  $\Omega_m$ 
7:   else
8:     Train a GP model on  $D_T$ ;
9:     Select a candidate  $\mathbf{x}_{T,t}$  in  $\Omega_m$  by optimizing an acquisition function
10:  end if
11:  Evaluate  $\mathbf{x}_{T,t}$  to get  $y_{T,t}$ , and let  $D_T = D_T \cup \{(\mathbf{x}_{T,t}, y_{T,t})\}$ ;
12:  Calculate  $Distance(D_i, D_T)$ , and sort them in ascending order to get the ranks  $r_i$ ;
13:  For each source task  $f_i$ , update its weight  $w_i$  by Eq. (4), where  $i = 1, \dots, K$ ;
14:  For each node  $n$ , update its potential value  $p_n$  by Eq. (3);
15:  if  $m.isSplitable$  then
16:    Expand  $m$ 
17:  end if
18:  Back-propagate and update the visit times and samples of nodes on the path to  $m$ ;
19:   $\mathcal{T} \leftarrow \text{Treeify}(\mathcal{T})$ 
20: end for
```

make full use of the information from the source tasks, we take the similarity between a source task and the target task into account, and try to assign higher weights to more similar source tasks, as introduced in Section 3.2. With new samples added in D_T , the distance between each source task D_i and the target task D_T may change. Thus, the distance, rank and weight of each source task are re-calculated in lines 12–13 according to Eq. (4), and the potential value of each node in the tree is re-evaluated in line 14 according to Eq. (3). After that, the node m is expanded in lines 15–17 if it is splitalle, i.e., contains more than θ samples of the target task and can be clustered into two clusters. Note that for node expansion, only samples of the target task are considered. In line 18, the back-propagation step is performed to refresh the information of each node on the path to m , including the visit times and the contained samples. Finally, we will check whether the updated tree conforms to the property that the left child of a node has a larger potential value than its right node. For any node violating this property, its sub-tree will be reconstructed. This process is accomplished by employing the Treeify procedure in Algorithm 2 in Appendix A.

In MCTS-transfer, the tree structure exhibits several properties that align with the requirements of search space transfer learning. 1) Tree structure is natural to model the search space, where the root represents the entire search space, and the node expansion corresponds to the space partition with each child node representing a subspace. 2) The node potential evaluation allows for the extraction of multiple irregularly shaped promising search subspaces represented by multiple leaf nodes. 3) The tree can be adjusted by updating and expanding nodes on the basis of the original tree, so that new information can be absorbed while historical information can also be retained, i.e., the tree structure is inheritable. This enables the knowledge of space partition from source task data to be transferred to the target task.

4 Experiments

In this section, we introduce a simple case to highlight the features of existing search space transfer algorithms and assess the effectiveness of MCTS-transfer. We also conduct experiments across a variety of tasks, including the synthetic benchmark BBOB, real-world problems, Design-Bench and the hyper-parameter optimization benchmark HPOB. The details of the problems, data and experimental results can be seen in Appendix B.2, B.3 and E.

Compared methods. Our baselines consist of two non-transfer BO, i.e., basic GP, and LA-MCTS [43], three search space transfer algorithms, i.e., Box-GP [23], Ellipsoid-GP [23], Supervised-

GP [17]), and one state-of-the-art surrogate model transfer BO algorithm PFN [21]. Detailed configurations and settings of the hyper-parameters for each algorithm are provided in Appendix B.1.

Basic Settings. In our experimental settings, we explore two types of transfer learning—similar and mixed transfer—to demonstrate the advantages of MCTS-transfer in handling complex source tasks. Similar transfer involves using data from similar source tasks, while mixed transfer uses a combination of similar and dissimilar source task data for pre-training. Given the challenges in real-world scenarios of determining task similarity and selecting appropriate tasks, our focus will be on evaluating MCTS-transfer’s performance in the more complex scenario, i.e., mixed transfer.

4.1 Motivating Cases: Sphere2D

We first conduct a simple but important experiment on Sphere2D to verify our motivation and evaluate the performance of existing search space transfer methods. The Sphere2D function is defined as $f(\mathbf{x}) = (\mathbf{x} - \mathbf{x}^*)^2$, with \mathbf{x}^* representing the optimal solution. We generate three source task datasets $D_{(5,5)}$, $D_{(5,-5)}$, and $D_{(-5,-5)}$, each containing 100 samples, by assigning $(5, 5)$, $(5, -5)$, and $(-5, -5)$ to \mathbf{x}^* and sampled by standard BO. The target task is defined as $f(\mathbf{x})$ with $\mathbf{x}^* = (4, 4)$.

Unlike the basic setting, we use the mixed setting and the dissimilar setting. Dissimilar setting uses only $D_{(5,-5)}$ and $D_{(-5,-5)}$ for pre-training, which are most dissimilar to the target task. This setting is to test MCTS-transfer’s ability to handle extreme cases without similar source tasks.

MCTS-transfer-GP is compared against three search space transfer algorithms—Box-GP, Ellipsoid-GP, and Supervised-GP—in both mixed and dissimilar settings, as shown in Figure 2. Box-GP and Ellipsoid-GP ignore task similarity and limit the search space to geometric areas encompassing all source task optima. In the dissimilar transfer setting, these methods fail because the target task’s optimal solution $\mathbf{x}^* = (4, 4)$ lies outside the defined subspace of source task optima. Supervised-GP takes task similarity to guide the selection and combination of promising region with a voting system. Although it pays attention to the importance of task similarity and can be applicable in dissimilar transfer, it tends to be an entire-space optimization algorithm when source tasks exhibit high variance or are dissimilar to the target task. MCTS-transfer recursively divides the search space, dynamically adjusts task weights and evaluates the subspace potential based on the task similarity. In addition, the method tends to increasingly rely on target task data due to decay factor γ . Therefore, MCTS-transfer achieves successful convergence with low variance in mixed and dissimilar transfer. The right part of Figure 2 displays the weight change curves for the three source tasks, confirming the anticipated results. Higher weights are assigned to datasets $D_{(5,5)}$ and $D_{(-5,5)}$, indicating that our method effectively recognizes and prioritizes tasks more similar to the target task.

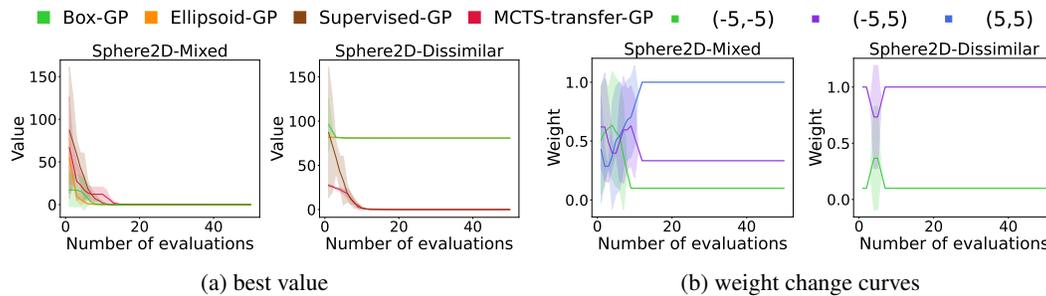


Figure 2: Comparison of MCTS-transfer and other algorithms on Sphere2D

4.2 Synthetic Functions

BBOB [9] is a popular benchmark for BBO. It offers 24 synthetic functions tailored for the continuous domain. We randomly select one function from each of the five function classes of BBOB as our target task, i.e., GriewankRosenbrock, Lunacek, Rastrigin, Rosenbrock, and SharpRidge.

As shown in Figure 3a, compared with LA-MCTS, we find that the results can indeed be greatly improved after space transfer, proving the effectiveness of MCTS-transfer. In mixed transfer, MCTS-transfer can still maintain warm-start, thanks to its ability to provide multiple compact promising

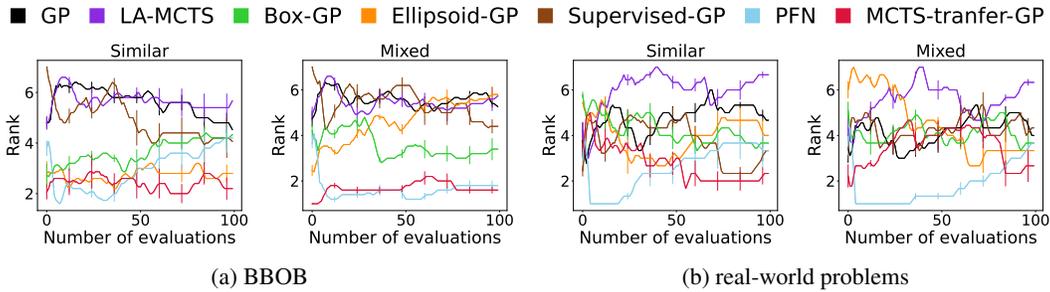


Figure 3: Comparison of MCTS-transfer and other algorithms on BBOB and real-world problems

subspaces by effective search space pre-learning. Throughout the optimization process, MCTS-transfer, combined with a simple GP, is able to reach or even surpass SOTA’s surrogate model transfer algorithm PFN, reflecting the superior performance.

4.3 Real-world Problems

To demonstrate the practical value of MCTS-transfer, we test MCTS-transfer with three real-world problems, including LunarLander, RobotPush, and Rover.

Figure 3b shows comparison on similar and mixed transfer. Based on the overall ranking, MCTS-transfer-GP outperforms other search space transfer baselines. In mixed transfer, although it’s surpassed by PFN initially, MCTS-transfer-GP is finally able to find better solutions even if there are dissimilar tasks misleading. The ability to adaptively correct the search space may come from the efficient utilization of source task data by reasonable node potential evaluation, node expansion, and tree reconstruction.

4.4 Design-Bench Problems

We further verify the performance of MCTS-transfer in more complex and high-dimensional problems from Design-Bench [39], including three continuous problems, Superconductor, Ant morphology and D’Kitty Morphology.

In similar and mixed transfer, MCTS-transfer-GP gets the best ranking after 40 iterations. The stable performance indicates that the algorithm can effectively transfer and utilize previous experience when facing different source tasks. This ability is particularly important for black-box problems because it allows the system to quickly adapt to unseen situations while maintaining high performance.

4.5 Hyper-parameter Optimization

Hyper-parameter optimization(HPO) is a common application of BBO transfer learning. The HPO-B benchmark [24], sourced from OpenML, includes 176 search spaces/algorithms and 196 datasets, totaling 6.4 million HPO evaluations. Details of Source task data selection and curves of each problem are placed in Appendix B.3 and E, respectively.

The HPOB benchmark has several challenges. Firstly, for these problems, algorithms are easy to converge on a certain area, leading to rapid convergence. Secondly, the source tasks are complex, presenting significant variability in data distribution and scale. To tackle the difficulties, the search algorithm should be equipped with ability to handle complex source tasks and robust optimization capabilities.

Experimental results are shown in Figure 4b, which represent the mean ranks and the variance of all algorithms on 39 test problems. In both transfer settings, the baseline ranks are very close, because there are only tiny gaps among final convergence values, which can be seen in Appendix E. However, our method still have clear advantages. Although our method doesn’t have enough warm-start strength when the source tasks are diverse and complex, but it can still adjust the search tree structure and locate the better search area, exceeding existing search space transfer algorithms.

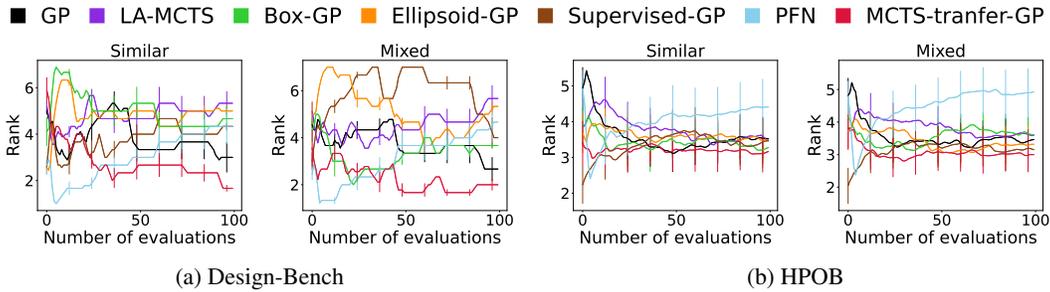


Figure 4: Comparison of MCTS-transfer and other algorithms on Design-Bench and HPOB

4.6 Runtime Analysis

To verify whether MCTS-transfer incurs excessive time overhead, we analyze the runtime proportion of each component of MCTS-transfer on BBOB and three real-world tasks (i.e., LunarLander, RobotPush, and Rover). We divide MCTS-transfer into three main components: evaluation, backpropagation and reconstruction. The evaluation time includes the time required for surrogate model fitting, candidate solution selection and evaluation, which is the common component shared by all optimization algorithms. Backpropagation and reconstruction constitute the two principal modules specific to MCTS-transfer.

As shown in Figure 5, the additional computational burden introduced by MCTS-transfer (i.e., backpropagation and reconstruction) represents a relatively minor fraction of the total runtime, particularly in the three expensive real-world problems. In expensive evaluation problems, MCTS-transfer can bring great improvement by introducing small additional computational overhead. Other details, such as the exact time cost of each component and the frequency MCTS subtree reconstruction, are provided in Appendix F.

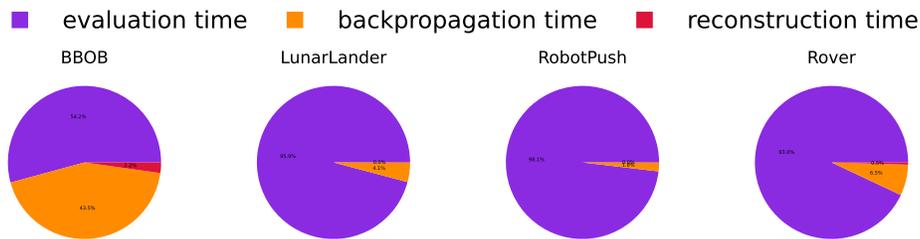


Figure 5: Time cost proportion of evaluation, backpropagation and reconstruction of MCTS-transfer

5 Conclusions and Limitations

In this paper, we propose MCTS-transfer to solve the expensive BBO problem. MCTS-transfer uses MCTS to perform search space transfer, extracting similar tasks during the optimization process and giving more accurate space partition results. Compared with other space transfer algorithms, our algorithm is more generalizable. Specifically, it can extract the most similar source tasks and give them higher weights to accelerate optimization. Besides, it is reliable because it can dynamically adjust the tree structure, improving the probability that the optimal solution falls in search space of the chosen node. Comprehensive experiments on BBOB, real-world problems, Design-Bench and HPOB demonstrate the effectiveness of our algorithm. However, there are still limitations, such as the inability of MCTS-transfer to handle search space transfer tasks for problems with different domains or different dimensions. Future work will focus on solving the heterogeneous space transfer [6], exploring more accurate similarity measures and trying to give the theoretical analysis of the effectiveness of MCTS-transfer.

Acknowledgments and Disclosure of Funding

We thank the reviewers for their insightful and valuable comments. This work was supported by the Science and Technology Project of the State Grid Corporation of China (5700-202440332A-2-1-ZX).

References

- [1] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.
- [2] Tianyi Bai, Yang Li, Yu Shen, Xinyi Zhang, Wentao Zhang, and Bin Cui. Transfer learning for Bayesian optimization: A survey. *arXiv:2302.05927*, 2023.
- [3] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13:281–305, 2012.
- [4] Bernd Bischl, Martin Binder, Michel Lang, Tobias Pielok, Jakob Richter, Stefan Coors, Janek Thomas, Theresa Ullmann, Marc Becker, Anne-Laure Boulesteix, Difan Deng, and Marius Lindauer. Hyperparameter optimization: Foundations, algorithms, best practices, and open challenges. *WIREs Data. Mining. Knowl. Discov.*, 13(2), 2023.
- [5] Cameron Browne, Edward Jack Powley, Daniel Whitehouse, Simon M. Lucas, Peter I. Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez Liebana, Spyridon Samothrakis, and Simon Colton. A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(1):1–43, 2012.
- [6] Zhou Fan, Xinran Han, and Zi Wang. Transfer learning for bayesian optimization on heterogeneous search spaces. *arXiv preprint arXiv:2309.16597*, 2023.
- [7] Matthias Feurer, Jost Tobias Springenberg, and Frank Hutter. Initializing Bayesian hyperparameter optimization via meta-learning. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI'15)*, pages 1128–1135, Austin, TX, 2015.
- [8] Peter I. Frazier. A tutorial on Bayesian optimization. *arXiv:1807.02811*, 2018.
- [9] Nikolaus Hansen, Steffen Finck, Raymond Ros, and Anne Auger. Real-parameter black-box optimization benchmarking 2009: Noiseless functions definitions. Technical Report RR-6829, INRIA, 2009.
- [10] Tomoharu Iwata. End-to-end learning of deep kernel acquisition functions for Bayesian optimization. *arXiv:2111.00639*, 2021.
- [11] Hadi Samer Jomaa, Sebastian Pineda Arango, Lars Schmidt-Thieme, and Josif Grabocka. Transfer learning for Bayesian HPO with end-to-end meta-features, 2022.
- [12] 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.
- [13] Tinu Theckel Joy, Santu Rana, Sunil Gupta, and Svetha Venkatesh. A flexible transfer learning framework for Bayesian optimization with convergence guarantee. *Expert Syst. Appl.*, 115: 656–672, 2019.
- [14] Beomjoon Kim, Kyungjae Lee, Sungbin Lim, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Monte carlo tree search in continuous spaces using voronoi optimistic optimization with regret bounds. In *The 34th AAAI Conference on Artificial Intelligence (AAAI'20)*, pages 9916–9924, New York, NY, 2020.
- [15] Harold J. Kushner. A new method of locating the maximum point of an arbitrary multippeak curve in the presence of noise. *Journal of Basic Engineering*, 86(1):97–106, 1964.
- [16] Ho Chung Leon Law, Peilin Zhao, Leung Sing Chan, Junzhou Huang, and Dino Sejdinovic. Hyperparameter learning via distributional transfer. In *Advances in Neural Information Processing Systems 32 (NeurIPS'18)*, pages 6801–6812, Vancouver, Canada, 2018.

- [17] Yang Li, Yu Shen, Huaijun Jiang, Tianyi Bai, Wentao Zhang, Ce Zhang, and Bin Cui. Transfer learning based search space design for hyperparameter tuning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'22)*, pages 967–977, Washington, DC, 2022.
- [18] Xingchen Ma and Matthew B. Blaschko. Additive tree-structured covariance function for conditional parameter spaces in bayesian optimization. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS'20)*, pages 1015–1025, Palermo, Italy, 2020.
- [19] Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, Joe W. J. Jiang, Ebrahim M. Songhori, Shen Wang, Young-Joon Lee, Eric Johnson, Omkar Pathak, Sungmin Bae, Azade Nazi, Jiwoo Pak, Andy Tong, Kavya Srinivasa, William Hang, Emre Tuncer, Anand Babu, Quoc V. Le, James Laudon, Richard Ho, Roger Carpenter, and Jeff Dean. Chip placement with deep reinforcement learning. *arXiv:2004.10746*, 2020.
- [20] Henry B. Moss, David S. Leslie, and Paul Rayson. Mumbo: Multi-task max-value Bayesian optimization. In *Proceedings of the 31th Machine Learning and Knowledge Discovery in Databases: European Conference (ECML/PKDD'20)*, pages 447–462, Ghent, Belgium, 2020.
- [21] Samuel Müller, Matthias Feurer, Noah Hollmann, and Frank Hutter. Pfn4bo: In-context learning for Bayesian optimization. In *Proceedings of the 40th International Conference on Machine Learning (ICML'23)*, pages 25444–25470, Honolulu, HI, 2023.
- [22] Rémi Munos. Optimistic optimization of a deterministic function without the knowledge of its smoothness. In *Advances in Neural Information Processing Systems 24 (NeurIPS'11)*, pages 783–791, Granada, Spain, 2011.
- [23] Valerio Perrone and Huibin Shen. Learning search spaces for Bayesian optimization: Another view of hyperparameter transfer learning. In *Advances in Neural Information Processing Systems 32 (NeurIPS'19)*, pages 12751–12761, Vancouver, Canada, 2019.
- [24] Sebastian Pineda-Arango, Hadi S. Jomaa, Martin Wistuba, and Josif Grabocka. HPO-B: A large-scale reproducible benchmark for black-box HPO based on openml. In *Advances in Neural Information Processing Systems 34 (NeurIPS'21)*, Virtual, 2021.
- [25] Anil Ramachandran, Sunil Gupta, Santu Rana, and Svetha Venkatesh. Selecting optimal source for transfer learning in Bayesian optimisation. In *Proceedings of the 15th Pacific Rim International Conference on Artificial Intelligence (PRICAI'18)*, pages 42–56, Nanjing, China, 2018.
- [26] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [27] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V. Le. Regularized evolution for image classifier architecture search. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI'19)*, pages 4780–4789, Honolulu, HI, 2019.
- [28] Nicolas Schilling, Martin Wistuba, and Lars Schmidt-Thieme. Scalable hyperparameter optimization with products of Gaussian process experts. In *Proceedings of the 27th Machine Learning and Knowledge Discovery in Databases: European Conference (ECML/PKDD'16)*, pages 199–214, Riva del Garda, Italy, 2016.
- [29] Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- [30] Yunqi Shi, Ke Xue, Lei Song, and Chao Qian. Macro placement by wire-mask-guided black-box optimization. In *Advances in Neural Information Processing Systems 36*, New Orleans, LA, 2023.
- [31] Alistair Shilton, Sunil Gupta, Santu Rana, and Svetha Venkatesh. Regret bounds for transfer learning in Bayesian optimisation. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS'17)*, pages 307–315, Fort Lauderdale, FL, 2017.

- [32] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Vedavyas Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy P. Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [33] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy P. Lillicrap, Karen Simonyan, and Demis Hassabis. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv:1712.01815*, 2017.
- [34] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy P. Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- [35] Lei Song, Ke Xue, Xiaobin Huang, and Chao Qian. Monte carlo tree search based variable selection for high dimensional bayesian optimization. In *Advances in Neural Information Processing Systems 35(NeurIPS'22)*, New Orleans, LA, 2022.
- [36] Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012.
- [37] Kevin Swersky, Jasper Snoek, and Ryan Prescott Adams. Multi-task Bayesian optimization. In *Advances in Neural Information Processing Systems 26 (NeurIPS'13)*, pages 2004–2012, Lake Tahoe, NV, 2013.
- [38] Petru Tighineanu, Kathrin Skubch, Paul Baireuther, Attila Reiss, Felix Berkenkamp, and Julia Vinogradska. Transfer learning with Gaussian processes for Bayesian optimization. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS'22)*, pages 6152–6181, Virtual, 2022.
- [39] Brandon Trabucco, Xinyang Geng, Aviral Kumar, and Sergey Levine. Design-bench: Benchmarks for data-driven offline model-based optimization. In *Proceedings of the 39th International Conference on Machine Learning (ICML'22)*, pages 21658–21676, Baltimore, MD, 2022.
- [40] Michael Volpp, Lukas P. Fröhlich, Kirsten Fischer, Andreas Doerr, Stefan Falkner, Frank Hutter, and Christian Daniel. Meta-learning acquisition functions for transfer learning in Bayesian optimization. In *Proceedings of the 8th International Conference on Learning Representations (ICLR'20)*, Addis Ababa, Ethiopia, 2019.
- [41] Linnan Wang, Saining Xie, Teng Li, Rodrigo Fonseca, and Yuandong Tian. Sample-efficient neural architecture search by learning action space. *arXiv:1906.06832*, 2019.
- [42] Linnan Wang, Yiyang Zhao, Yuu Jinnai, Yuandong Tian, and Rodrigo Fonseca. Alphax: exploring neural architectures with deep neural networks and monte carlo tree search. *arXiv:1903.11059*, 2019.
- [43] Linnan Wang, Rodrigo Fonseca, and Yuandong Tian. Learning search space partition for black-box optimization using monte carlo tree search. In *Advances in Neural Information Processing Systems 33 (NeurIPS'20)*, pages 19511–19522, Virtual, 2020.
- [44] Zi Wang, Clement Gehring, Pushmeet Kohli, and Stefanie Jegelka. Batched large-scale Bayesian optimization in high-dimensional spaces. In *Proceedings of 21st International Conference on Artificial Intelligence and Statistics (AISTATS'18)*, pages 745–754, Playa Blanca, Spain, 2018.
- [45] Ying Wei, Peilin Zhao, and Junzhou Huang. Meta-learning hyperparameter performance prediction with neural processes. In *Proceedings of the 38th International Conference on Machine Learning (ICML'21)*, pages 11058–11067, Virtual, 2021.
- [46] James T. Wilson, Riccardo Moriconi, Frank Hutter, and Marc Peter Deisenroth. The reparameterization trick for acquisition functions. *arXiv:1712.00424*, 2017.

- [47] Martin Wistuba and Josif Grabocka. Few-shot Bayesian optimization with deep kernel surrogates. In *Proceedings of the 9th International Conference on Learning Representations (ICLR'21)*, Virtual, 2021.
- [48] Martin Wistuba, Nicolas Schilling, and Lars Schmidt-Thieme. Learning hyperparameter optimization initializations. In *Proceedings of the 2nd IEEE International Conference on Data Science and Advanced Analytics (DSAA'15)*, pages 1–10, Paris, France, 2015.
- [49] Martin Wistuba, Nicolas Schilling, and Lars Schmidt-Thieme. Hyperparameter search space pruning - A new component for sequential model-based hyperparameter optimization. In *Proceedings of the 26th Machine Learning and Knowledge Discovery in Databases: European Conference (ECML/PKDD'15)*, pages 104–119, Porto, Portugal, 2015.
- [50] Martin Wistuba, Nicolas Schilling, and Lars Schmidt-Thieme. Two-stage transfer surrogate model for automatic hyperparameter optimization. In *Proceedings of the 27th Machine Learning and Knowledge Discovery in Databases: European Conference (ECML/PKDD'16)*, pages 199–214, Riva del Garda, Italy, 2016.
- [51] Martin Wistuba, Nicolas Schilling, and Lars Schmidt-Thieme. Scalable Gaussian process-based transfer surrogates for hyperparameter optimization. *Machine Learning*, 107:43–78, 2017.
- [52] Jian Wu and Peter Frazier. Practical two-step lookahead bayesian optimization. In *Advances in Neural Information Processing Systems 32 (NeurIPS'19)*, 2019.
- [53] Shangda Yang, Vitaly Zankin, Maximilian Balandat, Stefan Scherer, Kevin T. Carlberg, Neil Walton, and Kody J. H. Law. Accelerating look-ahead in bayesian optimization: Multilevel monte carlo is all you need. In *Proceedings of the 41st International Conference on Machine Learning (ICML'24)*, Vienna, Austria, 2024.
- [54] Xin-She Yang and Suash Deb. Eagle strategy using lévy walk and firefly algorithms for stochastic optimization. In *Proceedings of the 4th Nature Inspired Cooperative Strategies for Optimization (NICSO'10)*, pages 101–111, Granada, Spain, 2010.
- [55] Quanming Yao, Mengshuo Wang, Hugo Jair Escalante, Isabelle Guyon, Yi-Qi Hu, Yu-Feng Li, Wei-Wei Tu, Qiang Yang, and Yang Yu. Taking human out of learning applications: A survey on automated machine learning. *arXiv:1810.13306*, 2018.
- [56] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. In *Proceedings of the 5th International Conference on Learning Representations (ICLR'17)*, Toulon, France, 2017.

A Treeify Pseudocode

To maintain the property that the potential of the left child should exceed the potential of the right child for each node, tree reconstruction should be implemented in each iteration. The process of adjusting the tree structure, i.e., deleting and reconstructing subtrees that violate the desired property, is called Treeify. The pseudocode is provided below, as shown in Algorithm 2.

Algorithm 2 Treeify

Input: a MCTS model \mathcal{T}

Output: an updated MCTS model \mathcal{T}

```
1: Initialize empty queues  $\mathcal{Q}, \mathcal{N}$ ;  
2:  $\mathcal{Q} \leftarrow \mathcal{Q}.\text{Enqueue}(\text{ROOT of } \mathcal{T})$ ;  
3: while not  $\mathcal{Q}.\text{isEmpty}$  do  
4:    $node \leftarrow \mathcal{Q}.\text{Dequeue}$ ;  
5:   if not  $node.\text{isLeaf}$  then  
6:     if  $p_{node.left} < p_{node.right}$  then  
7:       Delete the subtree of  $node$ ;  
8:        $\mathcal{N} \leftarrow \mathcal{N}.\text{Enqueue}(node)$   
9:     else  
10:       $\mathcal{Q} \leftarrow \mathcal{Q}.\text{Enqueue}(node.left)$ ;  
11:       $\mathcal{Q} \leftarrow \mathcal{Q}.\text{Enqueue}(node.right)$   
12:    end if  
13:  end if  
14: end while  
15: for each  $node$  in  $\mathcal{N}$  do  
16:   if  $node.\text{isSplittable}$  then  
17:     Expand  $node$   
18:   end if  
19: end for
```

B Detailed settings

B.1 Algorithms

We use the authors' reference implementations for LA-MCTS³, Box-GP⁴, Ellipsoid-GP⁵, Supervised-GP⁶ and PFNs4BO⁷. The hyper-parameters of all compared algorithms are summarized as follows:

- **GP-EI.** We use the GP model in Scikit-learn⁸ and the ExpectedImprovement acquisition [46]. The kernel is set to be ConstantKernel(1.0)*Matern(length_scale=1.0, nu=2.5). To optimize the acquisition function, we generate 10000 points randomly and then select the one with the maximum expected improvement to be evaluated, which is similar to LA-MCTS [43].
- **LA-MCTS** [43]. We set $C_p = 0.1, \theta = 10$. We adopt global GP as the modeling approach. SVM with rbf kernel is used for space division in Sphere2D, BBOB, and HPOB, while Logistic Regression is applied for real-world problems. In the sampling process, we sample the entire space 10,000 times, retaining qualifying candidate points, and repeat this process three times. The parameters of GP are consistent with GP-EI. All other parameters are set to their default values.
- **Supervised-GP.** The sampling process is consistent with LA-MCTS. The parameters of GP are consistent with GP-EI. All other parameters are set to their default values.

³<https://github.com/facebookresearch/LaMCTS>

⁴<https://github.com/PKU-DAIR/SSD/tree/master>

⁵<https://github.com/PKU-DAIR/SSD/tree/master>

⁶<https://github.com/PKU-DAIR/SSD/tree/master>

⁷<https://github.com/automl/pfns4bo>

⁸<https://github.com/scikit-learn/scikit-learn>

- **Box-GP.** The sampling process is consistent with LA-MCTS. The parameters of GP are consistent with GP-EI. All other parameters are set to their default values.
- **Ellipsoid-GP.** The sampling process is consistent with LA-MCTS. The parameters of GP are consistent with GP-EI. All other parameters are set to their default values.
- **MCTS-transfer-GP.** We set $\gamma = 0.99$, $\alpha = 0.5$. We adopt 5 best solutions distance as similarity measure and linear-change strategy as weight assignment. Other parameters and sampling method is consistent with LA-MCTS. The parameters of GP are consistent with GP-EI.
- **PFN.** We set epochs = 5000, batch size=256, learning rate=3e-4, warm up steps = 100, model dimension = 256, head = 4, layer number = 4, hidden layer size = 256, dropout=0.1.

B.2 Problems

We consider the following three real-world problems in our experiments.

LunarLander.⁹ This problem is to learn the parameters of a controller for a lunar lander, which is implemented in OpenAI gym.¹⁰ It involves a 12-dimensional continuous input space depicting the lander’s actions. Our goal is to enhance the control algorithm to maximize the mean terminal reward across a consistent batch of 50 randomly generated landscapes, incorporating varying initial positions and velocities.

RobotPush.¹¹ This function computes the distance between a predefined target location and two objects that are manipulated by a pair of robotic appendages. The movement trajectory of these objects is governed by a set of 14 parameters, which encapsulate attributes such as position, orientation, speed, and direction of motion. We need to control the robot to push items to a designated location. This is implemented with a physics engine Box2D¹².

Rover.¹³ The task optimizes 2D trajectories for a rover by defining start and goal positions and a cost function over the state space [44]. The trajectory costs $c(x)$ is computed for solutions x within a 60-dimensional unit hypercube. We need to design a reasonable trajectory to minimize the cost.

To verify the performance of MCTS-transfer in more complex and high-dimensional cases, we consider the 3 continuous problems from Design-Bench¹⁴ [39].

Superconductor. It’s a critical temperature maximization for superconducting materials. This task is taken from the domain of materials science, where the goal is to design the chemical formula for a superconducting material that has a high critical temperature. The search space is a continuous space with 86 dimensions.

Ant morphology. It’s a robot morphology optimization. The goal is to optimize the morphological structure of Ant from OpenAI Gym¹⁵ to make this quadruped robot to run as fast as possible. The search space is a continuous space with 60 dimensions.

D’Kitty Morphology. It’s robot morphology optimization. The goal is to optimize the morphology of D’Kitty robot to navigate the robot to a fixed location. The search space is a continuous space with 56 dimensions.

B.3 Source Task Data Construction for Similar and Mixed Transfer

Data collection The source task data for pre-training is generated by seven optimization algorithms: Random Search [3], Shuffled Grid Search, Hill Climbing, Regularized Evolution [27], Eagle Strategy [54], and GP-EI [12]. These methods span heuristic, evolutionary, and BO techniques, providing a diverse set of optimization behaviors for our analysis.

⁹https://github.com/yucenli/bnn-bo/blob/main/test_functions/lunar_lander.py

¹⁰<https://gym.openai.com/envs/LunarLander-v2>

¹¹https://github.com/zi-w/Ensemble-Bayesian-Optimization/blob/master/test_functions/push_function.py

¹²<https://box2d.org/>

¹³https://github.com/zi-w/Ensemble-Bayesian-Optimization/blob/master/test_functions/rover_function.py

¹⁴<https://github.com/brandontrabucco/Design-Bench>

¹⁵<https://github.com/openai/gym>

BBOB In order to obtain similar source task data of function f , we perform some similar transformations on f to obtain function family \mathcal{F}_f with seed from 0 to 499. Then we use the sampling algorithms mentioned above with randomly selected seed to collect 300 samples on each function in \mathcal{F}_f . In this section, we randomly select 20 datasets from datasets of \mathcal{F}_f for function f . In the similar setting, We only use selected datasets of \mathcal{F}_f . In mixed setting, we use selected datasets of all 5 synthetic functions.

Real-world problems and Design-Bench problems Since it is difficult to find data with the same dimensions in real-world problems, we choose to use the artificial function *Sphere* of the same dimension to generate similar data and dissimilar data and mix it into the source task data set. Among them, to simulate similar functions, we set the optimal point \mathbf{x}^* of the sphere function to be the optimal point of all these algorithms found so far; the optimal point of the *Sphere* function that simulates dissimilar functions is set to $1 - \mathbf{x}^*$. In the similar setting, we randomly select 20 trajectories from the problem with different seeds and function transformations, combined with data generated by similar sphere function, as the dataset. In the mixed setting, we further add 7 trajectories from dissimilar sphere function. In this section, it is a problem with significantly more similar source tasks than dissimilar tasks.

HPOB We choose search spaces from HPO-B-v3 and divide them into 4 groups based on dimensions. Specailly, we obtain group 5860 and 5970 with dimension 2, group 5859 and 5889 with dimension 5, group 7607 and 7609 with dimension 9, group 5806 and 5971 with dimension 16. In the similar setting, we only use the training data from the search space itself for pre-learning. In the mixed setting, we use all training data of search spaces in the same group.

C Discussions of methods

C.1 Task Similarity

Measuring the similarity of two tasks based on the dataset of a source task D_i and the dataset of the target task D_T is a complex task. We define task distance as a measure of similarity. Smaller distances between tasks represent greater similarity. Here we mainly focus on two types of methods: point-based similarity measures and distribution-based similarity measures. The former includes optimal solutions distance, best N solutions mean distance, best N percent solutions mean distance; and the latter includes Kendall coefficient of datasets, Kullback-Leibler Divergence(KL Divergence) of distributions. Detailed description are as follows:

Optimal solutions distance Let \mathbf{x}_T^* be the best solution of target task, and \mathbf{x}_i^* be the best solution of target task. The task distance is calculated as $Distance(\mathbf{x}_T^*, \mathbf{x}_i^*)$.

Best N (or N percent) solutions mean distance Similar to optimal solutions distance, we replace \mathbf{x}_T^* and \mathbf{x}_i^* with the mean of best N (or N percent) of target task and source task i , denoted as $\bar{\mathbf{x}}_T^*$ and $\bar{\mathbf{x}}_i^*$. The task distance is $Distance(\bar{\mathbf{x}}_T^*, \bar{\mathbf{x}}_i^*)$. Intuitively, it's more robust than directly using the best solution.

Kendall coefficient Kendall coefficient is a measure of rank correlation, which can be used to evaluate the consistency of two datasets. We first build a surrogate model (usually GP) M_i on D_i , and then predict all sample values in D_T . According the ranking results of all these data, we get the similarity of the two tasks by Eq (5), where \mathbb{I} is the indicator function and \oplus is the exclusive-nor operation, in which the statement value is true only if the two sub-statements return the same value.

$$Distance(D_T, D_i) = \frac{2 \sum_{j=1}^{|D_T|} \sum_{k=j+1}^{|D_T|} \mathbb{I}((M_i(x_{T,j}) < M_i(x_{T,k})) \oplus (y_{T,j} < y_{T,k}))}{|D_T| * (|D_T| - 1)} \quad (5)$$

KL divergence of distribution KL divergence is a measure of the dissimilarity between two probability distributions, usually expressed as Eq. (6). For D_T and D_i , we first use Guassian KDE to estimate their density of distribution, denoted as p and q . Then we sample a set of $\{\mathbf{x}_i\}_{i=1}^n$ in Ω and

evaluate them. KL divergence is calculated based on $\{p(x)\}_{i=1}^n$ and $\{q(x)\}_{i=1}^n$.

$$Distance(D_T, D_i) = KL(p||q) = \sum_x p(x) \log \frac{p(x)}{q(x)} \quad (6)$$

C.2 Weight Assignment

Based on the similarity between the source tasks and the target task, we can rank the source tasks, and the smaller rank r_i means higher weight assigned. We can choose different weight change strategies, for example, Linear Change Strategy in Eq. (7), Exponential Change Strategy in Eq. (8) or All One Strategy, i.e., all source task weights are set to be 1, to control the proportion of influential source tasks.

$$w_i = \begin{cases} 1.0 - \frac{r_i}{\alpha N_m} & \text{if } r_i < \alpha N_m \\ 0.1 & \text{otherwise} \end{cases}, \quad (7)$$

$$w_i = \beta^{r_i}, 0 \leq r_i \leq N_m - 1 \quad (8)$$

In Eq. (7), N_m is the number of source tasks that have solutions located in Ω_m , and $\alpha \in [0, 1]$ controls the proportion. In Eq. (8), $\beta \in [0, 1]$ is a hyper-parameter of weight change speed.

C.3 Discussion on Conditional search space.

MCTS-transfer can also be applied in conditional search space, which is common in hyper-parameter optimization. For conditional optimization, we consider the problem $\min_{\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d} f(\mathbf{x})$. Specifically, the search space is tree-structured, formulated as $\mathcal{T} = \{V, E\}$, where $v \in V$ is a node representing subspace and $e \in E$ is an edge representing condition. The objective function is also defined based on \mathcal{T} , formulated as $f_{\mathcal{T}}(\mathbf{x}) := f_{p_j, \mathcal{T}}(\mathbf{x}|l_j)$, where p_j is a condition and $\mathbf{x}|l_j$ is the restriction of \mathbf{x} to l_j [18]. In the pre-learning stage, it builds subtrees for each $v \in V$ and generates the MCTS model \mathcal{T}' based on \mathcal{T} . In each iteration, followed by UCB value, it finds the target node m located in the subtree of v with condition p_i , optimizes in Ω_m , selects and evaluates the candidate using $f_{p_i, \mathcal{T}}(\mathbf{x}|l_i)$. After that, it updates the task weights and node potential in the whole tree \mathcal{T}' and tries to reconstruct the tree. Note that the tree reconstruction only happens in each subtree of $v \in V$.

C.4 Discussion between the state value in MCTS of AlphaZero and MCTS-Transfer.

MCTS measures state value differently in reinforcement learning and black box optimization. For example, AlphaZero [33] is designed to master complex games through self-play without relying on human knowledge or guidance. MCTS plays a crucial role in AlphaZero's decision-making process, whose state value is used to predict the expected future reward from the current state to the end, rather than the historical information. Different from that a state's future reward of AlphaZero can be obtained through multi-step simulations, i.e., alternating decisions through self-play, evaluation values in BBO can only be obtained through actual evaluations. Consequently, MCTS-transfer utilizes historical information. We have observed that some recent look-ahead BO works [52, 53] have been used to predict the expected value of future steps in BBO problems, which have the potential to be applied in MCTS-transfer as estimates for state values to further improve the performance.

C.5 Discussion of combining MCTS-transfer with other advanced algorithms

MCTS-transfer is a general search space transfer learning framework, which can be combined with other advanced algorithms. We equip MCTS-transfer with PFN and test MCTS-transfer-PFN on mixed real-world problems. As shown in Figure 6, MCTS-transfer-PFN makes further improvements compared to MCTS-transfer-GP and PFN after around 80 iterations, and its ranking is stable and excellent throughout the optimization. This result demonstrates the versatility of MCTS-transfer combining with advanced BO algorithms.

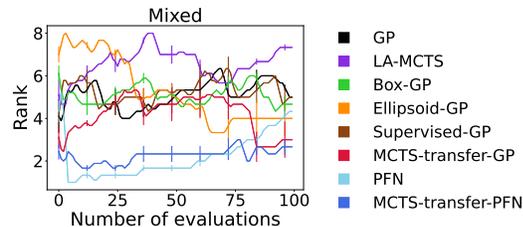


Figure 6: Experimental results on mixed real world problems (with MCTS-transfer-PFN)

D Sensitivity Analysis of Hyper-parameters of MCTS-transfer

Main hyper-parameters of MCTS-transfer includes modeling approach, similarity measures, weight assignment strategies, decay factor γ , important source task ratio α , exploration factor C_p , splitting threshold θ and the binary classifier. We conduct sensitivity analysis of these important parameters on real-world problems LunarLander, RobotPush, Rover under mixed setting.

Modeling approach In MCTS-transfer, we can choose to train GP on full D_T dataset or on the selected points in Ω_m . As shown in Figure 7, the global GP can utilize more global information, demonstrating enhanced search capabilities in these three problems.

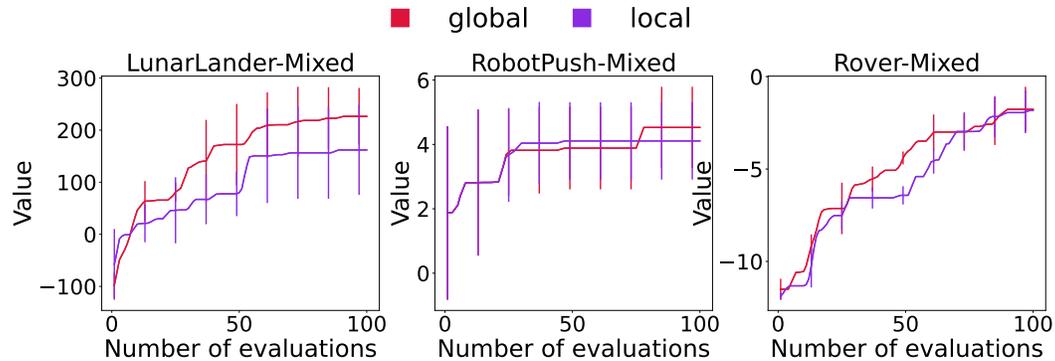


Figure 7: Sensitivity Analysis of Local and Global Modeling Approaches

Similarity measures We propose 5 similarity measures in this paper, including optimal solutions distance, best N or N percent solutions mean distance, Kendall coefficient and KL divergence, as C.1 displayed. The first three methods are point-based measures and the last three are distribution-based methods. Here we set $N=5$ and $N=30\%$ separately for best N solutions and best N percent solutions, and use Gaussian KDE to evaluate KL divergence of distributions. The results can be seen in Figure 8. The distribution-based methods give more precise measure of similarity, but the point-based methods are designed to focus on the optimal region. Each problem has its own suitable measure of similarity. In mixed transfer real-world problems, KL divergence is more appropriate to the problem.

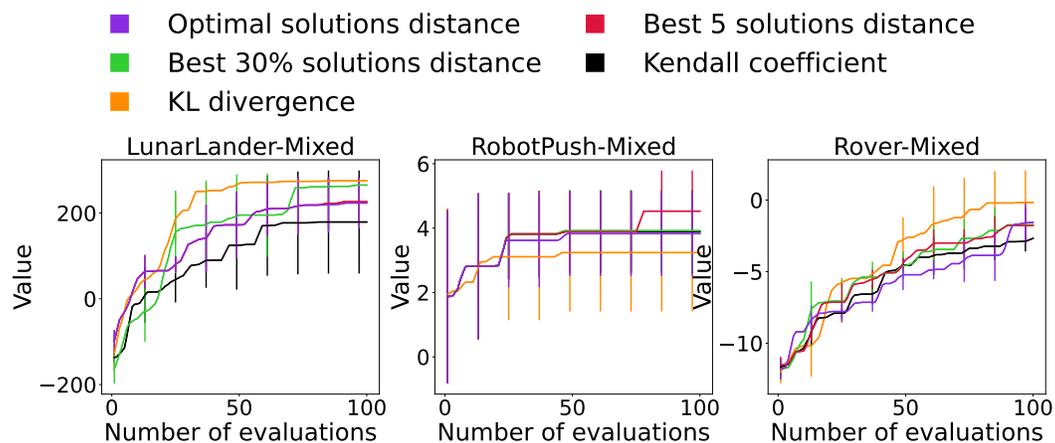


Figure 8: Sensitivity Analysis of Similarity Measures

Weight assignment To assign weights to source tasks, we consider three strategies: linear-change strategy (i.e., Eq. (7)), the exponential-change strategy (i.e., Eq. (8)) and all-one strategy. We set $\alpha = 0.5$ for linear-change strategy, $\beta = 0.5$ for exponential-change strategy. As shown in Figure 10,

the exponential-change strategy, since the weight decaying rapidly with the rank, can only effectively leverages the 1-3 source datasets. So it tends to have faster convergence speed if the similar tasks are effectively identified at the initial stage, as demonstrated in Rover. The all-one strategy are easy to be disturbed by dissimilar data but can fully utilize information, so it may have a higher convergence value at later stage. The linear-change strategy is relatively more stable because it takes the advantages of the above two strategies into account.

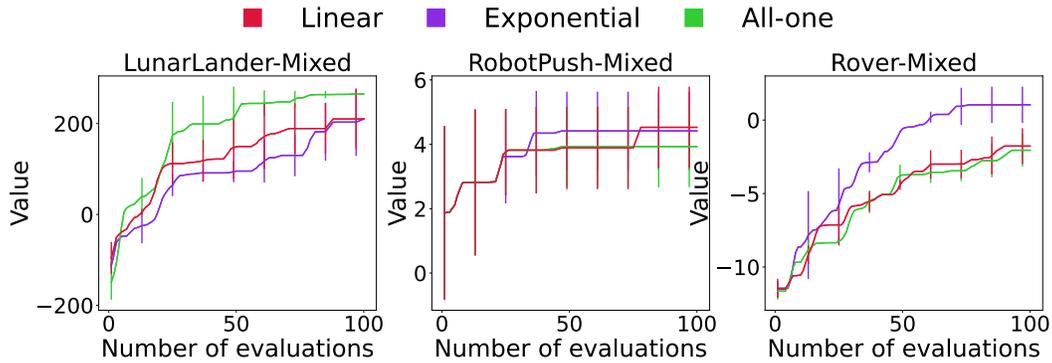


Figure 9: Sensitivity Analysis of Weight Change Strategy

Decay factor γ . The decay factor γ is used to control the influence of source tasks. A higher decay factor γ results in a quicker decay of source tasks' influence, thus making the node potential rely more on the target task's data. However, an excessively rapid decay of source task lead unstable evaluation outcomes and under-utilization of source task data. Specially, for the nodes preferred by source tasks, if they can't be selected as the sampling node at first, their potential will decrease rapidly and these nodes will be less likely to be selected. The analysis experiment of γ can be seen in figure 10, and the result is as expected. The curves of MCTS-transfer with $\gamma = 0.99$ and $\gamma = 1.0$ overlap, better than $\gamma = 0.1$ in most cases, due to data exploitation. And an appropriate decay factor will help to combine information from source and target tasks to accelerate optimization.

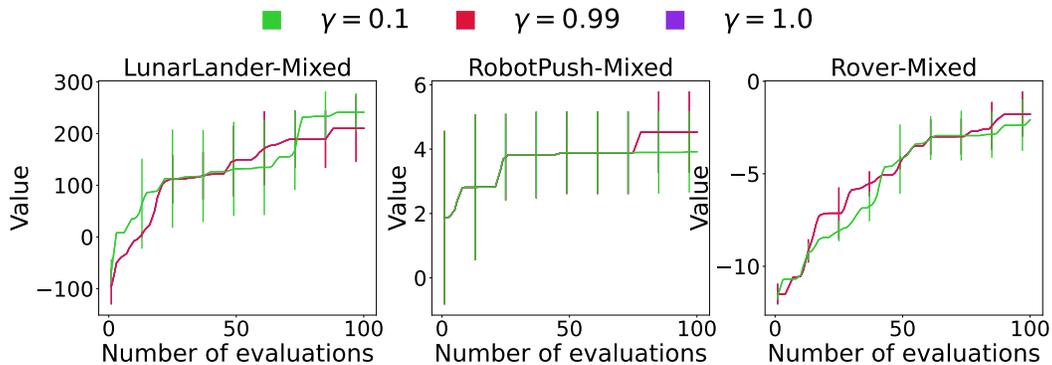


Figure 10: Sensitivity Analysis of Decay Factor γ

Important source task ratio α In linear-change strategy, we can set α to determine the ratio of important source tasks that have high weights. We choose $\alpha = 10\%$, 50% and 100% and the result is shown in Figure 11. A higher α means that more source tasks are influential on the evaluation of the tree node potential. However, this also increases the risk of the interference from dissimilar data. MCTS-transfer with a smaller α only trusts the most similar tasks and under-utilizing information from the source task, leading to slower convergence in early stage.

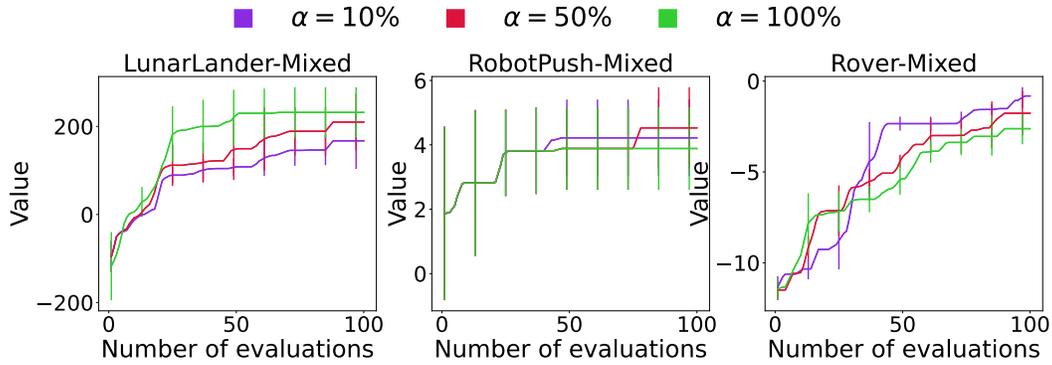


Figure 11: Sensitivity Analysis of Important Source Task Ratio α

Exploration factor C_p C_p controls the degree of exploration. A large C_p prompts the MCTS-transfer to select less promising regions, enhancing exploration. As shown in the Figure 12, too big C_p results in the reduction of exploit useful information, as the algorithm overexploits less promising areas, leading to slower convergence.

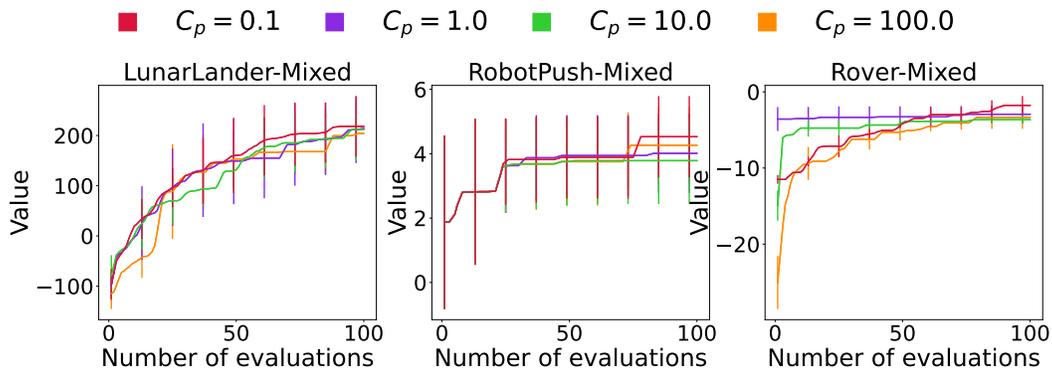


Figure 12: Sensitivity Analysis of Exploration Factor C_p

The splitting threshold θ The threshold θ controls the depth of the tree: a node is only allowed to be further divided when it contains more solutions than θ . A smaller θ leads to a deeper tree. In our experiment (Figure 13), $\theta = 100$ means the tree nodes won't be splitted in the optimization, so the suggesting search space will be bigger, which is less efficient.

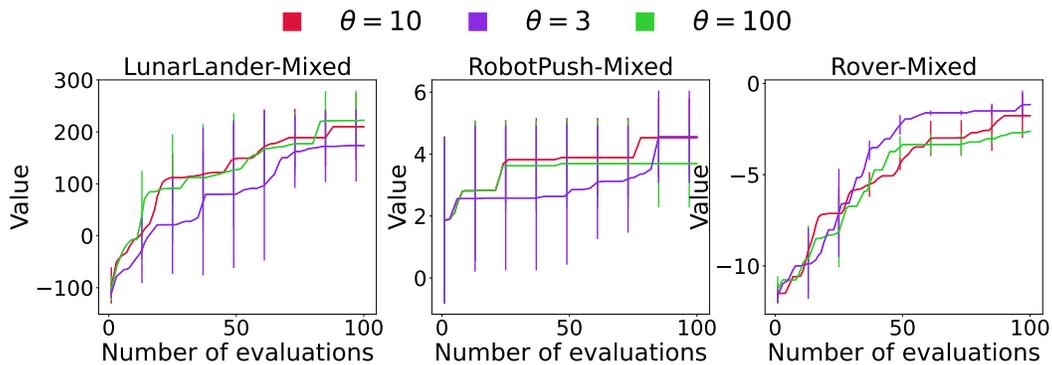


Figure 13: Sensitivity Analysis of The Splitting Threshold θ

Binary classifier The binary classifier decides the boundary of "good" and "bad" clusters. We try the following classifiers: Logistic Regression and SVM (with rbf, linear, or poly kernel). According to the results in Figure 14, SVM with linear kernel and Logistic Regression give more effective search space partition. We can choose a binary classifier with higher partition efficiency according to specific scenarios.

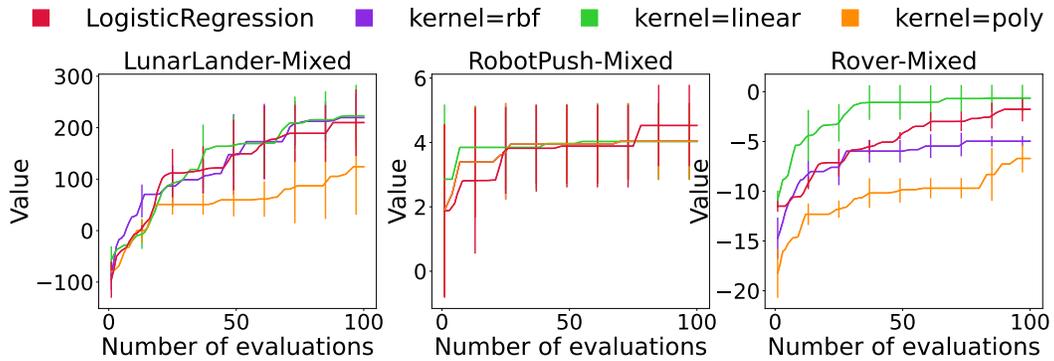


Figure 14: Sensitivity Analysis of Binary Classifier

E Detailed experimental results

The evaluation result curves on BBOB, real-world problems and Design-Bench are summarized in figures 15 and 16. The curves of HPOB in similar setting and mixed setting are shown in figure 18 and figure 19, and the format of subfigure titles is {search space id}-{dataset id}-{Similar/Mixed}.

■ GP ■ LA-MCTS ■ Box-GP ■ Ellipsoid-GP ■ Supervised-GP ■ PFN ■ MCTS-transfer-GP

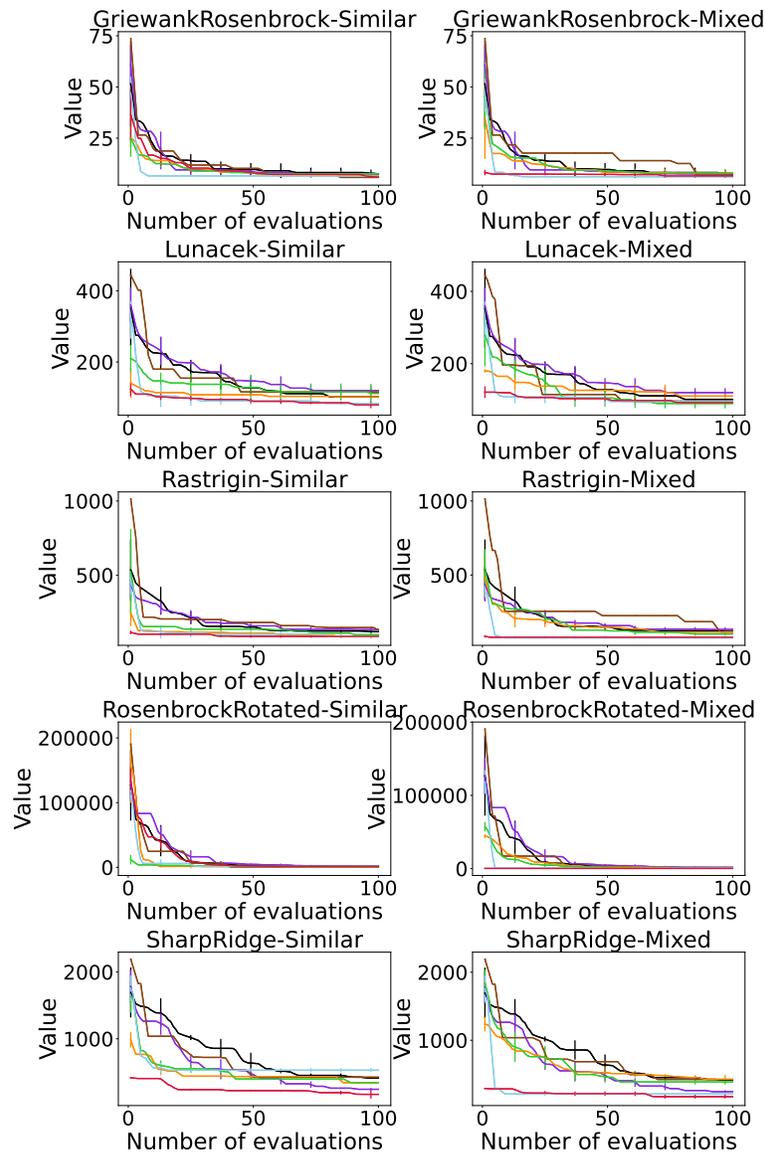


Figure 15: Evaluations of MCTS-transfer and other algorithms on BBOB

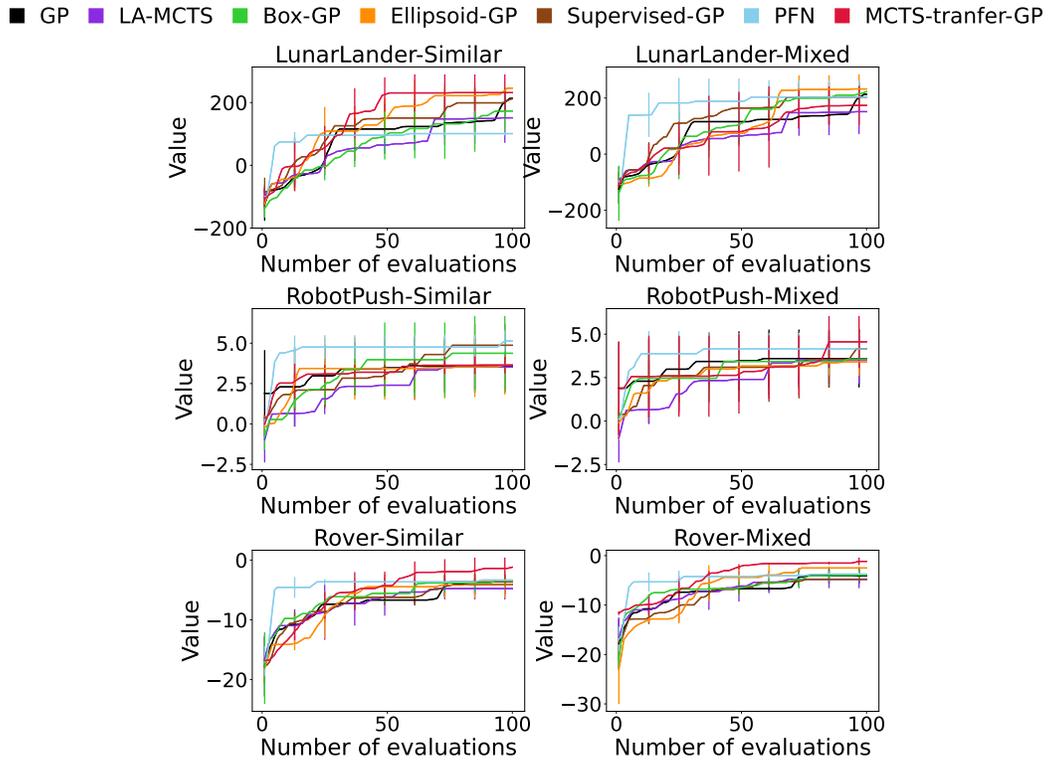


Figure 16: Evaluations of MCTS-transfer and other algorithms on Real-world Problems

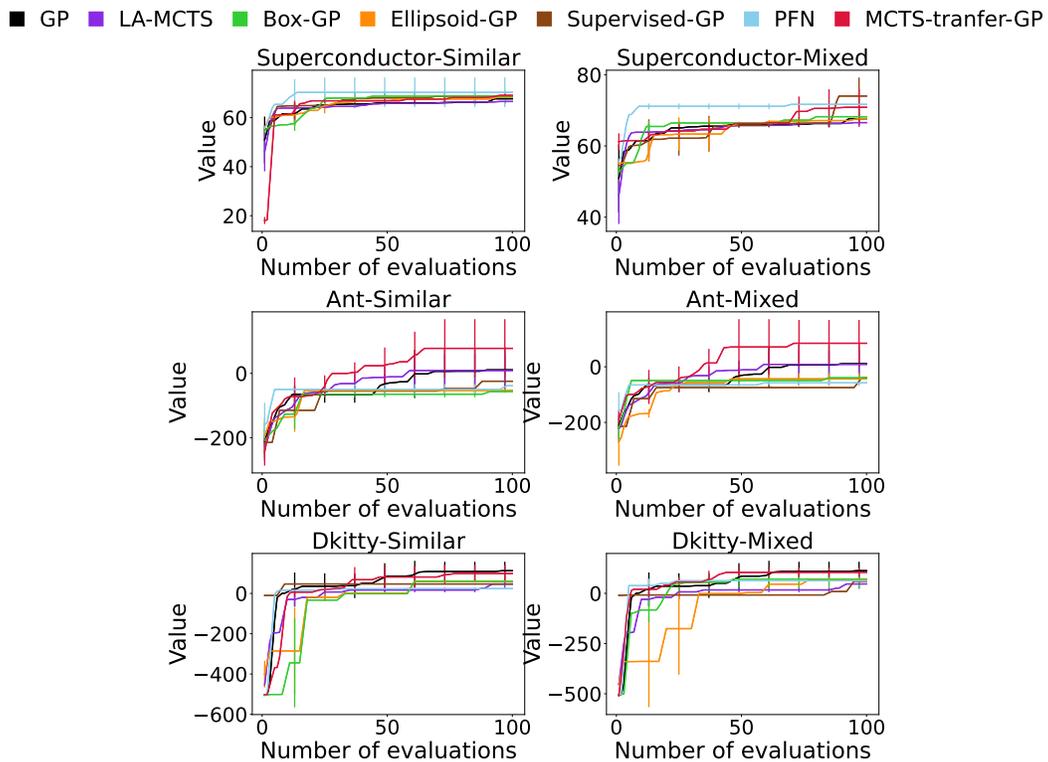
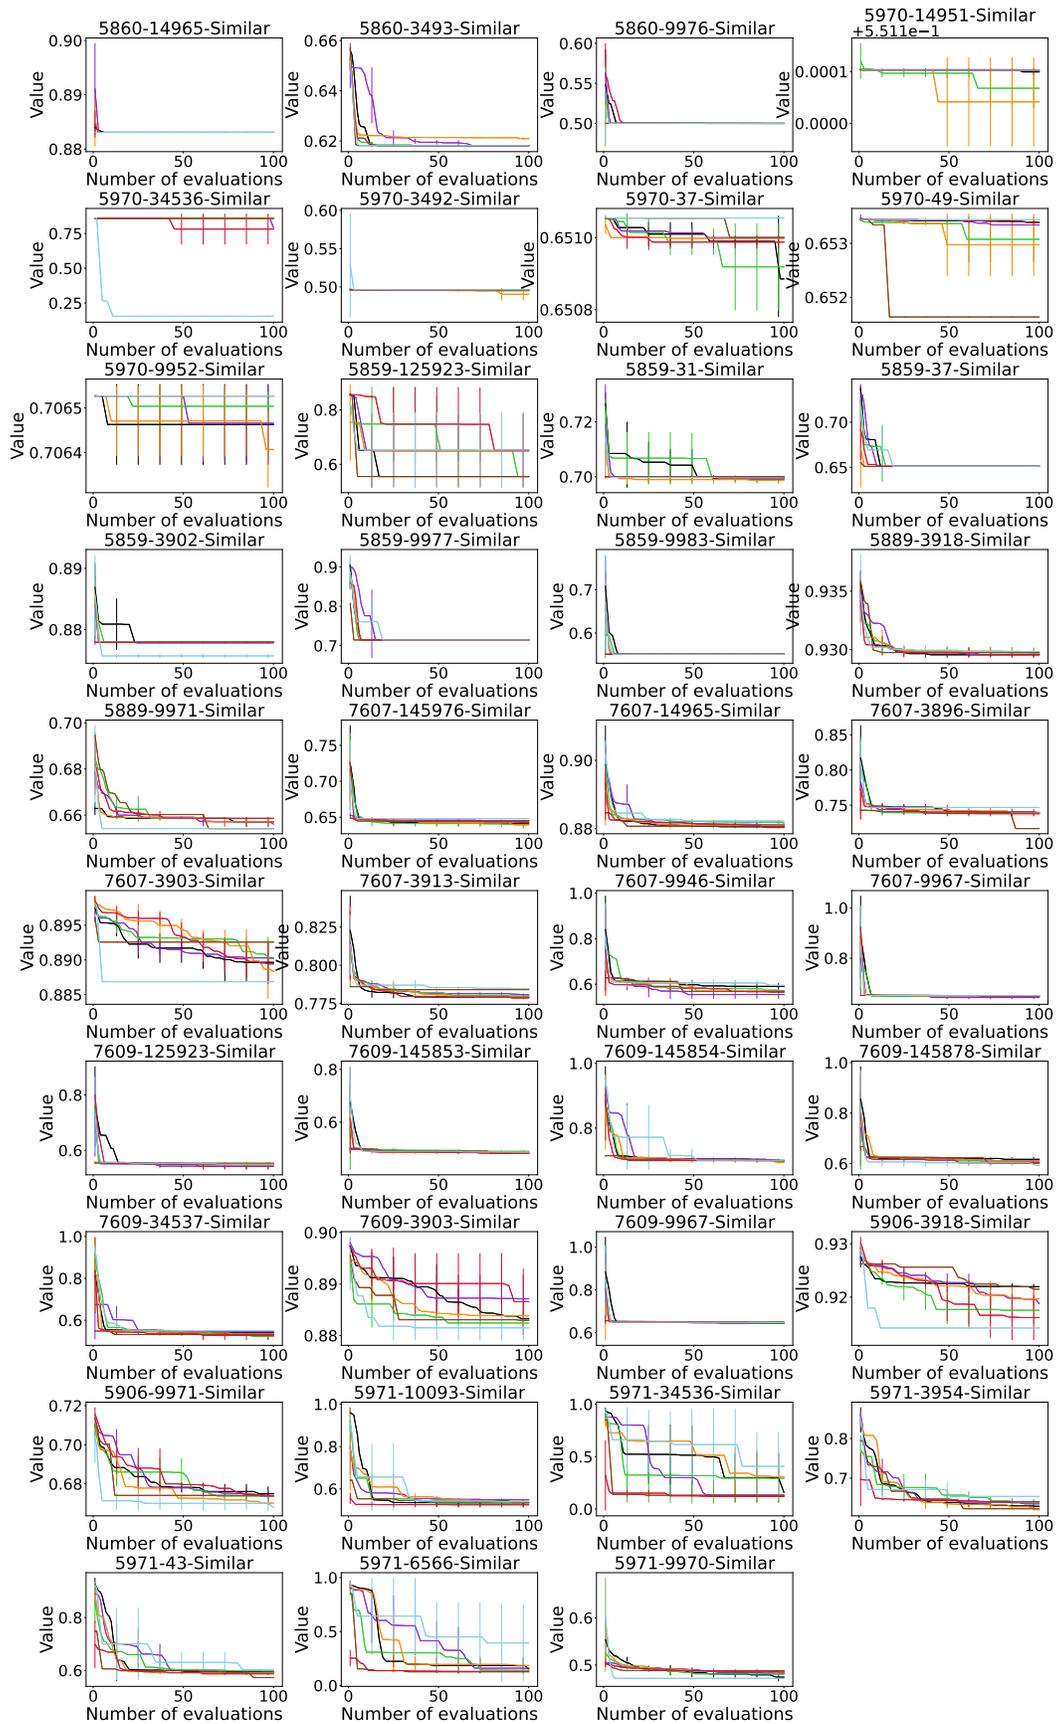
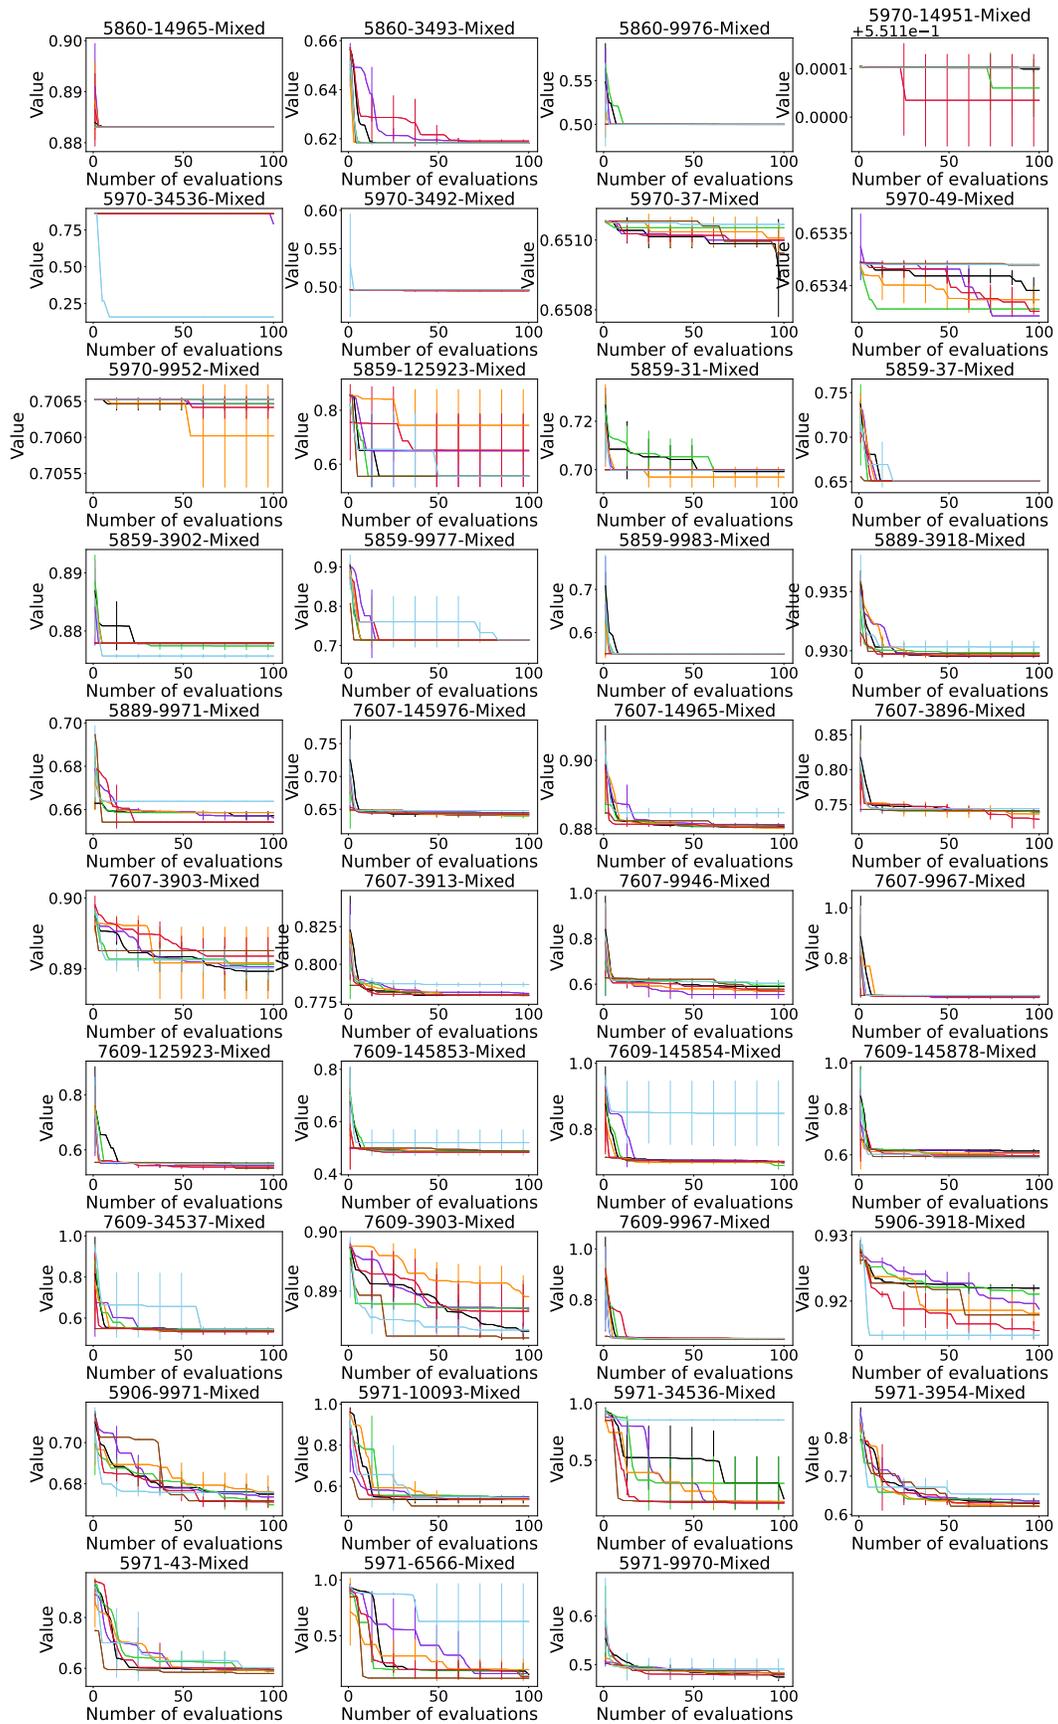


Figure 17: Evaluations of MCTS-transfer and other algorithms on Design-Bench

■ GP
 ■ LA-MCTS
 ■ Box-GP
 ■ Ellipsoid-GP
 ■ Supervised-GP
 ■ PFN
 ■ MCTS-transfer-GP



■ GP
 ■ LA-MCTS
 ■ Box-GP
 ■ Ellipsoid-GP
 ■ Supervised-GP
 ■ PFN
 ■ MCTS-transfer-GP



F Details of runtime analysis

In order to analysis the additional computational overhead introduced by MCTS-transfer, we calculate the time cost and the corresponding proportion of each component, and record the frequency of subtree reconstruction in each iteration.

Specifically, we divide MCTS-transfer into three main components: evaluation, backpropagation and reconstruction. The evaluation component includes the time required for surrogate model fitting, candidate solution selection and evaluation, which is a common component shared by all compared optimization algorithms. Backpropagation and reconstruction components are the two principal modules specific to MCTS-transfer.

We test MCTS-transfer on BBOB benchmark and three real-world tasks (i.e., LunarLander, RobotPush, and Rover). Among them, the evaluation process of the real-world tasks is more time-consuming and the evaluation cost of BBOB are relatively cheaper.

As illustrated in Figure 20, the additional computational burden introduced by MCTS-transfer (i.e., backpropagation and reconstruction) represents a relatively minor fraction of the total runtime, particularly in the three real-world scenarios. These scenarios precisely exemplify the computationally intensive cases that transfer BO is designed to address, wherein MCTS-transfer demonstrates small additional computational overhead. Furthermore, our result reveals that the average frequency of tree reconstructions is low, with the corresponding reconstruction time being almost negligible when compared to the evaluation time.

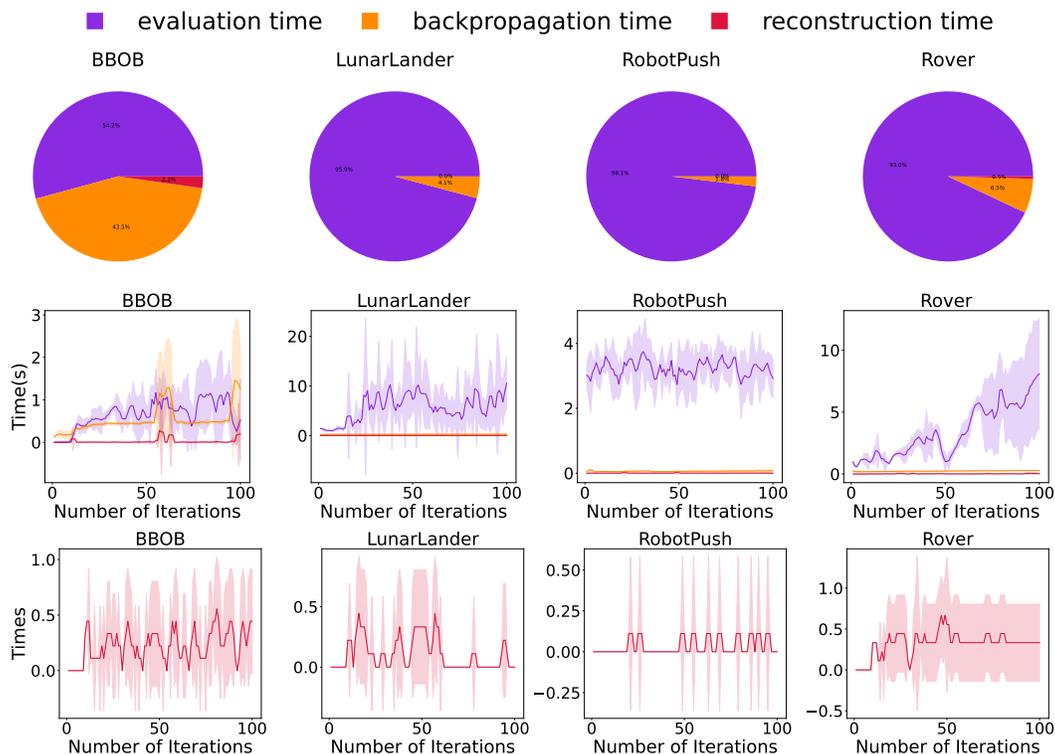


Figure 20: Time cost of different components in MCTS-transfer. The top row: the time proportion of components in optimization. The middle row: the time cost of components in each iteration (avg±std). The bottom row: the number of subtree reconstructions/reconstruction times in each iteration (avg±std).

G Visualization of Weight Change Curve of Source Tasks

Here we show the weight change curve of three mixed real-world problems. We set the weight assignment strategy to be the linear change strategy with $\alpha = 0.5$ and the decay factor to be $\gamma = 0.99$. Figure 21 shows that the weights of real-world problem and similar sphere problem exceed those of dissimilar sphere problem in most cases, regardless of any inconsistencies in initialization. The results prove that the weight change strategy can prioritize similar source task data, which will lead to more accurate node potential evaluation and search space partition.

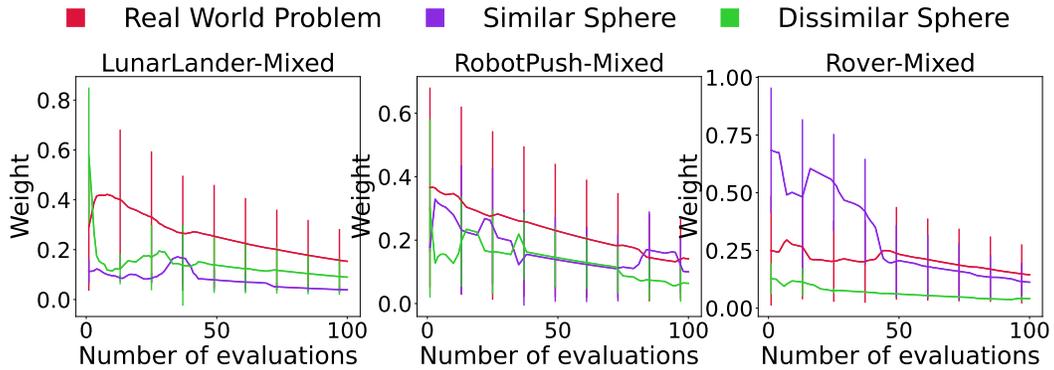


Figure 21: Weight Changes: MCTS-transfer

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Both the abstract and introduction accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitations are proposed in the Conclusion section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The detailed information is provided in Appendix B and C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We will submit code, data and instructions, and we will release the code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The detailed information is provided in Appendix B and C.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The error bars are already shown in figures in paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [No]

Justification: We didn't provide computing resource-related content in the article because that content is not relevant to the topic of the paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Our research conform with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We describe it in introduction part.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We correctly cite the origin paper and sources.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.