Additive Tree-Structured Conditional Parameter Spaces in Bayesian Optimization: A Novel Covariance Function and a Fast Implementation

Xingchen Ma and Matthew B. Blaschko

Abstract—Bayesian optimization (BO) is a sample-efficient global optimization algorithm for black-box functions which are expensive to evaluate. Existing literature on model based optimization in conditional parameter spaces are usually built on trees. In this work, we generalize the additive assumption to tree-structured functions and propose an additive tree-structured covariance function, showing improved sample-efficiency, wider applicability and greater flexibility. Furthermore, by incorporating the structure information of parameter spaces and the additive assumption in the BO loop, we develop a parallel algorithm to optimize the acquisition function and this optimization can be performed in a low dimensional space. We demonstrate our method on an optimization benchmark function, on a neural network compression problem and on pruning pre-trained VGG16 and ResNet50 models. Experimental results show our approach significantly outperforms the current state of the art for conditional parameter optimization including SMAC, TPE and Jenatton et al. (2017).

Index Terms—nonparametric statistics, global optimization, parameter learning.

1 INTRODUCTION

IN many applications, we are faced with the problem of optimizing an expensive black-box function and we wish to find its optimum using as few evaluations as possible. *Bayesian Optimization* (BO) [1] is a global optimization technique, which is specially suited for these problems. BO has gained increasing attention in recent years [2], [3], [4], [5], [6] and has been successfully applied to sensor location [2], hierarchical reinforcement learning [3], and automatic machine learning [7].

In the general BO setting, we aim to solve the following problem:

$$\min_{\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d} f(\mathbf{x}),$$

where $\mathcal{X}$ is a d -dimensional parameter space and f is a black-box function which is expensive to evaluate. Typically, the parameter space $\mathcal{X}$ is treated as structureless, however, for many practical applications, there exists a conditional structure in $\mathcal{X}$:

$$f(\mathbf{x} \mid \mathbf{x}_{I_1}) = f(\mathbf{x}_{I_2} \mid \mathbf{x}_{I_1}), \quad (1)$$

where I_1, I_2 are index sets and are subsets of $[d] := \{1, \dots, d\}$. Equation (1) means given the value of $\mathbf{x}_{I_1}$, the value of $f(\mathbf{x})$ remains unchanged after removing $\mathbf{x}_{[d] \setminus (I_1 \cup I_2)}$. Here we use set based subscripts to denote the restriction of $\mathbf{x}$ to the corresponding indices.

This paper investigates optimization problems where the parameter space exhibits such a conditional structure. In particular, we focus on a specific instantiation of the general conditional structure in Equation (1): Tree-structured parameter spaces, which are also studied in [8]. A realistic

example is neural network compression. If we compress different layers with different compression algorithms, the search space is a tree-structured parameter space. Another practical example is searching for activation functions of each layer in neural architecture search.

By exploring the properties of this tree structure, we design an *additive tree-structured* (Add-Tree) covariance function, which enables information sharing between different data points under the *additive assumption*, and allows GP to model f in a sample-efficient way. Furthermore, by including the tree structure and the additive assumption in the BO loop, we develop a parallel algorithm to optimize the acquisition function, making the overall execution faster. Our proposed method also helps to alleviate the curse of dimensionality through two advantages: (i) we avoid modeling the response surface directly in a high-dimensional space, and (ii) the acquisition optimization is also operated in a lower-dimensional space.

This article is an extended version of [9]. It has been expanded to include theoretical convergence results of our proposed Add-Tree covariance function in Section 4.3, additional practical details in Section 4.4, empirical results of pruning two large-scale neural networks in Section 5.3 and experiments on neural architecture search in Section 5.4.

In the next section, we will briefly review BO together with the literature related to optimization in a conditional parameter space. In Section 2, we formalize the family of objective functions that can be solved using our approach. We then present our Add-Tree covariance function in Section 3. In Section 4, we give the inference procedure and BO algorithm using our covariance function. We further present theoretical properties of an Add-Tree covariance function in Section 4.3. We then report a range of experiments in Section 5. Finally, we conclude in Section 6. All proofs and implementation details are provided in the supplementary

- X. Ma and M. B. Blaschko are with the Center for Processing Speech and Images, Department of Electrical Engineering, KU Leuven, Kasteelpark Arenberg 10, 3001 Leuven, Belgium.
E-mail: xingchen.ma@kuleuven.be

material.

1.1 Related Work

1.1.1 Bayesian Optimization

BO has two major components. The first one is a probabilistic regression model used to fit the response surface of f . Popular choices include GPs [3], random forests [4] and adaptive Parzen estimators [10]. We refer the reader to [11] for the foundations of Gaussian Processes. The second one is an acquisition function u_t which is constructed from this regression model and is used to propose the next evaluation point. Popular acquisition functions include the *expected improvement* (EI) [1], *knowledge gradient* (KG) [12], *entropy search* (ES) [13] and *Gaussian process upper confidence bound* (GP-UCB) [2].

One issue that often occurs in BO is, in high-dimensional parameter spaces, its performance may be no better than random search [14], [15]. This deterioration is because of the high uncertainty in fitting a regression model due to the curse of dimensionality [16, Ch. 2], which in turn leads to pure-explorational behavior of BO. This will further cause inefficiency in the acquisition function, making the proposal of the next data point behave like random selection. Standard GP-based BO ignores the structure in a parameter space, and fits a regression model in $\mathbb{R}^d$. By leveraging this structure information, we can work in a low-dimensional space $\mathbb{R}^m$ (recall Equation (1)) instead of $\mathbb{R}^d$. Another well-known issue of high-dimensional BO is the problem of optimizing an acquisition function is also performed in a high-dimensional space. This difficulty can also be sidestepped by using this structure information if GP-UCB is used as our acquisition function.

1.1.2 Conditional Parameter Spaces

Sequential Model-based Algorithm Configuration (SMAC) [4] and Tree-structured Parzen Estimator Approach (TPE) [10] are two popular non-GP based optimization algorithms that are aware of the conditional structure in $\mathcal{X}$, however, they lack favorable properties of GPs: uncertainty estimation in SMAC is non-trivial and the dependencies between dimensions are ignored in TPE. Additionally, neither of these methods have a particular sharing mechanism, which is valuable in the low-data regime.

In the category of GP-based BO, which is our focus in this paper, [17] proposed a covariance function that can explicitly employ the tree structure and share information at those categorical nodes. However, their specification for the parameter space is too restrictive and they require the shared node to be a categorical variable. By contrast, we allow shared variables to be continuous (see Section 3). [18] applied the idea of [17] in a BO setting, but their method still inherits the limitations of [17]. Another covariance function to handle tree-structured dependencies is presented in [19]. In that case, they force the similarity of two samples from different condition branches to be zero and the resulting model can be transformed into several independent GPs. We perform a comparison to an independent GP baseline in Section 5.1. In contrast to Add-Tree, the above approaches either have very limited applications, or lack a sharing mechanism. [8] presented another GP-based BO approach, where

they handle tree-structured dependencies by introducing a weight vector linking all sub-GPs, and this introduces an explicit sharing mechanism. Although [8] overcame the above limitations, the enforced linear relationships between different paths make their semi-parametric approach less flexible compared with our method. We observe in our experiments that this can lead to a substantial difference in performance.

2 PROBLEM FORMULATION

We begin by listing notation used in this paper. Vectors, matrices and sets will be denoted by bold lowercase letters, bold uppercase letters and regular uppercase letters respectively. Let $\mathcal{T} = (V, E)$ be a tree, in which V is the set of vertices, E is the set of edges, $P = \{p_i\}_{1 \leq i \leq |P|}$ be the set of leaves and r be the root of $\mathcal{T}$ respectively, $\{l_i\}_{1 \leq i \leq |P|}$ be the ordered set of vertices on the path from r to the i -th leaf p_i , and h_i be the number of vertices along l_i (including r and p_i). To distinguish an objective function defined on a tree-structured parameter space from a general objective function, we use $f_{\mathcal{T}}$ to indicate our objective function. In what follows, we will call $f_{\mathcal{T}}$ a *tree-structured function*. An example of such a tree-structured function is depicted in Figure 1, where each node of $\mathcal{T}$ is associated with a bounded continuous variable and each edge represents a specific setting of a categorical variable.

Definition 1 (Tree-structured parameter space). A *tree-structured parameter space* $\mathcal{X}$ is associated with a tree $\mathcal{T} = (V, E)$. For any $v \in V$, v is associated with a bounded continuous variable of $\mathcal{X}$; for any $e \in E$, the set of outgoing edges E_v of v represent one categorical variable of $\mathcal{X}$ and each element of E_v represents a specific setting of the corresponding categorical variable.

Definition 2 (Tree-structured function). A *tree-structured function* $f_{\mathcal{T}} : \mathcal{X} \rightarrow \mathbb{R}$ is defined on a d -dimensional tree-structured parameter space $\mathcal{X}$. The i -th leaf p_i is associated with a function $f_{p_i, \mathcal{T}}$ of the variables associated with the vertices along l_i . $f_{\mathcal{T}}$ is called *tree-structured* if for every leaf of the tree-structured parameter space

$$f_{\mathcal{T}}(\mathbf{x}) := f_{p_j, \mathcal{T}}(\mathbf{x}|_{l_j}), \quad (2)$$

where p_j is selected by the categorical values of $\mathbf{x}$ and $\mathbf{x}|_{l_j}$ is the restriction of $\mathbf{x}$ to l_j .

To aid in the understanding of a tree-structured function (and subsequently our proposed Add-Tree covariance function), we depict a simple tree-structured function in Figure 1. The categorical variable in this example is $t \in \{1, 2\}$ and its values are shown around the outgoing edges of r . The continuous variables are $v_r \in [-1, 1]^2$, $v_{p_1} \in [-1, 1]^2$, $v_{p_2} \in [-1, 1]^3$ and they are displayed inside vertices r, p_1, p_2 respectively. Here v_r is a shared continuous parameter, as evidenced by it appearing in both leaves. Two functions $f_{p_1, \mathcal{T}}$ and $f_{p_2, \mathcal{T}}$ associated with leaves p_1 and p_2 respectively are also shown in Figure 1. In Definition 2, the restriction of one input to a path means we collect variables associated with the vertices along that path and concatenate them using a fixed ordering. For example, in Figure 1, let $\mathbf{x} \in \mathcal{X}$ be an 8-dimensional input, then the restriction of $\mathbf{x}$ to path l_1 is

Fig. 1. A simple tree-structured function. The dimension of $\mathcal{X}$ is $d = 8$, and the effective dimension at p_1 and p_2 are 4 and 5 respectively.

a 4-dimensional vector. The function illustrated in Figure 1 can be compactly written down as:

$$f_{\mathcal{T}}(\mathbf{x}) = \mathbb{1}_{t=1} f_{p_1, \mathcal{T}}(\mathbf{x}|_{l_1}) + \mathbb{1}_{t=2} f_{p_2, \mathcal{T}}(\mathbf{x}|_{l_2}), \quad (3)$$

where $\mathbf{x}$ is the concatenation of $(\mathbf{v}_r, \mathbf{v}_{p_1}, \mathbf{v}_{p_2}, t)$, $\mathbb{1}$ denotes the indicator function, $f_{p_1, \mathcal{T}}(\mathbf{x}|_{l_1}) = \|\mathbf{v}_r\|^2 + \|\mathbf{v}_{p_1}\|^2$ and $f_{p_2, \mathcal{T}}(\mathbf{x}|_{l_2}) = \|\mathbf{v}_r\|^2 + \|\mathbf{v}_{p_2}\|^2$.

A tree-structured function $f_{\mathcal{T}}$ is actually composed of several smaller functions $\{f_{p_i, \mathcal{T}}\}_{1 \leq i \leq |P|}$, given a specific setting of the categorical variables, $f_{\mathcal{T}}$ will return the associated function at the i -th leaf. To facilitate our description in the following text, we define the *effective dimension* in Definition 3.

Definition 3 (Effective dimension). The *effective dimension* of a tree-structured function $f_{\mathcal{T}}$ at the i -th leaf p_i is the sum of dimensions of the variables associated with the vertices along l_i .

Remark 4. The effective dimension of $f_{\mathcal{T}}$ can be much smaller than the dimension of $\mathcal{X}$.

Particularly, if $\mathcal{T}$ is a perfect binary tree, in which each vertex is associated with a 1-dimensional continuous variable, the effective dimension of $f_{\mathcal{T}}$ at every leaf is the depth h of $\mathcal{T}$, while the dimension of $\mathcal{X}$ is $3 \cdot 2^{h-1} - 2$. If the tree structure information is thrown away, we have to work in a much higher dimensional parameter space. It is known that in high dimensions, BO behaves like random search [14], [15], which violates the entire purpose of model based optimization.

Now we have associated the parameter space $\mathcal{X}$ with a tree structure, which enables us to work in the low-dimensional effective space. How can we leverage this tree structure to optimize $f_{\mathcal{T}}$? Recalling $f_{\mathcal{T}}$ is a collection of $|P|$ functions $\{f_{p_i, \mathcal{T}}\}_{1 \leq i \leq |P|}$, a trivial solution is using $|P|$ independent GPs to model each $f_{p_i, \mathcal{T}}$ separately. In BO settings, we are almost always in low data regime because black-box calls to $f_{\mathcal{T}}$ are expensive (e.g. the cost of training and evaluating a machine learning model). Modeling $f_{\mathcal{T}}$ using a collection of GPs is obviously not an optimal way because we discard the correlation between $f_{p_i, \mathcal{T}}$ and $f_{p_j, \mathcal{T}}$ when $i \neq j$. How to make the most of the observed data, especially how to share information between data points coming from different leaves remains a crucial question. In this paper, we assume additive structure within each $f_{p_i, \mathcal{T}}$ for $i = 1, \dots, |P|$. More formally, $f_{p_i, \mathcal{T}}$ can be decomposed

in the following form:

$$f_{p_i, \mathcal{T}}(\mathbf{x}) = \sum_{j=1}^{h_i} f_{ij}(v_{ij}) \quad (4)$$

where v_{ij} is the associated variable on the j -th vertex along l_i . The additive assumption has been extensively studied in the GP literature [20], [21], [22], [23] and is a popular method of dimension reduction [16]. We note the tree-structured function discussed in this paper is a generalization of the objective function presented in these publications and our additive assumption in Equation (4) is also a generalization of the additive structure considered previously. For example, the additive function discussed in [21] can be viewed as a tree-structured function the associated tree of which has a branching factor of 1, i.e. $|P| = 1$. Our generalized additive assumption will enable an efficient sharing mechanism, as we develop in Section 3.

3 THE ADD-TREE COVARIANCE FUNCTION

In this section, we describe how we use the tree structure and the additive assumption to design a covariance function, which is sample-efficient in low-data regime. We start with the definition of the Add-Tree covariance function (Definition 5), then we show the intuition (Equation (9)) behind this definition and present an algorithm (Algorithm 1) to automatically construct an Add-Tree covariance function from the specification of a tree-structured parameter space, finally a proof of the validity of this covariance function is given.

Definition 5 (Add-Tree covariance function). For a tree-structured function $f_{\mathcal{T}}$, let $\mathbf{x}_{i'}$ and $\mathbf{x}_{j'}$ be two inputs of $f_{\mathcal{T}}$, p_i and p_j be the corresponding leaves, a_{ij} be the lowest common ancestor (LCA) of p_i and p_j , l_{ij} be the path from r to a_{ij} . A covariance function $k_{f_{\mathcal{T}}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is said to be an *Add-Tree covariance function* if for each $\mathbf{x}_{i'}$ and $\mathbf{x}_{j'}$

$$\begin{aligned} k_{\mathcal{T}}(\mathbf{x}_{i'}, \mathbf{x}_{j'}) &:= k_{l_{ij}}(\mathbf{x}_{i'}|_{l_{ij}}, \mathbf{x}_{j'}|_{l_{ij}}) \\ &= \sum_{v \in l_{ij}} k_v(\mathbf{x}_{i'}|_v, \mathbf{x}_{j'}|_v) \end{aligned} \quad (5)$$

where $\mathbf{x}_{i'}|_{l_{ij}}$ is the restriction of $\mathbf{x}_{i'}$ to the variables along the path l_{ij} , and k_v is any positive semi-definite covariance function on the continuous variables appearing at a vertex v on the path l_{ij} . We note the notation l_{ij} introduced here is different from the notation l_i introduced in the beginning of Section 2.

We take the tree-structured function illustrated in Figure 1 (Equation (3)) as an example.¹ Suppose there are n_1 and n_2 data points in l_1 and l_2 respectively, and let $X_1 \in \mathbb{R}^{n_1 \times d_1}$ and $X_2 \in \mathbb{R}^{n_2 \times d_2}$ be the inputs from l_1 and l_2 , where $d_1 = 2 + 2$ and $d_2 = 2 + 3$ are the effective dimensions of $f_{\mathcal{T}}$ at p_1 and p_2 respectively. Denote the latent variables associated to the decomposed functions² at r , p_1 and p_2 by $\mathbf{f}_r \in \mathbb{R}^{n_1+n_2}$, $\mathbf{f}_1 \in \mathbb{R}^{n_1}$ and $\mathbf{f}_2 \in \mathbb{R}^{n_2}$, respectively.

¹ To simplify the presentation, we use a two-level tree structure in this example. The covariance function, however, generalizes to tree-structured functions of arbitrary depth (Algorithm 1).

² On functions and latent variables, one can refer to [11, chap. 2].

Reordering and partition $\mathbf{f}_r$ into two parts corresponding to p_1 and p_2 , so that

$$\mathbf{f}_r = \begin{bmatrix} \mathbf{f}_r^{(1)} \\ \mathbf{f}_r^{(2)} \end{bmatrix}, \mathbf{f}_r^{(1)} \in \mathbb{R}^{n_1}, \mathbf{f}_r^{(2)} \in \mathbb{R}^{n_2}.$$

Let the Gram matrix corresponding to $\mathbf{f}_r, \mathbf{f}_1, \mathbf{f}_2$ be $\mathbf{K}_r \in \mathbb{R}^{(n_1+n_2) \times (n_1+n_2)}$, $\mathbf{K}_1 \in \mathbb{R}^{n_1 \times n_1}$, $\mathbf{K}_2 \in \mathbb{R}^{n_2 \times n_2}$. W.l.o.g, let the means of $\mathbf{f}_r, \mathbf{f}_1, \mathbf{f}_2$ be $\mathbf{0}$. By the additive assumption in Equation (4), the latent variables corresponding to the associated functions at p_1 and p_2 are $\mathbf{f}_r^{(1)} + \mathbf{f}_1$ and $\mathbf{f}_r^{(2)} + \mathbf{f}_2$, we have:

$$\begin{bmatrix} \mathbf{f}_r^{(1)} + \mathbf{f}_1 \\ \mathbf{f}_r^{(2)} + \mathbf{f}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{f}_r^{(1)} \\ \mathbf{f}_r^{(2)} \end{bmatrix} + \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \end{bmatrix}, \begin{bmatrix} \mathbf{f}_r^{(1)} \\ \mathbf{f}_r^{(2)} \end{bmatrix} \sim \mathcal{N}(\mathbf{0}, \mathbf{K}_r). \quad (6)$$

Due to $\mathbf{f}_1 \perp \mathbf{f}_2$, where $\perp$ denotes $\mathbf{f}_1$ is independent of $\mathbf{f}_2$, we have:

$$\begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{K}_2 \end{bmatrix}\right), \quad (7)$$

furthermore, because of the additive assumption in Equation (4),

$$\begin{bmatrix} \mathbf{f}_r^{(1)} \\ \mathbf{f}_r^{(2)} \end{bmatrix} \perp \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \end{bmatrix}. \quad (8)$$

Combining Equations (6) to (8), we arrive at our key conclusion:

$$\begin{bmatrix} \mathbf{f}_r^{(1)} + \mathbf{f}_1 \\ \mathbf{f}_r^{(2)} + \mathbf{f}_2 \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}_r^{(11)} + \mathbf{K}_1 & \mathbf{K}_r^{(12)} \\ \mathbf{K}_r^{(21)} & \mathbf{K}_r^{(22)} + \mathbf{K}_2 \end{bmatrix}\right), \quad (9)$$

where $\mathbf{K}_r$ is decomposed as follows:

$$\mathbf{K}_r = \begin{bmatrix} \mathbf{K}_r^{(11)} & \mathbf{K}_r^{(12)} \\ \mathbf{K}_r^{(21)} & \mathbf{K}_r^{(22)} \end{bmatrix}$$

in which $\mathbf{K}_r^{(11)} \in \mathbb{R}^{n_1 \times n_1}$, $\mathbf{K}_r^{(12)} \in \mathbb{R}^{n_1 \times n_2}$, $\mathbf{K}_r^{(21)} \in \mathbb{R}^{n_2 \times n_1}$, $\mathbf{K}_r^{(22)} \in \mathbb{R}^{n_2 \times n_2}$. The observation in Equation (9) is crucial in two aspects: firstly, we can use a single covariance function and a global GP to model our objective, secondly and more importantly, this covariance function allows an efficient sharing mechanism between data points coming from different paths, although we cannot observe the decomposed function values at the shared vertex r , we can directly read out this sharing information from $\mathbf{K}_r^{(12)}$.

We summarize the construction of an Add-Tree covariance function in Algorithm 1. We note Algorithm 1 is used to construct a function object of an Add-Tree covariance function, and during this construction stage, no observations are involved in the computation. The 7-th line in Algorithm 1 needs further explanation. If $\mathbf{x}$ and $\mathbf{x}'$ don't intersect, at least one of $\mathbf{x}_{si \leq i < ei}$ and $\mathbf{x}'_{si \leq i < ei}$ is empty (it is possible that both of them are empty). In this case, we let k_c map an empty set to any non-negative scalar, and a common choice is 0. In practice, when $\mathbf{x}$ and $\mathbf{x}'$ don't contain vertex v at the same time, the value of $k_c(\mathbf{x}_{si \leq i < ei}, \mathbf{x}'_{si \leq i < ei})$ is irrelevant because it is not needed.

Proposition 6. *The Add-Tree covariance function defined by Definition 5 is positive semi-definite for all tree-structured functions defined in Definition 2 with the additive assumption satisfied.*

Algorithm 1: Add-Tree Covariance Function

Input : The associated tree $\mathcal{T} = (V, E)$ of $f_{\mathcal{T}}$

Output: Add-Tree covariance function $k_{\mathcal{T}}$

```

1  $k_{\mathcal{T}} \leftarrow 0$ 
2 for  $v \leftarrow V$  do
3    $vi \leftarrow \text{Index}(v)$ 
4    $k_v^d \leftarrow k_v^{\delta}(\mathbf{x}, \mathbf{x}') \quad /* \quad k_v^{\delta}(\mathbf{x}, \mathbf{x}') = 1 \text{ iff } \mathbf{x} \text{ and } \mathbf{x}' \text{ both contain } v \text{ in their paths } */$ 
5    $si \leftarrow vi + 1 \quad /* \text{ start index of } v */$ 
6    $ei \leftarrow vi + 1 + \text{Dim}(v) \quad /* \text{ end index of } v */$ 
7    $k_v^c \leftarrow k_c(\mathbf{x}_{si \leq i < ei}, \mathbf{x}'_{si \leq i < ei}) \quad /* \quad k_c \text{ is any p.s.d. covariance function } */$ 
8    $k_v \leftarrow k_v^d \times k_v^c$ 
9    $k_{\mathcal{T}} \leftarrow k_{\mathcal{T}} + k_v$ 

```

4 BO FOR TREE-STRUCTURED FUNCTIONS

In this section, we first describe how to perform the inference with our proposed Add-Tree covariance function, then we present a parallel algorithm to optimize the acquisition function. At the end of this section, we give the convergence rate of an Add-Tree covariance function.

4.1 Inference with Add-Tree

Given noisy observations $D_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$, we would like to obtain the posterior distribution of $f_{\mathcal{T}}(\mathbf{x}_*)$ at a new input $\mathbf{x}_* \in \mathcal{X}$. We begin with some notation. Let p_* be the leaf selected by the categorical values of $\mathbf{x}_*$, l_* be the path from the root r to p_* . All n inputs are collected in the design matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$, where the i -th row represents $\mathbf{x}_i \in \mathbb{R}^d$, and the targets and observation noise are aggregated in vectors $\mathbf{y} \in \mathbb{R}^n$ and $\boldsymbol{\sigma} \in \mathbb{R}^n$ respectively. Let $\boldsymbol{\Sigma} = \text{diag}(\boldsymbol{\sigma})$ be the noise matrix, where $\text{diag}(\boldsymbol{\sigma})$ denotes a diagonal matrix containing the elements of vector $\boldsymbol{\sigma}$, $S = \{i \mid l_* \cap l_i \neq \emptyset\}$, $\mathbf{I} \in \mathbb{R}^{n \times n}$ be the identity matrix, $\mathbf{Q} \in \mathbb{R}^{|S| \times n}$ be a selection matrix, which is constructed by removing the j -th row of $\mathbf{I}$ if $j \notin S$, $\mathbf{X}' = \mathbf{Q}\mathbf{X} \in \mathbb{R}^{|S| \times d}$, $\mathbf{y}' = \mathbf{Q}\mathbf{y} \in \mathbb{R}^{|S|}$, $\boldsymbol{\Sigma}' = \mathbf{Q}\boldsymbol{\Sigma}\mathbf{Q}^T \in \mathbb{R}^{|S| \times |S|}$. We note it is possible that a vertex is not associated with any continuous variable. In particular, if there is no continuous variable associated with root r , $|S|$ is smaller than n . Then we can write down the joint distribution of $f_{\mathcal{T}}(\mathbf{x}_*)$ and $\mathbf{y}'$ as:

$$\begin{bmatrix} f_{\mathcal{T}}(\mathbf{x}_*) \\ \mathbf{y}' \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} k_{\mathcal{T}}(\mathbf{x}_*, \mathbf{x}_*) & \mathbf{K}_{\mathcal{T}}(\mathbf{x}_*) \\ \mathbf{K}_{\mathcal{T}}(\mathbf{x}_*)^T & \mathbf{K}_{\mathcal{T}} + \boldsymbol{\Sigma}' \end{bmatrix}\right),$$

where $\mathbf{K}_{\mathcal{T}} := \mathbf{K}_{\mathcal{T}}(\mathbf{X}') = [k_{\mathcal{T}}(\mathbf{x}, \mathbf{x}')]_{\mathbf{x}, \mathbf{x}' \in \mathbf{X}'}$, $\mathbf{K}_{\mathcal{T}}(\mathbf{x}_*) = [k_{\mathcal{T}}(\mathbf{x}_*, \mathbf{x}')]_{\mathbf{x} \in \mathbf{X}'}$. We note that this joint distribution has the same standard form [11] as in all GP-based BO, but that it is made more efficient by the selection of $\mathbf{X}'$ based on the tree structure. The predictive distribution of $f_{\mathcal{T}}(\mathbf{x}_*)$ is:

$$f_{*\mathcal{T}} \mid \mathbf{X}', \mathbf{y}', \mathbf{x}_* \sim \mathcal{N}(\mu_{\mathcal{T}}(\mathbf{x}_*), \sigma_{\mathcal{T}}^2(\mathbf{x}_*)) \quad (10)$$

where

$$\begin{aligned} \mu_{\mathcal{T}}(\mathbf{x}_*) &= \mathbf{K}_{\mathcal{T}}(\mathbf{x}_*) \mathbf{K}_{\mathcal{T}, \mathcal{T}}^{-1} \mathbf{y}, \\ \sigma_{\mathcal{T}}^2(\mathbf{x}_*) &= k_{\mathcal{T}}(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{K}_{\mathcal{T}}(\mathbf{x}_*) \mathbf{K}_{\mathcal{T}, \mathcal{T}}^{-1} \mathbf{K}_{\mathcal{T}}(\mathbf{x}_*)^T, \\ \mathbf{K}_{\mathcal{T}, \mathcal{T}} &= \mathbf{K}_{\mathcal{T}} + \boldsymbol{\Sigma}'. \end{aligned}$$

Black-box calls to the objective function usually dominate the running time of BO, and the time complexity of fitting GP is of less importance in BO settings. Details on the time complexity analysis of an Add-Tree covariance function can be found in [9], and we only sketch the idea here. The time complexity of inference using an Add-Tree covariance function can be derived by recursion. W.l.o.g, let $\mathcal{T}$ be a binary tree. If root r of $\mathcal{T}$ is associated with a continuous parameter, in this worst case, the complexity is $\mathcal{O}(n^3)$, where n is the number of observations. Otherwise, if r is not associated to any continuous parameter, let n_l and n_r be the number of samples falling into the left and the right path respectively, we have $T(r) = T(r_l) + T(r_r)$, where r_l and r_r are the left and the right child of r respectively, and $T(r), T(r_l), T(r_r)$ are the running time at nodes r, r_l, r_r respectively. Because the worst-case running time at nodes r_l and r_r is $\mathcal{O}(n_l^3)$ and $\mathcal{O}(n_r^3)$ respectively, we have $T(r) = \mathcal{O}(n_l^3 + n_r^3)$.

4.2 Acquisition Function Optimization

In BO, the acquisition function $u_{t-1}(\mathbf{x}|D_{t-1})$, where t is the current step of optimization, is used to guide the search for the optimum of an objective function. By trading off the exploitation and exploration, it is expected we can find the optimum using as few calls as possible to the objective. To get the next point at which we evaluate an objective function, we solve $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} u_{t-1}(\mathbf{x}|D_{t-1})$. For noisy observations, GP-UCB [2] is proved to be no-regret and has explicit regret bounds for many commonly used covariance functions, and in this work, we will use GP-UCB, which is defined as:

$$u_{t-1}(\mathbf{x}|D_{t-1}) = \mu_{t-1}(\mathbf{x}) + \beta_{t-1}^{1/2} \sigma_{t-1}(\mathbf{x}),$$

where β_{t-1} are suitable constants which will be given later in Section 4.3, $\mu_{t-1}(\mathbf{x})$ and $\sigma_{t-1}(\mathbf{x})$ are the predictive posterior mean and standard deviation at $\mathbf{x}$ from Equation (10).

A naïve way to obtain the next evaluation point for a tree-structured function is to independently find $|P|$ optima, each one corresponding to the optimum of the associated function at a leaf, and then choose the best candidate across these optima. This approach is presented in [8] and the authors there already pointed out this is too costly. Here we develop a much more efficient algorithm, which is dubbed as Add-Tree-GP-UCB, to find the next point and we summarise it in Algorithm 2. By explicitly utilizing the associated tree structure $\mathcal{T}$ of $f_{\mathcal{T}}$ and the additive assumption, the first two nested for loops can be performed in parallel. Furthermore, as a by-product, each acquisition function optimization routine is now performed in a low dimensional space whose dimension is even smaller than the effective dimension. We note the Add-Tree covariance function developed here can be combined with any other acquisition function, however, Algorithm 2 may not be applicable.

4.3 Theoretical Results

In this section, we give theoretical results regarding our proposed Add-Tree covariance function. Let $r_t = f_{\mathcal{T}}(\mathbf{x}^*) - f_{\mathcal{T}}(\mathbf{x}_t)$ be the instantaneous regret, where $f_{\mathcal{T}}(\mathbf{x}^*) = \max_{\mathbf{x} \in \mathcal{X}} f_{\mathcal{T}}(\mathbf{x})$, and let $R_t = \sum_{i=1}^t r_i$ be the cumulative regret after t iterations [24]. Similar to [25], we show our

Algorithm 2: Add-Tree-GP-UCB

Input : The associated tree $\mathcal{T} = (V, E)$ of $f_{\mathcal{T}}$,
Add-Tree covariance function $k_{\mathcal{T}}$ from
Algorithm 1, paths $\{l_i\}_{1 \leq i \leq |P|}$ of $\mathcal{T}$

```

1  $D_0 \leftarrow \emptyset$ 
2 for  $t \leftarrow 1, \dots$  do
3   for  $v \leftarrow V$  do
4      $\mathbf{x}_t^v \leftarrow \arg \max_{\mathbf{x}} \mu_{t-1}^v(\mathbf{x}) + \sqrt{\beta_{t-1}^v} \sigma_{t-1}^v(\mathbf{x})$ 
5      $u_t^v \leftarrow \mu_{t-1}^v(\mathbf{x}_t^v) + \sqrt{\beta_{t-1}^v} \sigma_{t-1}^v(\mathbf{x}_t^v)$ 
6   for  $i \leftarrow 1, \dots, |P|$  do
7      $U_t^{l_i} \leftarrow \sum_{v \in l_i} u_t^v$  /* additive
      assumption from Equation (4) */
8    $j \leftarrow \arg \max_i \{U_t^{l_i} \mid i = 1, \dots, |P|\}$ 
9    $\mathbf{x}_t \leftarrow \bigcup \{\mathbf{x}_t^v \mid v \in l_j\}$ 
10   $y_t \leftarrow f_{\mathcal{T}}(\mathbf{x}_t)$ 
11   $D_t \leftarrow D_{t-1} \cup \{(\mathbf{x}_t, y_t)\}$ 
12  Fitting GP using  $D_t$  to get  $\{(\mu_t^v, \sigma_t^v)\}_{v \in V}$ 

```

proposed Add-Tree covariance function has sublinear regret bound if each covariance function defined on the vertices of $\mathcal{T}$ is a squared exponential kernel or a Matérn kernel. We further show without taking the structure of the parameter space into consideration, the cumulative regret has an exponential dependence with the dimension d of $\mathcal{X}$, while our method has an exponential dependence with the highest dimension $\tilde{d}$ of functions defined in vertices of $\mathcal{T}$. As stated earlier, $\tilde{d}$ is smaller than the effective dimension of a tree-structured function and the effective dimension is usually much smaller than d .

4.3.1 Adapting the Kernel Hyper-Parameters

The GP-UCB algorithm proposed in [2] is shown to have an explicit sublinear regret bound $R_t = \mathcal{O}(\sqrt{t \beta_t \gamma_t})$ with high probability for commonly used covariance functions, including the squared exponential kernel and the Matérn kernel. Here γ_t is the maximum mutual information gain after t rounds and is defined to be:

$$\gamma_t := \max_{A \subset \mathcal{X}, |A|=t} I(\mathbf{y}_A; f), \quad (11)$$

where $I(\mathbf{y}_A; f) = 1/2 \log |\mathbf{I} + \sigma^{-2} \mathbf{K}_A|$ is the mutual information between observations in A and prior of f under a GP model, $\mathbf{K}_A$ is the Gram matrix $[k_{\theta}(\mathbf{x}, \mathbf{x}')]_{\mathbf{x}, \mathbf{x}' \in A}$ and is dependent on kernel k 's hyperparameters θ . The squared exponential kernel is defined as $k(\mathbf{x}, \mathbf{x}') = \exp(-\frac{1}{2\theta^2} \|\mathbf{x} - \mathbf{x}'\|_2^2)$, where θ is the lengthscale parameter. The Matérn kernel is defined as $k(\mathbf{x}, \mathbf{x}') = \frac{2^{1-\nu}}{\Gamma(\nu)} r^{\nu} B_{\nu}(r)$, where $r = \frac{\sqrt{2\nu}}{\theta} \|\mathbf{x} - \mathbf{x}'\|_2$, B_{ν} is the modified Bessel function with $\nu > 1$, Γ is the gamma function.

One practical issue is [2] requires information about the RKHS norm bound of f and the correct hyperparameters of the GP's covariance function, which are usually unavailable. In this work, we follow the adaptive GP-UCB algorithm (A-GP-UCB) proposed in [25], which can adapt the norm bound

of an RKHS and hyperparameters of a covariance function. The adapting mechanism in [25] is by setting:

$$\beta_t^{1/2} = B_t + 4\sigma\sqrt{I_{\theta_t}(\mathbf{y}_t; f) + 1 + \ln(1/\delta)} \quad (12)$$

$$B_t = b(t)g(t)^d B_0, \quad \theta_t = \theta_0/g(t) \quad (13)$$

where θ_0 and B_0 are initial guess for the lengthscales of kernel k and RKHS norm bound of f respectively, $g(t) : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ and $b(t) : \mathbb{N} \rightarrow \mathbb{R}_{>0}$, with $b(0) = g(0) = 1$ are monotonically increasing functions.

The motivation of adapting the norm bound and the lengthscales is shown in Figure 2. As pointed out in [25], it is always possible that a local optimum is located in a local bump area and zero observation falls into that region. GP-UCB with misspecified lengthscales θ and RKHS bound B cannot fully explore the input space and this causes BO to terminate too early. For example, in the first column of Figure 2, if we maximize the GP-UCB acquisition function (the red line), the algorithm will terminate prematurely and we will never have the chance of exploring the region containing the global maximum because of misspecified hyperparameters. Intuitively, decreasing the lengthscales of a covariance function and increasing the norm bound have the effect of expanding the function class, which results in additional exploration, and eventually the true function is contained in the considered space.

Fig. 2. Misspecified lengthscales and norm bound cause BO terminate too early. Scaling hyperparameters expands the function class and results in additional exploration. $\beta^{1/2}$ is a proxy to the RKHS bound B and the relationship between $\beta^{1/2}$ and B is shown in Equation (12). Crosses indicate the observations so far, θ shown on the left side indicates the lengthscales of an exponential squared kernel, the dashed line is the posterior mean $\mu(x)$ of a GP model using this covariance function, the shaded area is the confidence interval $\mu(x) \pm \beta^{1/2}\sigma(x)$, where $\sigma(x)$ is the posterior standard deviation, the solid line indicates the true function, the red line is the upper bound of this confidence interval, which is the GP-UCB acquisition function.

4.3.2 Bounds of Cumulative Regrets

A high-level summary of this section is as follows. Firstly, we use Theorem 7 to bound the cumulative regret of our proposed Add-Tree covariance function in terms of the norm bound and the maximum information gain. Secondly, we show a tree-structured function has bounded RKHS norm in an RKHS with an Add-Tree kernel under suitable assumptions using Proposition 8. Thirdly, to upper bound this information gain quantity, we use Theorem 9 to connect it with the tail sum of the operator spectrum of an Add-Tree covariance function. Fourthly, we derive the upper bound of this tail sum when each covariance function defined on the vertices of $\mathcal{T}$ is a squared exponential kernel or a Matérn kernel using proposition 14 and Proposition 10. Finally, we combine all these together to obtain the cumulative regret of an Add-Tree covariance function.

Theorem 7 (Theorem 1 in [25]). Assume that f has bounded RKHS norm $\|f\|_{k_\theta}^2 \leq B$ in a RKHS that is parametrized by a stationary kernel $k_\theta(\mathbf{x}, \mathbf{x}')$ with unknown lengthscales θ . Based on an initial guess, θ_0 and B_0 , define monotonically increasing functions $g(t) > 0$ and $b(t) > 0$ and run A-GP-UCB with $\beta_t^{1/2} = b(t)g(t)^d B_0 + 4\sigma\sqrt{I_{\theta_t}(\mathbf{y}_t; f) + 1 + \ln(1/\delta)}$ and GP lengthscales $\theta_t = \theta_0/g(t)$. Then, with probability at least $1 - \delta$, we obtain a regret bound of:

$$R_t \leq 2B \max \left(g^{-1}(\max_i \frac{[\theta_0]_i}{[\theta]_i}), b^{-1}(\frac{B}{B_0}) \right) + \sqrt{C_1 t \beta_t I_{\theta_t}(\mathbf{y}_t; f)} \quad (14)$$

where I_{θ_t} is the mutual information based on a GP model with lengthscales θ_t and $C_1 = 8/\log(1 + \sigma^{-2})$.

Before we turn to our main theoretical results, we recall and extend some existing notation. Let $f_{\mathcal{T}}$ be a tree-structured function defined in Definition 2 and $f_{\mathcal{T}}$ satisfies the additive decomposition in Equation (4). Let M be the height of $\mathcal{T}$, v_{ij} be the j -th vertex on the i -th level of $\mathcal{T}$, where $i \in \{1, \dots, |P|\}$, $j \in \{1, \dots, M\}$. Let k_{ij} be a covariance function defined on the continuous variables appearing at vertex v_{ij} , $\mathcal{H}_{ij}$ be an RKHS with kernel k_{ij} , f_{ij} be a function defined on the continuous variables on vertex v_{ij} and assume $f_{ij} \in \mathcal{H}_{ij}$. Let $\tilde{d}$ be the highest dimension of functions in $\{f_{ij}\}$, $f_i := \sum_{j=1}^{\tilde{d}} f_{ij}$ denote the sum of functions defined on the i -th path of $\mathcal{T}$, and let k_i be the sum of all covariance functions on the i -th path of $\mathcal{T}$. We note the collection $\{f_{ij}\}$ is a multiset, since different paths can share variables.

To utilize Theorem 7, we need to ensure a tree-structured function $f_{\mathcal{T}}$ has bounded RKHS norm in an RKHS with an Add-Tree covariance function $k_{\mathcal{T}}$.

Proposition 8. Assume every f_{ij} has bounded RKHS norm $\|f_{ij}\|_{k_{ij}}^2 \leq \tilde{B}$ in an RKHS with a kernel k_{ij} , and let $k_{\mathcal{T}}$ be an Add-Tree covariance function constructed from k_{ij} , then $f_{\mathcal{T}}$ has bounded RKHS norm in an RKHS with a kernel $k_{\mathcal{T}}$:

$$\|f_{\mathcal{T}}\|_{k_{\mathcal{T}}}^2 \leq M\tilde{B}. \quad (15)$$

Theorem 9 (Theorem 8 in [2]). Suppose that $D \subset \mathbb{R}^d$ is compact, and $k(\mathbf{x}, \mathbf{x}')$ is a covariance function for which the additional assumption of Theorem 2 in [2] holds. Moreover, let $B_k(T_*) = \sum_{s>T_*} \lambda_s$, where $\{\lambda_s\}$ is the operator spectrum of k

with respect to the uniform distribution over D . Pick $\tau > 0$, and let $n_T = C_4 T^\tau \log T$, with C_4 depends only on τ and D . Then the following bound holds true:

$$\begin{aligned} \gamma_T &\leq \frac{1/2}{1 - e^{-1}} \max_{r \in \{1, \dots, T\}} \left(T_* \log \left(\frac{r n_T}{\sigma^2} \right) \right. \\ &\quad \left. + C_4 \sigma^2 \left(1 - \frac{r}{T} \right) (B_k(T_*) T^{\tau+1} + 1) \log T \right) \\ &\quad + \mathcal{O}(T^{1-\tau/d}) \end{aligned} \quad (16)$$

for any $T_* \in \{1, \dots, n_T\}$.

Using Theorem 9, we can bound the maximum information gain in terms of $B_k(T_*)$, the tail sum of the operator spectrum of a kernel k . Our first main theoretical result is shown in the following Proposition 10. Comparing with the results in [25], our bounds of γ_t for two commonly used covariance functions have a linear dependence with the total dimension d and exponential dependence with $\tilde{d}$, which is usually much smaller than d in a tree-structured function.

Proposition 10. *Let k_T be an Add-Tree covariance function defined in Definition 5 and is parametrized by lengthscales θ . Scaling lengthscales of k_T and norm bound at iteration t by Equation (13).*

- *If every covariance function k_{ij} defined on vertex v_{ij} is a squared exponential kernel, then*

$$\gamma_t = \mathcal{O}(d \tilde{d}^{\tilde{d}} (\log t)^{\tilde{d}+1} g(t)^{2\tilde{d}})$$

- *If every covariance function k_{ij} defined on vertex v_{ij} is a Matérn kernel, then*

$$\gamma_t = \mathcal{O}(M^2 g(t)^{2\nu+\tilde{d}} (\log t) t^{\frac{\tilde{d}(\tilde{d}+1)}{2\nu+\tilde{d}(\tilde{d}+1)}})$$

Proposition 11. *Under the assumptions of Proposition 10, $b(t)$ and $g(t)$ grow unbounded, then we obtain the following high-probability bounds using our proposed Add-Tree covariance function:*

- *Squared exponential kernel:*

$$\begin{aligned} R_t &= \mathcal{O} \left(\sqrt{t} g(t)^{2\tilde{d}} b(t) M \sqrt{d \tilde{d}^{\tilde{d}} (\log T)^{\tilde{d}+1}} \right. \\ &\quad \left. + \sqrt{t} g(t)^{2\tilde{d}} d \tilde{d}^{\tilde{d}} (\log T)^{\tilde{d}+1} \right) \end{aligned}$$

- *Matérn kernel:*

$$\begin{aligned} R_t &= \mathcal{O} \left(M^2 b(t) \sqrt{t g(t)^{2\nu+3\tilde{d}} (\log t) t^{1-\tau\tilde{d}}} \right. \\ &\quad \left. + M^2 g(t)^{2\nu+\tilde{d}} (\log t) t^{1-\tau\tilde{d}} \sqrt{t} \right) \end{aligned}$$

where $\tau = 2\nu\tilde{d}/(2\nu + \tilde{d}(1 + \tilde{d}))$.

Proposition 11 should be compared with the results given in [25], where bounds of R_t have an exponential dependence with d , while our bounds have an exponential dependence with $\tilde{d}$. This improvement is a result of the assumed additive structure. Since we are adapting the norm bound and the lengthscale, and we use the additive assumption, our regret bounds in Proposition 11 are not directly comparable to the results stated in [2]. Fortunately, insights can be gained by comparing the forms of Proposition 11

and the results in [2]. In the initial stage, the true function is not contained in the current space, since the misspecified norm bound B_t or the RKHS space associated with k_{θ_t} are too small. In this case, we can bound the regret by $2B$ [25]. Since $g(t)$ and $b(t)$ grow unbounded, B_t can be larger and θ_t can be smaller than their optimal values. Consider the case for a squared exponential kernel, the additional factor $g(t)^{2\tilde{d}} b(t)$ in the first term leads to additional exploration and is the price we pay for our ignorance of the true norm bound and the lengthscale parameter.

4.4 Practical Considerations

Although Proposition 11 holds for $b(t)$ and $g(t)$ growing unbounded, what we really want is an asymptotically no-regret algorithm:

$$\lim_{t \rightarrow \infty} R_t/t = 0$$

To achieve this goal, we need to choose $b(t)$ and $g(t)$ from a narrower class and make $\sqrt{\beta_t \gamma_t}$ grow slower than $\sqrt{t}$. Several methods to choose $b(t)$ and $g(t)$ are given in [25]. In brief, (i) we first select a reference regret, reflecting our willingness to perform additional exploration. (ii) Then we match an estimator of the cumulative regret with this reference regret. (iii) finally we compute the $g(t)$ and $b(t)$.

Two natural questions are whether we can test the validity of the additive assumption, and what to do if the additive assumption doesn't hold. First, the cornerstone of BO is a probabilistic regression model and BO inherits all the limitations of this regression model. That means there is little we can do in a high-dimensional space, even with hundreds of thousands of observations [16], unless we use dimension reduction techniques. It is thus generally unreliable to test if the additive assumption is satisfied in BO settings, because we are almost always in a low-data regime. Additionally, automatically discovering an additive structure based on a few hundreds of observations in a high-dimensional space could lead to serious overfitting. The second question has nothing to do with the inference procedure proposed in this work, but rather a model assumption problem. It is always possible to fit a GP model using our proposed Add-Tree covariance function based on the observations collected in the initial phase of BO. If the mean squared error is sufficient, we can continue the BO.

5 EXPERIMENTS

In this section, we present empirical results using four sets of experiments. To demonstrate the efficiency of our proposed Add-Tree covariance function, we first optimize a synthetic function presented in [8] in Section 5.1, comparing with SMAC [4], TPE [10], random search [26], standard GP-based BO from GPyOpt [27], and the semi-parametric approach proposed in [8]. To facilitate our following description, we refer to the above competing algorithms as **smac**, **tpe**, **random**, **gpyopt**, and **tree** respectively. Our approach is referred as **add-tree**. To verify our Add-Tree covariance function indeed enables sharing between different paths, we compare it with independent GPs in a regression setting showing greater sample efficiency for our method. We then apply our method to compress a three-layer fully connected neural network, VGG16 and ResNet50 in Section 5.2 and

Section 5.3. Finally, in Section 5.4, we search for activation functions of ResNet20 and their corresponding hyperparameters. Experimental results show our method outperforms all competing algorithms.

For all GP-based BO, including **gpyopt**, **tree** and **add-tree**, we use the squared exponential covariance function. To optimize the parameters of Add-Tree, we maximize the marginal log-likelihood function of the corresponding GP and set $\theta_t = \min(\theta_t^{\text{MAP}}, \theta_0/g(t))$ following [25]. As for the numerical routine used in fitting the GPs and optimizing the acquisition functions, we use multi-started L-BFGS-B, as suggested by [28].

In our experiments, we use official implementations of competing algorithms. For **tpe**, we use Hyperopt [29]. For **smac**, we use SMAC3³. We note the original code for [8] is unavailable, thus we have re-implemented their framework to obtain the results presented here. There are several hyperparameters in their algorithm which are not specified in the publication. To compare fairly with their method, we tune these hyperparameters such that our implementation has a similar performance on the synthetic functions to that reported by [8], and subsequently fix the hyper-parameter settings in the model compression task in Section 5.2.

5.1 Synthetic Experiments

In our first experiment, we optimize the synthetic tree-structured function depicted in Figure 3. This function is originally presented in [8]. In this example, continuous variables include $x_4, x_5, x_6, x_7, r_8, r_9$, and r_8, r_9 of them are shared variables. x_4, x_5, x_6, x_7 are defined in $[-1, 1]$ and r_8, r_9 are bounded in $[0, 1]$. Categorical variables include x_1, x_2, x_3 and they are all binary.

Fig. 3. A synthetic function from [8]. The dimension of this function is $d = 9$ and the effective dimension at any leaf is 2.

Figure 4a shows the optimization results of six different methods. The x-axis shows the iteration number and the y-axis shows the $\log_{10}$ distance between the best minimum achieved so far and the known minimum value, which in this case is 0.1. It is clear from Figure 4a that our method has a substantial improvement in performance compared with other algorithms. After 60 iterations, the $\log_{10}$ distance of **tree** is still higher than -4, while our method can quickly obtain a much better performance in less than 20 iterations. Interestingly, our method performs substantially better than independent GPs, which will be shown later, while in [8], their algorithm is inferior to independent GPs. We note **gpyopt**⁴ performs worst, and this is expected (recall

Section 1.1.1). **gpyopt** encodes categorical variables using a one-hot representation, thus it actually works in a space whose dimension is $d' = d + c = 12$, which is relatively high considering we have less than 100 data points. In this case, **gpyopt** behaves like random search, but in a 12-dimensional space instead of the 9-dimensional space of a naïve random exploration.

To show that Add-Tree allows efficient information sharing across different paths, we compare it with independent GPs and **tree** in a regression setting and results are shown in Figure 4b. The training data is generated from the synthetic function in Figure 3: categorical values are generated from a Bernoulli distribution with $p = 0.5$, continuous values are uniformly generated from their domains. The x-axis shows the number of generated training samples and the y-axis shows the $\log_{10}$ of Mean Squared Error (MSE) on a separate randomly generated data set with 50 test samples. From Figure 4b, we see that Add-Tree models the response surface better even though independent GPs have more parameters. For example, to obtain a test performance of 10^{-4} , Add-Tree needs only 24 observations, while independent GPs require 44 data points. If we just look at the case when we have 20 training samples, the absolute MSE of independent GPs is 10^{-1} , while for Add-Tree, it is 10^{-3} . The reason for such a huge difference is when training data are rare, some paths will have few data points, and Add-Tree can use the shared information from other paths to improve the regression model. This property of Add-Tree is valuable in BO settings.

5.2 Model Compression

Neural network compression is essential when it comes to deploying models on devices with limited memory and low computational resources. For parametric compression methods, like Singular Value Decomposition (SVD) and unstructured weight pruning (P-L1), it is necessary to tune their parameters in order to obtain the desired trade-off between model size and performance. Existing publications on model compression usually determine parameters for a single compression method, and do not have an automated selection mechanism over different methods. By encoding this problem using a tree-structured function, different compression methods can now be applied to different layers and this formulation is more flexible than the current literature.

In this experiment, we apply our method to compress a three-layer fully connected network FC3 defined in [30]. FC3 has 784 input nodes, 10 output nodes, 1000 nodes for each hidden layer and is pre-trained on the MNIST dataset. For each layer, we optimize a discrete choice for the compression method, which is among SVD and P-L1, then optimize either the rank of the SVD or the pruning threshold of P-L1. We only compress the first two layers, because the last layer occupies 0.56% of total weights. The rank parameters are constrained to be in $[10, 500]$ and the pruning threshold parameters are bounded in $[0, 1]$. The tree structure of this problem has a depth of 3, and there is no continuous parameter associated with the root. Following [30], the objective function used in compressing FC3 is:

$$\gamma \mathcal{L}(\tilde{f}_\theta, f^*) + R(\tilde{f}_\theta, f^*), \quad (17)$$

3. <https://github.com/automl/SMAC3>

4. GPyOpt [27] is a state-of-the-art open source Bayesian optimization software package with support for categorical variables.

Fig. 4. Comparison results on minimizing a synthetic function and compressing FC3 using different algorithms. The solid line is the incumbent mean of 10 independent runs, and the shaded region is twice the standard deviation (95 % confidence interval for mean) of these 10 runs. Figure 4a and Figure 4c shows six different algorithms on optimizing the synthetic function and compressing a simple 3-layer neural network respectively. Figure 4b shows the performance comparison in a regression setting, indicating our proposed Add-Tree covariance function indeed enables information sharing between different paths.

TABLE 1
Wilcoxon Signed-Rank Test

Experiment	Iter	smac	tpe	random	gpyopt	tree
synthetic function	40	0.003	0.030	0.005	0.003	0.018
	60	0.003	0.003	0.003	0.003	0.003
	80	0.003	0.003	0.003	0.003	0.003
model compression	40	0.101	0.023	0.101	0.003	0.014
	60	0.037	0.018	0.011	0.003	0.008
	80	0.166	0.005	0.003	0.003	0.006

where f^* is the original FC3, $\tilde{f}_\theta$ is the compressed model using parameter θ , $R(\tilde{f}_\theta, f^*)$ is the compression ratio and is defined to be the number of weights in the compressed network divided by the number of weights in the original network. $\mathcal{L}(\tilde{f}_\theta, f^*) := \mathbb{E}_{x \sim P}(\|\tilde{f}_\theta(x) - f^*(x)\|_2^2)$, where P is an empirical estimate of the data distribution. Intuitively, the R term in Equation (17) prefers a smaller compressed network, the $\mathcal{L}$ term prefers a more accurate compressed network and γ is used to trade off these two terms. In this experiment, γ is fixed to be 0.01, and the number of samples used to estimate $\mathcal{L}$ is 50 following [30]. Figure 4c shows the results of our method (Add-Tree) compared with other methods. For this experiment, although **smac**, **tpe** and **tree** all choose SVD for both layers at the end, our method converges significantly faster, once again demonstrating our method is more sample-efficient than other competing methods. We note **gpyopt** has the worst performance among all other competing methods in this experiment.

Table 1 shows the results of pairwise Wilcoxon signed-rank tests⁵ for the above two objective functions at different iterations. In Table 1, **addtree** almost always performs significantly better than other competing methods (significance level $\alpha = 0.05$), while no method is significantly better than ours.

5. We used the Wilcoxon signed rank implementation from `scipy.stats.wilcoxon` with option `(alternative == 'greater')`.

5.3 Model Pruning

In this section, we perform model pruning on pre-trained VGG16 [31] and ResNet50 [32] on the CIFAR-10 dataset [33] and prune the weights of convolutional layers in VGG16 and ResNet50. Similar to Section 5.2, for each layer, the pruning method is automatically chosen among the unstructured weight pruning (P-L1) as in Section 5.2, and the L2 structured weight pruning (P-L2), which prunes output channels with the lowest L2 norm. For ResNet50, we don't prune its first block because it occupies only 0.91% of total convolutional weights. A block in ResNet50 is composed of several bottlenecks, and to prune a bottleneck in a block, we use 4 continuous parameters, corresponding to 4 convolutional layers in this bottleneck. To prune a block, we repeat this process for each bottleneck in this block. In this experiment, these 4 continuous parameters are shared for all bottlenecks in one block, and there are 7 categorical parameters, 56 continuous parameters.⁶ For VGG16, we group all convolutional layers into 4 blocks in a similar fashion with ResNet50. The first block contains the first 4 convolutional layers and later blocks contain 3 consecutive convolutional layers. We don't prune the first block because it takes up only 1.77% of total convolutional weights. To prune each block, we use 3 continuous parameters, and there are 7 categorical parameters, 42 continuous parameters in this problem. The tree structure of pruning VGG16 has a depth of 4, and each node in this tree except the root is associated with a 3-dimensional continuous parameter. The tree of pruning ResNet50 has a similar structure to the tree of pruning VGG16, except each node except the root is associated with a 4-dimensional continuous variable.

In Section 5.1 and Section 5.2, it has been shown **gpyopt** doesn't perform well in high-dimensional problems, and **tree** isn't competitive with other algorithms, even worse than **random** in a model compression problem. Thus in this section, we omit the comparison with these two algorithms.

The objective function used in this experiment is a trade-off between the model sparsity and the top-1 accu-

6. We use the parameter counting method in SMAC. Using different counting method could lead to different number of categorical parameters, but it doesn't make a difference in practice.

racy on a separate validation set: $S(\tilde{f}_\theta, f^*) + T(\tilde{f}_\theta)$, where $S(\tilde{f}_\theta, f^*) = 1 - R(\tilde{f}_\theta, f^*)$, $T(\tilde{f}_\theta)$ is the top-1 accuracy of the compressed model after finetuning for two epochs. This objective function is bounded above by 2. In a concrete application, it is possible to add a trade-off parameter between $S(\tilde{f}_\theta, f^*)$ and $T(\tilde{f}_\theta)$, depending on whether we care about the size or the accuracy of the pruned model. We emphasize that the focus of this work is the comparison with other algorithms in conditional parameter spaces, instead of the comparison with other existing pruning methods.

We run all algorithms for 300 iterations and the first 50 iterations are random initialization. The results of pruning VGG16 and ResNet50 are shown in Figure 5. We note the public **smac** implementation cannot handle problems with more than 40 hyper-parameters. To make a comparison possible, we fix some continuous parameters to **smac**'s default value 0.5. This is of course a biased choice, and changing these fixed parameters to other values may give a better or worse performance. Choosing different hyper-parameters to freeze could also have a huge impact on the performance. From Figures 5a and 5d, it is clear our proposed Add-Tree performs much better than other competing algorithms. The distribution of optimized objective values from 10 independent runs are shown in Figures 5b and 5e, which are zoomed-in versions of Figures 5a and 5d at the 300-th iteration, respectively. In Figures 5c and 5f, we show the blockwise compression analysis of pruning VGG16 and ResNet50 and we see later blocks have much more verbosity than earlier blocks. At the 300-th iteration, the average running time of different methods is shown in Table 2. The reason **random** takes longer time than **tpe** is the implementation of **random** comes from SMAC3 and their exist some computational overheads. Regarding the average time of running all algorithms for 300 iterations, **addtree** is comparable to **smac** and **tpe** is comparable to **random**. Since **addtree** can achieve a comparable performance at the 100-th iteration, we also show the average running time of **addtree** at the 100-th iteration in parentheses.

TABLE 2
Average Running Time (s) at the 300-th Iteration

Net	addtree	random	smac	tpe
VGG16	20282 (7391)	15889	20153	15303
ResNet50	65494 (20205)	58477	68824	55877

As stated earlier, TPE doesn't model the dependencies between parameters, and effectively performs multiple non-parametric one-dimensional regressions, each regression corresponding to one continuous parameter. The results shown here should be compared with the synthetic experiment in Section 5.1 and the simple model compression in Section 5.2. In Figure 4a, **tpe** performs well among all competing algorithms and this is because parameters of the synthetic function in Figure 3 are indeed independent with each other. In Figure 4c, **tpe** is also very competitive, possibly because for every layer in FC3, there is only one continuous parameter. While in pruning VGG16 and ResNet50, **tpe** performs much worse than **addtree**, and this indicates there are complex interactions between parameters in one block.

In Table 3, we show the optimized discrete choices of pruning methods for each block. M1, M2, M3 indicate the selected pruning methods for the second, the third, and the forth block respectively. From Table 3, it is clear that using a common compression method across all layers is not an optimal choice.

TABLE 3
Optimized Pruning Methods at the End of Optimization

Algo	M1	M2	M3	Obj	Model
addtree	L1	L2	L1	1.88	VGG16
addtree	L1	L1	L1	1.86	ResNet50
random	L1	L1	L1	1.78	VGG16
random	L1	L1	L1	1.77	ResNet50
smac	L1	L2	L2	1.81	VGG16
smac	L1	L1	L1	1.75	ResNet50
tpe	L1	L2	L1	1.84	VGG16
tpe	L2	L1	L1	1.81	ResNet50

5.4 Neural Architecture Search

Neural architecture search (NAS) tries to automate the process of designing novel and efficient neural network architectures. In this section, we apply our method to search for activation functions of ResNet20 and their corresponding hyper-parameters on CIFAR-10 dataset. ResNet20 is composed of three main blocks, each of which can further be decomposed into several smaller basic blocks, and every basic block has two ReLU activation functions. An interesting question is whether using other activation functions can improve the performance. In this experiment, we limit the activation function to be exponential linear unit (ELU) [34] or Leaky ReLU (LReLU) [35], and all basic blocks in one main block share the activation function and corresponding hyper-parameters. The final search space is parametrized by: (1) discrete choices of activation functions, which could be ELU or LReLU; (2) corresponding hyper-parameters (slope parameter for LReLU, α for ELU), which are 2 continuous parameters constrained in $[0, 1]$ in each main block. For this problem, there are 7 categorical parameters and 28 continuous parameters. The tree structure of this problem has a depth of 4, and each node except the root is associated with a 2-dimensional continuous parameter.

We run all algorithms for 100 iterations and the first 50 iterations are random initialization. Every iteration in the hyper-parameter optimization is as follows. First we use the sampled (initialization stage) or optimized (optimization stage) hyper-parameters to construct a ResNet20. Then we train this network on CIFAR-10. For learning, we use stochastic gradient descent with momentum with mini-batches of 128 samples for 10 epochs. The start learning rate is set to be 0.1 and is decreased down to be 0.01 and 0.001 in the beginning of the 7-th epoch and 9-th epoch respectively. The momentum is fixed to be 0.9. Last we test this trained network and use its accuracy on the validation set of CIFAR-10 as our objective.

The results of optimizing ResNet20 is shown in Figure 6. We see our proposed Add-Tree is significantly better than **tpe** and **random**. Since there is a overlapping between **smac** and **addtree** in Figure 6b, we cannot claim our method performs significantly better than **smac**. Figure 6c shows

Fig. 5. Comparison results on pruning VGG16 and ResNet50 using different algorithms. For VGG16, there are 49 hyper-parameters in total, 42 of which are continuous parameters lying in $[0, 1]$. For ResNet50, there are 63 hyper-parameters, 56 of which are continuous parameters lying in $[0, 1]$. In the first column (Figures 5a and 5d), the solid line is the incumbent median of 10 independent runs, and the shaded region is 95% confidence interval of median. The second column (Figures 5b and 5e) shows the best results among all iterations. The third column shows the per-block compression ratio for VGG16 and ResNet50 of **addtree** using the best parameters.

the running time of different algorithms at the 100-th iteration. In this experiment, the running time of all methods is comparable, and this is because we run all algorithms for only 100 iterations and evaluating the expensive objective function dominates the inference time of the random forest, the GP and the Parzen estimator in **smac**, **addtree** and **tpe** respectively. The existence of several extreme outliers in Figure 6c is due to we running these algorithms in a computer cluster managed by a software called HTCondor, which can assign jobs to very fast or very slow machines in a totally random way.

TABLE 4
Optimized Hyper-parameters Using Our Method

Block	Act	#1 Param	#2 Param
Block 1	ELU	0.01	0.75
Block 2	LReLU	0.01	0.99
Block 3	ELU	0.01	0.01

Table 4 shows the optimized hyper-parameters using our algorithm. The first row of Table 4 means for every basic block in the first main block in ResNet20, ELU is

selected, and the hyper-parameters α of the first ELU and the second ELU are 0.01 and 0.7484 respectively. During the optimization, we train ResNet20 for only 10 epochs in order to speed up this process, and the accuracy shown in Figure 6a underestimate the true accuracy. Thus we construct a new ResNet20 based on the optimized hyper-parameters shown in Table 4, and train this network for 200 epochs. We decrease the learning to 0.1 and 0.001 in the beginning of the 100-th epoch and 150-th epoch respectively and also add a L2-weight decay regularization term, which is set to be 0.0001. The final test error is 7.74%, compared with the original test error (8.27%), our optimized ResNet20 performs much better.

6 CONCLUSION

In this work, we have designed a covariance function that can explicitly utilize the problem structure, and demonstrated its efficiency on a range of problems. In the low data regime, our proposed Add-Tree covariance function enables an explicit information sharing mechanism, which in turn makes BO more sample-efficient compared with other model based optimization methods. Contrary to other

Fig. 6. Comparison results of NAS on ResNet20 using different algorithms. There are 35 hyper-parameters in total, 28 of which are continuous parameters lying in $[0, 1]$. In Figure 6a, the solid line is the incumbent median of 10 independent runs, and the shaded region is 95% confidence interval of median. Figure 6b shows the best results among all iterations. In Figure 6c, we show the running time of different methods.

GP-based BO methods, we do not impose restrictions on the structure of a conditional parameter space, greatly increasing the applicability of our method. We also directly model the dependencies between different observations under the framework of Gaussian Processes, instead of placing parametric relationships between different paths, making our method more flexible. In addition, we incorporate this structure information and develop a parallel algorithm to optimize the acquisition function. For both components of BO, our proposed method allows us to work in a lower dimensional space compared with the dimension of the original parameter space.

Empirical results on an optimization benchmark function, a simple neural network compression problem, on pruning VGG16 and ResNet50, and on searching for activation functions of ResNet20 show our method significantly outperforms other competing model based optimization algorithms in conditional parameter spaces, including SMAC, TPE and [8].

ACKNOWLEDGMENTS

Xingchen Ma is supported by Onfido. This research received funding from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme.

REFERENCES

- [1] D. R. Jones, M. Schonlau, and W. J. Welch, “Efficient Global Optimization of Expensive Black-Box Functions,” *Journal of Global Optimization*, vol. 13, no. 4, pp. 455–492, Dec. 1998.
- [2] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, “Gaussian Process Optimization in the Bandit Setting: No Regret and Experimental Design,” *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, Jun. 2010.
- [3] E. Brochu, V. M. Cora, and N. de Freitas, “A Tutorial on Bayesian Optimization of Expensive Cost Functions, with Application to Active User Modeling and Hierarchical Reinforcement Learning,” *CoRR*, Dec. 2010.
- [4] F. Hutter, H. H. Hoos, and K. Leyton-Brown, “Sequential Model-Based Optimization for General Algorithm Configuration,” in *Learning and Intelligent Optimization*, ser. Lecture Notes in Computer Science, C. A. C. Coello, Ed. Springer Berlin Heidelberg, 2011, pp. 507–523.

- [5] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, “Taking the Human Out of the Loop: A Review of Bayesian Optimization,” *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, Jan. 2016.
- [6] P. I. Frazier, “A Tutorial on Bayesian Optimization,” *arXiv:1807.02811 [cs, math, stat]*, Jul. 2018.
- [7] A. Klein, S. Falkner, S. Bartels, P. Hennig, and F. Hutter, “Fast Bayesian Optimization of Machine Learning Hyperparameters on Large Datasets,” in *Artificial Intelligence and Statistics*, Apr. 2017, pp. 528–536.
- [8] R. Jenatton, C. Archambeau, J. González, and M. Seeger, “Bayesian Optimization with Tree-structured Dependencies,” in *International Conference on Machine Learning*, Jul. 2017, pp. 1655–1664.
- [9] X. Ma and M. B. Blaschko, “Additive tree-structured covariance function for conditional parameter spaces in Bayesian optimization,” in *Artificial Intelligence and Statistics (AISTATS)*, 2020.
- [10] J. S. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, “Algorithms for Hyper-Parameter Optimization,” in *Advances in Neural Information Processing Systems 24*, J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2011, pp. 2546–2554.
- [11] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, ser. Adaptive Computation and Machine Learning. Cambridge, Mass: MIT Press, 2006.
- [12] P. Frazier, W. Powell, and S. Dayanik, “The Knowledge-Gradient Policy for Correlated Normal Beliefs,” *INFORMS Journal on Computing*, vol. 21, no. 4, pp. 599–613, May 2009.
- [13] P. Hennig and C. J. Schuler, “Entropy Search for Information-efficient Global Optimization,” *J. Mach. Learn. Res.*, vol. 13, no. 1, pp. 1809–1837, Jun. 2012.
- [14] Z. Wang, M. Zoghi, F. Hutter, D. Matheson, and N. De Freitas, “Bayesian Optimization in High Dimensions via Random Embeddings,” in *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, ser. IJCAI ’13. AAAI Press, 2013, pp. 1778–1784.
- [15] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, “Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization,” *arXiv:1603.06560 [cs, stat]*, Mar. 2016.
- [16] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk, *A Distribution-Free Theory of Nonparametric Regression*, ser. Springer Series in Statistics. New York, NY: Springer New York, 2002.
- [17] F. Hutter and M. A. Osborne, “A Kernel for Hierarchical Parameter Spaces,” *arXiv:1310.5738 [cs, stat]*, Oct. 2013.
- [18] K. Swersky, D. Duvenaud, J. Snoek, F. Hutter, and M. A. Osborne, “Raiders of the Lost Architecture: Kernels for Bayesian Optimization in Conditional Parameter Spaces,” *arXiv:1409.4011 [stat]*, Sep. 2014.
- [19] J. Lévesque, A. Durand, C. Gagné, and R. Sabourin, “Bayesian optimization for conditional hyperparameter spaces,” in *2017 International Joint Conference on Neural Networks (IJCNN)*, May 2017, pp. 286–293.

- [20] D. Duvenaud, H. Nickisch, and C. E. Rasmussen, "Additive gaussian processes," in *Proceedings of the 24th International Conference on Neural Information Processing Systems*, ser. NIPS'11. USA: Curran Associates Inc., 2011, pp. 226–234. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2986459.2986485>
- [21] K. Kandasamy, J. Schneider, and B. Poczos, "High Dimensional Bayesian Optimisation and Bandits via Additive Models," in *International Conference on Machine Learning*, Jun. 2015, pp. 295–304.
- [22] J. Gardner, C. Guo, K. Weinberger, R. Garnett, and R. Grosz, "Discovering and Exploiting Additive Structure for Bayesian Optimization," in *Artificial Intelligence and Statistics*, Apr. 2017, pp. 1311–1319.
- [23] P. Rolland, J. Scarlett, I. Bogunovic, and V. Cevher, "High-dimensional bayesian optimization via additive models with overlapping groups," in *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Storkey and F. Perez-Cruz, Eds., vol. 84. Playa Blanca, Lanzarote, Canary Islands: PMLR, 09–11 Apr 2018, pp. 298–307. [Online]. Available: <http://proceedings.mlr.press/v84/rolland18a.html>
- [24] N. Cesa-Bianchi and G. Lugosi, *Prediction, Learning, and Games*. Cambridge ; New York: Cambridge University Press, 2006, oCLC: ocm62493289.
- [25] F. Berkenkamp, A. P. Schoellig, and A. Krause, "No-Regret Bayesian Optimization with Unknown Hyperparameters," *Journal of Machine Learning Research*, vol. 20, no. 50, pp. 1–24, 2019.
- [26] J. Bergstra and Y. Bengio, "Random Search for Hyper-parameter Optimization," *J. Mach. Learn. Res.*, vol. 13, pp. 281–305, Feb. 2012.
- [27] The GPyOpt authors, "GPyOpt: A Bayesian optimization framework in python," 2016. [Online]. Available: <http://github.com/SheffieldML/GPyOpt>
- [28] J. Kim and S. Choi, "On Local Optimizers of Acquisition Functions in Bayesian Optimization," *arXiv:1901.08350 [cs, stat]*, Jan. 2019.
- [29] J. Bergstra, D. Yamins, and D. Cox, "Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures," in *International Conference on Machine Learning*, Feb. 2013, pp. 115–123.
- [30] X. Ma, A. Rannen Triki, M. Berman, C. Sagonas, J. Cali, and M. B. Blaschko, "A Bayesian optimization framework for neural network compression," in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019.
- [31] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," *arXiv:1409.1556 [cs]*, Apr. 2015.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," *arXiv:1512.03385 [cs]*, Dec. 2015.
- [33] A. Krizhevsky, "Learning multiple layers of features from tiny images," *Tech. Rep.*, 2009.
- [34] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, "Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)," *arXiv:1511.07289 [cs]*, Feb. 2016.
- [35] A. L. Maas, "Rectifier nonlinearities improve neural network acoustic models," 2013.
- [36] M. J. Wainwright, *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, 1st ed. Cambridge University Press, Feb. 2019.
- [37] W. Fulton, "Eigenvalues, invariant factors, highest weights, and Schubert calculus," *arXiv:math/9908012*, Mar. 2000.
- [38] M. W. Seeger, S. M. Kakade, and D. P. Foster, "Information Consistency of Nonparametric Gaussian Process Methods," *IEEE Transactions on Information Theory*, vol. 54, no. 5, pp. 2376–2382, May 2008.

Xingchen Ma is a PhD student in the Electrical Engineering department at KU Leuven.

Matthew Blaschko is a faculty member in the Electrical Engineering department at KU Leuven.