
On Uncertainty Estimation by Tree-based Surrogate Models in Sequential Model-based Optimization

Jungtaek Kim
POSTECH

Seungjin Choi
BARO AI

Abstract

Sequential model-based optimization sequentially selects a candidate point by constructing a surrogate model with the history of evaluations, to solve a black-box optimization problem. Gaussian process (GP) regression is a popular choice as a surrogate model, because of its capability of calculating prediction uncertainty analytically. On the other hand, an ensemble of randomized trees is another option and has practical merits over GPs due to its scalability and easiness of handling continuous/discrete mixed variables. In this paper we revisit various ensembles of randomized trees to investigate their behavior in the perspective of prediction uncertainty estimation. Then, we propose a new way of constructing an ensemble of randomized trees, referred to as *BwO forest*, where bagging with oversampling is employed to construct bootstrapped samples that are used to build randomized trees with random splitting. Experimental results demonstrate the validity and good performance of BwO forest over existing tree-based models in various circumstances.

1 INTRODUCTION

Sequential model-based optimization (SMO) (Brochu et al., 2010; Hutter et al., 2011) constructs a statistical surrogate model in order to estimate a function value and its uncertainty – both estimates are employed to balance a trade-off between exploitation and exploration. To determine where next to evaluate carefully, a surrogate model is one of key components in SMO (Bodin et al., 2020). In a common

setting of SMO including the formulation of Bayesian optimization (Srinivas et al., 2010; Azimi et al., 2010; Snoek et al., 2012), Gaussian process (GP) regression (O’Hagan, 1978; Williams and Rasmussen, 1996) is a popular choice as a surrogate model due to its flexibility and expressibility (Rasmussen and Williams, 2006). However, it requires the assumption on smoothness, which can induce a *mismatch on the smoothness degree* for an objective of interest (Schulz et al., 2016), and moreover its exact complexity over the number of the query points already evaluated scales *cubically* (Rasmussen and Williams, 2006).

Instead of GP regression, random forest regression (Breiman, 2001) is another option and has practical merits as a surrogate model; sequential model-based algorithm configuration (Hutter et al., 2011) shows its strength in various real-world applications such as automated machine learning (Feurer et al., 2015), neural architecture search (Ying et al., 2019), and water distribution system (Candelieri et al., 2018). In particular, it *inherently deals with a categorical variable*, because a randomized tree – a base estimator of random forest – is capable of defining a split criterion for categorical variables without any complex techniques. Compared to GPs for categorical variables (e.g., using the Aitchison and Aitken kernel (Aitchison and Aitken, 1976) or using a random embedding to lower-dimensional space (Wang et al., 2016)),¹ it provides an easy-to-use implementation as well as reliable performance. Furthermore, a tree-based surrogate model tends to be robust in a *high-dimensional search space*, in comparison with GP regression. These strengths mentioned above let us pursue in-depth and thorough studies on randomized tree-based models.

In this paper, we investigate SMO with a random forest-based model and provide a new understanding of tree-based models. Under such an understanding, we suggest our strategies with sophisticated ensemble models of trees, e.g., Bayesian additive regression trees (BART) (Chipman et al., 2010), Mondrian forest (Lak-

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

¹Many GP-based approaches to solving this topic have been studied, but they are not the scope of this work.




Figure 1: Results with GP regression and tree-based surrogate models such as random forest (RF), extremely randomized trees (ERTs), BART, Mondrian forest (MF), NGBoost, and BwO forest (ours). We randomly sample 5 points (green \times), where a true function is the sine function (blue) and observation noises exist. Mean function (orange) and region with ± 1.96 standard deviation is plotted. Qualitative analyses are presented in Table 2.

shminarayanan et al., 2016), and NGBoost (Duan et al., 2020), and then propose a new tree-based surrogate model, named bagging with oversampling (BwO) forest. To be clear, the goal of this work is not to outperform sequential GP-based optimization in every task. We aim to re-examine the random forest-based approach in terms of prediction uncertainty estimation, discuss about SMO strategies with diverse tree-based surrogate models, and propose a new method that inherits *the nature of the ensemble of randomized trees* and also follows *the underlying intuition about uncertainties*.

To briefly present our main analyses on prediction uncertainty estimation by tree-based surrogate models, we demonstrate 1D examples using GP, random forest, extremely randomized trees (Geurts et al., 2006), BART, Mondrian forest, NGBoost, and BwO forest, as shown in Figure 1. While the results by random forest, BART, Mondrian forest, and NGBoost are distinct from the result by GP regression, our BwO forest yields satisfactory uncertainty estimation, which represents small variance on the region that the previous decisions have already been evaluated and large variance on an unexplored region. By the definition of uncertainties (Gal, 2016), we can interpret that random forest, BART, (Mondrian forest), and NGBoost rely on aleatoric uncertainty, which comes from the uncertainty of data such as a noise in data, and on the contrary, GP, (Mondrian forest), and BwO forest capture epistemic uncertainty as well.² Finally, such intuitive results lead us to obtain the global optimization results that BwO forest tends to be more beneficial than the other tree-based models in various circumstances defined on continuous, high-dimensional binary, and mixed search spaces. In addition to such results, our BwO forest consistently enjoys the advantage of computational efficiency, similar to some of tree-based surrogate models such as random forest and extremely randomized trees.

²The reason why Mondrian forest is included in both groups is that the results (e.g., Figure 3(b)) show that Mondrian forest also works well by capturing the uncertainty derived from the amount of the knowledge of data.

Our contributions are summarized as follows:

- (i) We investigate the characteristics of tree-based surrogate models (e.g., random forest, extremely randomized trees, BART, Mondrian forest, and NGBoost) in terms of prediction uncertainties;
- (ii) We propose a new ensemble of randomized trees, named BwO forest, which can elaborate uncertainty estimation and yield the satisfactory results that follow the intuition about prediction uncertainties;
- (iii) We employ various tree-based surrogate models including BwO forest as a component of SMO in solving diverse global optimization problems.

2 PREDICTION UNCERTAINTY ESTIMATION BY TREE-BASED SURROGATE MODELS

In this section, we begin by introducing the notation of tree-based surrogate models. Denote that a decision tree $\mathcal{T} = (\boldsymbol{\tau}, \boldsymbol{\delta}, \boldsymbol{\eta})$, where $\boldsymbol{\tau} = \{\boldsymbol{\tau}_r, \boldsymbol{\tau}_d, \boldsymbol{\tau}_l\}$ is the nodes of tree including the root node $\boldsymbol{\tau}_r$, all the decision nodes $\boldsymbol{\tau}_d$, and all the leaf nodes $\boldsymbol{\tau}_l$, $\boldsymbol{\delta}$ is a collection of all the split dimensions of the parent nodes, and $\boldsymbol{\eta}$ is a collection of the split locations thereof. A tree-based model $\hat{f} : \mathbb{R}^d \rightarrow \mathbb{R}$ is an ensemble of B decision trees $\{\mathcal{T}_b\}_{b=1}^B$, where N points $\mathbf{X} \in \mathbb{R}^{N \times d}$ and their evaluations $\mathbf{y} \in \mathbb{R}^N$ are given. For example, if we define a surrogate output as the average of all the outputs of base decision trees, the surrogate is defined as

$$\hat{f}(\mathbf{x}; \mathbf{X}, \mathbf{y}) = \frac{1}{B} \sum_{b=1}^B g(\mathbf{x}; \mathcal{T}_b, \mathbf{X}, \mathbf{y}), \quad (1)$$

where a function g guides a route to a leaf node over \mathbf{x} :

$$g(\mathbf{x}; \mathcal{T}, \mathbf{X}, \mathbf{y}) = \sum_{\boldsymbol{\tau} \in \boldsymbol{\tau}_l} \frac{\sum_{i=1}^N y_i \mathbf{1}_{\mathbf{x}_i \in \boldsymbol{\tau}}}{\sum_{i=1}^N \mathbf{1}_{\mathbf{x}_i \in \boldsymbol{\tau}}} \mathbf{1}_{\mathbf{x} \in \boldsymbol{\tau}}, \quad (2)$$

and $\mathbf{1}_{\mathbf{x} \in \boldsymbol{\tau}}$ is 1 if $\mathbf{x} \in \boldsymbol{\tau}$ is true; otherwise, it is 0. For brevity, we denote $\frac{\sum_{i=1}^N y_i \mathbf{1}_{\mathbf{x}_i \in \boldsymbol{\tau}}}{\sum_{i=1}^N \mathbf{1}_{\mathbf{x}_i \in \boldsymbol{\tau}}}$ by $\mu_{\boldsymbol{\tau}}$, which can be

preemptively computed using \mathbf{X} and \mathbf{y} . In addition to μ_τ , without loss of generality, the variance of node τ , σ_τ^2 can also be defined with μ_τ . As will be discussed, the definition of tree-based surrogate model can differ as to how we define its formulation; however we first describe the surrogate that defines as the form of *sum-of-trees model* (1).

2.1 Sum-of-Trees Models

Generic tree-based surrogate models such as bootstrap aggregating (bagging) (Breiman, 1996) with decision trees and random forest (Breiman, 2001), however, do not model a posterior predictive distribution over real-valued variables $p(y|\mathbf{x}, \mathbf{X}, \mathbf{y})$. To define a function prediction with its uncertainty for the sum-of-trees model, under the assumption that a joint distribution over all random variables is a multivariate normal distribution, the posterior predictive distribution is defined as

$$p(y|\mathbf{x}, \mathbf{X}, \mathbf{y}) = \mathcal{N}(y|\mu(\mathbf{x}; \{\mathcal{T}_b\}_{b=1}^B, \mathbf{X}, \mathbf{y}), \sigma^2(\mathbf{x}; \{\mathcal{T}_b\}_{b=1}^B, \mathbf{X}, \mathbf{y})), \quad (3)$$

where

$$\begin{aligned} \mu(\mathbf{x}; \{\mathcal{T}_b\}_{b=1}^B, \mathbf{X}, \mathbf{y}) &= \frac{1}{B} \sum_{b=1}^B \mu_b(\mathbf{x}) \\ &= \frac{1}{B} \sum_{b=1}^B \sum_{\tau \in \tau_{b,l}} \mu_\tau \mathbf{1}_{\mathbf{x} \in \tau}, \end{aligned} \quad (4)$$

$$\begin{aligned} \sigma^2(\mathbf{x}; \{\mathcal{T}_b\}_{b=1}^B, \mathbf{X}, \mathbf{y}) &= \frac{1}{B} \sum_{b=1}^B (\sigma_b^2(\mathbf{x}) + \mu_b^2(\mathbf{x})) - \mu(\mathbf{x}; \{\mathcal{T}_b\}_{b=1}^B, \mathbf{X}, \mathbf{y})^2 \\ &= \frac{1}{B} \sum_{b=1}^B \left(\left(\sum_{\tau \in \tau_{b,l}} \sigma_\tau \mathbf{1}_{\mathbf{x} \in \tau} \right)^2 + \left(\sum_{\tau \in \tau_{b,l}} \mu_\tau \mathbf{1}_{\mathbf{x} \in \tau} \right)^2 \right) \\ &\quad - \left(\frac{1}{B} \sum_{b=1}^B \mu_b(\mathbf{x}) \right)^2, \end{aligned} \quad (5)$$

by the law of total variance, as described in (Hutter et al., 2014). Note that $\tau_{b,l}$ is a set of leaf nodes for tree b . Hutter et al. (2011) have applied the formulation (3) in SMO, and it is straightforwardly employed to estimate a prediction uncertainty using BART and Mondrian forest.

The uncertainty of such an ensemble model is derived from the randomness of individual trees, which is achieved by one or more of these techniques:

- (i) bagging: it samples a bootstrap sample from \mathbf{X} with replacement and then aggregates base estimators;

- (ii) random feature selection: this technique randomly selects δ from a set of dimensions;
- (iii) random selection of split locations: it randomly selects η between lower and upper bounds of the selected dimension;
- (iv) random tree sampling: this strategy randomly samples a tree under the assumption on a prior distribution over trees.

As shown in Figure 1, these techniques are effective in estimating an uncertainty. However, compared to the result by GP regression, the results by random forest, BART, and Mondrian forest (see Figure 1(b), Figure 1(d), and Figure 1(e), respectively) tend not to follow the underlying property of epistemic uncertainty, which has a small uncertainty on the region that the previous decisions by SMO have already been evaluated and a large uncertainty on the region that has not been explored yet. This property of epistemic uncertainty is required to explore an unseen region effectively. Before explaining why it occurs, we specify all the algorithms based on their respective original papers; random forest employs (i) and (ii), BART employs (i), (ii), and (iv), and Mondrian forest employs (i) and (iii); see the Leo Breiman’s seminal work and the corresponding original references for the details of these algorithms.

By (5), the results that do not follow the underlying property of epistemic uncertainty imply that two adjacent points have the same variance; formally, given two adjacent points \mathbf{x}, \mathbf{x}' where $\|\mathbf{x} - \mathbf{x}'\| < \varepsilon$ for $0 < \varepsilon \ll 1$, the following equation $|\sigma^2(\mathbf{x}; \{\mathcal{T}_b\}_{b=1}^B, \mathbf{X}, \mathbf{y}) - \sigma^2(\mathbf{x}'; \{\mathcal{T}_b\}_{b=1}^B, \mathbf{X}, \mathbf{y})| = 0$ is satisfied almost everywhere. This consequence is mainly induced due to the deterministic selection of split locations. Although there exist a large enough number of distinct bootstrap samples – we can choose B bootstrap samples among $\binom{2N-1}{N-1}$ bootstrap samples (Holmes, 2004), for example, if $N = 10$, there exist 92,378 bootstrap samples, possible aggregations of base estimators are finite and the aforementioned equation is satisfied almost everywhere. However, in addition to this statement, we need to explain the result by Mondrian forest, which uses the technique, random selection of split location (Geurts et al., 2006) but tend not to follow the intuition about uncertainties; also see Figure 7. An appropriate explanation is that this outcome is derived from a bootstrapping technique, which makes a surrogate model underfit due to random sampling with replacement. These understandings lead us to propose BwO forest, as will be presented in Section 3.

Before proposing our method, we first review a recent study on a sophisticated tree-based surrogate model.

2.2 Gradient Boosting Models

Compared to a class of surrogate models described in Section 2.1, a more direct approach to estimating parameters has recently been proposed (Duan et al., 2020). This approach updates parameters θ (e.g., mean and variance) using their gradients in terms of the objective of parametric distribution (e.g., likelihood function or continuous ranked probability score). For example, one of potential objectives, a log likelihood function can be used to find its maximizer:

$$\mathcal{L}(\theta; \mathbf{X}, \mathbf{y}) = \sum_{i=1}^N \log p(y_i | \theta(\mathbf{x}_i)). \quad (6)$$

In particular, in (Duan et al., 2020), natural gradients (Amari, 1998) are used in updating θ in order to consider an appropriate distance between two parameter vectors, which is capable of representing the gradient direction in Riemannian space, and a gradient boosting machine (Friedman, 2001) with respect to the parameters is built. Duan et al. (2020) show that the numerical results with the gradient boosting machine updated by natural gradients, dubbed NGBoost, outperform the results by generic gradients.

This gradient boosting method over parameters is undoubtedly a reasonable approach to estimating parameters as a probabilistic regression model. However, such a multi-parameter boosting algorithm is not robust in a high-dimensional space, and relies on aleatoric uncertainties (Malinin et al., 2020).

3 ELABORATING UNCERTAINTY ESTIMATION BY TREE-BASED SURROGATE MODELS

Before elaborating uncertainty estimation by tree-based surrogate models, we introduce a context of how our proposed model is motivated. According to Section 2.1, if a split location is deterministic and bootstrapping is applied, an uncertainty is not estimated properly. Similar to these observations, Tang et al. (2018) have discussed that either no or severe subsampling leads to inconsistent forest construction, while they have not mentioned about the uncertainty of tree-based surrogate model. Here, we verify this issue with the property related to the number of unique original elements in a bootstrap sample. As pointed out in the work (Mendelson et al., 2016), the expectation and variance of an indicator for the existence of \mathbf{x}_i in

Algorithm 1 Training BwO Forest

Input: Size of ensemble model B , training data $\mathbf{X} \in \mathbb{R}^{N \times d}$, training function evaluations $\mathbf{y} \in \mathbb{R}^N$, size of bootstrap sample $M = \alpha N$ for the rate of over-sampling $\alpha > 1$.

Output: Set of decision trees $\{\mathcal{T}_b\}_{b=1}^B$

- 1: Initialize a set of decision trees.
 - 2: **for** $b = 1, \dots, B$ **do**
 - 3: Sample a bootstrap sample $\mathbf{B}_b \in \mathbb{R}^{M \times d}$ from \mathbf{X} ; the evaluations of \mathbf{B}_b are also stored using the indices that have already been sampled to construct \mathbf{B}_b and \mathbf{y} .
 - 4: Set a root node τ_r that contains all the elements in \mathbf{B}_b , and τ_r is set as the current split node.
 - 5: **while** a stopping criterion has not been met **do**
 - 6: Randomly choose a fixed number of split dimensions δ from all the feature dimensions $\{1, \dots, d\}$.
 - 7: Determine the best split by uniformly sampling a split location η between lower and upper bounds of the values along the selected dimensions δ .
 - 8: Split the current split node into two decision nodes using δ and η .
 - 9: Update all the parameters of decision tree, $(\tau_b, \delta_b, \eta_b)$.
 - 10: Choose the next split node from decision nodes τ_d included in τ_b .
 - 11: **end while**
 - 12: Update $\tau_b = \{\tau_r, \tau_d, \tau_l\}$ by determining leaf nodes τ_l .
 - 13: Update a set of decision trees by adding $\mathcal{T}_b = (\tau_b, \delta_b, \eta_b)$.
 - 14: **end for**
 - 15: **return** A set of decision trees $\{\mathcal{T}_b\}_{b=1}^B$
-

a bootstrap sample \mathbf{B} are expressed as

$$\begin{aligned} \mathbb{E}[1_{\mathbf{x}_i \in \mathbf{B}}] &= p(1_{\mathbf{x}_i \in \mathbf{B}} = 1) \\ &= 1 - p(1_{\mathbf{x}_i \in \mathbf{B}} = 0) \\ &= 1 - \left(1 - \frac{1}{N}\right)^M, \end{aligned} \quad (7)$$

$$\begin{aligned} \text{Var}[1_{\mathbf{x}_i \in \mathbf{B}}] &= p(1_{\mathbf{x}_i \in \mathbf{B}} = 0)p(1_{\mathbf{x}_i \in \mathbf{B}} = 1) \\ &= \left(1 - \frac{1}{N}\right)^M - \left(1 - \frac{1}{N}\right)^{2M}, \end{aligned} \quad (8)$$

where N is the size of \mathbf{X} and M is the size of a bootstrap sample. By (7) and (8), the distribution of unique original elements in a bootstrap sample \mathbf{B}

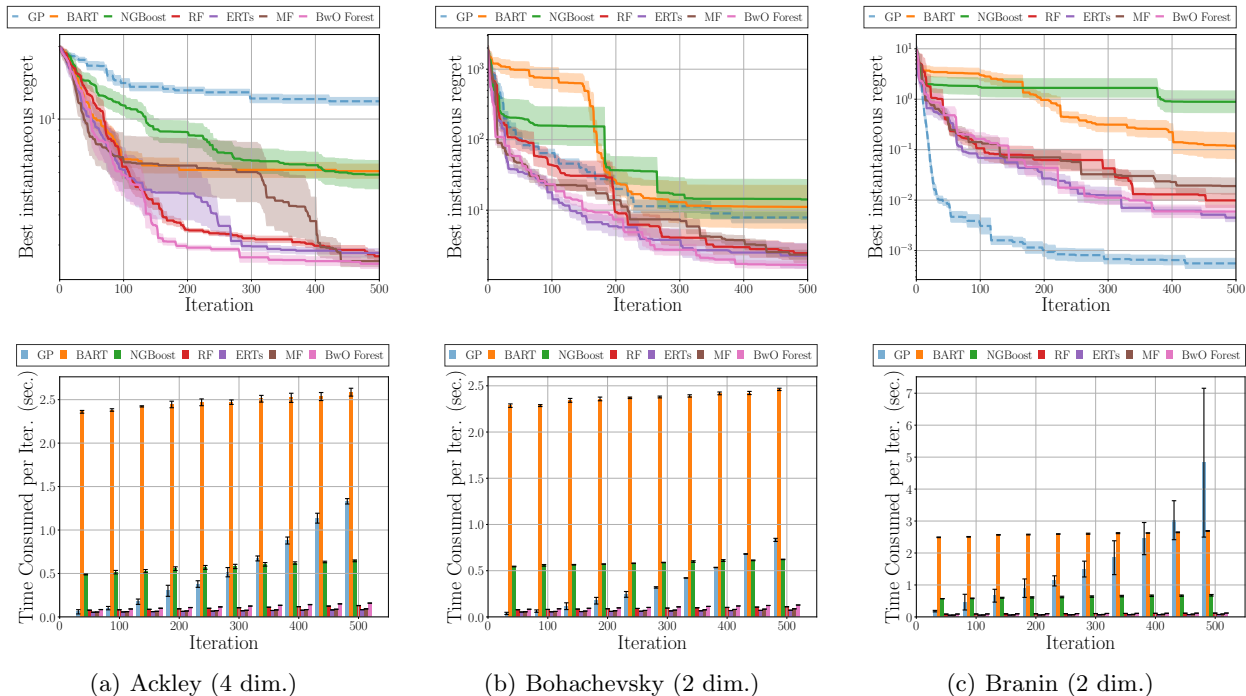


Figure 2: Results on various benchmark functions defined on continuous search spaces. The best instantaneous regret (left panel of each figure) and time consumed per iteration (right panel of each figure) versus iterations are plotted. All the runs are repeated 10 times. For the results on regrets, y -axis is set as a log-10 scale and relative log error bars are presented. For brevity, the results on runtime are plotted every 50 iterations.

can be described:

$$\mathbb{E}[|\text{unique}(\mathbf{B})|] = \mathbb{E}\left[\sum_{i=1}^N 1_{\mathbf{x}_i \in \mathbf{B}}\right] = N - \frac{(N-1)^M}{N^{M-1}}, \quad (9)$$

$$\begin{aligned} \text{Var}[|\text{unique}(\mathbf{B})|] &= \text{Var}\left[\sum_{i=1}^N 1_{\mathbf{x}_i \in \mathbf{B}}\right] \\ &= (N-1)\frac{(N-2)^M}{N^{M-1}} + \frac{(N-1)^M}{N^{M-1}} \\ &\quad - \frac{(N-1)^{2M}}{N^{2M-2}}, \end{aligned} \quad (10)$$

where $\text{unique}(\mathbf{B})$ filters duplicates and leaves unique original elements. For example, if $N = M = 5$, (9) and (10) return 3.362 and 0.509, respectively, and if $N = 5$ and $M = 20$, they are 4.942 and 0.055. Consequently, the distribution specified by (9) and (10) implies that combining two techniques, bagging and oversampling can help a tree construction process to fit well in \mathbf{X} .

From now, we propose our tree-based surrogate model, named BwO forest, which elaborates prediction uncertainty estimation by applying the technique, bagging with oversampling. As described in Algorithm 1, our BwO forest is trained by following a general pipeline of forest construction, given the size of ensemble model B , a training dataset (\mathbf{X}, \mathbf{y}) , and the size of bootstrap

samples $M = \alpha N$ for the rate of oversampling $\alpha > 1$. Note that a stopping criterion is satisfied when some pre-defined conditions such as maximum depth or the minimum number of elements in a node are encountered. To accommodate a page limit, we highlight the main components of our model. BwO forest utilizes random selection of split dimensions (Line 6) and random selection of split locations (Line 7), as well as bagging with oversampling (Line 3). Interestingly, as will be discussed in Section 5, all the components are important for appropriately estimating a function value and its uncertainty. As a prediction procedure, BwO forest estimates a function value and its uncertainty by computing (3), (4), and (5), where a set of trained decision trees $\{\mathcal{T}_b\}_{b=1}^B$ is given.

Finally, BwO forest is utilized as a surrogate model in the process of SMO; see Section B for the details. Compared to SMO with GP regression, it does not need a step for optimizing a kernel hyperparameter of GP, which is one of the most time-consuming steps in the corresponding procedure. Therefore, our method is more efficient than the GP-based approach, and besides it produces more satisfactory uncertainty estimation than the models with other tree-based surrogate models. More detailed empirical analyses can be found in the subsequent sections.

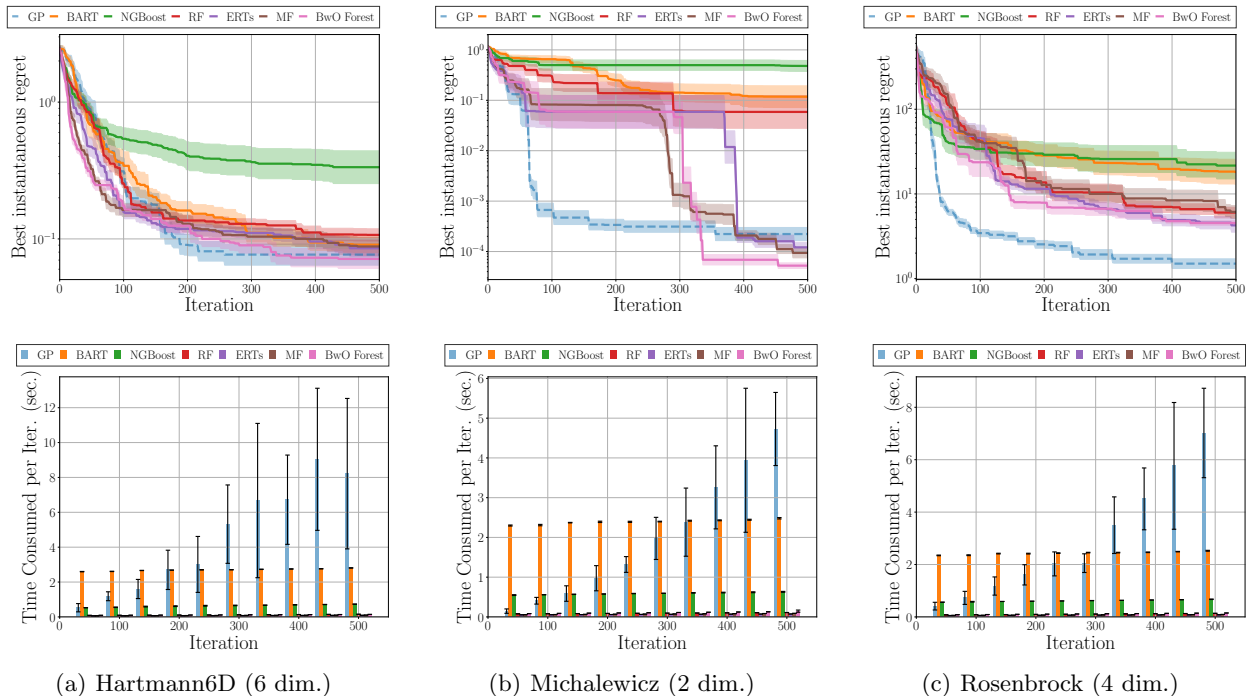


Figure 3: Results on various benchmark functions defined on continuous search spaces. All the experimental settings follow the settings described in Figure 2.

4 EXPERIMENTAL RESULTS

In this section, we show the experimental results on continuous, high-dimensional binary, and mixed search spaces using SMO strategies with diverse tree-based surrogate models and GP regression. Experimental setup is described below.

Experimental setup. We use the implementation of decision trees and ensemble methods, included in scikit-learn (Pedregosa et al., 2011), and employ them in the implementation of our BwO forest. To fairly compare the results, we set the size of ensemble model B as 100 and the rate of oversampling α as 4. All the tree-based surrogate models employ random feature selection as the square root of the feature dimensionality. In addition, for simplicity of BwO implementation, we duplicate a training dataset β times, where β is larger than α , and then pick a bootstrap sample that contains α/β of the duplicated dataset, e.g., if $\beta = 16$, a quarter of the duplicated dataset is sampled by bootstrapping for every bootstrap sample.

GP regression with Matérn 5/2 kernel is used, and its hyperparameters are optimized by marginal likelihood maximization. Because the Cholesky decomposition is applied to compute a marginal likelihood and a posterior distribution, it is slightly faster than the vanilla GP regression model. To focus on the tree-based sur-

rogate models, we do not apply more sophisticated techniques to speed up GP regression. Nevertheless, we include the results with the aforementioned GP regression, in order to briefly compare it to the results with tree-based surrogate models.

For the SMO setting, we use the expected improvement (Moćkus et al., 1978) as an acquisition function. To optimize an acquisition function for tree-based surrogate models, a fixed number of points are sampled to compute acquisition function values; we sample 50,000 points using the Sobol’ sequence. Additionally, for GP regression, L-BFGS-B with multiple initializations is used. Note that every set of initial points is fixed across surrogate models, so that they are started from same regret values, and 5 initial points are given for every run.

All the computations are conducted on a system with CPU, and each experiment set is tested on the same machine in order to measure wall-clock time precisely. In addition, all the missing details are described in the supplementary material.

4.1 Continuous Search Spaces

We test popular benchmark functions such as Ackley (4 dim.), Bohachevsky (2 dim.), Branin (2 dim.), Hartmann6D (6 dim.), Michalewicz (2 dim.), and Rosenbrock (4 dim.) functions. As shown in Figure 2 and

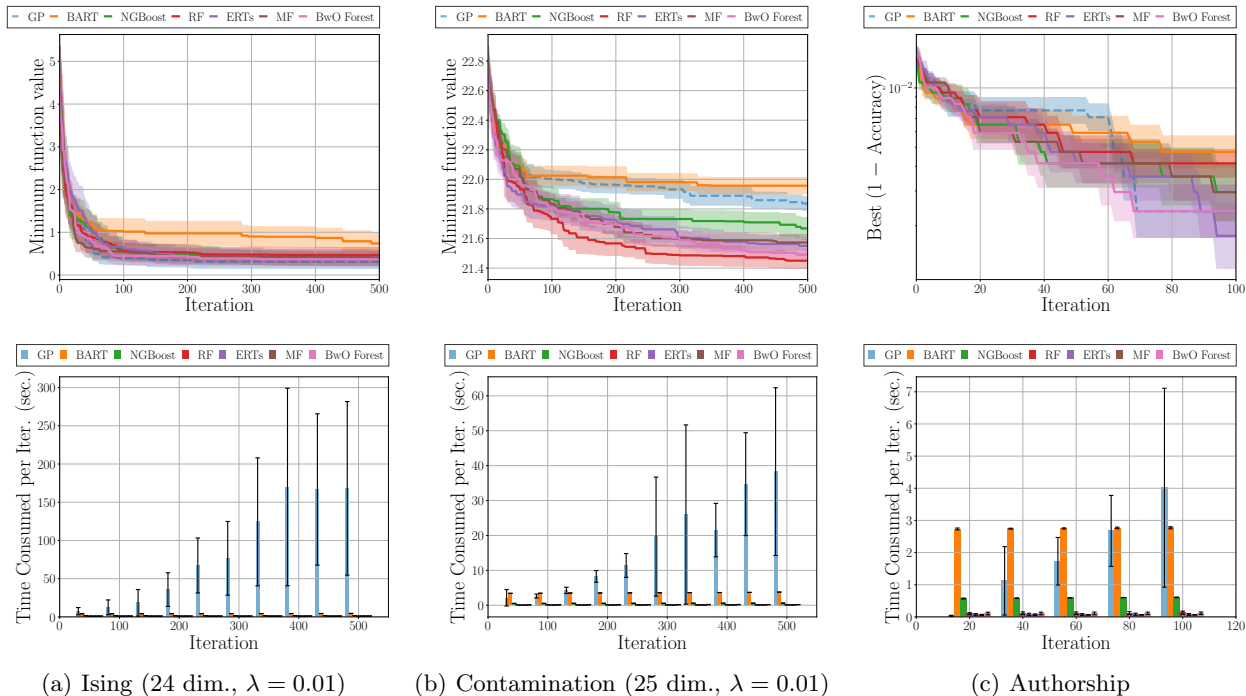


Figure 4: Results on two functions defined on high-dimensional binary search spaces and automated machine learning for the Authorship dataset. All the runs are repeated 10 times.

Figure 3, our method with BwO forest works well, compared to SMO strategies with the other tree-based models.

4.2 High-Dimensional Binary Search Spaces

We conduct the methods studied in this work on two high-dimensional binary problems such as Ising (24 dim.) and contamination (25 dim.) problems, as in the first two columns of Figure 4. For the case of contamination problem, the SMO method with random forest regression shows better results than other methods; see Figure 4(b). The details of these problems are described in the supplementary material.

4.3 Mixed Search Spaces

We carry out the experiments of automated machine learning, which are defined on a mixed search space; see Table 1 for the detailed description of search space. We choose one of machine learning algorithms through a categorical variable and simultaneously tune their numerical or ordinal hyperparameters using SMO. To optimize a categorical variable, we employ a one-hot encoding, which can be defined as a simplex. Four datasets such as Authorship, Breast Cancer, Digits, and Phoneme are used to train and test an automated machine learning model. The tendency of the results by our method is better than the other methods, as

Table 1: Details of automated machine learning. AB, GB, DT, ET, and RF indicate AdaBoost, Gradient-Boosting, Decision Tree, ExtraTrees, Random Forest classifiers, respectively, and Cat. and Int. stand for categorical and integer variables.

Hyperparam.	Range	Type
Algorithm type	{AB, GB, DT, ET, RF}	Cat.
Ensemble size	[10, 200]	Int.
Max. depth	[2, 8]	Int.
Max. features	[0.2, 1.0]	Float
log learning rate	[-3, 1]	Float

presented in Figure 4(c) and Figure 5.

5 DISCUSSION AND LIMITATIONS

Here, we provide the discussion on tree-based models as well as BwO forest, which is about oversampling, tree construction techniques, epistemic uncertainty, and cheap required computation. Furthermore, we introduce a future direction of this research to resolve the limitations of tree-based surrogate models, which are related to the choice of tree-based surrogate models and extrapolation.

Oversampling. While this technique is widely used in solving an imbalanced data problem (Yap et al.,

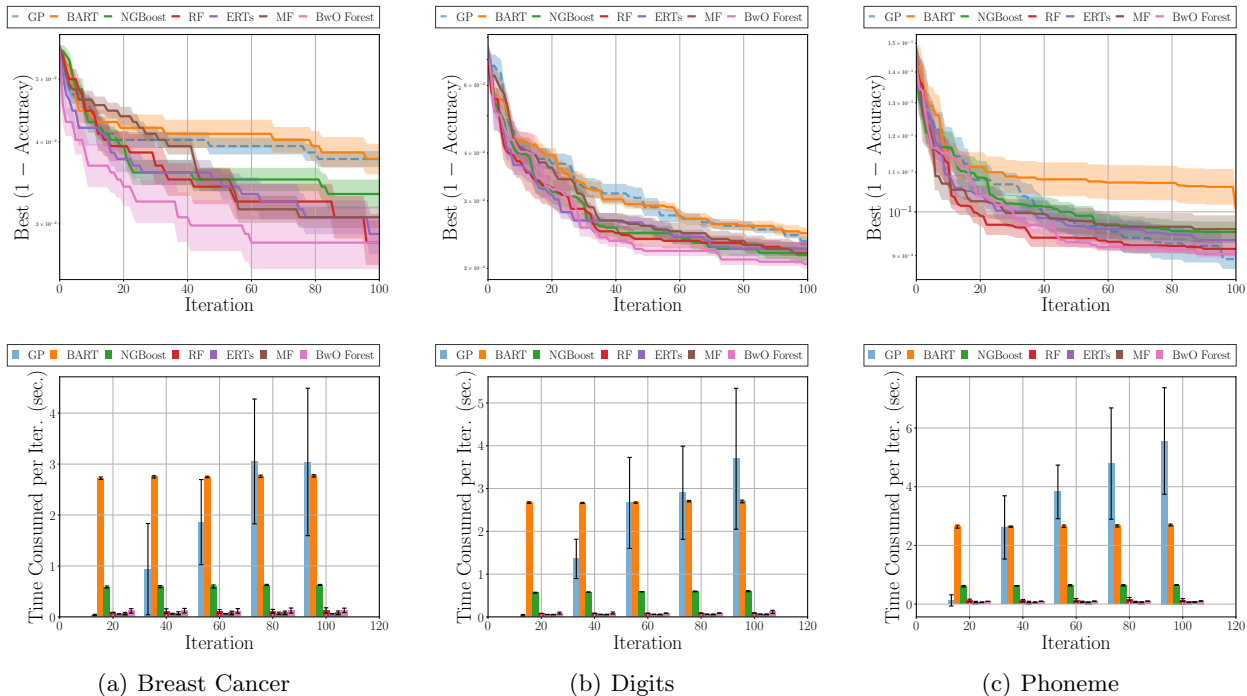


Figure 5: Results on automated machine learning for three datasets, Breast Cancer, Digits, and Phoneme. Note that the result for the Authorship dataset is presented in Figure 4(c). All the runs are repeated 10 times.

2014; Pérez-Ortiz et al., 2015; Yan et al., 2019), it is not popular for bagging. Since it samples a bootstrap sample from an original set of data points \mathbf{X} , it increases the number of unique original elements as described in Section 3, but it is a solid method to construct a bootstrap sample. To sum up, bagging with oversampling prevents a surrogate model underfitting to \mathbf{X} . Moreover, it is more effective in the case that N is relatively small, which is common in SMO, than the case N is large, as in Figure 1, Figure 7, and Table 2.

Tree construction techniques. We show the effects of tree construction techniques such as bagging (denoted as B), oversampling (denoted as O), and random sampling of split location (denoted as R), as presented in Figure 6. The prediction uncertainty estimation by B + O is abnormal compared to other results, because it overfits to the duplicated training data and split locations are not properly determined. The R + B result is similar to the result by Mondrian forest, but it tends to underfit to training data. The result in Figure 6(b) looks similar to Figure 6(d), however the uncertainty on the region out of the range of training data goes to zero. We provide more diverse results by individual trees in Figure 9, Figure 10, and Figure 11.

Qualitative analysis on regression. To compare regression results qualitatively, we measure Kullback–Leibler divergence from GP to a result by each

Table 2: Kullback–Leibler divergence from GP to respective results by tree-based surrogate models.

Method	Figure 1	Figure 7	Figure 8
RF	1.38	0.00	0.00
ERTs	<u>0.06</u>	0.07	0.09
BART	0.93	<u>0.01</u>	<u>0.01</u>
MF	0.93	<u>0.01</u>	0.00
NGBoost	0.21	0.05	0.00
BwO forest	0.04	<u>0.01</u>	<u>0.01</u>

tree-based surrogate model, which shows how similar a result is to GP. As shown in Table 2, our BwO forest tends to follow the GP appropriately, compared to random forest, extremely randomized trees, BART, Mondrian forest, and NGBoost.

Epistemic uncertainty by tree-based surrogate models. Epistemic uncertainty is reducible if we collect more data or identify a model well (Gal, 2016). This type of uncertainty is induced from the randomness of model or bootstrapping. Thus, our BwO forest can be considered that the randomness of model is maximized with the tree construction techniques such as random feature selection, random selection of split locations as well as bagging with oversampling. Especially, the bootstrap sample constructed by our technique, bagging with oversampling, contains more

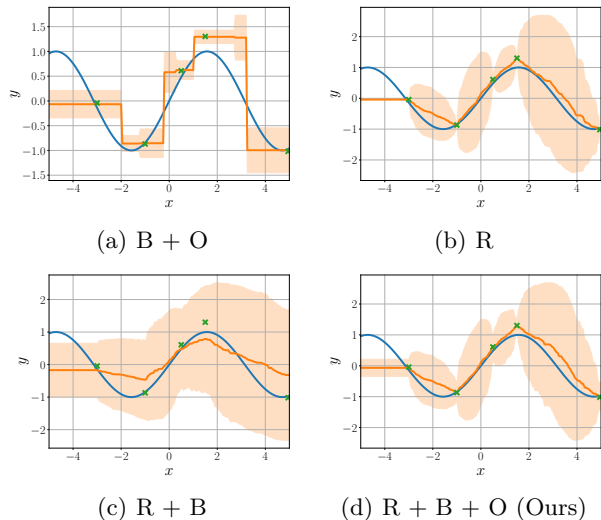


Figure 6: Results on tree construction techniques. B, O, and R indicate bagging, oversampling, and random split location, respectively. The case of only B is shown in Figure 1(b) as the result by random forest.

unique original elements to enhance the degree of fitting to training data than a standard bagging method.

Cheap required computation. In general, tree-based surrogate models are more time-consuming in predicting an output than other types of estimators, because it has to compute the outputs of all the base estimators. However, fortunately the setup of SMO does not assume a large number of data points, which implies that the size of ensemble model can be maintained as a small size. On the other hand, as shown in Section 4, tree-based models are consistently faster than GP regression in many cases, and they are able to be directly applied to speed up the overall procedure of SMO with comparable performance.

Choice of tree-based surrogate models. Inevitably, a specific SMO method successfully optimizes some specific functions, and simultaneously fails to optimize other functions, depending on the characteristics of target functions. As shown in Section 4, our method is robust in many cases, but either random forest or Mondrian forest can be a good option in some cases, e.g., Figure 3(b) and Figure 4(b).

Extrapolation. Lastly, tree-based surrogate models are vulnerable in predicting a function value and its corresponding uncertainty out of the range of training data. At least, even if a function estimate is not correct, an uncertainty estimate should become larger than the results shown in the paper. In order to expand the usage of tree-based models in SMO, improv-

ing the ability to extrapolate is left to a future work. According to our preliminary experiments, injecting a noise in the duplicates is likely to improve the extrapolation ability.

6 RELATED WORK

From now, we briefly review tree-based estimators and SMO with tree-based surrogate models.

Tree-based estimators. An estimator that aggregates a set of decision trees is attractive to many machine learning practitioners in both classification and regression tasks, since it shows reliable performance despite its relatively efficient training and test procedures (Biau, 2012; Natekin and Knoll, 2013; Louppe, 2014). In particular, compared to a deep neural network, such a randomized tree-based approach has a practical strength in real-world problems, which is shown by the popularity of gradient boosting machines such as XGBoost (Chen and Guestrin, 2016) and LightGBM (Ke et al., 2017). As mentioned in the previous sections, Breiman (1996, 2001); Friedman (2001); Geurts et al. (2006); Chipman et al. (2010); Lakshminarayanan et al. (2014, 2016); Duan et al. (2020) have proposed the foundations and theories of the approaches used in this work; see the work by Dietterich (2000); Zhou (2012); Louppe (2014) for the details. Moreover, while we omit it in this work, Mentch and Hooker (2016) suggest a method to estimate uncertainties using subbagging for random forest.

Sequential model-based optimization with tree-based surrogate models. Hutter et al. (2011) propose SMO with random forest by applying itself in optimizing an algorithm configuration. This method is widely adopted in many applications (Feurer et al., 2015; Candelieri et al., 2018; Ying et al., 2019) including Auto-sklearn (Feurer et al., 2020). Importantly, it is beneficial in certain circumstances presented and tested in this paper.

7 CONCLUSION

We re-examined sequential random forest-based optimization and suggested our methods defined with diverse randomized tree-based surrogate models including extremely randomized trees, BART, Mondrian forest, and NGBoost. Then, we proposed a new tree-based surrogate model, named BwO forest, which uses an ensemble construction technique, bagging with oversampling. The empirical analyses on such methods help us to understand the tree-based models thoroughly and provide a future research direction of SMO with tree-based surrogate models.

References

- Aitchison, J. and Aitken, C. G. G. (1976). Multivariate binary discrimination by the kernel method. *Biometrika*, 63(3):413–420.
- Amari, S. (1998). Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276.
- Azimi, J., Fern, A., and Fern, X. Z. (2010). Batch Bayesian optimization via simulation matching. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 23, pages 109–117, Vancouver, British Columbia, Canada.
- Biau, G. (2012). Analysis of a random forests model. *Journal of Machine Learning Research*, 13(1):1063–1095.
- Bodin, E., Kaiser, M., Kazlauskaitė, I., Dai, Z., Campbell, N., and Ek, C. H. (2020). Modulating surrogates for Bayesian optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 970–979, Virtual.
- Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24(2):123–140.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1):5–32.
- Brochu, E., Cora, V. M., and de Freitas, N. (2010). A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*.
- Candelieri, A., Perego, R., and Archetti, F. (2018). Bayesian optimization of pump operations in water distribution systems. *Journal of Global Optimization*, 71(1):213–235.
- Chen, T. and Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 785–794, San Francisco, California, USA.
- Chipman, H. A., George, E. I., and McCulloch, R. E. (2010). BART: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298.
- Choi, S.-C. T., Hickernell, F. J., Jagadeeswaran, R., McCourt, M., and Sorokin, A. G. (2021). Quasi-Monte Carlo software. *arXiv preprint arXiv:2102.07833*.
- Dietterich, T. G. (2000). Ensemble methods in machine learning. In *Proceedings of the First International Workshop on Multiple Classifier Systems*, pages 1–15, Cagliari, Italy.
- Dua, D. and Graff, C. (2019). UCI machine learning repository. <http://archive.ics.uci.edu/ml>.
- Duan, T., Avati, A., Ding, D. Y., Thai, K. K., Basu, S., Ng, A., and Schuler, A. (2020). NGBoost: Natural gradient boosting for probabilistic prediction. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 2690–2700, Virtual.
- Feurer, M., Eggensperger, K., Falkner, S., Lindauer, M., and Hutter, F. (2020). Auto-Sklearn 2.0: Hands-free AutoML via meta-learning. *arXiv preprint arXiv:2007.04074*.
- Feurer, M., Klein, A., Eggensperger, K., Springenberg, J. T., Blum, M., and Hutter, F. (2015). Efficient and robust automated machine learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 28, pages 2962–2970, Montreal, Quebec, Canada.
- Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *The Annals of Statistics*, 29:1189–1232.
- Gal, Y. (2016). *Uncertainty in deep learning*. PhD thesis, University of Cambridge.
- Geurts, P., Ernst, D., and Wehenkel, L. (2006). Extremely randomized trees. *Machine Learning*, 63(1):3–42.
- Holmes, S. P. (2004). How many different bootstrap samples are there? <http://statweb.stanford.edu/~susan/courses/s208/node11.html>. Last accessed on 25/05/2021.
- Hutter, F., Hoos, H. H., and Leyton-Brown, K. (2011). Sequential model-based optimization for general algorithm configuration. In *Proceedings of the International Conference on Learning and Intelligent Optimization (LION)*, pages 507–523, Rome, Italy.
- Hutter, F., Xu, L., Hoos, H. H., and Leyton-Brown, K. (2014). Algorithm runtime prediction: Methods & evaluation. *Artificial Intelligence*, 206:79–111.
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, pages 3146–3154, Long Beach, California, USA.
- Lakshminarayanan, B., Roy, D. M., and Teh, Y. W. (2014). Mondrian forests: Efficient online random forests. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 27, pages 3140–3148, Montreal, Quebec, Canada.
- Lakshminarayanan, B., Roy, D. M., and Teh, Y. W. (2016). Mondrian forests for large-scale regression when uncertainty matters. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1478–1487, Cadiz, Spain.

- Louppe, G. (2014). Understanding random forests: From theory to practice. *arXiv preprint arXiv:1407.7502*.
- Malinin, A., Prokhorenkova, L., and Ustimenko, A. (2020). Uncertainty in gradient boosting via ensembles. *arXiv preprint arXiv:2006.10562*.
- Mendelson, A. F., Zuluaga, M. A., Hutton, B. F., and Ourselin, S. (2016). What is the distribution of the number of unique original items in a bootstrap sample? *arXiv preprint arXiv:1602.05822*.
- Mentch, L. and Hooker, G. (2016). Quantifying uncertainty in random forests via confidence intervals and hypothesis tests. *Journal of Machine Learning Research*, 17(1):841–881.
- Moćkus, J., Tiesis, V., and Žilinskas, A. (1978). The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, 2:117–129.
- Natekin, A. and Knoll, A. (2013). Gradient boosting machines, a tutorial. *Frontiers in neurorobotics*, 7:21.
- Oh, C., Tomczak, J., Gavves, E., and Welling, M. (2019). Combinatorial Bayesian optimization using the graph cartesian product. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, pages 2914–2924, Vancouver, British Columbia, Canada.
- O’Hagan, A. (1978). Curve fitting and optimal design for prediction. *Journal of the Royal Statistical Society B*, 40(1):1–42.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Pérez-Ortiz, M., Gutiérrez, P. A., Tino, P., and Hervás-Martínez, C. (2015). Oversampling the minority class in the feature space. *IEEE Transactions on Neural Networks and Learning Systems*, 27(9):1947–1961.
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT Press.
- Schulz, E., Speekenbrink, M., Hernández-Lobato, J. M., Ghahramani, Z., and Gershman, S. J. (2016). Quantifying mismatch in Bayesian optimization. In *Neural Information Processing Systems Workshop on Bayesian Optimization (BayesOpt)*, Barcelona, Spain.
- Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 25, pages 2951–2959, Lake Tahoe, Nevada, USA.
- Srinivas, N., Krause, A., Kakade, S., and Seeger, M. (2010). Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 1015–1022, Haifa, Israel.
- Tang, C., Garreau, D., and von Luxburg, U. (2018). When do random forests fail? In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, pages 2987–2997, Montreal, Quebec, Canada.
- Vanschoren, J., van Rijn, J. N., Bischl, B., and Torgo, L. (2013). OpenML: networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60.
- Wang, Z., Hutter, F., Zoghi, M., Matheson, D., and de Freitas, N. (2016). Bayesian optimization in a billion dimensions via random embeddings. *Journal of Artificial Intelligence Research*, 55:361–387.
- Williams, C. K. I. and Rasmussen, C. E. (1996). Gaussian processes for regression. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 8, pages 514–520, Denver, Colorado, USA.
- Yan, Y., Tan, M., Xu, Y., Cao, J., Ng, M., Min, H., and Wu, Q. (2019). Oversampling for imbalanced data via optimal transport. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 5605–5612, Honolulu, Hawaii, USA.
- Yap, B. W., Abd Rani, K., Abd Rahman, H. A., Fong, S., Khairudin, Z., and Abdullah, N. N. (2014). An application of oversampling, undersampling, bagging and boosting in handling imbalanced datasets. In *Proceedings of the First International Conference on Advanced Data and Information Engineering (DaEng)*, pages 13–22, Kuala Lumpur, Malaysia.
- Ying, C., Klein, A., Christiansen, E., Real, E., Murphy, K., and Hutter, F. (2019). NAS-Bench-101: Towards reproducible neural architecture search. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 7105–7114, Long Beach, California, USA.
- Zhou, Z.-H. (2012). *Ensemble methods: foundations and algorithms*. CRC Press.

Supplementary Material: On Uncertainty Estimation by Tree-based Surrogate Models in Sequential Model-based Optimization

In this material, we describe the examples and contents that are missing in the main article.

A 1D EXAMPLES

Two 1D examples: (i) $y = \sin(x) + \epsilon$, (ii) $y = x^3 + \epsilon$, where ϵ is an observation noise, are demonstrated.

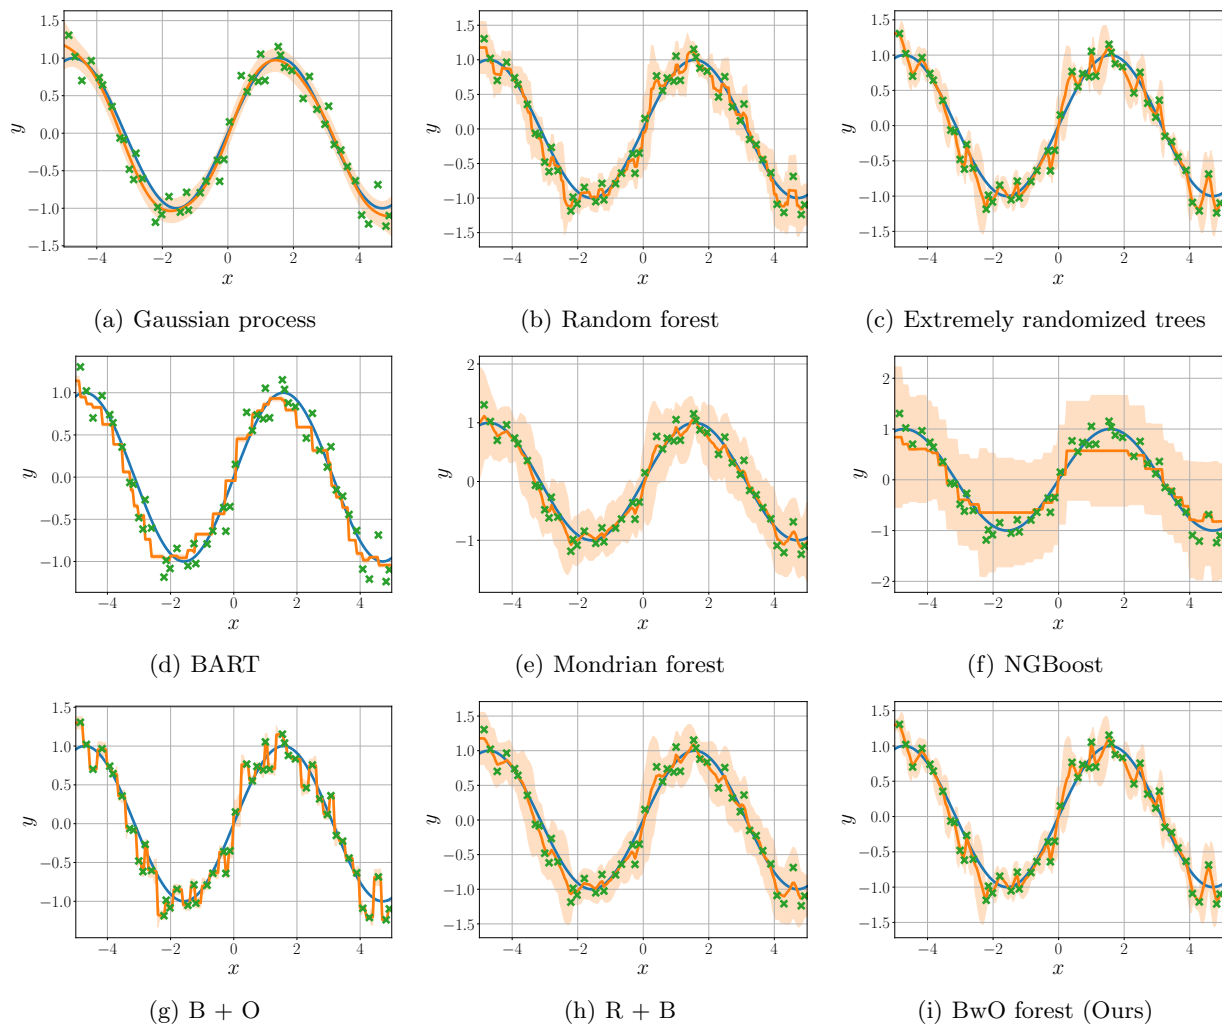


Figure 7: Examples on the sine with 50 points.

The examples by diverse tree-based surrogate models as well as Gaussian process surrogate are shown in Figure 7 and Figure 8. Such examples present the characteristics of surrogate models, as described in the main article.

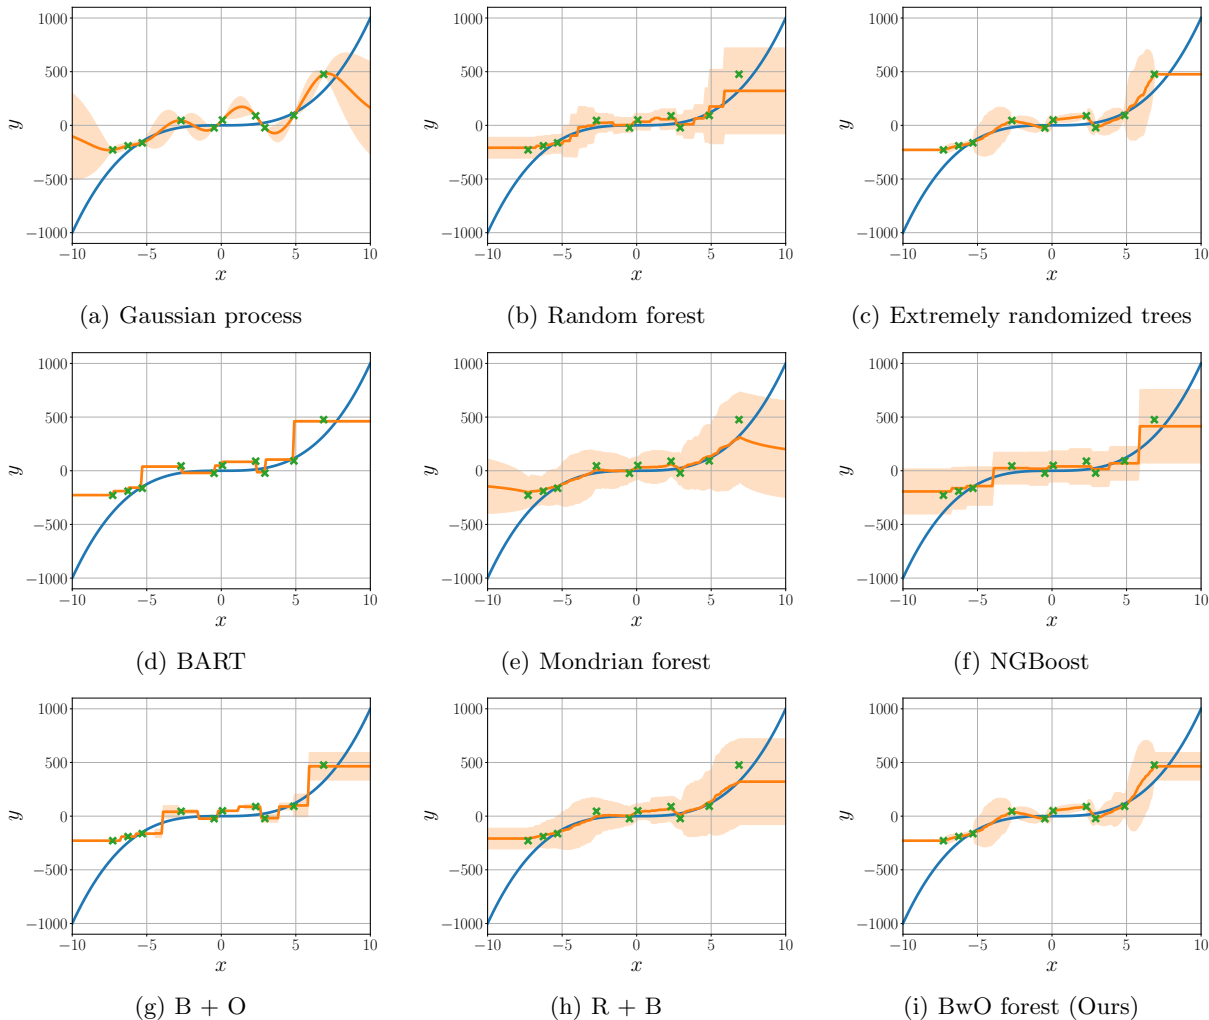


Figure 8: Examples on a cubic equation with 10 points.


B SEQUENTIAL MODEL-BASED OPTIMIZATION WITH OUR PROPOSED SURROGATE MODEL

In this section, we present the sequential model-based optimization procedure with BwO forests. It follows generic steps of sequential model-based optimization (Brochu et al., 2010).


As shown in Algorithm 2, we are given initial points \mathbf{X}_0 , their corresponding evaluations \mathbf{y}_0 , the number of iterations T , and the size of ensemble model B . Sequentially, we acquire a query point and evaluate it every iteration by fitting BwO forest and optimizing an acquisition function. Finally, the best query point \mathbf{x}_{best} among \mathbf{X}_T is determined, by considering the evaluations \mathbf{y}_T .

C REGRESSION RESULTS BY INDIVIDUAL TREES


We visualize the regression results by individual trees for the cases shown in Figure 1, Figure 7, and Figure 8. These results presented in Figure 9, Figure 10, and Figure 11 help us to understand the consequences by tree-based ensemble models.




(a) B (originally proposed as random forest)




(b) R (originally proposed as extremely randomized trees)



(c) B + O




(d) R + B




(e) R + B + O (i.e., BwO forest)


Figure 9: Results by individual trees for the case shown in Figure 1. For brevity, each result is randomly sampled.




(a) B (originally proposed as random forest)




(b) R (originally proposed as extremely randomized trees)



(c) B + O




(d) R + B




(e) R + B + O (i.e., BwO forest)


Figure 10: Results by individual trees for the case shown in Figure 7.




(a) B (originally proposed as random forest)




(b) R (originally proposed as extremely randomized trees)



(c) B + O



(d) R + B



(e) R + B + O (i.e., BwO forest)

Figure 11: Results by individual trees for the case shown in Figure 8.

Algorithm 2 Sequential Model-based Optimization with BwO Forests

Input: Initial points and their evaluations $(\mathbf{X}_0, \mathbf{y}_0)$, the number of iterations T , and the size of ensemble model B .

Output: The best query point \mathbf{x}_{best} .

- 1: **for** $t = 1, \dots, T$ **do**
 - 2: Fit BwO forest $\{\mathcal{T}_b\}_{b=1}^B$ using \mathbf{X}_{t-1} and \mathbf{y}_{t-1} .
 - 3: Acquire a query point by optimizing an acquisition function, $\mathbf{x}_t = \arg \max_{\mathbf{x}} a(\mathbf{x}; \{\mathcal{T}_b\}_{b=1}^B)$.
 - 4: Evaluate $y_t = f(\mathbf{x}_t) + \epsilon$ where ϵ is an observation noise.
 - 5: Update $\mathbf{X}_t \leftarrow [\mathbf{X}_{t-1}; \mathbf{x}_t]$ and $\mathbf{y}_t \leftarrow [\mathbf{y}_{t-1}; y_t]$.
 - 6: **end for**
 - 7: **return** The best query point \mathbf{x}_{best} among \mathbf{X}_T
-

D DETAILS OF EXPERIMENTS

In this section, we describe the detailed setup of experiments. As described in the main article, to implement our BwO forest, we sufficiently duplicate an original dataset and then subsample part of duplicated datasets, where $\alpha = 4$ and $\beta = 16$. As part of our implementation, we utilize scikit-learn (Pedregosa et al., 2011) and QMCpy (Choi et al., 2021). For the implementation of BART, Mondrian forest, and NGBoost, we use the following open-source projects: <https://github.com/JakeColtman/bartpy>, <https://github.com/scikit-garden/scikit-garden>, and <https://github.com/stanfordmlgroup/ngboost>, respectively.

High-Dimensional Binary Search Spaces. We adopt the experimental setup suggested by Oh et al. (2019). Ising sparsification is to optimize the KL-divergence between two probability mass functions with the regularization technique controlled by λ . Contamination control in food supply chain is a problem that optimizes food contamination with minimum prevention cost. It is also regularized by a balancing hyperparameter λ . See the work (Oh et al., 2019) for the details of experiments.

Mixed Search Spaces. Our automated machine learning problem selects one of such classifiers: (i) AdaBoost, (ii) GradientBoosting, (iii) Decision Tree, (iv) ExtraTrees, and (v) Random Forest. This selection is represented by a one-hot encoding. In addition to this categorical variable, we optimize the following hyperparameters for classifiers: (i) the size of ensemble model, (ii) maximum depth, (iii) maximum features, and (iv) log learning rate. It is summarized in Table 1. We use scikit-learn (Pedregosa et al., 2011) to implement all the algorithms. Datasets used in this work: (i) Authorship (Vanschoren et al., 2013), (ii) Breast Cancer (Dua and Graff, 2019), (iii) Digits (Dua and Graff, 2019), and (iv) Phoneme (Vanschoren et al., 2013), are available at the respective references.