

Co-Learning Bayesian Optimization

Zhendong Guo, Yew-Soon Ong, *Fellow, IEEE*, Tiantian He, and Haitao Liu

Abstract—Bayesian optimization (BO) is well known to be sample-efficient for solving black-box problems. However, the BO algorithms may get stuck in suboptimal solutions even with plenty of samples. Intrinsically, such suboptimal problem of BO can attribute to the poor surrogate accuracy of the trained Gaussian process (GP), particularly that in the regions where the optimal solutions locate. Hence, we propose to build multiple GP models instead of a single GP surrogate to complement each other, thus resolving the suboptimal problem of BO. Nevertheless, according to the bias-variance tradeoff equation, the individual prediction errors can increase when increasing the diversity of models, which may lead to even worse overall surrogate accuracy. On the other hand, based on the theory of Rademacher complexity, it has been proven that exploiting the agreement of models on unlabeled information can reduce the complexity of hypothesis space, therefore achieving the required surrogate accuracy with fewer samples. Such value of model agreement has been extensively demonstrated for co-training style algorithms to boost model accuracy with a small portion of samples. Inspired by the above, we propose a novel BO algorithm labeled as co-learning BO (CLBO), which exploits both model diversity and agreement on unlabeled information to improve the overall surrogate accuracy with limited samples, therefore achieving more efficient global optimization. Through tests on five numerical toy problems and three engineering benchmarks, the effectiveness of proposed CLBO has been well demonstrated.

Index Terms—Bayesian optimization, co-training style semi-supervised learning, black-box problems, robotics control, aircraft design.

I. INTRODUCTION

BAYESIAN optimization is well-known to be sample efficient, which has drawn wide attentions for solving black-box problems [1]–[3], such as active control of robotics systems [4] and hyperparameter tuning [5]. The basic process of BO is as follows. First, it uses Gaussian process (GP) [6] to build a surrogate of the underlying objective function. Second, it employs an acquisition function [7]–[10], which makes use of the information extracted from GP to select the most promising solution candidate to query. The above process is iterated till the stopping condition is met.

Zhendong Guo is with the School of Energy and Power Engineering, Xi’an Jiaotong University, China, and also with Data Science and Artificial Intelligence Research Center (DSAIR), Nanyang Technological University, Singapore. (e-mail: guozhendong@xjtu.edu.cn; ericzhendong@163.com).

Yew-Soon Ong is the corresponding author of this paper, he is with School of Computer Science and Engineering, Nanyang Technological University, Singapore, and also with Agency for Science, Technology and Research (A*STAR), Singapore (e-mail: asysong@ntu.edu.sg; ongyewsoon@hq.a-star.edu.sg).

Tiantian He is with Center for Frontier AI Research (CFAR), Institute of High Performance Computing (IHPC), Agency for Science, Technology and Research (A*STAR), and also with Data Science and Artificial Intelligence Research Center (DSAIR), Nanyang Technological University, Singapore (email: He_Tiantian@ihpc.a-star.edu.sg).

Haitao Liu is with School of Energy and Power Engineering, Dalian University of Technology, 116024, Dalian, China (e-mail: htliu@dlut.edu.cn).

Though tremendous success has been achieved [11], [12], BO approaches sometimes can get stuck by poor solutions even with plenty of adaptively and sequentially selected samples [13]–[15]. *Intrinsically, such a suboptimal problem of BO can attribute to poor accuracy of the trained GP surrogate, particularly that in the region where the true optima locate* [16]. To improve the accuracy of GP surrogate in BO process, some studies resort to extracting information from the related source tasks [17], [18], or building multi-fidelity surrogate [19]–[21]. Nevertheless, neither the source task nor multi-fidelity samples can be always available. Alternatively, as no single model can be accurate for all the scenarios [22]–[25], some other studies propose to use multiple models [1], [26] instead of a single GP to improve the surrogate accuracy in BO process.

Note that, the idea of using multiple models has been extensively explored in machine learning community [27], [28], such as ensemble learning [29]. In particular, the overall accuracy by using multiple models at a single point site or over a region can be calculated by the following bias-variance tradeoff equation [30]:

$$\underbrace{(\hat{y} - y)^2}_E = \underbrace{\sum w_i (\hat{y}^{(i)} - y)^2}_{E_{ind}} - \underbrace{\sum w_i (\hat{y}^{(i)} - \hat{y})^2}_{Diversity} \quad (1)$$

$$\hat{y} = \sum_i w_i \hat{y}^{(i)}, \quad \sum_i w_i = 1$$

where $\hat{y}^{(i)}$ and \hat{y} denote the overall prediction and the individual model prediction, respectively, and w_i is the corresponding weight. The terms $(\hat{y} - y)^2$, $\sum w_i (\hat{y}^{(i)} - y)^2$, and $\sum w_i (\hat{y}^{(i)} - \hat{y})^2$ are respectively the errors of overall model (E) and individual models (E_{ind}), and the diversity among individual models (ambiguity term). Ideally, simultaneously increasing the model diversity and reducing the error of individual models (i.e., E_{ind}) can reduce the overall prediction error. However, *purely increasing the diversity among models may increase E_{ind} , which may contrarily result in an incremental overall error of prediction*. Hence, *both individual model accuracy and model diversity should be taken into account when building multiple models*.

On the other hand, according to Rademacher complexity [31], *the bound of prediction loss (error) is related to the complexity of hypothesis searching space* [32]:

$$E(L(Y, f(X))) \leq \hat{E}_n(L(Y, f(X))) + \underbrace{R_n(\tilde{L} \circ F)}_{Complexity} + \sqrt{\frac{8 \log(2/\delta)}{n}} \quad (2)$$

where $E(L(Y, f(X)))$ and $\hat{E}_n(L(Y, f(X)))$ are the expectations of real and empirical loss, $R_n(\tilde{L} \circ F)$ is the Rademacher

complexity of the hypothesis space, and n and δ are the number of training samples and the certainty in the success of the whole procedure, respectively.

As has been previously proven [32], [33], *the agreement of models on unlabeled information can reduce the complexity of hypothesis space (i.e., $R_n(\tilde{L} \circ F)$ in Eq.(2)), therefore helping to achieve the required model accuracy with fewer samples.* The significance of model agreement shown as above has been extensively demonstrated by co-training style algorithms [28], [34], [35], which achieved tremendous success in boosting the model accuracy with limited available samples [36], [37].

Inspired by the above, we propose a novel algorithm labeled as co-learning Bayesian Optimization (CLBO). The motivation behind CLBO is as follows. First, since no surrogate model can be accurate at all sites of the design space, we propose to build multiple GP models to improve the overall surrogate accuracy, particularly that in the region where the true optima locate. Second, based on the analysis with Eq.(2), we propose to enforce model agreement on unlabeled information to reduce the sample complexity of hypothesis space, therefore building reasonably accurate individual models with limited training samples. And further, the differences among these reasonably accurate individual models can help them to complement each other if they are combined together. Following the way shown above, the suboptimal problem of BO can be addressed by much better overall surrogate accuracy particularly that in the vicinity of real optima, and therefore better optimization results can be expected.

More specifically, we use two kinds of GP models in CLBO. On one hand, we use multi-output GP (MOGP) [38] with the subsets of training samples to obtain multiple GP predictions, wherein we treat the predictions as “subtasks” and enforce model agreement on curve bumpiness to help build individual prediction models with reasonable accuracy. On the other hand, though it cannot be accurate at every site of the design space, the surrogate trained with more samples generally can have relatively better average accuracy over space. Therefore we also build a single-output GP surrogate with the full training set, labeled as SOGP. Then, the combination of MOGP and SOGP models can be expected to have good surrogate accuracy over space, which helps to achieve the optimal solutions more efficiently.

The contributions of the paper can be summarized as follows:

- Inspired by bias-variance tradeoff equation and agreement-assisted co-training paradigm, we propose a novel BO framework with multiple GP models, labeled as co-learning BO (CLBO). The proposed framework can exploit both model diversity and agreement on unlabeled information to improve the overall surrogate accuracy in BO process, therefore achieving an more efficient global optimization.
- In doing so, we use multi-output GP to enforce model agreement on curve bumpiness to obtain individual surrogates of good accuracy. In the meantime, we build the additional single-output GP model with full samples to further improve the overall surrogate accuracy in BO process.

- Through tests on benchmark functions and engineering problems, and comparison against other BO algorithms, the effectiveness of CLBO is well demonstrated.

The remainder of this paper is organized as follows. Section II presents the preliminaries and related work of CLBO. Section III illustrates the details of proposed CLBO. Section IV then shows the experimental studies on the benchmark functions and engineering problems. And finally, we draw conclusions in Section V.

II. BACKGROUND AND RELATED WORK

In this section, we introduce the basics of BO, which motivate us to use multiple GP models instead of a single GP surrogate in CLBO. We then investigate the recent works that are closely related to the proposed framework. It should be noted that, we focus on BO algorithms adopting the expected improvement (EI) as the acquisition function in this paper, although our framework can be combined with various acquisition functions.

A. Basics and Issues of EI-based Bayesian Optimization

Since GP and acquisition function constitute the key ingredients of BO, we introduce them briefly in this subsection. In particular, the motivation why we use multiple GP models is illustrated by analyzing the “over-exploitation” and “over-exploration” issues of EI.

1) *Standard Gaussian Process*: GP is a stochastic process to approximate the input-output relation, which formulates the function prediction $y(\mathbf{x})$ as a Gaussian distribution [6] :

$$y(\mathbf{x}) \sim \text{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \quad (3)$$

where $m(\mathbf{x})$ and $k(\mathbf{x}, \mathbf{x}')$ denote the mean and covariance function, respectively. Without loss of generality, we take $m(\mathbf{x}) \equiv 0$ and $k(\mathbf{x}, \mathbf{x}')$ can be the popular squared exponential (SE) [39]:

$$k_{SE}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^T P^{-1}(\mathbf{x} - \mathbf{x}')\right) \quad (4)$$

where σ_f^2 models the function output scale, P is an $m \times m$ diagonal matrix with diagonal element $\{l_h^2\}_{1 \leq h \leq d}$, d is the dimension of design space, and l_h is the characteristic length that describes how quickly the function value changes as \mathbf{x}' moves away from \mathbf{x} along dimension h .

Given samples $\mathbf{y}(X) = [y(\mathbf{x}_1), y(\mathbf{x}_2) \cdots y(\mathbf{x}_n)]^T$, the mean prediction and corresponding mean squared prediction error σ^2 at an unknown site \mathbf{x} are expressed as:

$$\begin{aligned} \hat{y}(\mathbf{x}) &= \mathbf{k}(\mathbf{x}, X)(K + \sigma_n^2 \mathbf{I})^{-1} \mathbf{y} \\ \sigma^2(\mathbf{x}) &= \mathbf{k}(\mathbf{x}, \mathbf{x}) - \mathbf{k}(\mathbf{x}, X)(K + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}(X, \mathbf{x}) + \sigma_n^2 \end{aligned} \quad (5)$$

where σ_n^2 is the Gaussian noise, \mathbf{k} is the cross-correlation vector between \mathbf{x} and X , and K is the covariance matrix over the input features of samples X , which is formulated as:

$$\begin{aligned} k(\mathbf{x}, X) &= [k(\mathbf{x}, \mathbf{x}_1), \cdots, k(\mathbf{x}, \mathbf{x}_n)] \\ K &= \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & \cdots & k(\mathbf{x}_1, \mathbf{x}_n) \\ \vdots & \ddots & \vdots \\ \text{Sym} & & k(\mathbf{x}_n, \mathbf{x}_n) \end{bmatrix} \end{aligned} \quad (6)$$




Fig. 1. Plots of $\Phi(z)$ and $\phi(z)$ and $\Phi(z)/\phi(z)$.

where Sym in Eq.(6) means that K is a symmetric matrix.

Note that the form of \mathbf{k} and K in Eq.(6) can be adjusted in various forms to achieve even better surrogate accuracy in BO process. However, no single model can be guaranteed to be accurate for all the scenarios, therefore it is attractive to employ multiple models instead of a single GP surrogate. In the meantime, according to Eq. (1), purely increasing the model diversity can increase the error of individual models (i.e., E_{ind}), resulting in even worse overall surrogate accuracy. Therefore, inspired by the agreement-assisted co-training algorithms, we propose to exploit unlabeled information and enforce model agreement constraint to improve the overall surrogate accuracy in BO process.

2) *Over-Exploitation and Over-Exploration of EI*: Assuming that the GP prediction $Y(\mathbf{x})$ follows the distribution as $Y(\mathbf{x}) \sim N(\hat{y}(\mathbf{x}), \sigma^2(\mathbf{x}))$, and let $F_{Y(\mathbf{x})}$ denote the probability distribution of $Y(\mathbf{x})$, and $F'_{Y(\mathbf{x})}$ be the derivative of $F_{Y(\mathbf{x})}$, the expectation of the objective improvement can be derived as:

$$\begin{aligned} EI(\mathbf{x}) &= \int_{-\infty}^{\infty} (f_{min} - Y(\mathbf{x}), 0) dF_{y(\mathbf{x})} \\ &= \int_{-\infty}^{f_{min}} (f_{min} - \hat{y}) dF_{y(\mathbf{x})} - \int_{-\infty}^{f_{min}} (y - \hat{y}) dF_{y(\mathbf{x})} \quad (7) \\ &= (f_{min} - \hat{y}) F_{y(\mathbf{x})}(f_{min}) + \sigma F'_{y(\mathbf{x})}(f_{min}) \end{aligned}$$

In each optimization cycle, we select the sample to query by maximizing EI. The partial derivatives of EI with respect to \hat{y} and σ and the gradient of EI are:

$$\begin{aligned} \frac{\partial(EI)}{\partial \hat{y}} &= -F_{Y(\mathbf{x})}(f_{min}) = -\Phi\left(\frac{f_{min} - \hat{y}}{\sigma}\right) = -\Phi(z) \quad (8) \\ \frac{\partial(EI)}{\partial \sigma} &= F'_{Y(\mathbf{x})}(f_{min}) = \phi\left(\frac{f_{min} - \hat{y}}{\sigma}\right) = \phi(z) \end{aligned}$$

$$\begin{aligned} -\nabla EI(\mathbf{x}) &= -\left(\frac{\partial EI}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial \mathbf{x}} + \frac{\partial EI}{\partial \sigma} \frac{\partial \sigma}{\partial \mathbf{x}}\right) \quad (9) \\ &= \Phi(z) \frac{\partial \hat{y}}{\partial \mathbf{x}} - \phi(z) \frac{\partial \sigma}{\partial \mathbf{x}} \end{aligned}$$

As shown in Eq. (8), there is a balance between $\partial \hat{y} / \partial \mathbf{x}$ and $\partial \sigma / \partial \mathbf{x}$ to attain maximum EI, which can be varied according to the corresponding weights, i.e., $\Phi(z)$ and $\phi(z)$.

In particular, the influences of $\Phi(z)$ and $\phi(z)$ on the performance of EI are discussed with Fig.1. When $z > 3$, $\Phi(z)$ can be overwhelmingly greater than that of $\phi(z)$. Correspondingly, the following equation can be derived:

$$\begin{aligned} -\nabla EI(\mathbf{x}|z > 3) &= \Phi(z > 3) \frac{\partial \hat{y}}{\partial \mathbf{x}} - \phi(z > 3) \frac{\partial \sigma}{\partial \mathbf{x}} \quad (10) \\ &\approx \Phi(z > 3) \frac{\partial \hat{y}}{\partial \mathbf{x}} \end{aligned}$$

This means EI focuses sampling near the optima of the fitted GP surrogate when z takes large positive values, and thereby it may ignore exploring other regions of the optimization space and get stuck in a sub-optimal solution, which is the well-known issue of EI labeled as ‘‘over-exploitation’’.

On the other hand, as shown in Fig.1, $\phi(z)$ gradually becomes greater than $\Phi(z)$ when $z < 0$. In particular, we can ‘‘approximately’’ have the following equation when $z < -3$:

$$\begin{aligned} -\nabla EI(\mathbf{x}|z < -3) &= \Phi(z < -3) \frac{\partial \hat{y}}{\partial \mathbf{x}} - \phi(z < -3) \frac{\partial \sigma}{\partial \mathbf{x}} \quad (11) \\ &\approx -\phi(z < -3) \frac{\partial \sigma}{\partial \mathbf{x}} \end{aligned}$$

This happens at the later stage of the optimization process when no objective improvement can be expected directly from the GP surrogate. At such stage with $z < 0$, the issue of ‘‘over-exploration’’ can take place, i.e., EI is likely to query samples with maximum $\sigma(\mathbf{x})$ but poor $y(\mathbf{x})$ for a long period of time.

More specifically, the ‘‘over-exploration’’ issue can be explained from a historical view of the optimization process, with the following equation:

$$\begin{aligned} z = \frac{f_{min} - \hat{y}(\mathbf{x})}{\sigma(\mathbf{x})} > 0 &\Rightarrow \hat{y}(\mathbf{x}_{EI}^*) < f_{min} \quad (12) \\ z = \frac{f_{min} - \hat{y}(\mathbf{x})}{\sigma(\mathbf{x})} \leq 0 &\Rightarrow \text{No constraints on } \hat{y}(\mathbf{x}_{EI}^*) \end{aligned}$$

where \mathbf{x}_{EI}^* denotes the selected sample to query by maximizing EI, and $\hat{y}(\mathbf{x}_{EI}^*)$ is the related GP prediction. Note that the EI-based BO process often starts with $z > 0$. As shown in Eq. (12), the condition $z > 0$ confines $\hat{y}(\mathbf{x}_{EI}^*)$ to be queried in the areas with relatively better $\hat{y}(\mathbf{x})$. In the meantime, the GP surrogate can more or less capture the global trend of objective function. Therefore, as the iteration goes on with $z > 0$, the density of samples at regions with relatively better $y(\mathbf{x})$ becomes denser than those with poor $y(\mathbf{x})$. When $z \leq 0$, EI turns to exploring the optimization space globally by putting much greater weight on $\sigma(\mathbf{x})$ (see Eq. (11)). However, the condition $z \leq 0$ does not impose any constraints on $\hat{y}(\mathbf{x}_{EI}^*)$ (see Eq. (12)). Then, the ‘‘over-exploration’’ issue takes place, and EI is likely to query samples with maximum $\sigma(\mathbf{x})$ but poor $y(\mathbf{x})$ continuously and even forever [2].

To summarize, the parameter z , which is a function of $\hat{y}(\mathbf{x})$ and $\sigma(\mathbf{x})$, plays a crucial role in determining the behavior of EI and consequently the performance of BO. While increasing $\sigma(\mathbf{x})$ can help to address the issue of ‘‘over-exploitation’’ (i.e., the issue with $z > 0$), larger $\sigma(\mathbf{x})$ can also lead the algorithm get stuck in the ‘‘over-exploration’’ issue more easily, which takes place when $z < 0$.

Intrinsically, the issues of ‘‘over-exploitation’’ and ‘‘over-exploration’’ can attribute to the poor accuracy of the trained GP surrogate, particularly that in the region where the true optima locate. Therefore, we propose to use multiple GP models to improve the overall surrogate accuracy in CLBO. And further, we propose to make use of both model diversity and agreement on unlabeled information to manage the overall accuracy of multiple GP models in BO process, which are illustrated with more details in Section III.

B. Related Work

The proposed CLBO is related to the following categories of BO algorithms. First, CLBO is most close to the BO algorithms assisted by multiple GP kernels [1] or a set of surrogate techniques [26], as they all attempt to take advantage of model differences to complement each other and thus achieving better optimization results. Differently, in addition to model differences, CLBO also makes use of model consistency under the spirit of co-training to manage the overall prediction accuracy in the BO process. Second, CLBO shares similarity with batch BO algorithms, such as constant liar [40], batch BO via local penalization [41], etc., since they both query more than one samples in each optimization cycle. More importantly, the querying samples in each cycle of the above studies are obtained by maximizing a modified EI function. In other words, instead of using multiple GP surrogates, these studies can be regarded as the BO algorithms that assisted by multiple acquisition functions. Third, CLBO can also have relation to those BO algorithms that augment the role of $\sigma(\mathbf{x})$ [42], [43], or change the incumbent [44] of EI formulation, as they share the same purpose to resolve the “over-exploitation” issue of EI. However, the above studies may not be helpful to address the “over-exploration” issue of EI [2]. By exploiting both model agreement and disagreement (i.e., diversity) to improve the overall surrogate accuracy, our work can help to resolve the issues of “over-exploitation” and “over-exploration” simultaneously.

III. METHODOLOGY

To address the suboptimal problem of BO and therefore achieve an efficient global optimization, we propose CLBO which exploits both model diversity and agreement on unlabeled information to improve the overall surrogate accuracy in the optimization process. In this section, we show the details of the proposed CLBO.

A. General Framework of CLBO

The basic idea of CLBO is as follows: Firstly, while no surrogate can be guaranteed to be accurate at all sites of the design space, multiple surrogate models can complement each other to achieve a better approximation accuracy particular that in neighborhood of the true optimal solution, therefore we propose to build multiple GP predictions with the subsets of training samples in BO process. Secondly and more importantly, the GP models trained with a small portion of subset samples may suffer from poor individual accuracy and therefore resulting in even worse overall surrogate accuracy (see Eq.(1)). Nonetheless, enforcing model agreement on unlabeled information can reduce the number of samples necessary to achieve required model accuracy [32]. Hence, we propose to exploit model agreement to ensure reasonable accuracy of individual models. Then, these reasonably accurate individual models are combined to complement each other and thus achieving better overall approximation accuracy in the BO process. Along this idea, we show the framework of proposed CLBO in Algorithm 1, which is consisted of four components, i.e., (1) generation of diverse sample subsets; (2) building

Algorithm 1 General framework of CLBO

Input: Target optimization problem and the number of initial training samples

Output: Optimized solution of the target optimization problem

- 1: **Generation of Diverse Sample Subsets:** Use bootstrapping (see Algorithm 2) or other sampling strategy to obtain a set of sample sets, and evaluate the objective function value.
 - 2: **while** the termination condition is not satisfied **do**
 - 3: **Building of Agreement-Assisted GP Models:** Build multiple GP models, where agreement constraints are imposed on all or part of GP models in the training process to obtain predictions of reasonable accuracy.
 - 4: **Sample Search with Multiple GP Models:** Use acquisition functions such as EI to select the samples to query, and then evaluate the objective functions of all the new samples.
 - 5: **Sample Exchange between Multiple GP Models:** Add new samples to each training set of the multiple models.
 - 6: **end while**
-

of agreement-assisted GP models; (3) sample search with multiple GP models and (4) sample exchange with multiple GP models.

Note that the building of agreement-assisted GP models is the most important part of CLBO, and there can be many ways to exploit model agreement on unlabeled knowledge. In the following subsections, we present an instantiation of agreement-assisted GP models. After that, we illustrate the details of the other components of CLBO.

B. Agreement-Assisted GP Models

According to the bias-variance tradeoff equation (see Eq.(1)), it is desirable to combine individual models of reasonable accuracy to complement each, and therefore achieving better overall surrogate accuracy in BO process. To achieve it, we use two kinds of GP models to deal with the balance in between individual model accuracy and diversity among models. One is the single-output GP (SOGP) that trained with full sample set, which is believed to have relatively good mean accuracy over design space. The other kind of GP models are built by using multi-output GP (MOGP) [45], [46] with training subsets, wherein agreement constraint are imposed on curve bumpiness to ensure reasonable model accuracy. As we can train the SOGP by following the standard process shown in Section II.A, this section will be focused on illustrating the motivation and implementing details of agreement assisted MOGP.

Specifically, the agreement assisted MOGP built in this paper is inspired by the multi-task GP (MTGP) [38], which is originally proposed for transfer learning between distinct tasks [47], [48]. In this paper, by treating the subsets as related source data, we extend the MTGP to generate multiple GP models with subsets. The motivation behind is as follows. That is, as the surrogate models trained with subsets are built for the

same task, they should have similar curvature/curve bumpiness in predicting the underlying objective function in unexplored design space. To distinguish it from the multi-task case, we name it as multi-form GP (abbreviated as MFGP).

In MFGP, let Eq. (13) denote the training subsets to build MFGP, we use a correlation matrix to model the relations between them, where ρ_{ij} in Eq.(14) denotes the correlation between the outputs associated with X_i and X_j .

$$X = [X_1, \dots, X_j \dots X_m], X_1 \neq \dots \neq X_j \neq \dots \neq X_m \quad (13)$$

$$\mathbf{y} = [\mathbf{y}_1, \dots, \mathbf{y}_j, \dots \mathbf{y}_m]^T, \mathbf{y}_1 \neq \dots \neq \mathbf{y}_j \neq \dots \neq \mathbf{y}_m$$

$$K^f = \begin{bmatrix} \rho_{11} & \dots & \rho_{1m} \\ \vdots & \ddots & \vdots \\ \rho_{m1} & \dots & \rho_{mm} \end{bmatrix} \text{ where } \rho_{ij} = \rho_{ji} \quad (14)$$

Then, we reformulate the cross-correlation vector k_j and the covariance matrix K_M as:

$$K_M = \begin{bmatrix} \rho_{11}K_{11} & \dots & \rho_{1m}K_{1m} \\ \vdots & \ddots & \vdots \\ \rho_{m1}K_{m1} & \dots & \rho_{mm}K_{mm} \end{bmatrix} + D \otimes I \quad (15)$$

where, $\{K_{ij}\}_{1 \leq i, j \leq m}$ is the input correlation matrix associated with X_i and X_j , and D is an $m \times m$ diagonal matrix with diagonal elements $\{\sigma_{n,j}^2\}_{1 \leq j \leq m}$, $\sigma_{n,j}^2$ is the noise term of the GP associated with X_j . Then, the GP predictions at \mathbf{x} can be expressed as:

$$\hat{\mathbf{y}}(\mathbf{x}) = \begin{bmatrix} k_1(\mathbf{x}, X)K_M^{-1}\mathbf{y} \\ \vdots \\ k_m(\mathbf{x}, X)K_M^{-1}\mathbf{y} \end{bmatrix} \quad (16)$$

$$\sigma^2(\mathbf{x}) = \begin{bmatrix} k_1(\mathbf{x}, \mathbf{x}) - k_1(\mathbf{x}, X)K_M^{-1}k_1(X, \mathbf{x}) + \sigma_{n,1}^2 \\ \vdots \\ k_m(\mathbf{x}, \mathbf{x}) - k_m(\mathbf{x}, X)K_M^{-1}k_m(X, \mathbf{x}) + \sigma_{n,m}^2 \end{bmatrix}$$

Also note that, in original MTGP [38], the correlation coefficient ρ_{ij} also measures the differences between source tasks to avoid ‘‘negative transfer’’ [49], [50]. In our MFGP, we use ρ_{ij} to model the differences of function outputs associated with subsets, which plays a crucial role in generating different models. To ensure the diversity of GP models, we generate subsets by using the Algorithm 2, and we also take ρ_{ij} and $\sigma_{n,i}^2$ as hyper-parameters in the tuning process.

Equation (17) shows the hyper-parameters of MFGP to be tuned by the negative logarithm maximum likelihood (NLML) [45]. As contrast, Eqs. (18) and (19) show the GP models of subsets trained by SOGP, where $\sigma_{f,i}^2$ models the function output scale associated with the subset (X_i, \mathbf{y}_i) .

$$\arg \min NLML_{MFGP}(l_h^2, \sigma_{f,i}^2, \sigma_{n,i}^2, \rho_{i,j}) \quad (17)$$

$$\hat{\mathbf{y}}_i(\mathbf{x}_q) = \mathbf{k}(\mathbf{x}_q, X_i)(K(X_i) + \sigma_{n,i}^2 \mathbf{I})^{-1} \mathbf{y}_i$$

$$\hat{\sigma}_i^2(\mathbf{x}_q) = \mathbf{k}(\mathbf{x}_q, \mathbf{x}_q) \quad (18)$$

$$- \mathbf{k}(\mathbf{x}_q, X_i)(K(X_i) + \sigma_{n,i}^2 \mathbf{I})^{-1} \mathbf{k}(X_i, \mathbf{x}_q) + \sigma_{n,i}^2$$

$$\arg \min NLML_{SOGP}(l_{h,i}^2, \sigma_{f,i}^2, \sigma_{n,i}^2) \quad (19)$$

More specifically, l_h^2 in Eq.(17) models the curve bumpiness of all the subsets in MFGP. In contrast, in Eq. (19), $l_{h,i}^2$ is the length hyper-parameter of each subset. In other words,

Algorithm 2 Bootstrap sampling

Input: The full training set X , the number n of samples in X and the number m of subsets

Output: The subsets X_1, \dots, X_m

- 1: **for** $i \leq m$ **do**
 - 2: Bootstrapping n samples from X to obtain X_{sub}
 - 3: Use the unique function to delete the duplicate samples in X_{sub} to get X_i , i.e., $X_i = \text{unique}(X_{sub})$
 - 4: **end for**
-

Algorithm 3 Sample search with multiple GP models

Input: The trained SOGP and MFGP models, the number m of subsets

Output: The new query samples, i.e., X_{new}

- 1: Optimize the EI acquisition function with the SOGP model, i.e., $X_{new}^{SOGP} = \arg \max EI(\mathbf{x})$
 - 2: **for** $i \leq m$ **do**
 - 3: Optimize the EI with the i^{th} sub-surrogate in MFGP, i.e., $\mathbf{x}_{new} = \arg \max EI(\mathbf{x})$
 - 4: Check the minimum distance condition for \mathbf{x}_{new} , i.e., $d_{\min} = \min \|\mathbf{x}_{new} - X \cup X_{new}^{SOGP} \cup X_{new}^{MOGP}\|$
 - 5: **if** $d_{\min} < \varepsilon$ **then**
 - 6: $\mathbf{x}'_{new} = \arg \max PEI(\mathbf{x}, \mathbf{x}_{new})$;
 - 7: $\mathbf{x}_{new} = \mathbf{x}'_{new}$
 - 8: **end if**
 - 9: $X_{new}^{MFGP} = X_{new}^{MFGP} \cup \mathbf{x}_{new}$
 - 10: $X_{new} = X_{new}^{MFGP} \cup X_{new}^{SOGP}$
-

compared to the SOGP shown in Eq. (19), an agreement constraint on curve bumpiness has been imposed among the GP models in MFGP, which can be written as below:

$$\underbrace{l_{h,1}^2 = \dots = l_{h,i}^2 = \dots = l_{h,m}^2}_{\text{Agreement constraint on curve bumpiness in MFGP}} \equiv l_h^2 \quad (20)$$

Agreement constraint on curve bumpiness in MFGP

The effectiveness of the agreement constraint in MFGP will be verified in Section IV.

C. Sample Search with Multiple GP Models

After building multiple models using MFGP and SOGP, we use EI to search new samples. Besides, we set a minimum distance threshold (i.e., $d_{\min} < \varepsilon$) to avoid the ill-conditioning of the inversion of correlation matrix. To prevent samples being null due to the threshold in an optimization cycle, we use PEI [51] as a remedy for the search with subsets. The PEI can be expressed as:

$$PEI(\mathbf{x}, \mathbf{x}^*) = EI(\mathbf{x}) \times IF(\mathbf{x}, \mathbf{x}^*)$$

$$IF(\mathbf{x}, \mathbf{x}^*) = 1 - \exp\left(-\sum_{h=1}^d \|x_h - x_h^*\|^2 / 2l_h^2\right) \quad (21)$$

where, \mathbf{x}^* is the selected sample by $\max EI$ and l_h^2 is the characteristic length of covariance function in Eq.(4). According to our observations, PEI will only be called in rare cases. More details are shown in Algorithm 3.

Algorithm 4 Sample exchange and update of training set

Input: The new sample points as $X_{new} = \mathbf{x}_{new}^{SOGP} \cup X_{new}^{MFGP}$

Output: The updated training sets of SOGP and MFGP

- 1: Evaluate X_{new} in parallel with the real objective function
 - 2: Update the current best optimal solution,
 - i. $f_{\min} \leftarrow \min(f_{\min}, Y(X_{new}))$
 - ii. $\mathbf{x}_{\min} \leftarrow \mathbf{x} \in X : y(\mathbf{x}) = f_{\min}$
 - 3: Update the training set of SOGP
 - i. $X \leftarrow X \cup X_{new}$
 - ii. $Y \leftarrow Y \cup Y_{new}$
 - 4: Update the training set of MFGP
 - 5: **if** \mathbf{x}_{\min} comes from \mathbf{x}_{new}^{SOGP} **then**
 - 6: **for** $i \leq m$ **do**
 - 7: $X_i \leftarrow X_i \cup \mathbf{x}_{new}^{SOGP}$
 - 8: $Y_i \leftarrow Y_i \cup Y_{new}^{SOGP}$
 - 9: **end for**
 - 10: **else if** \mathbf{x}_{\min} comes from X_{new}^{MOGP} **then**
 - 11: $nbest = \arg \min(Y_{new}^{MOGP})$
 - 12: **for** $i \leq m$ & $i \neq nbest$ **do**
 - 13: $X_i \leftarrow X_i \cup X_{new}^{MFGP}(nbest)$
 - 14: $Y_i \leftarrow Y_i \cup Y_{new}^{MFGP}(nbest)$
 - 15: **end for**
 - 16: Randomly choose X_j
 - 17: $X_j \leftarrow X_j \cup \mathbf{x}_{new}^{SOGP}$
 - 18: $Y_j \leftarrow Y_j \cup Y_{new}^{SOGP}$
 - 19: **end if**
-

D. Sample Exchange between Multiple GP models

On one hand, sample exchanges will provide clues to each GP model about where the real optimum may locate; on the other hand, we should keep model diversity to gain maximum improvement by the co-training style process. Therefore, we randomly assign the new sample of SOGP model to a MFGP subset if it was inferior to the current best solution f_{\min} , otherwise it will be added to each MFGP subset. In the meantime, we directly add the new querying samples of MFGP models to the SOGP training set. The implementation details are shown in Algorithm 4.

IV. EXPERIMENTS

In this section, we test CLBO and compare it with both classical and recent BO algorithms on numerical and engineering benchmark problems. After that, we examine the influence of different GP models on CLBO performance.

A. Baselines for comparison

To show the effectiveness of CLBO, it is compared with the classical EGO [7] and the following three kinds of BO algorithms.

To verify the effectiveness of exploiting both model diversity and agreement on unlabeled information in surrogate management for BO, CLBO is compared with algorithms that purely take advantage of the differences between models to improve algorithm performance. In particular, an efficient global optimization algorithm assisted by multiple surrogate techniques (labeled as MSEGO [26]) is taken as a baseline.

In MSEGO, the predictions from surrogates such as radial basis function (RBF), support vector machine (SVM), etc. are combined with the uncertainty imported from GP to formulate new EI acquisition functions, and then a set of samples are queried in each optimization cycle by maximizing the new EI functions. Additionally, we also build an algorithm labeled as MSBO, which trains the surrogates with the full and subset samples using standard GP while remaining the other settings exactly the same as those of CLBO.

To show the effectiveness of improving GP surrogate accuracy other than modifying the acquisition function to address the suboptimal problem of BO, CLBO is compared against two variants of EI-based algorithms, i.e., generalized EI (GEI) [43] and Exploration Enhanced EI (E3I) [44]. Specifically, GEI proposes to address the “over-exploitation” issue of EI (see Eq.(10)) by dynamically adjusting the weight of $\sigma(\mathbf{x})$ with a global factor g ; and E3I proposes to use the optima of Thompson samples to replace the current best solution f_{\min} as the incumbent in EI, and therefore lowering the weight of $\partial \hat{y} / \partial \mathbf{x}$ to resolve the “over-exploitation” when $z > 0$. Additionally, the standard BO with MCMC hyperparameters [52] is selected as a baseline, labeled as BO-MCMC in the following tests.

Since CLBO evaluates multiple samples in each optimization cycle, it is also compared against a set of batch BO algorithms, such as constant liar (CL) [40], batch BO via local penalization (BBO-LP) [41] and pseudo EI (PEI) [51]. The common feature of these batch algorithms is that they all rely on a single GP surrogate to guide the optimization process.

B. Implementing Details

Except BO-MCMC, BBO-LP and MSEGO, CLBO and the other compared baselines are all built based on the GPML toolbox [6]. In particular, the MFGP in CLBO is programmed based on the MTGP toolbox¹, which is an extension of the GPML toolbox. Additionally, the GpyOpt [53] is used to run BO-MCMC and BBO-LP. The source code of MSEGO, which is built upon the SURROGATES toolbox², is also employed for the tests.

The number of initial training samples is set as six times of the problem dimension. For numerical benchmark functions, the sample budget is set as 30 times of the problem dimension. However, the sample evaluation for engineering benchmarks can be quite expensive, and therefore the number of total evaluations for engineering benchmarks is set as no greater than 200. Furthermore, since the best solutions of numerical benchmark functions are known, the regret [54] is used to measure the best solution obtained thus far, i.e., $R_t = f_{\min}^{(t)} - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$, wherein $\min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ is the real optimal solution of the benchmarks, and $f_{\min}^{(t)}$ denotes the best solution at the t^{th} iteration. For engineering benchmarks, the best objective function values obtained thus far are directly used to show the convergence history of compared algorithms.

By default, we use MFGP to obtain two GP predictions with the training subsets, and one single-output GP with the full samples in CLBO. The influence of different GP models

¹<https://www.rob.uni-luebeck.de/index.php?id=410&L=0>

²<https://sites.google.com/site/srgtstoolbox/>

on CLBO performance are further inspected in Subsection E. Accordingly, three samples are evaluated in each optimization cycle when running batch BO algorithms such as CL, PEI and BBO-LP. By following the demo shown in SURROGATE toolbox, the radial basis function (RBF) and linear Shepard function are used to combine with standard GP for MSEGO.

Although the algorithm efficiency can be boosted by evaluating a set of samples simultaneously, the main purpose of this paper is to address the suboptimal problem of BO. Therefore, the performance of sequential and batch BO algorithms are compared by the convergence history in terms of function calls within a fixed sample budget. And accordingly, the termination conditions for the compared algorithms are set as follows:

- (1) $n_{total} \leq n_T$, i.e., the total number of function calls n_{total} cannot be greater than the sample budget n_T .
- (2) $t_{total} \leq t_T$, i.e., the total number of iterations t_{total} should be no greater than a threshold t_T .

In particular, the condition $t_{total} \leq t_T$ is set to address the following issue. That is, to avoid the ill-conditioning of the correlation matrix a distance constraint $\min\|\mathbf{x} - \mathbf{x}_i\| > \varepsilon$ is imposed, where \mathbf{x}_i is any one of the training samples, and ε is set to be 0.001 in the following tests. Nevertheless, the constraint $\min\|\mathbf{x} - \mathbf{x}_i\| > \varepsilon$ can prevent an algorithm from querying more samples at later stage of the iteration process, and consequently the algorithm may get trapped in infinite loop but cannot meet the termination condition $n_{total} \leq n_T$. Therefore, we propose to terminate the algorithm when either the condition (1) or (2) is met.

C. Comparisons on Numerical Tests

Five numerical benchmark functions are selected from [55] to inspect the algorithm performance. Specifically, Hartman6 is a non-convex problem that frequently used as a benchmark for testing BO algorithms [1], [2]; Michalewicz is a multimodal function characterized by steep ridges; and Rastrigin and Ackley are functions featured by many local optima. The Michalewicz, Rastrigin and Ackley are tested over five-dimensional design spaces. The steepness parameter m of Michalewicz is set to be 10. While the variable ranges of Michalewicz and Rastrigin are set according to the descriptions in [55]; the range of each variable in Ackley is set as [-2, 2], by referring to [56]. Additionally, the 10-dimensional Trid function is also selected for the test. As the function values of Trid vary in several orders of magnitude, it cannot be easily fitted with a small portion of samples.

Figures 2 to 6 show the testing results on numerical benchmarks, where the shaded areas in Figs.2(a) to 6(a) and Figs. 2(b) to 6(b) show the results variation of different algorithms over 20 runs. The boxplots in Figs.2(c) to 6(c) exhibit the medians and the distributions of the final optimal solutions. Obviously, CLBO always achieves the best or the second best optimal solutions within sample budget.

When comparing BO-MCMC with CLBO and the other compared algorithms, it shows faster convergence rate at the early stage of the optimization process, especially when optimizing the Rastrigin and Ackley functions. However, at later stage of the optimization process, CLBO can have better

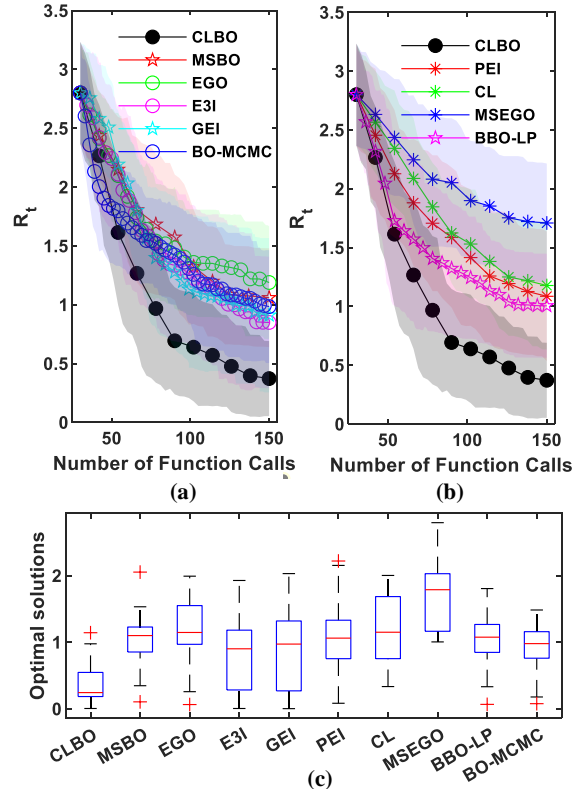


Fig. 2. Test results of five-dimensional Michalewicz function, (a) comparison with sequential BO algorithms, (b) comparison with batch BO algorithms, (c) boxplot of the regrets of final optimal solutions

results than those of BO-MCMC. Among the compared algorithms, E3I performs the best when optimizing the functions such as Trid and Hartman6 functions. However, it performs much poorer when testing the Michalewicz and Rastrigin functions. The reason behind can be explained as follows. To address the suboptimal problem of BO, E3I proposes to use the optima of Thompson samples to replace the current best solution f_{min} as the incumbent, and thereby the value of z becomes smaller in the EI formulation. In particular, the smaller $z > 0$ helps to lower the weight of $\partial\hat{y}/\partial\mathbf{x}$ (see Eq.(10)) and thus addressing the “over-exploitation” issue. However, the smaller $z < 0$ will increase the weight of $\partial\sigma/\partial\mathbf{x}$ (see Eq.(11)), which can make the algorithm get stuck in the “over-exploration” more easily. For the convex functions like Trid, the optimization stage with $z < 0$ can be delayed or even avoided when optimizing it with E3I. That is, by reducing the value of $z > 0$ to prevent too greedily sampling around the optima of the fitted surrogate, E3I is able to maintain the effective exploitation-exploration tradeoff with $z > 0$ and thus achieving better results than the other compared algorithms. However, $z < 0$ can take place many times for the functions with many local optima, therefore E3I may even increase the risk of “over-exploration” when $z < 0$. Hence, E3I performs much poorer when optimizing the Michalewicz and Rastrigin functions.

Similar situations can be observed for GEI, which proposes to resolve the “over-exploitation” issue by dynamically augmenting the role of $\sigma(\mathbf{x})$ with a factor g . However, adjusting

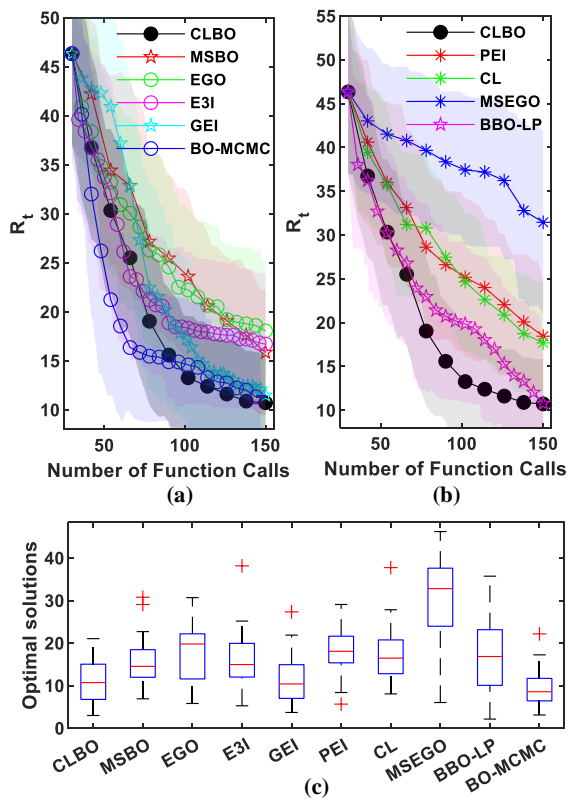


Fig. 3. Test results of five-dimensional Rastrigin function, (a) comparison with sequential BO algorithms, (b) comparison with batch BO algorithms, (c) boxplot of the regrets of final optimal solutions

$\sigma(\mathbf{x})$ also cannot address the “over-exploration” issue with $z < 0$. Moreover, the value of g are set empirically in [43], which may not adapt well to the real objective landscape. Hence, except Rastrigin function, GEI performs relatively poor when optimizing the benchmark functions.

With the above, it indicates that modifying the EI acquisition function may not be able to fully resolve the suboptimal problem of BO. Intrinsically, such suboptimal problem can attribute to the poor accuracy of the trained GP surrogate particularly that in the neighborhood of the real optimal solution. However, no surrogate can be accurate at every site of the design space. Therefore, an alternative way to address the suboptimal problem is to improve the overall surrogate accuracy by using multiple models, which can be promising to address both the issues of “over-exploitation” and “over-exploration”.

However, the performance of MSEGO and MSBO, which make use of multiple surrogate models, are rather poor when optimizing the numerical benchmarks. According to Eq.(1), the differences between surrogate models indeed can help them to complement each other, but the surrogates may suffer from incredibly large individual prediction errors, resulting in worse overall prediction accuracy. In MSEGO and MSBO, no strategy is used to take care of the individual prediction accuracy. Moreover, the information exchange between the surrogates is limited to exchanging the new queried samples in each optimization cycle. Hence, the overall surrogate accuracy particularly that in the vicinity of the real optimal solutions

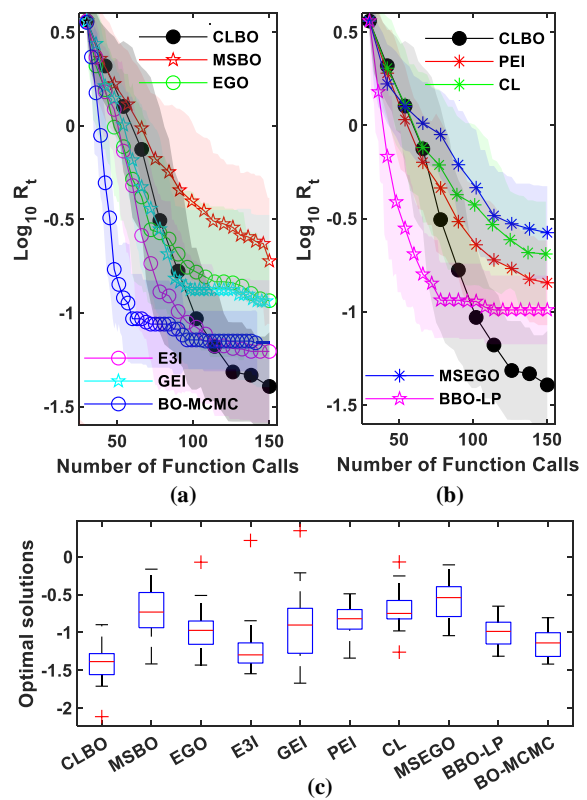


Fig. 4. Test results of five-dimensional Ackley function, (a) comparison with sequential BO algorithms, (b) comparison with batch BO algorithms, (c) boxplot of the regrets of final optimal solutions

may improve slowly as the iteration goes on, leading to the poor optimization results of MSEGO and MSBO.

Differently, good care of the individual model accuracy and model diversity are taken in CLBO. In particular, agreement constraint is imposed on curve bumpiness when training GP predictions with subset samples using MFGP. As has been proved in [32], the agreement constraint on unlabeled information helps to reduce the sample complexity of the effective hypothesis space. It means, with the same number of subset samples, the agreement constraint GP predictions can have smaller real prediction loss according to Eq.(2). Additionally, one more single-output GP is built with the full sample set, which helps to increase the model diversity with relatively good individual accuracy. And further, the combination of these reasonably accurate GP models can complement each other to achieve even better overall surrogate accuracy (see Eq.(1)) in the vicinity of the true optimal. Hence, our proposed CLBO always achieves the best or second best solutions with even faster convergence rate.

As the main purpose of batch BO algorithms such as BBO-LP is to enable parallel optimization without scarifying the global searching performance dramatically, the convergence performance of BBO-LP, PEI and CL look poorer than the sequential BO algorithms in most cases, when the function calls is concerned. By using multiple GP models and taking a balance in between individual model accuracy and model diversity, CLBO is the only batch BO algorithm which achieves better convergence rate than the sequential BO algorithms.

TABLE I
THE AVERAGED CPU TIME OF COMPARED ALGORITHMS WHEN TESTING THE FIVE BENCHMARK FUNCTIONS

Time (s)	Michalewicz	Rastrigin	Ackley	Hartman6	Trid
CLBO	274.4	247.1	236.1	228.0	888.7
MSBO	100.4	88.2	81.1	129.9	321.9
EGO	150.8	128.9	116.1	110.2	402.0
E3I	2511.5	2618.0	2579.4	1731.6	4595.9
GEI	127.0	126.6	108.6	171.6	383.2
BO-MCMC	6207.7	8681.0	4119.0	3426.3	16511.2
PEI	112.1	95.3	91.9	126.9	291.8
CL	164.1	149.9	125.4	124.1	480.4
MSEGO	693.5	653.1	727.1	1088.7	4505.9
BBO-LP	219.6	252.2	192.5	170.8	1030.2

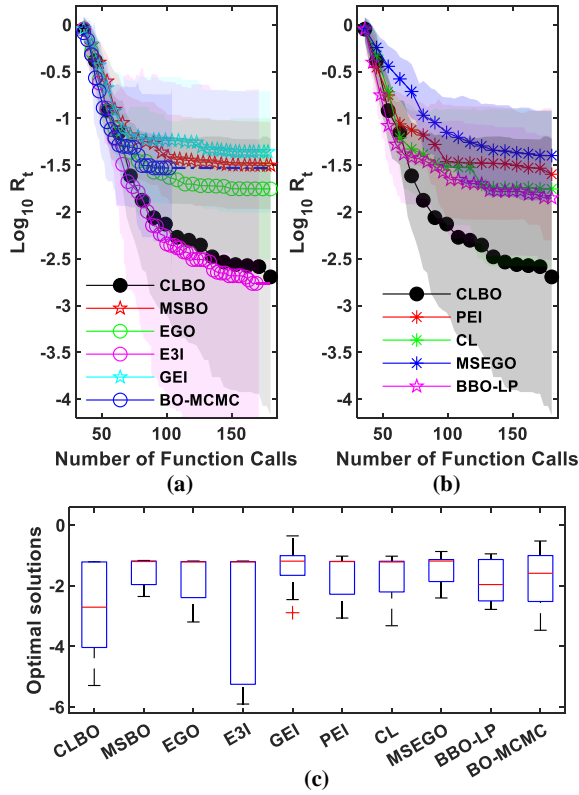


Fig. 5. Test results of Hartman6 function, (a) comparison with sequential BO algorithms, (b) comparison with batch BO algorithms, (c) boxplot of the regrets of final optimal solutions

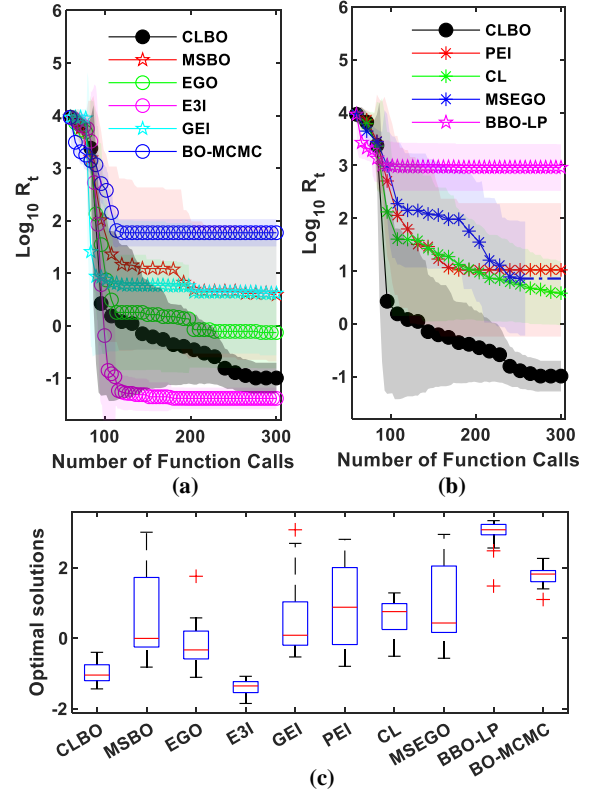


Fig. 6. Test results of 10-dimensional Trid function, (a) comparison with sequential BO algorithms, (b) comparison with batch BO algorithms, (c) boxplot of the regrets of final optimal solutions

To examine the computational costs of the compared algorithms, Table 1 lists the averaged CPU time over 20 runs. Since the dimensionality of the covariance matrix of MFGP (see Eq.(15)) in CLBO is larger than that of the standard GP (see Eq.(6)), the averaged CPU time of CLBO is larger than those of EGO, GEI and etc. In the meantime, due to the huge cost of Thompson sampling and MCMC, the computational cost of E3I and BO-MCMC can be one order higher than the other compared algorithms. Also note that, the BO algorithms are often applied to expensive problems, where a function call can take even much longer time to implement, the cost differences of BO algorithms may not be an issue.




Fig. 7. Schematic of a three-link bipedal robot [57]

D. Tests on Learning Problems

To further show the effectiveness of CLBO, we also test it on engineering benchmarks, such as the control of robotics




Fig. 8. Schematic of airfoil contour and the forces acting on it [2]

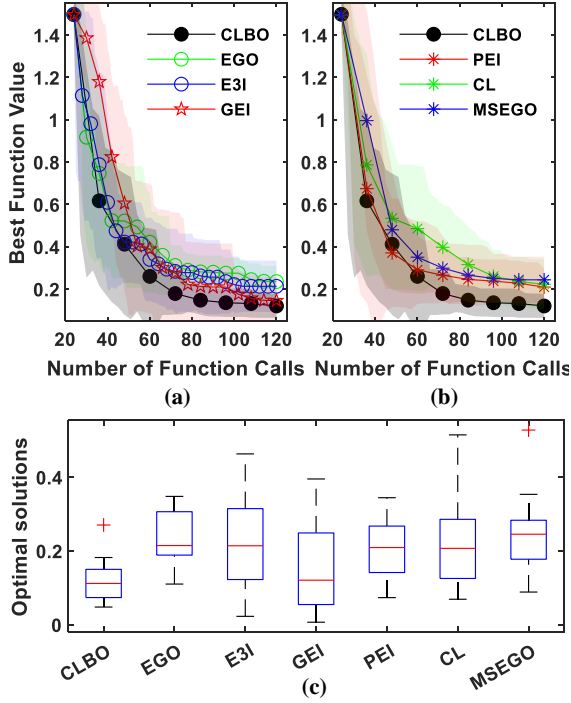


Fig. 9. Test results of active learning for robot pushing, (a) comparison with sequential BO algorithms, (b) comparison with batch BO algorithms, (c) boxplot of the regrets of final optimal solutions

systems and airfoil optimization [58], [59]. Note that simulation failures can be encountered when evaluating the samples for robotic systems and airfoil aerodynamic performance, which can get the algorithm stop early. Hence, we propose to evaluate a random sample instead to restart the optimization process once encountered, which needs code modifications in related BO algorithms. However, it may be not be easy for BBO-LP which implemented on a free-available software GPyOpt. On the other hand, as a contrast of CLBO, the poor performance of MSBO and the reason behind it have been discussed in the numerical tests. Hence, BBO-LP, BO-MCMC and MSBO are not tested for engineering benchmark problems. The detailed descriptions of the three selected engineering benchmarks are given as follows.

1) *Active Learning for Robot Pushing*: This task deals with the active learning for a pre-image learning problem for pushing [54]. The input is the pushing action of a robot, and the output is the distance of the pushed object to the expected location. The optimization objective is to find a good pre-image for pushing the object to the designated

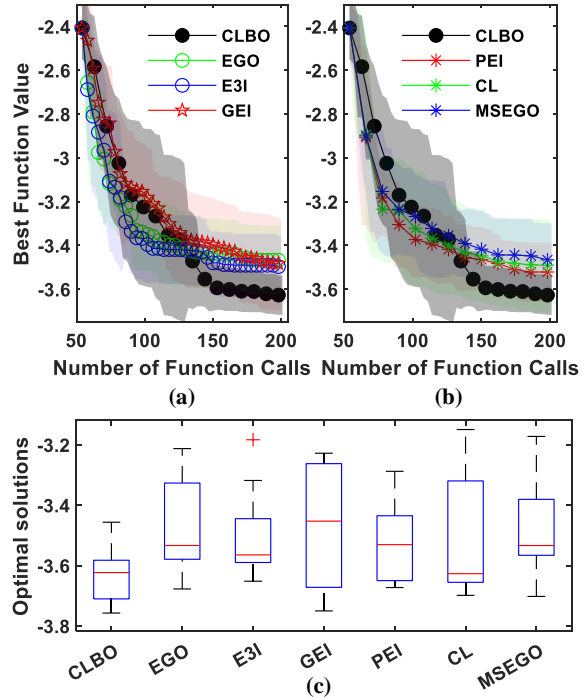


Fig. 10. Test results of a planar bipedal walker reinforcement learning task, (a) comparison with sequential BO algorithms, (b) comparison with batch BO algorithms, (c) boxplot of the regrets of final optimal solutions

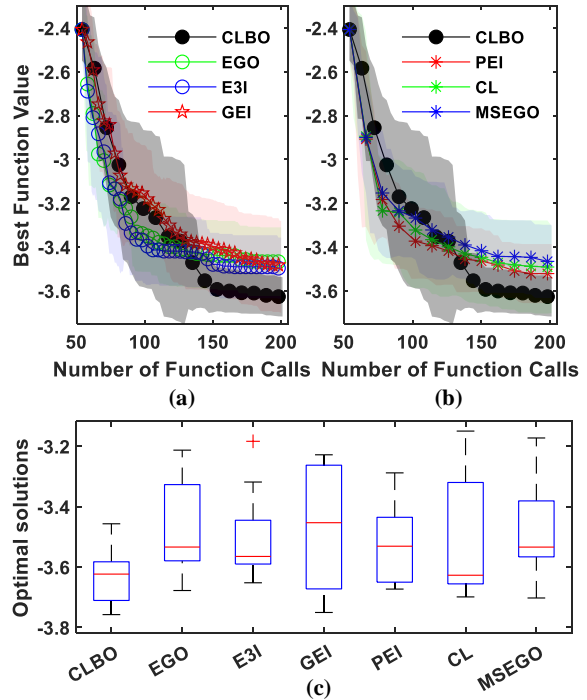


Fig. 11. Test results of airfoil aerodynamic optimization, (a) comparison with sequential BO algorithms, (b) comparison with batch BO algorithms, (c) boxplot of the regrets of final optimal solutions

location. Four variables including the robot location and angle as (r_x, r_y, r_θ) and the pushing duration (t) are considered for the optimization of this task.

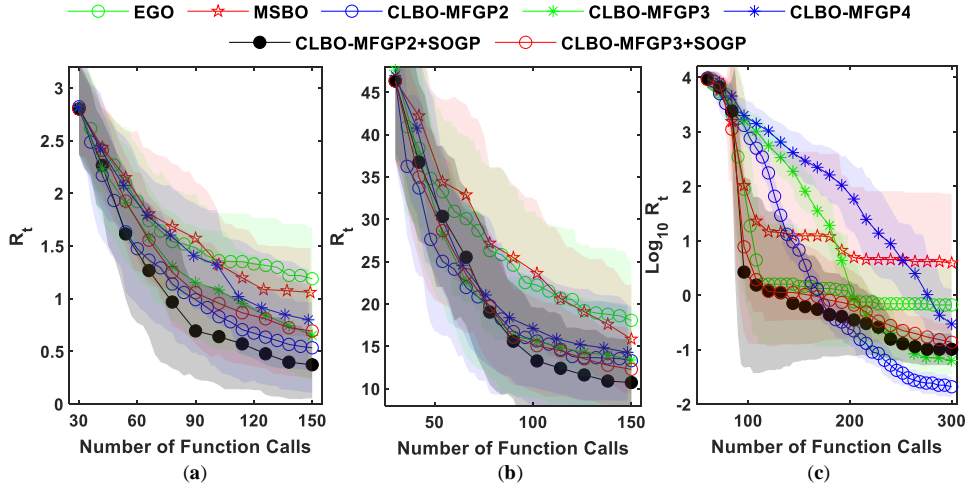


Fig. 12. Sensitivity analysis of GP models by testing on representative benchmark functions, (a) Michalewicz function, (b) Rastrigin function, (c) Trid function

2) *Reinforcement learning control of a planar bipedal walker*: This problem considers the walking pattern control of a three-link bipedal walker in [57], which is a reinforcement learning problem. Specifically, the controller modulates the walking pattern of a planar bipedal robot. The goal is to make the robot walking in a fast upright pattern, and the reward is the walking speed with a penalty for not maintaining the upright position. The controller has nine parameters as shown in Fig.7, where r is the length of the leg, l is the distance from the hips to the center of mass of the torso, $(\theta_1, \theta_2, \theta_3)$ are the absolute orientations of the various links, (q_1, q_2) are the relative angles and q_3 is the absolute orientation of the torso. In addition, there is one more parameter as $d\theta_1$, which is the input velocity of the stance leg.

3) *The airfoil design problem*: This problem considers the airfoil shape optimization in order to achieve excellent aerodynamic performance for aircraft wing. As shown in Fig.8, the forces acting on the airfoil can be decomposed into two components as the drag (D) and the lift (L). The drag and lift and their ratio L/D as well are the functions of the airfoil contour. In the meantime, the airfoil contours can be in various shapes according to different working conditions determined by the aerodynamic parameters such as the Mach number (Ma_∞), the Reynolds number (Re) and the angle of attack (AoA). In this study, we follow the settings in [2] for the optimization of a low-speed airfoil with 13 design variables. In particular, the design objective is to maximize the lift to drag ratio, i.e., L/D , with the operating conditions set as $Re = 1.8 \times 10^6$, $Ma = 0.01$ and $AoA = 0$ deg. The XFOIL [60] is used to calculate the objective performance L/D . The detailed settings of 13 design variables for airfoil optimization can refer to [2].

Figures 9 to 11 show the optimization results of engineering benchmark problems over 10 runs. Similar to the results shown in the numerical tests, CLBO achieves the best solutions within budget. For the 4-dimensional optimization problem as active learning for robot pushing, CLBO achieves better solutions almost from the start of the optimization process. However, when optimizing the 9-dimensional bipedal walker and 13-

dimensional airfoil design problems, CLBO progresses slower at the early stage. The slower convergence rate of CLBO can be attributed to the poor prediction accuracy built with a small portion of subset samples at the early stage. However, the situation is changed by imposing agreement constraint with sequentially added new samples, as seen by the faster convergence rate of CLBO at the later stage. In contrast, MSEGO, which also uses multiple surrogate models to guide the optimization process, obtains much worse results at the end of the optimization process.

More importantly, as shown in the boxplots in Figs.9(c) to 11(c), the variation of final solutions of CLBO is the smallest among the compared algorithms, and the medians of CLBO are also better. The above can be attributed to the exploiting of model agreement on unlabeled information. On one hand, the convergence of BO is highly related to the surrogate prediction accuracy in the vicinity of optimal solutions. On the other hand, the agreement constraint on curve bumpiness helps to reduce the sample complexity of the effective hypothesis space [32]. Then, with the same number of subset samples, the agreement constrained GP predictions can be expected to achieve better prediction accuracy in neighborhood of optimal solutions(see Eq.(2)). Additionally, the single-output GP trained with the full samples also help to increase the overall surrogate accuracy, by increasing the model diversity without scarifying the individual model error in BO process. With the above, CLBO outperforms both the sequential and batch BO algorithms.

E. Sensitivity Analysis on the Selection of GP Models

To address the suboptimal problem of BO, we propose to exploit both model diversity and agreement on unlabeled information to improve the overall surrogate accuracy in BO process. In doing so, the agreement-constraint MFGP is combined with SOGP to build CLBO, their influences on algorithm performance are examined in this subsection.

Specifically, three versions of CLBO that use MFGP alone are built to inspect the effect of agreement-constrained modeling, which are labeled as CLBO-MFGP2, CLBO-MFGP3

and CLBO-MFGP4, respectively. In particular, CLBO-MFGP2 means two GP predictions are built with bootstrapping based sub-sampling (see Algorithm 2) in CLBO, similar treatment for CLBO-MFGP3 and CLBO-MFGP4. Additionally, two more versions of CLBO that combines MFGP with SOGP are also tested, denoted by CLBO-MFGP2+SOGP and CLBO-MFGP3+SOGP, respectively. As contrast, the results of EGO and MSBO are also presented, as shown in Fig.12.

Among the five versions of CLBO, CLBO-MFGP2+SOGP achieves the best solutions when optimizing the Michalewicz and Rastrigin functions, and CLBO-MFGP2 performs the best on Trid. More interestingly, the convergence rates of CLBO-MFGP2+SOGP and CLBO-MFGP2 at the early stage of the optimization process are opposite to their final results. The reason behind can be explained as follows. For the multi-peak functions like Michalewicz and Rastrigin, the agreement-constrained MFGP predictions can have reasonable accuracy in promising areas where relatively better solutions may locate. Then, *similar to the local exploitation, the agreement on curve bumpiness confines the searching space of MFGP predictions to smaller yet promising areas, and hence CLBO-MFGP2 convergences more rapidly at the early stage of optimization process.* Nevertheless, the accuracy of MFGP predictions may be poor in the region where even better solutions locate, but the SOGP trained with full samples can complement the above inaccuracy of MFGP predictions. As a consequence, CLBO-MFGP2+SOGP finally achieves better solutions.

On the other hand, for the high-dimensional convex functions like Trid, the accuracy of MFGP predictions trained with subset samples can be too poor to capture the global trend of the high-dimensional objective landscape at the early stage, and hence CLBO-MFGP2 progresses slowly at this stage. As contrast, the SOGP trained with full samples can have reasonable accuracy over space, and hence CLBO-MFGP2+SOGP converges much faster at the early the stage when optimizing Trid. Nevertheless, at later stage when MFGP predictions have sufficient accuracy to capture the convex function trend, the agreement on curve bumpiness makes CLBO-MFGP2 focus exploiting promising areas. And then, CLBO-MFGP2 progresses much faster and finally achieves even better solutions for this convex function.

Also note that, the performance of CLBO becomes worse with the increase of the number of agreement-constrained MFGP predictions. The reason can be explained as follows. As pointed out by [32], *the extent of performance improvement by enforcing model agreement on unlabeled information depends on the diversity of the individual models.* In particular for our study, increasing MFGP predictions can reduce the differences between MFGP and SOGP predictions. As a consequence, the overall surrogate accuracy in promising areas can become poorer (see Eq.(1)), resulting in slower convergence rate and worse final solutions.

V. CONCLUSION

In this paper, we propose a co-learning Bayesian optimization framework (CLBO), which exploits both diversity and agreement of multiple GP models on unlabeled information to improve the surrogate accuracy in BO process,

and therefore achieving more efficient global optimization. Through tests on both numerical toy problems and various engineering benchmarks, it shows that both the agreement constraint and the diversity generated by multiple models contribute to the success of CLBO. Meanwhile, as CLBO can be run in parallel, the encouraging results of CLBO may also provide clues to resolve the exploitation-exploration challenge in batch BO. In future, we will further examine the advantage of CLBO with acquisition functions other than EI and parallel strategies that previously proposed under a single GP. More importantly, we will explore more efficient ways to use the model agreements/disagreements on unlabeled information to further boost the BO performance.


ACKNOWLEDGMENTS

This research is partially funded by Center for Frontier AI Research (CFAR), Agency for Science, Technology and Research (A*STAR), Data Science and Artificial Intelligence Research Center (DSAIR), Nanyang Technological University, Singapore, the National Science and Technology Major Project (2019-II-0008-0028) and the National Science Foundation of China (NO. 51936008, and NO. 61802317).

REFERENCES

- [1] R. Martinez-Cantin, "Funneled bayesian optimization for design, tuning and control of autonomous systems," *IEEE transactions on cybernetics*, no. 99, pp. 1–12, 2018.
- [2] Z. Guo, Y. Ong, and H. Liu, "Calibrated and recalibrated expected improvements for bayesian optimization," *Structural and Multidisciplinary Optimization*, 2021.
- [3] L. Song, Z. Guo, J. Li, and Z. Feng, "Research on metamodel-based global design optimization and data mining methods," *Journal of Engineering for Gas Turbines and Power*, vol. 138, no. 9, 2016.
- [4] U. Martinez-Hernandez, T. J. Dodd, and T. J. Prescott, "Feeling the shape: Active exploration behaviors for object recognition with a robotic hand," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 48, no. 12, pp. 2339–2348, 2017.
- [5] K. Swersky, J. Snoek, and R. P. Adams, "Multi-task bayesian optimization," in *Advances in neural information processing systems*, 2013, pp. 2004–2012.
- [6] C. E. Rasmussen and H. Nickisch, "Gaussian processes for machine learning (gpml) toolbox," *Journal of machine learning research*, vol. 11, no. Nov, pp. 3011–3015, 2010.
- [7] 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, 1998.
- [8] I. Couckuyt, D. Deschrijver, and T. Dhaene, "Fast calculation of multi-objective probability of improvement and expected improvement criteria for pareto optimization," *Journal of Global Optimization*, vol. 60, no. 3, pp. 575–594, 2014.
- [9] P. Hennig and C. J. Schuler, "Entropy search for information-efficient global optimization," *Journal of Machine Learning Research*, vol. 13, no. Jun, pp. 1809–1837, 2012.
- [10] J. Wu and P. Frazier, "The parallel knowledge gradient method for batch bayesian optimization," in *Advances in Neural Information Processing Systems*, 2016, pp. 3126–3134.
- [11] B. E. Stuckman and E. E. Easom, "A comparison of bayesian/sampling global optimization techniques," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 22, no. 5, pp. 1024–1032, 1992.
- [12] D. Huang, T. T. Allen, W. I. Notz, and N. Zeng, "Global optimization of stochastic black-box systems via sequential kriging meta-models," *Journal of global optimization*, vol. 34, no. 3, pp. 441–466, 2006.
- [13] A. D. Bull, "Convergence rates of efficient global optimization algorithms," *Journal of Machine Learning Research*, vol. 12, no. Oct, pp. 2879–2904, 2011.
- [14] C. Qin, D. Klabjan, and D. Russo, "Improving the expected improvement algorithm," in *Advances in Neural Information Processing Systems*, 2017, pp. 5381–5391.

- [15] M. J. Sasena, "Flexibility and efficiency enhancements for constrained global design optimization with kriging approximations," Ph.D. dissertation, University of Michigan., 2002.
- [16] 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, 2015.
- [17] T. T. Joy, S. Rana, S. Gupta, and S. Venkatesh, "A flexible transfer learning framework for bayesian optimization with convergence guarantee," *Expert Systems with Applications*, vol. 115, pp. 656–672, 2019.
- [18] A. T. W. Min, A. Gupta, and Y.-S. Ong, "Generalizing transfer bayesian optimization to source-target heterogeneity," *IEEE Transactions on Automation Science and Engineering*, 2020.
- [19] K. Kandasamy, G. Dasarathy, J. Schneider, and B. Póczos, "Multi-fidelity bayesian optimisation with continuous approximations," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 1799–1808.
- [20] Z. Guo, L. Song, C. Park, J. Li, and R. T. Haftka, "Analysis of dataset selection for multi-fidelity surrogates for a turbine problem," *Structural and Multidisciplinary Optimization*, vol. 57, no. 6, pp. 2127–2142, 2018.
- [21] Z. Guo, Q. Wang, L. Song, and J. Li, "Parallel multi-fidelity expected improvement method for efficient global optimization," *Structural and Multidisciplinary Optimization*, 2021.
- [22] A. Gupta, Y.-S. Ong, and L. Feng, "Insights on transfer optimization: Because experience is the best teacher," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 1, pp. 51–64, 2017.
- [23] Q. Zhang, C. Zhou, N. Xiong, Y. Qin, X. Li, and S. Huang, "Multimodel-based incident prediction and risk assessment in dynamic cybersecurity protection for industrial control systems," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 46, no. 10, pp. 1429–1444, 2015.
- [24] A. T. W. Min, Y.-S. Ong, A. Gupta, and C.-K. Goh, "Multiproblem surrogates: transfer evolutionary multiobjective optimization of computationally expensive problems," *IEEE Transactions on Evolutionary Computation*, vol. 23, no. 1, pp. 15–28, 2017.
- [25] Z. Zhou, Y. S. Ong, P. B. Nair, A. J. Keane, and K. Y. Lum, "Combining global and local surrogate models to accelerate evolutionary optimization," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 37, no. 1, pp. 66–76, 2006.
- [26] F. A. C. Viana, R. T. Haftka, and L. T. Watson, "Efficient global optimization algorithm assisted by multiple surrogate techniques," *Journal of Global Optimization*, vol. 56, no. 2, pp. 669–689, 2013.
- [27] T. He, Y. Liu, T. H. Ko, K. C. Chan, and Y.-S. Ong, "Contextual correlation preserving multiview featured graph clustering," *IEEE transactions on cybernetics*, 2019.
- [28] C. Xu, D. Tao, and C. Xu, "A survey on multi-view learning," *arXiv preprint arXiv:1304.5634*, 2013.
- [29] Z.-H. Zhou, "When semi-supervised learning meets ensemble learning," in *International Workshop on Multiple Classifier Systems*. Springer, 2009, pp. 529–538.
- [30] A. Krogh and J. Vedelsby, "Neural network ensembles, cross validation, and active learning," in *Advances in neural information processing systems*, 1995, pp. 231–238.
- [31] Luca, Oneto, Alessandro, Ghio, Sandro, Ridella, Davide, and Anguita, "Local rademacher complexity: Sharper risk bounds with and without unlabeled samples," *Neural Networks*, 2015.
- [32] B. Leskes and L. Torenvliet, "The value of agreement a new boosting algorithm," *Journal of Computer and System Sciences*, vol. 74, no. 4, pp. 557–586, 2008.
- [33] J. Farquhar, D. Hardoon, H. Meng, J. S. Shawe-Taylor, and S. Szedmak, "Two view learning: Svm-2k, theory and practice," in *Advances in neural information processing systems*, 2006, pp. 355–362.
- [34] A. Blum and T. Mitchell, "Combining labeled and unlabeled data with co-training," in *Proceedings of the eleventh annual conference on Computational learning theory*, 1998, pp. 92–100.
- [35] R. Navaratnam, A. W. Fitzgibbon, and R. Cipolla, "The joint manifold model for semi-supervised multi-valued regression," in *2007 IEEE 11th International Conference on Computer Vision*. IEEE, 2007, pp. 1–8.
- [36] Y. Zhou and S. Goldman, "Democratic co-learning," in *16th IEEE International Conference on Tools with Artificial Intelligence*. IEEE, 2004, pp. 594–602.
- [37] I. Muslea, S. Minton, and C. A. Knoblock, "Active learning with multiple views," *Journal of Artificial Intelligence Research*, vol. 27, pp. 203–233, 2006.
- [38] E. V. Bonilla, K. M. Chai, and C. Williams, "Multi-task gaussian process prediction," in *Advances in neural information processing systems*, 2008, pp. 153–160.
- [39] R. L. Riche and V. Picheny, "Revisiting bayesian optimization in the light of the coco benchmark," 2021.
- [40] D. Ginsbourger, R. Le Riche, and L. Carraro, "Kriging is well-suited to parallelize optimization," in *Computational intelligence in expensive optimization problems*. Springer, 2010, pp. 131–162.
- [41] J. González, Z. Dai, P. Hennig, and N. D. Lawrence, "Batch bayesian optimization via local penalization," *statistics*, 2015.
- [42] M. Schonlau and W. D. R. Jones, *Global versus local search in constrained optimization of computer models*. Institute of Mathematical Statistics, 1998.
- [43] W. Ponweiser, T. Wagner, and M. Vincze, "Clustered multiple generalized expected improvement: A novel infill sampling criterion for surrogate models," in *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*. IEEE, 2008.
- [44] J. Berk, V. Nguyen, S. Gupta, S. Rana, and S. Venkatesh, "Exploration enhanced expected improvement for bayesian optimization," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2018.
- [45] H. Liu, J. Cai, and Y.-S. Ong, "Remarks on multi-output gaussian process regression," *Knowledge-Based Systems*, vol. 144, pp. 102–121, 2018.
- [46] R. Dürichen, M. A. Pimentel, L. Clifton, A. Schweikard, and D. A. Clifton, "Multi-task gaussian process models for biomedical applications," in *IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)*. IEEE, 2014, pp. 492–495.
- [47] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2009.
- [48] J. Luo, A. Gupta, Y.-S. Ong, and Z. Wang, "Evolutionary optimization of expensive multiobjective problems with co-sub-pareto front gaussian process surrogates," *IEEE transactions on cybernetics*, vol. 49, no. 5, pp. 1708–1721, 2018.
- [49] P. Wei, R. Sagarna, Y. Ke, Y.-S. Ong, and C.-K. Goh, "Source-target similarity modelings for multi-source transfer gaussian process regression," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 3722–3731.
- [50] B. Da, A. Gupta, and Y.-S. Ong, "Curbing negative influences online for seamless transfer evolutionary optimization," *IEEE transactions on cybernetics*, vol. 49, no. 12, pp. 4365–4378, 2018.
- [51] D. Zhang, J. Qian, and Y. Cheng, "Pseudo expected improvement criterion for parallel ego algorithm," *Journal of Global Optimization*, vol. 68, no. 3, pp. 641–662, 2017.
- [52] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *Advances in neural information processing systems*, 2012, pp. 2951–2959.
- [53] T. G. authors, "GPyOpt: A bayesian optimization framework in python," <http://github.com/SheffieldML/GPyOpt>, 2016.
- [54] Z. Wang and S. Jegelka, "Max-value entropy search for efficient bayesian optimization," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 3627–3635.
- [55] S. Surjanovic and D. Bingham, "Virtual library of simulation experiments: Test functions and datasets," Retrieved June 10, 2021, from <http://www.sfu.ca/ssurjano>.
- [56] Y. Zhang, Z.-H. Han, and K.-S. Zhang, "Variable-fidelity expected improvement method for efficient global optimization of expensive functions," *Structural and Multidisciplinary Optimization*, vol. 58, no. 4, pp. 1431–1451, 2018.
- [57] E. R. Westervelt, J. W. Grizzle, C. Chevallereau, J. H. Choi, and B. Morris, *Feedback control of dynamic bipedal robot locomotion*. CRC press, 2018.
- [58] K. Chiu, M. Fuge *et al.*, "Airfoil design parameterization and optimization using b²ezier generative adversarial networks," *arXiv preprint arXiv:2006.12496*, 2020.
- [59] H. Wang, J. Doherty, and Y. Jin, "Hierarchical surrogate-assisted evolutionary multi-scenario airfoil shape optimization," in *2018 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, 2018, pp. 1–8.
- [60] M. Drela, "Xfoil: An analysis and design system for low reynolds number airfoils," in *Low Reynolds number aerodynamics*. Springer, 1989.




Zhendong Guo received the Ph.D. degree from the School of Energy and Power Engineering, Xi'an Jiaotong University (XJTU), China, in 2018. From 2019 to 2021, he was a research fellow with Data Science and AI Research (DSAIR) Center at Nanyang Technological University, Singapore. He is currently an assistant professor at Xi'an Jiaotong University, China. His main research interests include Bayesian optimization, generative model based optimization, multi-fidelity surrogate-based modeling and surrogate-based optimization and data mining for turbomachinery design.

ing for turbomachinery design.



Yew-Soon Ong (M'99-SM'12-F'18) received the Ph.D. degree in artificial intelligence in complex design from the University of Southampton, U.K., in 2003. He is President's Chair Professor in Computer Science at Nanyang Technological University (NTU) and Chief Artificial Intelligence Scientist of the Agency for Science, Technology and Research Singapore. At NTU, he serves as co-Director of the Singtel-NTU Cognitive & Artificial Intelligence Joint Lab. His Research interest is in artificial and computational intelligence. He is founding Editor-

in-chief of the IEEE Transactions on Emerging Topics in Computational Intelligence and associate editor of IEEE Transactions on Neural Networks & Learning Systems, IEEE Transactions on Cybernetics, IEEE Transactions on Artificial Intelligence.



Tiantian He received the Ph.D. degree in Computer Science from the Hong Kong Polytechnic University in 2017. Currently He is a Research Scientist in Center for Frontier AI Research (CFAR), Institute of High Performance Computing (IHPC), Agency for Science, Technology and Research (A*STAR). His research interests include artificial intelligence, machine learning, data-centric transfer optimization, and data mining.



Haitao Liu received the Ph.D. degree from the School of Energy and Power Engineering, Dalian University of Technology, China, in 2016. From 2016 to 2019, he was a research fellow with Rolls-Royce@NTU Corp Lab at Nanyang Technological University, Singapore. He is currently an associate professor at Dalian University of Technology, China. His main research interests include Gaussian process, representation learning, multi-task learning, active learning and intelligent design of turbomachinery. Dr. Liu is an Associate Editor of the IEEE

Transactions on Emerging Topics in Computational Intelligence.