Alleviating Search Bias in Evolutionary Bayesian Optimization with Many Heterogeneous Objectives

Xilu Wang, Yaochu Jin, Sebastian Schmitt, Markus Olhofer

2024

Preprint:

This is an accepted article published in IEEE Transactions on Systems, Man and Cybernetics: Systems. The final authenticated version is available online at: https://doi.org/10.1109/TSMC.2023.3306085

Alleviating Search Bias in Bayesian Evolutionary Optimization with Many Heterogeneous Objectives

Xilu Wang, Yaochu Jin, *Fellow*, *IEEE*, Sebastian Schmitt, and Markus Olhofer

***Abstract*—Multi-objective optimization problems whose objectives have different evaluation costs are commonly seen in the real world. Such problems are now known as multi-objective optimization problems with heterogeneous objectives (HE-MOPs). So far, however, only a few studies have been reported on addressing HE-MOPs, and most of them focus on bi-objective problems with one fast objective and one slow objective. In this work, we aim to deal with HE-MOPs having more than two black-box and heterogeneous objectives. To this end, we develop a multi-objective Bayesian evolutionary optimization approach to HE-MOPs that can alleviate search biases resulting from the different numbers of function evaluations allowed for the cheap and expensive objectives, which is achieved by designing a new acquisition function that penalizes the search bias towards the fast objectives, thereby achieving a balance between convergence and diversity. In addition, to make the best use of the different amounts of training data while avoiding increasing the computational cost, an ensemble consisting of two GPs is constructed for each cheap objective, one trained on the data collected before the Bayesian optimization starts, and the other on those evaluated during the Bayesian evolutionary optimization. Empirical studies on widely used multi-/many-objective benchmark problems whose objectives are assumed to be heterogeneously expensive demonstrate that the proposed algorithm is able to find high-quality solutions for HE-MOPs compared with the state-of-the-art methods.**

***Index Terms*—Multi/many-objective optimization, different evaluation costs, heterogeneous objectives, surrogate-assisted evolutionary algorithm, Bayesian optimization.**

I. Introduction

SURROGATE-ASSISTED evolutionary algorithms (SAEAs) are powerful tools for optimizing computationally expensive multi-objective problems (MOPs), where several conflicting objective functions must be simultaneously optimized and the evaluations of the objectives are highly time-consuming. While conventional multi-objective evolutionary algorithms (MOEAs) assume that candidate solutions can be accurately evaluated, SAEAs typically construct computationally efficient surrogate models to approximate the expensive real objective functions, and then the surrogates are used together with the real objective functions to guide the evolutionary optimization, known as model management [1]. Various classification or regression models are commonly used as surrogates in SAEAs, including support vector machines [2], radial basis function networks [3], and Gaussian processes (GPs), also known as Kriging which is a special case of GPs [4]. Among them, the GP is a popular choice for modelling expensive objective functions due to its ability of capturing the model's beliefs over the unknown objective function, providing both estimated objective values and the uncertainty of the estimations. The estimations provided by GPs can be utilized to design an acquisition function to select the next new data point to be evaluated by the real objective functions, guiding the search of the optimum. An SAEA with a GP as the surrogate model and an acquisition function to select new samples is known as Bayesian optimization BO [5].

SAEAs typically assume that the evaluation of each objective function takes the same period of time. Consequently, the selection operator of an MOEA can be conducted and the evolutionary search can proceed to the next generation. This assumption, however, can be violated in practice, e.g., the evaluation of aerodynamic and structural mechanics performance of an airplane wing design [6]–[8] or a car shape [9], [10] involves computationally intensive computational fluid dynamics, where several hours of evaluation time are typical for a single fitness evaluation. Additionally, some types of evaluations can be an order of magnitude slower than others, for example, crashworthiness assessment is much more resource- and time-consuming than computationally intensive computational fluid dynamics. Such MOPs exhibit so-called heterogeneous objectives [11] and, in particular, we consider heterogeneously expensive MOPs (HE-MOPs) in this paper, where non-uniform evaluation times in expensive MOPs give rise to the heterogeneity.

Most recently, new MOEAs have been proposed to effectively address HE-MOPs. Most of the algorithms, however, are limited to considering a class of bi-objective HE-MOPs having one computationally cheap (fast) objective function f^c and one computationally expensive objective function f^e (also called delayed or slow objective). Existing methods for handling HE-MOPs can be roughly categorized into two groups, i.e., non-surrogate based and surrogate based methods, which are briefly reviewed below.

Non-surrogate based methods: Allmendinger *et al.* [12] first introduced HE-MOPs and proposed a ranking-based MOEA to allow solutions with missing objective function values caused by f^e to guide the search. A missing objective value caused by a pending evaluation is assigned with a pseu-

Xilu Wang is with the Faculty of Technology, Bielefeld University, 33619 Bielefeld, Germany. Part of the work was done when she was with the Department of Computer Science, University of Surrey, Guildford, GU2 7XH, United Kingdom. Email: xilu.wang@uni-bielefeld.de.

Yaochu Jin is with the Faculty of Technology, Bielefeld University, 33619 Bielefeld, Germany. He is also with the Department of Computer Science, University of Surrey, Guildford, GU2 7XH, United Kingdom. Email: yaochu.jin@uni-bielefeld.de. (*Corresponding author*)

Sebastian Schmitt and Markus Olhofer are with the Honda Research Institute Europe GmbH, Carl-Legien-Strasse 30, D-63073 Offenbach/Main, Germany. Email: {sebastian.schmitt;markus.olhofer}@honda-ri.de.

dovalue. Subsequently, new selection strategies are proposed based on the ranking subject to missing objective values. In a follow-up work by Allmendinger *et al.* [13], MOPs with non-uniform latencies are defined more formally based on the framework of MOEAs, and three general schemes are proposed to handle heterogeneous objectives, including *Waiting*, *Fast-first* and *Interleaving schemes*. While *Waiting* directly applies an MOEA to HE-MOPs by waiting for the completion of expensive evaluations, *Fast-first* employs a single-objective evolutionary algorithm (SOEA) to consume the additional fitness evaluations available for f^c during the waiting of expensive evaluations.Unlike *Fast-first*, more elaborated strategies are introduced in *Interleaving schemes* (i.e. *brood* and *speculative interleaving*) to utilize the limited evaluations by integrating the search results of each objective. Although the non-surrogate based methods shed light on possible directions for handling HE-MOPs, a major remaining issue is that the obtained solutions may be still far from Pareto optimal due to the limited evaluation budget available. In addition, how well they can scale to more complex problem settings with more objectives has not been explored.

Surrogate based methods: More recently SAEAs have been extended to HE-MOPs, which is motivated by the fact that SAEAs have emerged as powerful methods for the optimization of MOPs with expensive evaluations. Chugh *et al.* [14] proposed a heterogeneous Kriging-assisted evolutionary algorithm, called HK-RVEA. HK-RVEA adopted an SOEA and genetic operators to generate solutions for f^c when the initial population and the new samples are submitted for evaluations on both objectives, respectively. To make use of the additional evaluations on f^c, Wang *et al.* [15] developed a parameter-based transfer learning strategy based on a GP-assisted evolutionary algorithm (T-SAEA). In T-SAEA, common decision variables related to both f^c and f^e are determined first using a filter-based feature selection, then the corresponding parameters in the GPs can be shared.In a follow-up work, Wang *et al.* [16] proposed an instance-based transfer learning method (Tr-SAEA) to address the heterogeneous bi-objective problems. Domain adaptation techniques are adopted to generate synthetic samples for f^e, and a GP-based co-training method is introduced to augment the training data for surrogate models of f^e using the unlabeled synthetic data. Unfortunately, Tr-SAEA only learns the mapping in the objective space and requires an additional optimization method to obtain the corresponding solutions in the search space. Alternatively, in [17] a co-surrogate is adopted to model the relationship between f^c and f^e. Subsequently, transferable instances are identified from the search of f^c to speed up the optimization of f^e. Interestingly, trust region methods, instead of evolutionary algorithms, with the use of surrogates have also been successfully applied to HE-MOPs [18]. Multi-objective heterogeneous trust region algorithms are easily scalable to any number of objectives; however, they hinge on strong assumptions: the objective functions are black-box and twice continuously differentiable, and the cheap ones are given analytically and derivatives are easily available [18].

This work extends SAEAs to HE-MOPs with more general and more practical problem settings, i.e., HE-MOPs with more than three objectives (called HE-many-objective problems, HE-MaOPs) and different combinations of computationally cheap and expensive objectives. We propose to reduce the search bias resulting from the heterogeneous evaluation time of the objectives and limit the computational time for training GPs with an increased amount of data on the cheap objectives. The key contributions of the proposed search bias penalized BO, termed as SBP-BO, can be summarized as follows:

1) To make use of the different amount of available data evaluated on cheap and expensive objectives, an ensemble consisting of two GPs is constructed for each cheap objective, while one GP is used to approximate each expensive objective. Before optimization starts, while the initial population is evaluated on all objective functions, the cheap objectives can be explored using an SOEA, resulting in abundant extra solutions on the cheap objectives. For each cheap objective, one GP is trained and updated with the solutions evaluated during the Bayesian evolutionary optimization, while the other is trained with the data evaluated during the initialization only without re-training, thus making full use of all data available while avoiding highly intensive computation.

2) To alleviate the search bias towards the fast objectives and achieve a good balance between exploitation and exploration, a new acquisition function is proposed by introducing a penalty of search bias. The acquisition function can not only evaluate the quality of a solution in terms of convergence and diversity, but also promotes the exploration on the expensive objectives.

The rest of paper is organized as follows. Section II provides a problem description, followed by an introduction to multi-objective BO including GPs and acquisition functions. Then, the proposed SBP-BO is introduced in Section III. Section IV provides details about the experimental settings and Section V presents the experimental results to demonstrate the effectiveness of SBP-BO. Finally, we draw conclusions and discuss future research.

II. Background

A. Problem Description

Before we tackle the challenges posed by the presence of computationally heterogeneous objectives, we present the problem description at first. We consider expensive multi- and many-objective optimization problems with heterogeneous objectives (HE-MOPs and HE-MaOPs) in the following form:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_m(\mathbf{x})) \\ \text{s.t.} \quad & \mathbf{x} \in X \end{aligned} \tag{1}$$

where $\mathbf{x} = (x_1, x_2, \cdots, x_d)$ is the decision vector with d decision variables, X denotes the decision space, the objective vector $\mathbf{f}$ consists of m $(m > 2)$ objectives and for MaOPs the number of objectives m is larger than 3. The evaluation time of each objective is denoted by $\mathbf{t} = (t_1, t_2, \cdots, t_m)$, where we assume the objectives are ordered in terms of their computational complexity, ranked from the fastest with $i = 1$ to the slowest $i = m$, i.e. $t_1 \leq t_2 \leq \cdots \leq t_m$. The objective functions are black-boxes that can be evaluated by either time-consuming numerical simulations, or costly physical experiments. Building surrogate models based on the data collected via numerical simulations or experiments has been shown to be an efficient approach to such black-box expensive optimization problems [1]. In this work, we assume that the evaluation of each objective function can be done in parallel, and the computation time for building surrogates and applying the genetic operators of the evolutionary algorithm is negligible compared to that for evaluating the true objective functions. Therefore, we characterize the heterogeneity of the objectives by the number of affordable evaluations of an objective f_i relative to the slowest objective, f_m, which can be calculated given the evaluation time as the ratio $r_i = \lfloor \frac{t_m}{t_i} \rfloor$ [11], [13]. Here, $\lfloor . \rfloor$ denotes the floor operation.

For convenience, we introduce a notation to divide the m objectives into two groups based on the value of r_i. The i-th objective is called *cheap*, denoted as f^c, if $r_i > r_{\text{thres}}$; the objective is called *expensive*, denoted as f^e, if $r_i \leq r_{\text{thres}}$, where r_{thres} is a threshold separating the cheap objectives from expensive objectives. For real-world problems that have a natural separation between cheap and expensive objectives [19], the threshold can be defined straightforwardly. In case there is no intuitive separation between cheap and expensive objectives, the threshold can be defined according to the user's preference. It should be pointed out that the partitioning of the objectives has no direct influence on the effectiveness of the proposed algorithm, which does not make any assumptions on the ratios. Together with the idea of constructing surrogates for all objectives, this makes the proposed algorithm generic and applicable to a wide range of problems.

B. Multi-objective Bayesian Optimization

Multi-objective Bayesian optimization (MOBO), an extension of BO to MOPs, has been successfully applied to simultaneously optimizing expensive black-box MOPs [14], [20]. MOBO typically first trains a GP using data collected from previous evaluations to approximate each objective function of an MOP. A brief introduction to GPs is presented in Section I in the Supplementary material. Based on the way in which BO and evolutionary algorithms work together, MOBO can be further divided into two groups, evolutionary Bayesian optimization (EBO) and Bayesian evolutionary optimization (BEO) [21]. In BEO the evolutionary algorithm is the basic framework where the acquisition function is adopted as a criterion for model management, while in EBO Bayesian optimization is the basic framework in which the acquisition function is optimized using an evolutionary algorithm. In the following, we will briefly review the existing work on MOBO.

A straightforward way to address expensive MOPs using BO is to decompose an MOP into multiple single-objective problems, so that existing acquisition functions for single-objective optimization can be directly applied to MOPs. Some decomposition-based MOEAs (e.g., MOEA/D [22] and RVEA [23]) that decompose an MOP into a set of single-objective subproblems, have been extended to address expensive MOPs with GPs for function approximation and the expected improvement as the acquisition function [14], [24]. Instead, some quality indicators that are originally developed to assess the quality of the approximation of Pareto front solution sets have been employed as a scalar measure to reduce an MOP to a single-objective optimization problem. An early and popular performance indicator based algorithm is $\mathcal{S}$-Metric-Selection-based Efficient Global Optimization (SMS-EGO) [25], which is based on the $\mathcal{S}$ metric or hypervolume (HV) metric [26]. Moreover, there have been ample extensions of information-based acquisition functions for tackling expensive MOPs based on the information theory. For example, predictive entropy search [27] has been extended to MOPs by maximally reducing the entropy of the posterior distribution over the Pareto set, called PESMO; however, it is computationally expensive to approximate and optimize PESMO.

While various many-objective evolutionary algorithms have been proposed [28], [29], only a few attempts have been made to address expensive MaOPs. A representative method, a Kriging-assisted RVEA (K-RVEA), was proposed to handle MaOPs [14]. With the help of RVEA, K-RVEA decides to prioritize the diversity or convergence according to the number of the active adaptive reference vectors. Recently, Liu *et al.* [4] adopted GPs as surrogate models and proposed an amplified upper confidence bound to emphasize the uncertainty. To achieve a good balance of exploration and exploitation, one set of fixed and evenly distributed reference vectors and one set of adaptive reference vectors are used to perform convergence-related and diversity-related optimization, respectively. Due to the scalability issue of GPs, Guo *et al.* [30] proposed an efficient dropout neural network to replace the GPs when solving expensive many-objective problems.

In this work, the proposed acquisition function is based on the lower confidence bound (LCB) acquisition function [5] $\boldsymbol{AF}_{LCB}$ to evaluate the quality of a solution, owing to its computational efficiency. LCB tries to balance exploration and exploitation by maximising the uncertainty and minimising the mean value simultaneously. Considering an m-objective minimisation problem $\{f_i(\mathbf{x})\}_{i=1}^m$, the predictions of a candidate solution $\mathbf{x}$ is derived from the GP of each objective. The vectors of the predicted mean and variance of $\mathbf{x}$ are denoted as $\boldsymbol{\sigma}(\mathbf{x}) = \{\sigma_i(\mathbf{x})\}_{i=1}^m$ and $\boldsymbol{\mu}(\mathbf{x}) = \{\mu_i(\mathbf{x})\}_{i=1}^m$, respectively. The LCB for the candidate solution $\mathbf{x}$ is defined as

$$\boldsymbol{AF}_{LCB}(\mathbf{x}) = \boldsymbol{\mu} - k \cdot \boldsymbol{\sigma} \tag{2}$$

where k is a parameter to manage the trade-off between exploration and exploitation. Note that $\boldsymbol{AF}_{LCB}(\mathbf{x})$ is a vector with a length of m and minimised by an MOEA.

III. Proposed Algorithm

In this section, the proposed search bias penalized Bayesian optimization (SBP-BO) algorithm for solving HE-MOPs and HE-MaOPs is introduced. In the following, we will introduce the notations and the framework step by step, and then present the details of the two key components, i.e., the construction of GPs and the design of acquisition function.

A. Basic Ideas and the Overall Framework

The framework of SBP-BO is presented in Fig. 1, where the differences between conventional MOBO and the SBP-BO are highlighted. The main steps of the SBP-BO are described

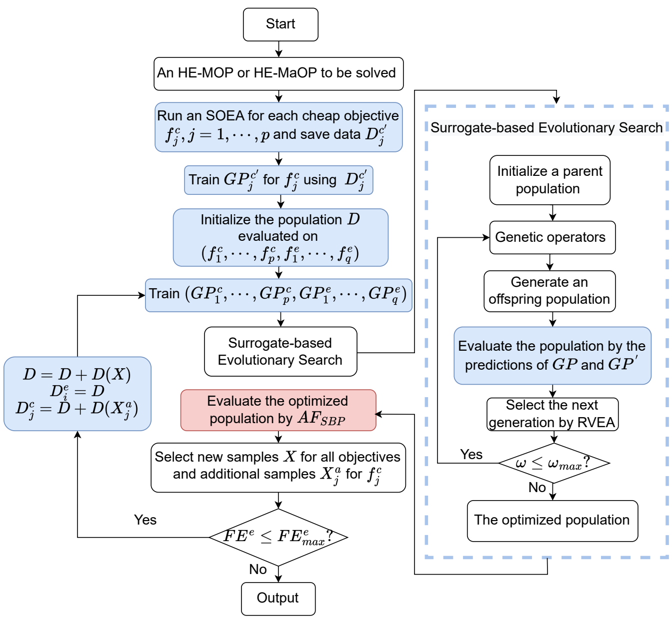

Fig. 1. The overall framework of SBP-BO.

below, and the pseudo code containing more details is given in Algorithm 1:

- Step 1: Initialization. An initial population $\mathbf{P}$ with n individuals is evaluated on m real objective functions using Latin Hypercube sampling that produces the data set $D = (\mathbf{P}, \mathbf{Y_P})$. During the evaluation of the expensive objectives, each cheap objective $f_j^c, j = 1, \cdots, p$ can be evaluated $n \cdot r_j$ times. This is done by using an SOEA to optimize each objective f_j^c, and the corresponding data is saved as $D_j^{c'}$.
- Step 2: Initial construction of GPs. SBP-BO constructs GPs $GP = [GP_1^c, \cdots, GP_p^c, GP_1^e, \cdots, GP_q^e]$ for each objective using the data D evaluated on all objectives. For each cheap objective $f_j^c, j = 1, \cdots, p$, an extra GP $GP_j^{c'}$ is trained on the additional data set $D_j^{c'}$, allowing us to construct an ensemble with two members GP_j^c and $GP_j^{c'}$.
- Step 3: Selection of new samples. Similar to the standard GP-assisted SAEA, SBP-BO adopts a baseline MOEA to optimize the HE-MOP/HE-MaOP for a certain number of generations, in which the ensemble surrogate predicts the value of each cheap objective $f_j^c, j = 1, \cdots, p$ for the candidate solutions, while GP_i^e predicts the value of the expensive objectives $f_i^e, i = 1, \cdots, q$. In our case, the RVEA [23] is adopted as the baseline MOEA, where a reference vector guided selection is introduced to select the next generation according to the angle-penalized distance. Subsequently, all individuals in the optimized population are evaluated by the proposed acquisition function that accounts for heterogeneous function evaluation times. Again, the reference vector guided selection is employed to identify a set of promising solutions, in which u new query points $\mathbf{X}$ and $u \cdot r_j - u$ additional query points $\mathbf{X}_j^a$ are randomly selected to be evaluated on all objectives and on cheap objectives f_j^c, respectively. Consequently, the newly evaluated solutions $\mathbf{X}$ are added to dataset D.
- Step 4: Update of GPs. SBP-BO follows a strategy used in [3], [31] to manage the training data: A predefined maximum number L of training data is set to $11d - 1 + 25$ [31], where d is the number of decision variables. When the number of data samples in D is less than L, e.g., in the beginning of the optimization, all solutions in D are used to train the GPs, i.e. $GP = [GP_1^c, \cdots, GP_p^c, GP_1^e, \cdots, GP_q^e]$. If D contains more samples than L, a subset D^t will be selected from the training data archive to limit the computation time, where the quality of the GPs and the optimization performance are considered. Since the solutions in D have been evaluated on all objectives, existing training data management methods for GP-assisted MOEAs, such as K-RVEA and MOEA/D-EGO, also can be used. While the GPs for f_i^e GP_i^e, $i = 1, \cdots, q$ are updated with the selected L training data samples D^t, each $GP_j^c, j = 1, \cdots, p$ is trained using both D^t and the extra new samples $\boldsymbol{X}_j^a$. Note that $GP_j^{c'}, j = 1, \cdots, p$ remain unchanged during the optimization, which are trained once with the offline data generated in Step 1 only, thus avoiding increasing the computational cost.
- Repeat Step 3 and Step 4 until the allowed computation budget is exhausted.

Algorithm 1 The framework of SBP-BO

Input: FE^e_{max}: the maximum number of the expensive objective function evaluations; $\boldsymbol{r}$: the ratio of the evaluation times between the expensive and cheap objectives; u: the number of new samples for updating the GPs; w_{max}: the maximum number of generations before updating GPs;
Output: Optimal solutions in D;

Initialization: Use the Latin Hybercube Sampling method to sample an initial population $\boldsymbol{P}$; $\boldsymbol{P}$ is evaluated on all objective functions, obtaining the objective values $\boldsymbol{Y_P}$; set $D = (\boldsymbol{P}, \boldsymbol{Y_P})$ and train GPs $GP = [GP_1^c, \cdots, GP_p^c, GP_1^e, \cdots, GP_q^e]$ for each objective using D that evaluated on all objectives; run an SOEA to optimize $f_j^c, j = 1, \cdots, p$ and save data in $D_j^{c'}$, then $GP_j^{c'}$ for each cheap objective f_j^c is trained on $D_j^{c'}$; set $FE = |D|$, $w = 1$ and $N_{\text{iter}} = 1$.
while $FE^e \leqslant FE^e_{max}$ **do**
  //Using the surrogate in the RVEA//
  Create the initial population;
  while $w \leqslant w_{max}$ **do**
    Generate offspring using the simulated binary crossover and polynomial mutation;
    Use the ensemble to predict cheap objective values and the $GP_1^e, \cdots, GP_q^e$ to predict the expensive objective values on the combined population;
    Use the reference vector guided selection to select the next generation;
    Perform the reference-vector-adaptation;
    $w = w + 1$;
  end while
  Use the proposed acquisition function to evaluate the optimized solutions found by RVEA;
  Use the reference vector guided selection to determine u new solutions $\boldsymbol{X}$ and $u \cdot r_j - u$ additional new solutions $\boldsymbol{X}_j^a, j = 1, \cdots, p$ to be evaluated on all objectives and on each $f_j^c, j = 1, \cdots, p$, respectively. Consequently, the newly solutions are saved in the corresponding datasets $D^{new} = (\boldsymbol{X}, \boldsymbol{Y})$ and $D_j^a = (\boldsymbol{X}_j^a, \boldsymbol{Y}_j^c)$, respectively;
  Add D^{new} to D and select training data D^t from data set D evaluated on all objectives;
  $GP = [GP_1^c, \cdots, GP_p^c, GP_1^e, \cdots, GP_q^e]$ is updated: $GP_i^e, i = 1, \cdots, q$ is updated with D^t, while $GP_j^c, j = 1, \cdots, p$ is updated with D^t and D_j^a.
  Update $FE^e = FE^e + u$, $N_{\text{iter}} = N_{\text{iter}} + 1$;
end while
Return the optimized solutions;

B. Ensemble Surrogate for Cheap Objectives

As described above, an ensemble surrogate including two GPs, GP_j^c and $GP_j^{c'}$, are constructed for each cheap objective $f_j^c, j = 1, \cdots, p$. The predicted mean value of f_j^c on a new sampled solution $\mathbf{x}$ provided by the ensemble is a weighted combination of the predictions of GP_j^c and $GP_j^{c'}$. Motivated by product of experts [32], confident predictions should have more influence on the combined prediction than the less confident ones. Hence, the weight is calculated based on the level of uncertainty of each GP's prediction, and the ensemble prediction is

$$
\begin{aligned}
\mu_j^c(\mathbf{x}) &= \alpha_j \, \mu_{GP_j^c}(\mathbf{x}) + \beta_j \, \mu_{GP_j^{c'}}(\mathbf{x}) \\
\alpha_j &= \frac{\sigma_{GP_j^{c'}}(\mathbf{x})}{\sigma_{GP_j^{c'}}(\mathbf{x}) + \sigma_{GP_j^c}(\mathbf{x})} \\
\beta_j &= \frac{\sigma_{GP_j^c}(\mathbf{x})}{\sigma_{GP_j^{c'}}(\mathbf{x}) + \sigma_{GP_j^c}(\mathbf{x})}
\end{aligned}
\tag{3}
$$

where $\mu_{GP_j^c}$ and $\sigma_{GP_j^c}$ are the predictions of GP_j^c, and $\mu_{GP_j^{c'}}$ and $\sigma_{GP_j^{c'}}$ are the predictions of $GP_j^{c'}$. As mentioned earlier, training two separate GPs on D and $D^{c'}$ makes the update of the GPs more efficient and computationally more effective. Specifically, while $GP = [GP_1^c, \cdots, GP_p^c, GP_1^e, \cdots, GP_q^e]$ are updated with newly sampled data selected according to the acquisition function, we do not re-train $GP^{c'} = [GP_1^{c'}, \cdots, GP_p^{c'}]$ during the optimization. Moreover, the method for selecting training samples used to retrain $GP = [GP_1^c, \cdots, GP_p^c, GP_1^e, \cdots, GP_q^e]$, as presented in Step 4, limits the maximum training time while allowing for building an effective model using relevant samples.

This training scheme makes it possible to adopt an existing strategy for selecting training data to update $GP = [GP_1^c, \cdots, GP_p^c, GP_1^e, \cdots, GP_q^e]$ in the context of HE-MOPs. Specifically, the number of training data is typically capped to limit the computational complexity of constructing the surrogates, which is a common practice in GP-based evolutionary algorithms [3], [14], [31]. For standard MOPs, it is desirable to select a subset so that the quality of the surrogates can be improved as much as possible, and the resulting surrogate-assisted search can maintain a balance between convergence and diversity. Note that each solution is evaluated on all objectives in standard MOPs, so that the balance between convergence and diversity can be estimated by selection criteria in MOEAs, such as nondominated sorting and the crowding distance. However, this is not the case for HE-MOPs since many solutions are partially evaluated. Consequently, it is difficult to select a subset that can balance convergence and diversity. To tackle this challenge, we train two GPs (GP_j^c and $GP_j^{c'}$) using solutions evaluated on all objectives and solutions evaluated on fast objectives, respectively. Hence, the existing strategies for selecting training data [3], [31] can be directly applied to HE-MOPs as we update GP_j^c only.

C. Search Bias Penalized Acquisition Function

In order to alleviate the search bias towards the cheap objectives resulting from the heterogeneous evaluation times, we propose to include a penalty term in the acquisition function for alleviating the search bias, which prioritizes the expensive objectives in minimizing the acquisition function. This penalty term is multiplied by the adaptive acquisition function reported in Eq. (2), resulting in a search bias penalized acquisition function ($\boldsymbol{AF}_{SBP}$). Having obtained the optimized population by RVEA using the GP ensemble, the mean and variance of the objective values of all individuals are predicted by the GPs at first. Given a candidate solution $\mathbf{x}$ in the optimized population, the proposed acquisition function can be computed analytically as follows:

$$
\boldsymbol{AF}_{SBP}(\mathbf{x}, N_{\text{iter}}) = \boldsymbol{AF}_{LCB}(\mathbf{x}) \circ \mathbf{SBP}(\mathbf{x}, N_{\text{iter}}) \tag{4}
$$

where N_{iter} is the current iteration number of the BO loop, $\mathbf{SBP}(\mathbf{x}, N_{\text{iter}})$ is the penalty term to be described below in detail, $\boldsymbol{AF}_{LCB}(\mathbf{x})$ denotes the acquisition function of Eq. (2), and $\circ$ denotes component-wise multiplication. As a result, each individual in the population can be evaluated according to $\boldsymbol{AF}_{SBP}$, obtaining a vector with length m. Hence, minimising $\boldsymbol{AF}_{SBP}$ is still an MOP, and therefore the reference vector guided selection in RVEA is adopted in this work.

Let $\boldsymbol{\mu}(\mathbf{x}) = \{\mu_i\}_{i=1}^m$ denote the predicted mean value of the objective vector with length m on the candidate solution $\mathbf{x}$ in the optimized population, and the maximum and minimum value of the predicted mean of the optimized population will be identified and denoted as $\boldsymbol{\mu}^{max} = \{\mu_i^{max}\}_{i=1}^m$ and $\boldsymbol{\mu}^{min} = \{\mu_i^{min}\}_{i=1}^m$, respectively. In order to calculate the penalty terms for $\mathbf{x}$, first we normalize the objective vector into the same range, i.e, [0,1]

$$
\bar{\boldsymbol{\mu}}(\mathbf{x}) = (\boldsymbol{\mu}(\mathbf{x}) - \boldsymbol{\mu}^{min})./(\boldsymbol{\mu}^{max} - \boldsymbol{\mu}^{min}), \tag{5}
$$

where again ./ indicates component-wise division. Hence, $\bar{\boldsymbol{\mu}}(\mathbf{x}) = \{\bar{\mu}_i\}_{i=1}^m$ is a vector with length m, and the corresponding penalty term $\mathbf{SBP} = \{SBP_i\}_{i=1}^m$ for the data sample $\mathbf{x}$ is also defined as a vector including the penalty term for each objective.

For the i-th objective, the penalty term SBP_i is calculated as

$$
SBP_i(\bar{\mu}_i, N_{\text{iter}}) = 1 - \pi\left(\bar{\mu}_i, N_{\text{iter}}\right) \tag{6}
$$

where $\pi\left(\bar{\mu}_i, N_{\text{iter}}\right)$ is calculated using an exponential distribution function,

$$
\pi\left(\bar{\mu}_i, N_{\text{iter}}\right) = \lambda\left(N_{\text{iter}}\right) e^{-\lambda(N_{\text{iter}})\bar{\mu}_i} \tag{7}
$$

with

$$
\lambda(N_{\text{iter}}) = \frac{1}{w_i N_{\text{iter}} + 1}, \tag{8}
$$

where $w_i = \frac{r_i}{\sum_i^m r_i}$ encodes the relative number of affordable evaluations of objective i. This penalty is intuitively motivated by the exponential distribution in modeling situations in which certain events occur at a constant probability per unit length [33]. Due to the fact that fast objectives can be explored more often than slow objectives in HE-MOPs, we construct different exponential distribution functions $\pi\left(\bar{\mu}_i, N_{\text{iter}}\right)$ for each objective function based on the evaluation times of different objectives to alleviate search biases. For example, the exploration on a fast objective is expected to occur at a lower probability than that on slow objectives. This is achieved by generating different $\lambda(N_{\text{iter}})$ values with respect to the evaluation times and the number of iterations. Hence, a large number of affordable function evaluations, i.e. a larger value of r_i, of an objective function f_i will result in a higher value of w_i, and accordingly a smaller value of λ. Since $0 \le \bar{\mu}_i \le 1$ and $0 \le \lambda \le 1$, this leads to more uniformly distributed and smaller values for π and therefore results in a larger penalty value on the corresponding fast objectives. Therefore, the acquisition function will prefer new samples that not only balance the local exploitation and global exploration, but also reduce the search bias by prioritizing the exploration for selecting slow functions. Similarly, as optimization progresses and N_{iter} increases, the penalty term will approach a value of 1 for all objectives, gradually reducing the disadvantage over the fast objectives in the acquisition function.

To take a closer look at the proposed SBP, we consider an example minimization HE-MOP with $\mathbf{r} = (5, 1)$ having a cheap and an expensive objective function (denoted as f^c and f^e, respectively). Contour plots of $h_1(\boldsymbol{Y}, N_{\text{iter}}) = \frac{SBP_c(\boldsymbol{Y}^c, N_{\text{iter}})}{SBP_e(\boldsymbol{Y}^e, N_{\text{iter}})}$ and $h_2(\boldsymbol{Y}, N_{\text{iter}}) = SBP_c(\boldsymbol{Y}^c, N_{\text{iter}}) + SBP_e(\boldsymbol{Y}^e, N_{\text{iter}})$ with respect to f^c, f^e and N_{iter} are given in Fig. 2. For $N_{\text{iter}} = 1$ shown in Figs. 2(a) and (b), the search bias penalty varies a lot in different regions of the objective space: (1) the penalty on f^e is always smaller than that on f^c; (2) there is a significant difference with respect to the penalty between the regions with smaller objective values and the regions with lager objective values. As the selection of new samples is a minimization MOP, this indicates that the SBP allows a smaller objective value to be preferred over a large one. Moreover, including such a penalty term into an acquisition function will guide the selection of new samples towards exploring f^e. As a result, the SBP will encourage the search towards f^e and mitigate the intrinsic search bias due to heterogeneous objectives. As the optimization proceeds, the difference between the cheap and expensive objectives will gradually shrink, as illustrated in Fig. 2(c) and 2(d), indicating a decreasing influence on the search bias. Finally, the penalty is almost equal across the whole objective space, as shown in Figs. 2(e) and 2(f). This allows SBP-BO to find out a set of satisfying solutions that cover the whole Pareto front.

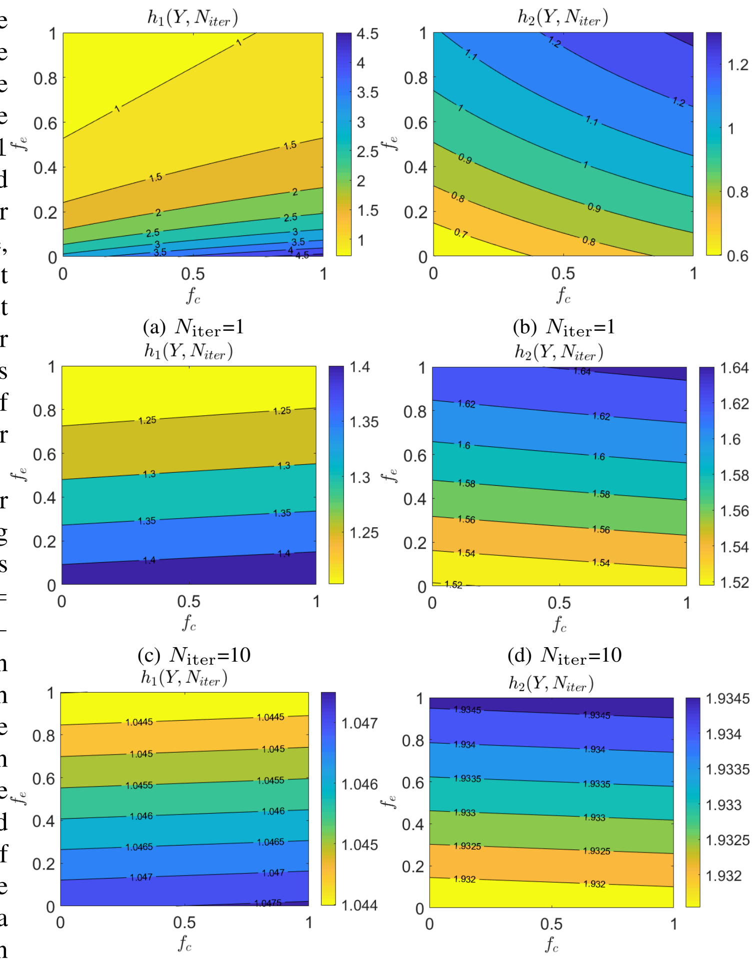

(a) N_{iter}=1 (b) N_{iter}=1

(c) N_{iter}=10 (d) N_{iter}=10

(e) N_{iter}=100 (f) N_{iter}=100

Fig. 2. Contour plots of $h_1(\boldsymbol{Y}, N_{\text{iter}}) = \frac{SBP_c(\boldsymbol{Y}^c, N_{\text{iter}})}{SBP_e(\boldsymbol{Y}^e, N_{\text{iter}})}$ ((a), (c) and (e)) and $h_2(\boldsymbol{Y}, N_{\text{iter}}) = SBP_c(\boldsymbol{Y}^c, N_{\text{iter}}) + SBP_e(\boldsymbol{Y}^e, N_{\text{iter}})$ ((b), (d) and (f)) with $\mathbf{r} = (5, 1)$.

IV. Experimental Studies

A. Experimental Settings

1) *Test Problems*: Although there are no test problems available that have inherently heterogeneous objectives, any existing multi-/many-objective benchmark problem can be adopted as HE-MOPs/MaOPs, assuming the evaluation times of the objectives are substantially different. Therefore, we have selected three widely used test suites of scalable multi-objective test problems, i.e., the DTLZ [34], UF [35] and WFG [36] test suites, and extend them to simulate HE-MOPs and HE-MaOPs. For all the test instances used in the experimental studies, the number of decision variables is set to 10. We set the last objective as the most expensive objective and vary the ratio r_i of the remaining objectives and the threshold r_{thres} to generate different HE-MOPs and HE-MaOPs.

2) *Performance Indicators*: The modified inverted generational distance (IGD) [37], the IGD$^+$ indicator [38], and the hypervolume (HV) [39] are adopted as the performance indicator to evaluate the quality of the obtained non-dominated solutions in terms of convergence and diversity. Let $Z = \{\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_{|Z|}\}$ be a given reference solution set, where $|Z|$ is the number of reference solutions, and $A = \{\boldsymbol{a}_1, \boldsymbol{a}_2, \ldots, \boldsymbol{a}_{|A|}\}$ be an obtained approximation to the Pareto front. The HV calculates the volume of the objective space dominated by an approximation set A, and the larger the HV value is, the better the quality of the approximation set. IGD^+ is calculated as follows:

$$IGD^+(A, Z) = \frac{1}{|Z|} \sum_{j=1}^{|Z|} \min_{\boldsymbol{a}_i \in A} d^+(\boldsymbol{a}_i, \boldsymbol{z}_j). \tag{9}$$

The distance d^+ is the distance between a reference solution $\boldsymbol{z} = (z_1, z_2, \cdots, z_m)$ and an objective vector $\boldsymbol{a} = (a_1, \cdots, a_m)$. Here, m is the number of objectives. d^+ is defined as:

$$d^+(\boldsymbol{a}, \boldsymbol{z}) = \sqrt{\sum_{k=1}^{m} (\max\{z_k - a_k, 0\})^2} \tag{10}$$

The smaller the IGD$^+$ value is, the better the quality of the non-dominated solution set.

Each algorithm under comparison is performed on each benchmark problem for 20 independent runs. The Wilcoxon rank sum test at a significance level of 0.05 is adopted to compare the results obtained by SBP-BO and other algorithms under comparison. To reduce the probability of making a type I error, the Holm-Bonferroni correction is adopted. The corresponding statistical results are presented in Tables I-V and Tables SI-SVIII in the Supplementary material, where symbols "(+)", "(–)", and "(≈)" indicate that the compared algorithm performs significantly better than, significantly worse than, or as well as the proposed algorithm, respectively. Note that we use notation 'S' to indicate tables and figures in the Supplementary material in order to avoid confusion.

3) *Algorithms Under Comparison*: To the best of authors knowledge, SBP-BO is the first algorithm designed for addressing HE-MOPs and HE-MaOPs. For comparison, one of state-of-the-art heterogeneity-handling methods, HK-RVEA [40], is slightly adapted to the proposed problem setting. Specifically, the SOEA in HK-RVEA optimizes each cheap objective using the different number of additional evaluations in the initialization. The genetic operators are used to generate additional samples for each cheap objectives while waiting for the expensive evaluation on new samples.As existing surrogate-assisted heterogeneity handling methods, e.g., T-SAEA [15], Tr-SAEA [16] and TC-SAEA [17], cannot address HE-MOPs or HE-MaOPs with more than two objectives, we compared them with SBP-BO on heterogeneous bi-objective optimization problems reported in [17]. Note that we did not include non-surrogate assisted algorithms for two reasons: 1) it has been shown that surrogate-assisted methods outperform non-surrogate assisted methods [14], [31], [41]; 2) most existing non-surrogate assisted methods work only for bi-objective problems with one fast and one slow objective. Since SBP-BO is based on GP assisted RVEA, a representative GP-assisted MOEA, K-RVEA [14], is also adopted as a surrogate assisted *Waiting* method to examine the performance of the proposed algorithm.

To further investigate the efficacy of the proposed surrogate ensemble and the search bias penalized acquisition function, we perform the following ablation studies:

4) *Parameter settings*: We use RVEA as the MOEA and a real-coded genetic algorithm that uses the simulated binary crossover and polynomial mutation as the SOEA. In addition, we use the DACE toolbox [42] to construct the GPs. All experiments are performed in MATLAB R2019a on an Intel Core i7-8750H with 2.21 GHz CPU. The parameter settings used in the experiments are summarized as follows:

- The initial population size for all the compared algorithms is $11d-1$ [31] where d is the number of decision variables.
- The maximum number of generations before updating the GPs (w_{max}) is set to 20 [14], [43].
- The number of new solutions selected for evaluations on all objectives at each BO iteration is set to $u = 3$ [16].
- The parameter in the LCB is set to $k = 2$ [44].
- The maximum number of function evaluations for the slow objectives (FE^e_{max}) is set to 200 for heterogeneous bi-objective problems, and 300 for HE-MOPs/MaOPs [41].

B. Experimental Results

1) *Comparison with state-of-the-art methods*: Each algorithm is performed on the test problems with $m = 3, 5, 10$ objectives, respectively. For simplicity, only one slow objective is included while all other objectives are considered as fast objectives with the same evaluation time, resulting in $\mathbf{r} = (r^c, r^c, \cdots, r^c, 1)$. Tables I-II and Tables SI-SII present the experimental results in terms of the IGD$^+$ and HV values obtained on each test instance with $r^c = 5$ and $r^c = 10$, respectively. The threshold r_{thres} is set to 1 for simplicity, and we will test different $\mathbf{r}$ and r_{thres} in the next subsection.

Firstly, the results presented in Table I show that the proposed SBP-BO significantly outperforms K-RVEA and HK-RVEA on 28 out of 48 test instances, respectively, indicating the effectiveness of the proposed strategies for handling HE-MOPs/HE-MaOPs. Secondly, it is interesting to see that both K-RVEA and HK-RVEA show significantly better performance than SBP-BO on DTLZ7 with $m = 10$. A possible explanation for this might be that in K-RVEA and HK-RVEA a fixed reference set is used to evaluate whether diversity or convergence should be prioritized in the selection of new samples. In this way, the exploration can be guaranteed, so that K-RVEA and HK-RVEA are able to account for the disconnected Pareto front of DTLZ7. According to the results on WFG1-WFG9 summarized in Table I, SBP-BO achieves the best performance in terms of IGD$^+$ metric on 21 test instances, followed by K-RVEA with 4 best results. Note that K-RVEA show better performance than the proposed algorithm on WFG1 with $m = 10$ and WFG9 with $m = 10$. Recall that both WFG1 and WFG9 are difficult for optimization algorithms to achieve a good diversity. WFG1 is designed by using the most complex transformation function to add complexity to a underlying problem, making it hard an optimization algorithm to converge to the true Pareto front. Similarly, WFG9 features a troublesome transformation function and is also a multi-model and nonseparable problem. Lastly, by comparing K-RVEA using the *Waiting* method with HK-RVEA and SBP-BO we can confirm the effectiveness of using a GP ensemble and the proposed acquisition function. While HK-RVEA generally shows similar performance with K-RVEA, SBP-BO exhibits better performance than K-RVEA on most test problems. Similarly, the results in terms of the HV values in Table II show that Tr-SAEA significantly outperforms K-RVEA and HK-RVEA on 24 and 29 test instances, respectively. Interestingly, the HV values on DTLZ1, DTLZ3 and DTLZ6 indicate that all the algorithms fail to satisfactorily approximate the Pareto front due to the ruggedness and complexity of the fitness landscape.

The results presented in Table SI are consistent with those in Table I, further supporting the benefit of using a GP ensemble to make use of the additional data on the fast objectives and the search bias penalized acquisition function. It is noteworthy that with r^c increasing from 5 to 10, the proposed algorithm maintains its advantage for solving HE-MOPs and HE-MaOPs by properly making use of the information obtained from the optimization of the fast objectives. This observation can be validated by the HV values in Table SII, where SBP-BO significantly outperforms K-RVEA on 29 test instances. Note, however, that the proposed algorithm may fail to maintain a good diversity of the obtained solutions on some test problems, such as DTLZ5 and WFG3 which have degenerate Pareto fronts. Hence, although the proposed SBP acquisition function promotes the search towards the slow objectives, it may lead to a poor diversity of the solutions on some problems.

To gain a deeper insight into the quality of the final solution sets obtained by K-RVEA, HK-RVEA and SBP-BO, Figs. S1-S3 in the Supplementary material show the nondominated solution set with the median IGD$^+$ value among 20 runs obtained by K-RVEA, HK-RVEA and SBP-BO on DTLZ2, DTLZ4, WFG2 and WFG6 with $m = 3, 5, 10$, respectively. Take Fig. SI as an example, the proposed SBP-BO shows promising performance in terms of both convergence and diversity on the selected three-objective test problems, compared with K-RVEA and HK-RVEA. For example, it is clear that SBP-BO finds a better approximation of the true Pareto front on DTLZ2 when compared with the other algorithms, indicating its good balance between diversity and convergence. It is worth noting that while HK-RVEA covers a small part of the true Pareto front, SBP-BO is able to achieve a set of well distributed solutions. This observation further supports the advantage of the proposed SBP for reducing the search bias towards to the cheap objectives. Regarding HE-MaOPs, similar observations can be made, as illustrated in Figs. S2-S3, where the solutions on the estimated Pareto front are shown in the parallel coordinate plots. These observations can be explained from the perspective of the use of additional data and the reduction of search bias by the SBP acquisition function.

To explore the performance of SBP-BO as the evolution proceeds, the IGD$^+$ values obtained by each algorithm over the number of real fitness evaluations (FEs) on test problems with $r^c = \{5, 10\}$ and the corresponding statistically significant differences are summarised in Tables SIII-SIV, respectively, in the Supplementary material. As can be seen from the tables, the proposed SBP-BO shows significantly better performance than K-RVEA on DTLZ2, WFG3, WFG5 and WFG6 with $FE^e = 150$, indicating the fast convergence of SBP-BO. Although HK-RVEA shows better performance than K-RVEA, SBP-BO can significantly outperform HK-RVEA on DTLZ2, WFG4, WFG6 and WFG7 with $FE^e = 200$. Subsequently, Figs. S4-S5 plot the boxplots of the IGD$^+$ values obtained by each algorithm on DTLZ2, DTLZ4, WFG2 and WFG6 with different number of FEs over 20 runs, confirming the fast convergence achieved by SBO-BO. Moreover, we demonstrate the search process of each algorithm in terms of IGD$^+$ values in Figs. S6-S7, where the error bars indicate the variance of IGD$^+$ values over 20 runs. According to Figs. S5-S6, similar conclusion can be made.

2) *Influence of different* $\mathbf{r}$ *and* r_{thres} *on the optimization performance*: Each objective in an HE-MOP/HE-MaOP requires a distinct period of time to be evaluated, it is therefore expected to test the proposed algorithm on problems with different $\mathbf{r} = (r_1^c, r_2^c, \cdots, r_p^c, r_1^e, \cdots, r_q^e)$, where p and q are the number of cheap and expensive objectives. In this subsection, the heterogeneity handling ability of SBP-BO is

TABLE I

MEAN (STANDARD DEVIATION) IGD$^+$ VALUES OBTAINED BY K-RVEA, HK-RVEA, AND SBP-BO WITH $FE^e_{max} = 300$ AND $r = (r^c, \ldots, r^c, 1)$ WHERE $r^c = 5$

Problem	m	K-RVEA	HK-RVEA	SBP-BO
DTLZ1	3	7.88e+1 (1.16e+1) ≈	9.07e+1 (1.42e+1) –	6.60e+1 (1.10e+1)
	5	3.96e+1 (1.21e+1) –	4.41e+1 (3.50e+0) –	2.45e+1 (8.88e+0)
	10	3.31e-1 (1.42e-1) –	1.97e-1 (8.40e-2) –	1.58e-1 (2.31e-2)
DTLZ2	3	7.52e-2 (1.02e-2) –	5.98e-2 (1.05e-2) ≈	3.75e-2 (1.27e-3)
	5	1.79e-1 (1.88e-2) –	1.38e-1 (9.34e-3) –	9.76e-2 (6.07e-3)
	10	2.38e-1 (1.32e-2) ≈	2.73e-1 (4.02e-2) –	2.04e-1 (9.70e-3)
DTLZ3	3	2.13e+2 (3.73e+1) ≈	2.35e+2 (3.55e+1) ≈	1.82e+2 (4.57e+1)
	5	1.25e+2 (3.49e+1) –	1.64e+2 (3.70e+1) –	9.17e+1 (2.53e+1)
	10	7.96e-1 (3.52e-1) ≈	8.05e-1 (3.58e-1) ≈	6.12e-1 (1.48e-1)
DTLZ4	3	2.59e-1 (7.57e-2) –	3.41e-1 (1.30e-1) –	1.71e-1 (1.29e-1)
	5	2.73e-1 (5.59e-2) ≈	3.26e-1 (6.29e-2) –	2.89e-1 (9.04e-2)
	10	2.58e-1 (1.77e-2) ≈	2.68e-1 (3.08e-2) ≈	2.57e-1 (2.60e-2)
DTLZ5	3	6.76e-2 (1.15e-2) –	6.52e-2 (9.38e-3) –	2.78e-2 (2.87e-3)
	5	2.90e-2 (6.77e-3) –	1.90e-2 (3.57e-3) ≈	1.88e-2 (3.55e-3)
	10	6.22e-3 (7.70e-4) ≈	7.31e-3 (1.43e-3) ≈	7.22e-3 (9.88e-4)
DTLZ6	3	3.03e+0 (6.09e-1) ≈	3.15e+0 (4.09e-1) ≈	2.85e+0 (4.83e-1)
	5	1.77e+0 (3.18e-1) ≈	1.90e+0 (3.09e-1) ≈	1.03e+0 (4.89e-1)
	10	3.85e-2 (7.26e-3) –	2.69e-2 (8.73e-3) –	2.47e-2 (7.67e-3)
DTLZ7	3	1.09e-1 (2.63e-2) –	6.64e-2 (1.04e-2) –	5.85e-2 (1.89e-2)
	5	4.79e-1 (2.94e-1) –	3.12e-1 (7.50e-2) ≈	3.12e-1 (3.48e-1)
	10	9.08e-1 (3.88e-2) +	9.36e-1 (2.65e-2) +	9.54e-1 (1.70e-1)
WFG1	3	1.64e+0 (4.10e-2) ≈	1.74e+0 (1.15e-1) ≈	1.66e+0 (1.34e-1)
	5	2.11e+0 (7.70e-2) +	2.22e+0 (8.27e-2) ≈	2.23e+0 (6.82e-2)
	10	2.81e+0 (1.32e-1) ≈	2.81e+0 (1.26e-1) ≈	2.79e+0 (1.46e-1)
WFG2	3	2.91e-1 (2.74e-2) –	2.20e-1 (2.22e-2) –	1.46e-1 (2.81e-2)
	5	3.93e-1 (6.16e-2) –	2.82e-1 (3.07e-2) –	1.90e-1 (2.89e-2)
	10	3.96e-1 (1.38e-1) –	3.98e-1 (1.45e-1) –	3.11e-1 (8.55e-2)
WFG3	3	4.12e-1 (5.12e-2) –	4.53e-1 (5.97e-2) –	2.10e-1 (3.02e-2)
	5	4.29e-1 (8.76e-2) –	3.43e-1 (3.57e-2) ≈	3.48e-1 (8.35e-2)
	10	5.59e-1 (6.59e-2) –	5.35e-1 (7.26e-2) ≈	5.36e-1 (7.64e-2)
WFG4	3	3.92e-1 (2.82e-2) –	3.84e-1 (2.41e-2) –	3.22e-1 (2.28e-2)
	5	7.73e-1 (4.30e-2) –	8.55e-1 (6.66e-2) –	6.86e-1 (3.89e-2)
	10	3.25e+0 (8.96e-1) –	3.45e+0 (8.52e-1) –	2.21e+0 (3.82e-2)
WFG5	3	3.78e-1 (6.53e-2) –	2.59e-1 (1.69e-2) ≈	2.29e-1 (5.32e-2)
	5	7.96e-1 (6.80e-2) –	7.12e-1 (3.17e-2) –	6.19e-1 (4.10e-2)
	10	2.12e+0 (4.55e-1) –	1.97e+0 (4.58e-1) ≈	1.78e+0 (3.69e-1)
WFG6	3	6.77e-1 (5.71e-2) –	5.04e-1 (7.46e-2) –	3.16e-1 (8.37e-2)
	5	1.18e+0 (1.41e-1) –	9.28e-1 (8.25e-2) –	7.41e-1 (4.45e-2)
	10	1.29e+0 (3.37e-2) ≈	1.28e+0 (5.73e-2) ≈	1.03e+0 (2.96e-2)
WFG7	3	4.98e-1 (4.06e-2) –	5.44e-1 (3.21e-2) –	4.25e-1 (2.11e-2)
	5	8.74e-1 (6.09e-2) ≈	1.03e+0 (8.34e-2) –	8.28e-1 (5.22e-2)
	10	3.49e+0 (4.68e-1) ≈	3.63e+0 (5.24e-1) –	3.33e+0 (4.32e-1)
WFG8	3	6.57e-1 (5.50e-2) –	5.64e-1 (2.90e-2) –	4.73e-1 (5.86e-2)
	5	1.49e+0 (4.32e-2) –	1.40e+0 (6.33e-2) –	1.20e+0 (2.85e-2)
	10	1.47e+0 (2.47e-1) ≈	3.57e+0 (1.10e+0) –	1.58e+0 (6.99e-1)
WFG9	3	5.76e-1 (7.03e-2) ≈	5.72e-1 (1.15e-1) ≈	5.91e-1 (1.36e-1)
	5	1.13e+0 (2.54e-1) –	1.31e+0 (1.95e-1) –	8.74e-1 (2.54e-1)
	10	3.71e+0 (8.56e-1) +	5.16e+0 (6.40e-1) ≈	4.69e+0 (8.73e-1)
+/–/≈		3/28/17	1/28/19	

tested on three-objective and ten-objective problems with different numbers of expensive objectives (q), different ratios of function evaluation times between objectives ($\mathbf{r}$), and different threshold values (r_{thres}). Firstly, SBP-BO is examined on three-objective test functions with two sets of $\mathbf{r}$ values, i.e., one set having $\mathbf{r} = (5, 5, 1)$, $\mathbf{r} = (7, 3, 1)$, and $\mathbf{r} = (9, 1, 1)$, and the other having $\mathbf{r} = (15, 5, 1)$, $\mathbf{r} = (10, 10, 1)$, and $\mathbf{r} = (19, 1, 1)$. The statistical results in terms of IGD$^+$ are presented in Tables SV-SVI. Secondly, ten-objective problems with $\mathbf{r} = (10, 9, 8, 7, 6, 5, 4, 3, 2, 1)$ and $r_{\text{thres}} = 1, 3, 5$ are used to test the impact of r_{thres} on the optimization performance, and the statistical results are presented in Table III. Lastly, the impact of different $\mathbf{r}$ is further investigated on ten-objective problems with $r_{\text{thres}} = 3$, where $\mathbf{r}$ is set to $\mathbf{r1} = (10, 8, 8, 7, 5, 4, 3, 2, 2, 1)$, $\mathbf{r2} = (10, 9, 8, 6, 3, 2, 2, 2, 1, 1)$,

TABLE II
MEAN (STANDARD DEVIATION) HV VALUES OBTAINED BY K-RVEA, HK-RVEA, AND SBP-BO WITH $FE^e_{max} = 300$ AND $r = (r^c, \ldots, r^c, 1)$ WHERE $r^c = 5$

Problem	m	K-RVEA	HK-RVEA	SBP-BO
DTLZ1	3	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0)
	5	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0)
	10	3.38e-1 (2.90e-1) –	5.57e-1 (5.42e-2) –	6.32e-1 (9.90e-2)
DTLZ2	3	4.52e-1 (2.80e-2) –	5.01e-1 (4.75e-2) –	5.32e-1 (6.36e-3)
	5	6.28e-1 (1.71e-2) –	6.46e-1 (8.58e-2) –	7.68e-1 (8.05e-4)
	10	8.71e-1 (1.28e-2) –	8.24e-1 (8.95e-2) ≈	9.62e-1 (3.28e-3)
DTLZ3	3	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0)
	5	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0)
	10	1.12e-1 (3.28e-3) –	8.08e-2 (0.00e+0) –	3.06e-1 (8.28e-2)
DTLZ4	3	5.93e-2 (6.35e-2) –	1.25e-2 (4.94e-2) –	2.44e-1 (1.06e-1)
	5	3.64e-1 (1.43e-1) +	3.33e-2 (3.43e-2) –	2.44e-1 (3.10e-2)
	10	8.50e-1 (3.08e-2) ≈	7.58e-1 (3.05e-2) –	8.43e-1 (1.64e-2)
DTLZ5	3	1.31e-1 (8.71e-3)≈	1.27e-2 (3.17e-2) ≈	1.46e-1 (1.01e-2)
	5	1.11e-1 (3.43e-3) ≈	1.14e-1 (2.93e-2) ≈	1.10e-1 (1.14e-3)
	10	9.76e-2 (5.36e-3) ≈	9.69e-2 (1.71e-2) ≈	9.81e-2 (6.75e-4)
DTLZ6	3	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0)
	5	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0)
	10	9.00e-2 (4.94e-4) ≈	9.48e-2 (4.94e-4) ≈	9.46e-2 (7.81e-4)
DTLZ7	3	2.49e-1 (3.85e-3) –	2.54e-1 (3.09e-3) –	2.63e-1 (1.70e-3)
	5	2.17e-1 (7.59e-3) ≈	2.21e-1 (9.13e-4) ≈	2.22e-1 (9.71e-5)
	10	2.49e-1 (3.85e-3) +	1.76e-1 (3.60e-3) ≈	1.70e-1 (3.70e-3)
WFG1	3	2.11e-1 (3.16e-2) –	2.12e-1 (1.89e-2) –	2.38e-1 (2.14e-2)
	5	2.24e-1 (3.88e-2) –	1.48e-1 (4.96e-2) –	2.57e-1 (3.88e-2)
	10	2.11e-1 (3.16e-2) –	2.12e-1 (1.89e-2) –	2.38e-1 (2.14e-2)
WFG2	3	7.76e-1 (2.12e-2) ≈	7.30e-1 (4.60e-2) –	7.92e-1 (4.14e-2)
	5	7.05e-1 (5.45e-2) –	7.77e-1 (3.30e-2) –	8.70e-1 (3.47e-2)
	10	6.14e-1 (1.13e-2) –	6.28e-1 (3.49e-2) –	6.96e-1 (4.79e-2)
WFG3	3	2.22e-1 (1.63e-2) –	1.89e-1 (1.52e-2) –	3.72e-1 (1.85e-2)
	5	6.54e-3 (6.66e-3) –	6.07e-3 (7.73e-3) –	7.47e-3 (8.47e-3)
	10	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0) ≈	0.00e+0 (0.00e+0)
WFG4	3	3.78e-1 (1.51e-2) ≈	3.57e-1 (1.39e-2) –	3.95e-1 (1.26e-2)
	5	5.35e-1 (4.21e-2) –	4.75e-1 (1.92e-2) –	5.93e-1 (2.35e-2)
	10	4.80e-1 (2.08e-2) ≈	4.79e-1 (2.98e-2) ≈	4.82e-1 (4.82e-2)
WFG5	3	3.97e-1 (3.59e-2) ≈	3.69e-1 (2.21e-2) ≈	3.95e-1 (4.01e-2)
	5	4.50e-1 (2.89e-2) –	4.59e-1 (2.28e-2) ≈	4.76e-1 (2.47e-2)
	10	4.15e-1 (1.15e-2) ≈	4.10e-1 (2.16e-2) ≈	4.28e-1 (3.60e-2)
WFG6	3	2.26e-1 (1.68e-2) –	2.15e-1 (9.52e-3) –	2.68e-1 (2.97e-2)
	5	3.06e-1 (5.71e-2) –	3.29e-1 (2.31e-2) –	3.81e-1 (3.41e-2)
	10	2.54e-1 (2.13e-2) –	3.62e-1 (1.86e-2) –	3.99e-1 (3.04e-2)
WFG7	3	3.07e-1 (1.64e-2) ≈	2.85e-1 (9.16e-3) –	3.17e-1 (2.19e-2)
	5	4.38e-1 (3.17e-2) –	4.23e-1 (2.15e-2) –	4.59e-1 (2.30e-2)
	10	4.76e-1 (4.56e-2) ≈	4.32e-1 (2.29e-2) –	4.98e-1 (3.08e-2)
WFG8	3	2.81e-1 (1.54e-2) –	2.63e-1 (1.52e-2) –	3.18e-1 (1.80e-2)
	5	3.28e-1 (1.31e-2) –	3.49e-1 (2.12e-2) –	3.78e-1 (1.82e-2)
	10	4.24e-1 (1.39e-2) –	3.80e-1 (1.39e-2) –	4.76e-1 (2.90e-2)
WFG9	3	2.68e-1 (3.84e-2) +	2.34e-1 (3.88e-2) ≈	2.48e-1 (1.95e-2)
	5	3.43e-1 (3.12e-2) –	3.24e-1 (3.12e-2) –	3.93e-1 (2.76e-2)
	10	3.72e-1 (4.01e-2) ≈	3.21e-1 (1.04e-2) –	3.63e-1 (2.15e-2)
+/–/≈		3/24/21	0/29/19	

and $\mathbf{r3} = (9, 7, 3, 3, 3, 2, 2, 2, 1, 1)$, respectively. The results are summarized in Table IV. Accordingly, the results in terms of the HV values are presented in Tables SVII-SVIII.

Although it is unclear what influence the exact form of $\mathbf{r}$ and the exact value of r_{thres} will have on the optimization process, we can make the following observations. First, the instance with $\mathbf{r} = (9, 1, 1)$ will cause a strong search bias towards the objective whose $r_1^c = 9$ compared with $\mathbf{r} = (5, 5, 1)$ or $\mathbf{r} = (7, 3, 1)$, making the problem more difficult to solve. Consequently, the instance with $\mathbf{r} = (19, 1, 1)$ will render the strongest search bias among all $\mathbf{r}$ situations for three-objective problems studied in this work. Secondly, regarding the ten-objective test instances, a larger r_{thres} will result in a stronger search bias towards the cheap objectives. Consequently, problems with $r_{\text{thres}} = 5$ become more challenging to solve than those with $r_{\text{thres}} = 1$. Moreover, we have a similar expectation that an optimization algorithm can achieve better performance on ten-objective problems with a smaller q (i.e., $\mathbf{r1}$) than those with a larger one (i.e., $\mathbf{r3}$).

The IGD$^+$ values in Tables SV-SVI achieved by HK-RVEA and the proposed algorithm accord with our earlier observations that compared with HK-RVEA, SBP-BO shows similar or significantly better performance on all test problems except DTLZ7, with different $\mathbf{r}$. This observation provides evidence for confirming SBP-BO's ability for solving various HE-MOPs. Besides, it is interesting to note that changing $\mathbf{r}$ from $(5, 5, 1)$ to $(9, 1, 1)$ generally causes a reduced performance for both HK-RVEA and SBP-BO on most test problems, which is expected. Similar conclusions can be drawn from Table IV, confirming the effectiveness of SBP-BO in handling HE-MOPs/HE-MaOPs. According to the results in Table III, SBP-BO significantly outperforms HK-RVEA on seven test instances, while HK-RVEA always shows the best performance on DTLZ1. We can also observe that the performance of SBP-BO and HK-RVEA degrades as $\mathbf{r}_{\text{thres}}$ increases. These observations can be confirmed by the HV values in Tables SVII-SVIII.

To further illustrate the influence of different $\mathbf{r}$ on the performance of the optimization algorithms considered in this work, the approximations of the true Pareto front on DTLZ5 with $\mathbf{r} = (5, 5, 1)$ and $\mathbf{r} = (9, 1, 1)$ obtained by HK-RVEA and SBP-BO are shown in Figs. 3-4, where FE^e_{max} is set to 300 and 1000, respectively. From these results, our observations can be summarized as follows:

- Consistent with the aforementioned hypothesis, the search bias is more likely to occur on problems when $\mathbf{r} = (9, 1, 1)$ in the different function evaluation ratios considered between objectives, rendering an MOEA inefficient. An illustrative example is given in Fig. 3 when $FE^e_{max} = 300$. SBP-BO can achieve better diversity on DTLZ5 with $\mathbf{r} = (5, 5, 1)$ compared with that on DTLZ5 with $\mathbf{r} = (9, 1, 1)$. Similar observations can be made from Fig. 4, where $FE^e_{max} = 1000$.
- The proposed SBP-BO can find a set of solutions with good quality on all considered test instances, while HK-RVEA suffers from the heterogeneous objectives. As depicted in Fig. 3, the solution set obtained by HK-RVEA only covers some subregions, especially those around the two end points of the true Pareto front of DTLZ5, which is a curve for that problem. We note that the proposed SBP-BO shows its advantage for handling HE-MOPs, which is achieved by an efficient use of the additional data on the cheap objectives with the help of the ensemble GP surrogate and the alleviation search bias by means of the proposed acquisition function. It is clear that SBP-BO finds a satisfying Pareto front approximation with respect to both diversity and convergence when $FE^e_{max} = 1000$, confirming that the proposed SBP acquisition function can reduce the search bias introduced by heterogeneous objectives.

3) *Results on bi-objective heterogeneous problems*: In this subsection, we compare SBP-BO with three transfer-learning (TL) based heterogeneity-handling methods, i.e., T-SAEA

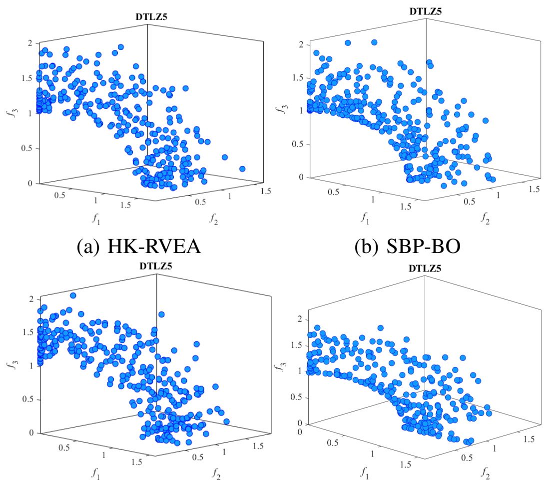

(a) HK-RVEA (b) SBP-BO

(c) HK-RVEA (d) SBP-BO

Fig. 3. The final solution set with the median IGD$^+$ values found by HK-RVEA and SBP-BO on DTLZ5 with $FE^s_{max} = 300$ and $\mathbf{r} = (5, 5, 1)$ ((a) and (b)) and $\mathbf{r} = (9, 1, 1)$ ((c) and (d))
, respectively.

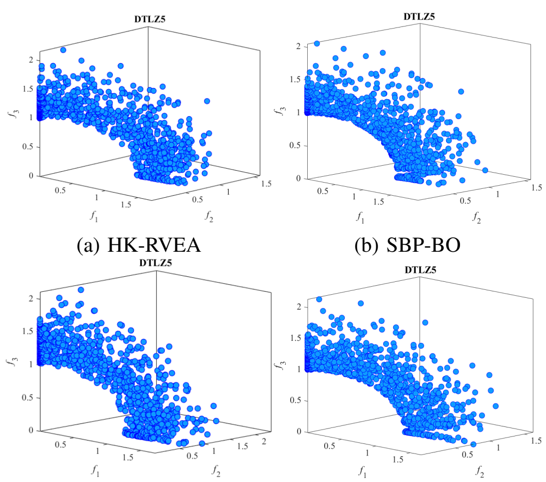

(a) HK-RVEA (b) SBP-BO

(c) HK-RVEA (d) SBP-BO

Fig. 4. The final solution set with the median IGD$^+$ values found by HK-RVEA and SBP-BO on DTLZ5 with $FE^s_{max} = 1000$ and $\mathbf{r} = (5, 5, 1)$ ((a) and (b)) and $\mathbf{r} = (9, 1, 1)$ ((c) and (d)), respectively.

TABLE III
MEAN (STANDARD DEVIATION) IGD$^+$ VALUES OBTAINED BY HK-RVEA, AND SBP-BO ON TEN-OBJECTIVE PROBLEMS WITH $r = (10, 9, 8, 7, 6, 5, 4, 3, 2, 1)$ AND DIFFERENT r_{thres} AND $FE^e_{max} = 300$

Problem	r_{thres}	HK-RVEA	SBP-BO
DTLZ1	1	2.37e-1 (7.60e-2) +	5.54e-1 (8.84e-2)
	3	3.17e-1 (1.61e-1) +	6.15e-1 (2.12e-1)
	5	3.26e-1 (1.42e-1) +	6.58e-1 (1.39e-1)
DTLZ4	1	2.52e-1 (1.73e-2) ≈	2.51e-1 (1.09e-2)
	3	2.64e-1 (1.65e-2) –	2.59e-1 (1.78e-2)
	5	2.71e-1 (3.94e-2) ≈	2.71e-1 (2.25e-2)
WFG2	1	3.45e-1 (2.26e-1) –	1.56e-1 (2.81e-2)
	3	6.18e-1 (1.96e-1) –	2.01e-1 (9.94e-2)
	5	1.53e+0 (9.21e-1) –	2.15e-1 (1.58e-1)
WFG6	1	1.21e+0 (5.28e-2) –	1.94e-1 (4.23e-2)
	3	1.49e+0 (1.89e-1) –	2.28e-1 (2.06e+0)
	5	2.77e+0 (1.66e+0) –	2.71e-1 (9.56e-1)
+/–/ ≈		3/7/2	

TABLE IV
MEAN (STANDARD DEVIATION) IGD$^+$ VALUES OBTAINED BY HK-RVEA, AND SBP-BO ON TEN-OBJECTIVE PROBLEMS WITH $r_{thres} = 3$, $\mathbf{r1} = (10, 8, 8, 7, 5, 4, 3, 2, 2, 1)$, $\mathbf{r2} = (10, 9, 8, 6, 3, 2, 2, 2, 1, 1)$ AND $\mathbf{r3} = (9, 7, 3, 3, 3, 2, 2, 2, 1, 1)$, AND $FE^e_{max} = 300$

Problem	r	HK-RVEA	SBP-BO
DTLZ1	**r1**	3.56e-1 (1.33e-1) –	2.73e-1 (1.00e-1)
	r2	2.66e-1 (1.01e-1) ≈	2.51e-1 (1.16e-1)
	r3	3.29e-1 (1.65e-1) ≈	2.91e-1 (9.60e-2)
DTLZ4	**r1**	2.79e-1 (1.70e-2) –	2.37e-1 (2.15e-2)
	r2	2.67e-1 (2.70e-2) ≈	2.57e-1 (2.36e-2)
	r3	2.80e-1 (4.18e-2) ≈	2.71e-1 (2.70e-2)
WFG2	**r1**	1.41e+0 (7.11e-1) –	8.35e-1 (6.09e-1)
	r2	7.30e-1 (4.76e-1) +	8.16e-1 (5.26e-1)
	r3	1.48e+0 (5.47e-1) –	8.23e-1 (3.60e-1)
WGF6	**r1**	3.44e+0 (1.66e+0) –	2.66e-1 (1.45e-1)
	r2	2.71e+0 (9.07e-1) –	3.18e-1 (3.76e-2)
	r3	4.01e+0 (9.13e-1) –	2.75e+0 (2.56e-2)
+/–/ ≈		1/7/4	

[15], Tr-SAEA [16] and TC-SAEA [17], on the same bi-objective heterogeneous problems with the same parameter setting reported in [17]. Specifically, as presented in the Supplementary material, DTLZ1 to DTLZ7 and two modified counterparts (DTLZ1a and DTLZ3a) of DTLZ1 and DTLZ3, and UF1 to UF7 from the UF test suite [35], are used as test instances. The statistical results in terms of the IGD$^+$ and HV values obtained by each algorithm on test instances with $r^c = 5$ and $r^c = 10$ are summarized in Table V and Tables SIX-SXI. According to the IGD$^+$ results, at least one of the TL-based methods significantly outperforms SBP-BO on 10 out of 16 problems for $r^c = 5, 10$, while SBP-BO is only significantly better than all TL approaches on three problems for $r^c = 5$ and two problems for $r^c = 10$. Similar observations can be made from the HV values, while all algorithms fail to approximate the true Pareto front on the hard-to-convergence instances, i.e., DTLZ1 and DTLZ3. The performance difference is understandable since in the TL-based approaches information on the correlation of the two objectives is acquired, which allows for an estimation of the expensive objective from the search experience on the cheap objective. Such information is not used in SBP-BO and the selection of new samples is only guided by the heterogeneous evaluation times in SBP-BO.

However, the extension of the TL-based approaches to problems with more than two objectives is nontrivial. Multiple models for transferring knowledge between each pair of objectives will need to be trained, which increases the computational complexity substantially. More importantly, it is an open question how to utilize the information from these models in a consistent manner, as, for example, information on one objective will be provided by multiple models that might contain contradicting information.

4) *Ablation studies* Further experiments are performed here to provide a deeper understanding of the performance of SBP-BO by testing the effectiveness of each component. The IGD$^+$ values obtained by SBP-BO and its variants are presented in Tables SXII and SXIII in the Supplementary material. The following observations can be made:

- According to Tables SXII and SXIII, we can see that the proposed algorithm yields the best IGD$^+$ values on 23 out

TABLE V
MEAN (STANDARD DEVIATION) IGD^{+} VALUES OBTAINED BY T-SAEA, TR-SAEA, TC-SAEA AND SBP-BO FOR BI-OBJECTIVE PROBLEMS WITH $FE^{e}_{max} = 200$ AND $r = (r^{c}, 1)$ WHERE $r^{c} = 5$

Problem	T-SAEA	Tr-SAEA	TC-SAEA	SBP-BO
DTLZ1	21.7 (11.9) –	20.7 (5.38) –	20.1 (8.16) –	17.2 (4.52)
DTLZ1a	1.06 (1.00) –	0.21 (0.07) –	0.36 (0.04) –	0.15 (0.03)
DTLZ2	0.05 (0.03) ≈	0.03 (0.01) ≈	0.02 (0.00) ≈	0.04 (0.01)
DTLZ3	203 (100) ≈	327 (82.1) ≈	132 (79.3) ≈	214 (62.3)
DTLZ3a	5.34 (37.5) +	3.39 (1.87) +	2.30 (0.66) +	13.1 (5.56)
DTLZ4	0.60 (0.13) –	0.16 (0.07) +	0.44 (0.13) –	0.36 (0.12)
DTLZ5	0.05 (0.02) ≈	0.03 (0.03) ≈	0.03 (0.00) ≈	0.04 (0.00)
DTLZ6	2.56 (1.21) +	0.72 (0.09) +	2.62 (1.95) +	5.36 (0.43)
DTLZ7	1.15 (0.91) +	0.03 (0.01) +	0.05 (0.08) +	5.41 (0.56)
UF1	0.19 (0.02) +	0.19 (0.01) +	0.19 (0.02) ≈	1.12 (0.14)
UF2	0.14 (0.02) +	0.12 (0.01) +	0.13 (0.02) +	0.57 (0.03)
UF3	0.19 (0.08) +	0.49 (0.01) ≈	0.42 (0.03) +	1.01 (0.05)
UF4	0.23 (0.02) –	0.22 (0.00) –	0.19 (0.01) –	0.17 (0.00)
UF5	2.49 (0.44) +	2.43 (0.28) +	2.42 (0.38) +	4.91 (0.36)
UF6	1.01 (0.25) +	1.32 (0.39) +	0.81 (0.19) +	5.43 (0.69)
UF7	0.37 (0.06) +	0.32 (0.11) +	0.33 (0.05) +	1.12 (0.08)
+/–/≈	9/4/3	9/3/4	7/4/4	

of 48 test instances for $r^{c} = 5$ and $r^{c} = 10$, confirming the effectiveness of the ensemble model and the search bias penalized acquisition function.

- The effectiveness of the proposed way of utilizing the additional data in SBP-BO can be validated by comparing SBP-BO with SBP-BO-C and SBP-BO-R. From Table SVII, SBP-BO significantly outperforms SBP-BO-C and SBP-BO-R on 37 test problems. This is consistent with the findings in [15]: how to utilize the additional data obtained from the search of cheap objectives plays a vital role in the optimization of HE-MOPs/HE-MaOPs. One weakness of the commonly used methods for training data selection based on clustering or randomly selection is that they cannot use all available data. This issue becomes more challenging for HE-MOPs where abundant training data are available for the fast objectives, making the algorithm inefficient for addressing problems with heterogeneous objectives. Similar observations can be made from Table SXIII. Note that there is no limitation on the available FEs for the relatively cheap objectives and thousands of cheap FEs are consumed in SBP-NoGP^{c}. It is interesting to see that SBP-BO is able to significantly outperform SBP-NoGP^{c} on 30 and 31 out of 48 test instances for $r^{c} = 5$ and $r^{c} = 10$, respectively. The comparison between SBP-BO and SBP-NoGP^{c} indicates that the algorithm can benefit from the use of surrogates on the cheap objectives. A possible explanation is that surrogates may smooth out some local optima and thus accelerate the search, which was discussed intuitively in [45] and empirically verified in [46].
- Compared with BO-LCB, the proposed algorithm shows significantly better performance on 22 out of 48 test instances, and similar performance on the remaining test problems, according to the results in Table SXI. It is worthy of noting that for HE-MOPs/HE-MaOPs with $r^{c} = 10$, the advantage of SBP-BO becomes a little less clear compared with BO-LCB, as can be observed from the results in Table SXI. SBP-BO is worse than BO-LCB on one test instance, but it only outperforms BO-AFF on 15 out of 48 instances. The results indicate that the algorithm can benefit from the use of the search bias penalty on some problems. However, since it is highly tricky to measure the search bias, it is challenging to apply an appropriate degree of penalty. This is might be the reason why SBP-BO and BO-LCB show similar performance on most test problems.

V. CONCLUSION

In this paper, we address heterogeneously expensive multi-/many-objective optimization problems, which have not received much attention in the evolutionary optimization community. We focus on exploiting the different amounts of data for the cheap and expensive objectives in constructing surrogates and reducing the search bias towards the cheap objectives within the Bayesian optimization framework. Specifically, to make full use of the available data for the cheap objectives while avoiding increasing the computational cost, an ensemble of GPs is constructed for each cheap objective to make use of both the solutions evaluated on all objectives and on the cheap objectives only. To reduce the bias towards the cheap objectives, we introduce a penalty term based on the heterogeneity in computational complexity of the objectives into the acquisition function, guiding the selection of new samples by taking the search bias into consideration. Although MOPs and MaOPs with heterogeneous objectives are ubiquitous in real-world applications, little work on BEO considering the computational heterogeneity has been reported in the literature. Different from most state-of-the-art algorithms that are limited to bi-objective optimization problems, the proposed algorithm is more generic in that it is applicable to problems with more than two objectives, where each objective can have a different evaluation time. Thus, the proposed work constitutes a valuable step forward towards solving real-world problems.

Encouraged by the promising results of the present work, we are interested in further investigating the efficient use of additional data on the cheap objectives, e.g., by properly guiding the single objective search. Meanwhile, the experimental results of the current work suggest that the proposed algorithm is less effective on nonseparable, multi-modal and disconnected problems, implying that more powerful search operators are required. Finally, this work adopts a simplified way to measure the search bias resulting from heterogeneous objectives, which leaves much room for further improvement in alleviating the search bias.

ACKNOWLEDGEMENT

This work was supported in part by the Honda Research Institute Europe. YJ is funded by an Alexander von Humboldt Professorship for AI endowed by the German Federal Ministry of Education and Research.

REFERENCES

[1] Y. Jin, H. Wang, T. Chugh, D. Guo, and K. Miettinen, "Data-driven evolutionary optimization: An overview and case studies," *IEEE Transactions on Evolutionary Computation*, vol. 23, no. 3, pp. 442–458, 2019.

[2] F. Kang, Q. Xu, and J. Li, "Slope reliability analysis using surrogate models via new support vector machines with swarm intelligence," *Applied Mathematical Modelling*, vol. 40, no. 11-12, pp. 6105–6120, 2016.

[3] D. Guo, Y. Jin, J. Ding, and T. Chai, "Heterogeneous ensemble-based infill criterion for evolutionary multiobjective optimization of expensive problems," *IEEE Transactions on Cybernetics*, vol. 49, no. 3, pp. 1012–1025, 2018.

[4] Q. Liu, R. Cheng, Y. Jin, M. Heiderich, and T. Rodemann, "Reference vector-assisted adaptive model management for surrogate-assisted many-objective optimization," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2022.

[5] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, "Taking the human out of the loop: A review of Bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2015.

[6] Z. Guo, Y.-S. Ong, T. He, and H. Liu, "Co-learning Bayesian optimization," *IEEE Transactions on Cybernetics*, vol. 52, no. 9, pp. 9820–9833, 2022.

[7] G. Li and Q. Zhang, "Multiple penalties and multiple local surrogates for expensive constrained optimization," *IEEE Transactions on Evolutionary Computation*, vol. 25, no. 4, pp. 769–778, 2021.

[8] H. Wang, L. Feng, Y. Jin, and J. Doherty, "Surrogate-assisted evolutionary multitasking for expensive minimax optimization in multiple scenarios," *IEEE Computational Intelligence Magazine*, vol. 16, no. 1, pp. 34 – 48, 2021.

[9] S. S. Thomas, J. Palandri, M. Lakehal-Ayat, P. Chakravarty, F. Wolf-Monheim, and M. B. Blaschko, "Kinematics design of a MacPherson suspension architecture based on Bayesian optimization," *IEEE Transactions on Cybernetics*, vol. 53, no. 4, pp. 2261 – 2274, 2023.

[10] Y. Morita, S. Rezaeiravesh, N. Tabatabaei, R. Vinuesa, K. Fukagata, and P. Schlatter, "Applying Bayesian optimization with Gaussian process regression to computational fluid dynamics problems," *Journal of Computational Physics*, vol. 449, p. 110788, 2022.

[11] R. Allmendinger and J. Knowles, "Heterogeneous objectives: state-of-the-art and future research," *arXiv preprint arXiv:2103.15546*, 2021.

[12] R. Allmendinger and J. Knowles, "Hang On a Minute: Investigations on the effects of delayed objective functions in multiobjective optimization," in *International Conference on Evolutionary Multi-Criterion Optimization*, pp. 6–20, Springer, 2013.

[13] R. Allmendinger, J. Handl, and J. Knowles, "Multiobjective optimization: When objectives exhibit non-uniform latencies," *European Journal of Operational Research*, vol. 243, no. 2, pp. 497–513, 2015.

[14] T. Chugh, Y. Jin, K. Miettinen, J. Hakanen, and K. Sindhya, "A surrogate-assisted reference vector guided evolutionary algorithm for computationally expensive many-objective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 22, no. 1, pp. 129–142, 2018.

[15] X. Wang, Y. Jin, S. Schmitt, and M. Olhofer, "Transfer learning for Gaussian process assisted evolutionary bi-objective optimization for objectives with different evaluation times," in *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, pp. 587–594, 2020.

[16] X. Wang, Y. Jin, S. Schmitt, M. Olhofer, and R. Allmendinger, "Transfer learning based surrogate assisted evolutionary bi-objective optimization for objectives with different evaluation times," *Knowledge-Based Systems*, p. 107190, 2021.

[17] X. Wang, Y. Jin, S. Schmitt, and M. Olhofer, "Transfer learning based co-surrogate assisted evolutionary bi-objective optimization for objectives with non-uniform evaluation times," *Evolutionary computation*, pp. 1–27, 2021.

[18] J. Thomann and G. Eichfelder, "A trust-region algorithm for heterogeneous multiobjective optimization," *SIAM Journal on Optimization*, vol. 29, no. 2, pp. 1017–1047, 2019.

[19] D. Hernández-Lobato, J. Hernandez-Lobato, A. Shah, and R. Adams, "Predictive entropy search for multi-objective Bayesian optimization," in *International Conference on Machine Learning*, pp. 1492–1501, PMLR, 2016.

[20] R. Allmendinger, M. T. Emmerich, J. Hakanen, Y. Jin, and E. Rigoni, "Surrogate-assisted multicriteria optimization: Complexities, prospective solutions, and business case," *Journal of Multi-Criteria Decision Analysis*, vol. 24, no. 1-2, pp. 5–24, 2017.

[21] S. Qin, C. Sun, Y. Jin, and G. Zhang, "Bayesian approaches to surrogate-assisted evolutionary multi-objective optimization: a comparative study," in *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 2074–2080, IEEE, 2019.

[22] Q. Zhang and H. Li, "MOEA/D: A multiobjective evolutionary algorithm based on decomposition," *IEEE Transactions on Evolutionary Computation*, vol. 11, no. 6, pp. 712–731, 2007.

[23] R. Cheng, Y. Jin, M. Olhofer, and B. Sendhoff, "A reference vector guided evolutionary algorithm for many-objective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 20, no. 5, pp. 773–791, 2016.

[24] Q. Zhang, W. Liu, E. Tsang, and B. Virginas, "Expensive multiobjective optimization by MOEA/D with gaussian process model," *IEEE Transactions on Evolutionary Computation*, vol. 14, no. 3, pp. 456–474, 2009.

[25] W. Ponweiser, T. Wagner, D. Biermann, and M. Vincze, "Multiobjective optimization on a limited budget of evaluations using model-assisted $\mathcal{S}$-metric selection," in *International Conference on Parallel Problem Solving from Nature*, pp. 784–794, Springer, 2008.

[26] N. Beume, B. Naujoks, and M. Emmerich, "SMS-EMOA: Multiobjective selection based on dominated hypervolume," *European Journal of Operational Research*, vol. 181, no. 3, pp. 1653–1669, 2007.

[27] J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani, "Predictive entropy search for efficient global optimization of black-box functions," *arXiv preprint arXiv:1406.2541*, 2014.

[28] W. Qiu, J. Zhu, G. Wu, H. Chen, W. Pedrycz, and P. N. Suganthan, "Ensemble many-objective optimization algorithm based on voting mechanism," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2022.

[29] M. Elarbi, S. Bechikh, A. Gupta, L. B. Said, and Y.-S. Ong, "A new decomposition-based nsga-ii for many-objective optimization," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 48, no. 7, pp. 1191–1210, 2018.

[30] D. Guo, X. Wang, K. Gao, Y. Jin, J. Ding, and T. Chai, "Evolutionary optimization of high-dimensional multiobjective and many-objective expensive problems assisted by a dropout neural network," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 52, no. 4, pp. 2084–2097, 2021.

[31] J. Knowles, "ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems," *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 1, pp. 50–66, 2006.

[32] Y. Cao and D. J. Fleet, "Generalized product of experts for automatic and principled fusion of Gaussian process predictions," *arXiv preprint arXiv:1410.7827*, 2014.

[33] M. Abdolshah, A. Shilton, S. Rana, S. Gupta, and S. Venkatesh, "Cost-aware multi-objective Bayesian optimisation," *arXiv preprint arXiv:1909.03600*, 2019.

[34] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler, "Scalable multi-objective optimization test problems," in *Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No. 02TH8600)*, vol. 1, pp. 825–830, IEEE, 2002.

[35] Q. Zhang, A. Zhou, S. Zhao, P. N. Suganthan, W. Liu, S. Tiwari, *et al.*, "Multiobjective optimization test instances for the CEC 2009 special session and competition," tech. rep., University of Essex, Colchester, UK and Nanyang technological University, 2008.

[36] S. Huband, P. Hingston, L. Barone, and L. While, "A review of multiobjective test problems and a scalable test problem toolkit," *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 5, pp. 477–506, 2006.

[37] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. G. Da Fonseca, "Performance assessment of multiobjective optimizers: An analysis and review," *IEEE Transactions on Evolutionary Computation*, vol. 7, no. 2, pp. 117–132, 2003.

[38] H. Ishibuchi, H. Masuda, Y. Tanigaki, and Y. Nojima, "Modified distance calculation in generational distance and inverted generational distance," in *International Conference on Evolutionary Multi-criterion Optimization*, pp. 110–125, Springer, 2015.

[39] L. While, P. Hingston, L. Barone, and S. Huband, "A faster algorithm for calculating hypervolume," *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 1, pp. 29–38, 2006.

[40] T. Chugh, R. Allmendinger, V. Ojalehto, and K. Miettinen, "Surrogate-assisted evolutionary biobjective optimization for objectives with non-uniform latencies," in *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 609–616, 2018.

[41] F. Li, L. Gao, and W. Shen, "Surrogate-assisted multi-objective evolutionary optimization with Pareto front model-based local search method," *IEEE Transactions on Cybernetics*, 2022.

[42] S. N. Lophaven, H. B. Nielsen, J. Søndergaard, *et al.*, *DACE: a Matlab kriging toolbox*, vol. 2. Citeseer, 2002.

[43] L. Pan, C. He, Y. Tian, H. Wang, X. Zhang, and Y. Jin, "A classification-based surrogate-assisted evolutionary algorithm for expensive many-objective optimization," *IEEE Transactions on Evolutionary Computation*, vol. 23, no. 1, pp. 74–88, 2019.

[44] M. T. Emmerich, K. C. Giannakoglou, and B. Naujoks, "Single-and multiobjective evolutionary optimization assisted by Gaussian random field metamodels," *IEEE Transactions on Evolutionary Computation*, vol. 10, no. 4, pp. 421–439, 2006.

[45] M. Hüsken, Y. Jin, and B. Sendhoff, "Structure optimization of neural networks for evolutionary design optimization," *Soft Computing*, vol. 9, no. 1, pp. 21–28, 2005.

[46] D. Lim, Y. Jin, Y.-S. Ong, and B. Sendhoff, "Generalizing surrogate-assisted evolutionary computation," *IEEE Transactions on Evolutionary Computation*, vol. 14, no. 3, pp. 329–355, 2009.

Xilu Wang received the B.Sc. and M.Sc. from Harbin Institute of Technology in 2016 and Xidian University in 2018, respectively, and received the Ph.D. degree at University of Surrey, UK, in 2022.

She is currently a Research Assistant with Faculty of Technology, Bielefeld University, Germany. Her research interests include evolutionary algorithms, multi-objective optimization, surrogate-assisted evolutionary optimization, Bayesian optimization and federated optimization.

Yaochu Jin (Fellow, IEEE) received the B.Sc., M.Sc., and Ph.D. degrees from Zhejiang University, Hangzhou, China, in 1988, 1991, and 1996, respectively, and the Dr.-Ing. degree from Ruhr University Bochum, Germany, in 2001.

He is an Alexander von Humboldt Professor for Artificial Intelligence endowed by the German Federal Ministry of Education and Research, with the Faculty of Technology, Bielefeld University, Germany. He is also a Distinguished Chair, Professor in Computational Intelligence, Department of Computer Science, University of Surrey, Guildford, U.K. He was a “Finland Distinguished Professor” of University of Jyväskylä, Finland, “Changjiang Distinguished Visiting Professor”, Northeastern University, China, and “Distinguished Visiting Scholar”, University of Technology Sydney, Australia. His main research interests include multi-objective and data-driven evolutionary optimization, evolutionary multi-objective learning, trustworthy AI, and evolutionary developmental AI.

Prof Jin is presently the President-Elect of the IEEE Computational Intelligence Society and the Editor-in-Chief of Complex & Intelligent Systems. He was the Editor-in-Chief of the IEEE TRANSACTIONS ON COGNITIVE AND DEVELOPMENTAL SYSTEMS, an IEEE Distinguished Lecturer in 2013-2015 and 2017-2019, and the Vice President for Technical Activities of the IEEE Computational Intelligence Society (2015-2016). He is the recipient of the 2018 and 2021 IEEE Transactions on Evolutionary Computation Outstanding Paper Award, and the 2015, 2017, and 2020 IEEE Computational Intelligence Magazine Outstanding Paper Award. He was named by the Web of Science as “a Highly Cited Researcher” from 2019 to 2022 consecutively. He is a Member of Academia Europaea.

Sebastian Schmitt received a Ph.D. degree in theoretical solid-state physics from the Technical University of Darmstadt, Darmstadt, Germany, in 2008. After a postdoc period from 2009 to 2011 with the Technical University of Dortmund, Dortmund, Germany, he joined the Honda Research Institute Europe in 2012. In his role as a Principal Scientist, he evaluates recent technologies in optimization, data analytics, machine learning and quantum computing for their use in real-world engineering settings relevant to R&D departments.

Markus Olhofer is Chief Scientist at the Honda Research Institute Europe GmbH since 2010 and responsible for the Complex Systems Optimisation and Analysis Group of the institute. Markus Olhofer graduated in Electrical Engineering in 2000 and received the Dr.-Ing. degree from the Institute for Neural Computation from Ruhr-Universität Bochum, Germany. He joined the Future Technology Research Division at Honda R&D Europe (Deutschland) GmbH in 1998 and since 2001 he works at the Honda Research Institute Europe GmbH..