

BI-OBJECTIVE BAYESIAN OPTIMIZATION WITH TRANSFORMED ADDITIVE GAUSSIAN PROCESSES

Caroline M. Kerfonta¹, Qiong Zhang¹ and Margaret M. Wiecek¹

¹School of Mathematical and Statistical Sciences, Clemson University, Clemson, SC, USA

ABSTRACT

Simulation is often used to optimize the performance of an engineering or scientific process. The input-output relation of a simulation model can be black-box functions and expensive to evaluate. Bayesian optimization, popularly used to improve efficiency in searching for optimal input settings, has been extended to bi-objective optimization problems. However, the common challenge in bi-objective Bayesian optimization is the computational challenge of maximizing the expected hypervolume improvement (EHI) to search for the input point to evaluate. In this paper, we utilize the transformed additive Gaussian process to simplify the objectives to additive functions for each dimension of the input space. Under this model assumption, the maximization of EHI over the entire decision space is reduced to the maximization of the resulting EHI functions for each dimension separately. We demonstrate the performances of the proposed method through numerical comparisons of the bi-objective Bayesian optimization with the traditional Gaussian process model.

1 INTRODUCTION

Simulation is a popular approach used to optimize engineering and scientific design (Lee et al. 2007; Zhang 2008; Fliege and Xu 2011; Hunter and McClosky 2016; Hunter et al. 2019; de Castro et al. 2022). For complex engineering design problems, running the corresponding simulation codes is often time-consuming. Also, the input and output relationship of the simulation model may not be expressed in a closed form to enable gradient-based optimization approaches. Thus, Bayesian optimization is often used to solve these problems by efficiently guiding new experiments through running the simulation models (Frazier 2018). Bayesian optimization approaches contain two key components: a computationally inexpensive surrogate model of the simulation and an acquisition function to select input points for the follow-up experiments. In terms of surrogate models, Gaussian process is a common choice (e.g., Jones et al. (1998)). A popular choice of the acquisition function is the expected improvement (EI, e.g. Qin et al. (2017)) to the target optimization problem over the input space. Then the input point for the new evaluation is given as a maximizer of EI over the input space.

Some engineering design problems involve two or more objectives. For multi-objective problems, the solution improvement can not be simply defined as for the single objective problems. In the literature of multi-objective optimization, the quality of the solution given by an optimization approach can be quantified by a variety of performance indicators, see various concepts given by Audet et al. (2021). Among them, the hypervolume indicator is a popular performance indicator to measure the quality of the solutions to multi-objective optimization problems (Knowles et al. 2003). Hence, in the literature on multi-objective Bayesian optimization (Hunter et al. 2019), a popular choice for the acquisition function is the expected hypervolume improvement (EHI), which is the expected difference between the hypervolume indicators computed after and before adding a new simulation evaluation (Emmerich et al. 2011). Even though EHI is widely used, there are some computational challenges to maximizing this acquisition function. When the dimension of the decision space increases, the optimization of EHI becomes much more complex, requires a longer computing time, and takes more memory space to complete (Hupkens et al. 2015). There are a number of existing approaches proposed to reduce the complexities of maximizing EHI. Yang et al. (2019)

consider bi-objective Bayesian optimization and provide an algorithm to evaluate the gradient function of EHI to enhance the optimization of EHI for searching new input points. A truncated form of the EHI calculation to optimize computing time is also proposed by Yang et al. (2016). This approach uses prior knowledge about the objective function values to more efficiently compute EHI. Since the gradient of EHI can be extremely complex and difficult to compute, Daulton et al. (2020) provide a method for computing EHI by using the gradients of the Monte Carlo estimator via auto-differentiation. This method simplifies the calculation of EHI by using first-order and quasi-second-order methods.

In this paper, we provide an algorithm for bi-objective Bayesian optimization to solve the computational challenges of maximizing EHI. Differently from perspectives in current literature, we approximate the objective functions to additive functions for each dimension of the decision space using the transformed additive Gaussian processes (Lin and Joseph 2020). Then the maximization of EHI over the entire decision space becomes the maximization of multiple univariate EHI functions, i.e., the number of EHI functions is equal to the number of dimensions of the input space. In our implementation, we maximize univariate EHI functions for each dimension of the input space through enumerating on a set of discrete points. We compare the numerical performance of the proposed approach with the bi-objective Bayesian optimization under the classical Gaussian process model.

This paper is organized as follows. Section 2 gives a review of related background, including single-objective Bayesian optimization with Gaussian processes and EI, bi-objective optimization and the hypervolume indicator, and bi-objective Bayesian optimization with Gaussian processes and EHI. Section 3 proposes our method, which uses transformed additive Gaussian process (TAG) models as the surrogate models and derives the EHI expression under this model assumption. Section 4 provides a numerical study that demonstrates the advantages of the proposed method. Section 5 concludes the paper and points out directions for future improvement of the proposed method.

2 BACKGROUND

In this section, we review related background approaches including Bayesian optimization with Gaussian processes and EI, bi-objective optimization, the hypervolume indicator, and bi-objective Bayesian optimization with classical Gaussian processes and EHI.

2.1 Bayesian Optimization

Let $y(\mathbf{x})$ be a scalar deterministic black-box function with inputs $\mathbf{x} = (x_1, \dots, x_d)^\top \in \mathcal{X}^d$. The goal of Bayesian optimization is to solve problems in the form:

$$\min_{\mathbf{x} \in \mathcal{X}} y(\mathbf{x}), \quad (1)$$

where \mathcal{X} is the domain of the decision. Bayesian optimization (Frazier 2018) is comprised of two parts: the surrogate statistical model and the acquisition function. The surrogate model is fitted by the data collected from evaluating the function $y(\mathbf{x})$ and informs the posterior distribution of $y(\mathbf{x})$ given the data. Under the distribution of $y(\mathbf{x})$, the acquisition function is optimized to select a new input point from \mathcal{X} . A generic Bayesian optimization algorithm is shown in Algorithm 1 for a budget of N function evaluation. The acquisition function is optimized in step 3, and the surrogate model is updated in step 4. Bayesian optimization algorithms could differ in different target optimization problems (e.g., single or multiple objectives), surrogate statistical models and acquisition functions.

We illustrate Bayesian optimization in detail using the Gaussian process as the surrogate model and the EI as the acquisition function. The Gaussian process model is a popular choice of surrogate models (Sacks et al. 1989). We assume that the function value $y(\mathbf{x})$ is a realization of the stochastic process

$$Y(\mathbf{x}) = \mu + Z(\mathbf{x}), \quad (2)$$

where μ is the unknown deterministic mean value, and $Z(\mathbf{x})$ with $\mathbf{x} \in \mathcal{X}$ is mean zero Gaussian process with variance σ^2 and correlation function $R(\cdot)$, e.g., the Gaussian correlation function (Sacks et al. 1989)

Algorithm 1 Generic Bayesian Optimization Algorithm

- 1: Provide a prior distribution of $y(\mathbf{x})$
 - 2: **for** $n = 0, 1, \dots, N - 1$ **do**
 - 3: Compute acquisition function based on the distribution of $y(\mathbf{x})$ and find \mathbf{x}_{n+1} that maximizes the acquisition over $\mathbf{x} \in \mathcal{X}$
 - 4: Collect a new data point by evaluating $y(\mathbf{x})$ at \mathbf{x}_{n+1} and update the distribution of $y(\mathbf{x})$ as the posterior distribution of $y(\mathbf{x})$ given all existing data.
 - 5: **end for**
- Return the input point with the smallest posterior mean as the optimal decision.
-

for two inputs points \mathbf{x} and \mathbf{x}' is given by

$$R(\mathbf{x} - \mathbf{x}') = \exp\left(-\sum_{i=1}^d \theta_i (x_i - x'_i)^2\right), \quad (3)$$

where θ_i 's are the correlation parameters. Under this model assumption, given two input points $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, we have that

$$Y(\mathbf{x}) \sim N(\mu, \sigma^2) \quad \text{and} \quad \text{Cov}(Y(\mathbf{x}), Y(\mathbf{x}')) = \sigma^2 R(\mathbf{x} - \mathbf{x}').$$

Assume that we have evaluated the function $y(\mathbf{x})$ at n input points $\mathbf{x}_1, \dots, \mathbf{x}_n$, and collected the outputs $\mathbf{y}_n = (y(\mathbf{x}_1), \dots, y(\mathbf{x}_n))^\top$. Under the Gaussian process assumption, the conditional distribution of $Y(\mathbf{x})$ given \mathbf{y}_n at a new input point \mathbf{x} (see for example, Jones et al. (1998)) is expressed by :

$$Y(\mathbf{x}) | \mathbf{y}_n \sim N(\hat{y}(\mathbf{x}), s^2(\mathbf{x})), \quad (4)$$

with mean

$$\hat{y}(\mathbf{x}) = \hat{\mu} + \mathbf{r}(\mathbf{x})^\top \mathbf{R}^{-1} (\mathbf{y}_n - \mathbf{1} \hat{\mu}) \quad (5)$$

and variance

$$s^2(\mathbf{x}) = \sigma^2 \left(1 - \mathbf{r}(\mathbf{x})^\top \mathbf{R}^{-1} \mathbf{r}(\mathbf{x}) - \frac{(1 - \mathbf{1}^\top \mathbf{R}^{-1} \mathbf{r}(\mathbf{x}))^2}{\mathbf{1}^\top \mathbf{R}^{-1} \mathbf{1}} \right) \quad (6)$$

where $\mathbf{1}$ is a vector of ones of size n , $\hat{\mu} = \mathbf{1}^\top \mathbf{R}^{-1} \mathbf{y}_n / \mathbf{1}^\top \mathbf{R}^{-1} \mathbf{1}$ is the best unbiased linear estimator of μ , $\mathbf{r}(\mathbf{x}) = (R(\mathbf{x} - \mathbf{x}_1), \dots, R(\mathbf{x} - \mathbf{x}_n))^\top$, and \mathbf{R} be the $n \times n$ correlation matrix with the (i, j) -th element $R(\mathbf{x}_i - \mathbf{x}_j)$.

EI, as a popular choice for the acquisition function (Jones et al. 1998), measures the expected improvement to the target optimization problem at input point \mathbf{x} . Let $y_n^{\min} = \min_{1 \leq i \leq n} y(\mathbf{x}_i)$ be the current minimum value at step n . Given by Jones et al. (1998), the expected improvement of input point \mathbf{x} is defined by:

$$\text{EI}_n(\mathbf{x}) = \text{E} \left[\max(y_n^{\min} - Y(\mathbf{x}), 0) \right], \quad (7)$$

where the expectation is taken with respect to the conditional distribution of $Y(\mathbf{x})$ given \mathbf{y}_n in (4). Then, EI can be expressed by:

$$\text{EI}_n(\mathbf{x}) = (y_n^{\min} - \hat{y}(\mathbf{x})) \Phi\left(\frac{y_n^{\min} - \hat{y}(\mathbf{x})}{s(\mathbf{x})}\right) + s(\mathbf{x}) \phi\left(\frac{y_n^{\min} - \hat{y}(\mathbf{x})}{s(\mathbf{x})}\right), \quad (8)$$

where $\hat{y}(\mathbf{x})$ and $s(\mathbf{x})$ are the mean and standard deviation of $Y(\mathbf{x}) | \mathbf{y}_n$ in (4), and $\Phi(\cdot)$ and $\phi(\cdot)$ are the cumulative distribution function and the probability density function of the standard normal distribution, respectively.

2.2 Bi-objective Optimization

Let $y_1(\mathbf{x})$ and $y_2(\mathbf{x})$ be two scalar-valued deterministic black-box functions with input $\mathbf{x} \in \mathcal{R}^d$. The bi-objective optimization problem is defined by

$$\begin{aligned} \min \quad & \mathbf{y}(\mathbf{x}) = [y_1(\mathbf{x}), y_2(\mathbf{x})] \\ \text{s.t.} \quad & \mathbf{x} \in \mathcal{X} \end{aligned} \tag{9}$$

where $\mathcal{X} \subset \mathcal{R}^d$ is a feasible set. Euclidean spaces \mathcal{R}^d and \mathcal{R}^2 are respectively referred to as the decision space and the objective (outcome) space. Given two input points $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$, the outcome $\mathbf{y}(\mathbf{x})$ is said to dominate the outcome $\mathbf{y}(\mathbf{x}')$, denoted as $\mathbf{y}(\mathbf{x}) \prec \mathbf{y}(\mathbf{x}')$, if and only if $y_l(\mathbf{x}) \leq y_l(\mathbf{x}')$ for $l = 1, 2$ and $y_l(\mathbf{x}) < y_l(\mathbf{x}')$ for $l = 1$ or $l = 2$. If $y_l(\mathbf{x}) \leq y_l(\mathbf{x}')$ for $l = 1, 2$, the outcome $\mathbf{y}(\mathbf{x})$ is said to weakly dominate the outcome $\mathbf{y}(\mathbf{x}')$, denoted as $\mathbf{y}(\mathbf{x}) \preceq \mathbf{y}(\mathbf{x}')$. An outcome point is said to be nondominated if there does not exist another outcome point dominating it. To solve problem (9), we apply the classical concept of Pareto-optimality, that is to identify the set of all input points whose images are nondominated outcomes (Ehrgott 2005). The Pareto set is defined by

$$\mathcal{P} = \{\mathbf{y}(\mathbf{x}) \in \mathcal{R}^2 : \mathbf{x} \in \mathcal{X} \mid \nexists \mathbf{x}' \in \mathcal{X} \text{ s.t. } \mathbf{y}(\mathbf{x}') \prec \mathbf{y}(\mathbf{x})\}. \tag{10}$$

Since the Pareto set is typically not available in a closed form even if the objective and constraint functions are available, various computational approaches and algorithms have been developed to provide approximation of different types (Ruzika and Wiecek 2005; Herzel et al. 2021). Due to the diversity of these methods and the resulting sets being exactly or only approximating the Pareto set, many ways to measure their quality have been proposed (Faulkenberg and Wiecek 2010; Audet et al. 2021). One way to measure the quality of the computed Pareto set is the hypervolume indicator (e.g., Knowles et al. (2003), Guerreiro et al. (2021)). Given a reference point $\mathbf{t} = (t_1, t_2) \in \mathcal{R}^2$, the hypervolume indicator $H(\mathcal{P}, \mathbf{t})$ is the two-dimensional Lebesgue measure Λ of the region weakly dominated by \mathcal{P} and bounded above by the reference point \mathbf{t} , i.e.,

$$H(\mathcal{P}, \mathbf{t}) = \Lambda(\{\mathbf{v} \in \mathcal{R}^2 \mid \nexists \mathbf{p} \in \mathcal{P} : \mathbf{p} \preceq \mathbf{v} \text{ and } \mathbf{v} \preceq \mathbf{t}\}). \tag{11}$$

The hypervolume improvement (HVI) (Hupkens et al. 2015) can be used to assess the effect of evaluating the vector-valued objective function $\mathbf{y}(\mathbf{x})$ at a new input point $\mathbf{x} \in \mathcal{X}$:

$$\text{HVI}(\mathbf{y}(\mathbf{x}), \mathcal{P}, \mathbf{t}) = H(\mathcal{P} \cup \{\mathbf{y}(\mathbf{x})\}, \mathbf{t}) - H(\mathcal{P}, \mathbf{t}). \tag{12}$$

2.3 Bi-objective Bayesian Optimization with EHI

In bi-objective Bayesian optimization, the expected value of the hypervolume improvement (EHI) in (12) is a commonly used acquisition function (Emmerich et al. 2011). Given that the two objectives $y_1(\mathbf{x})$ and $y_2(\mathbf{x})$ are modeled by two independent Gaussian processes $Y_1(\mathbf{x})$ and $Y_2(\mathbf{x})$ as in (2), we provide the expression of EHI.

Assume that we evaluated the two objectives at n input points $\mathbf{x}_1, \dots, \mathbf{x}_n$, and collected the outputs $\mathbf{y}_{n,1} = (y_1(\mathbf{x}_1), \dots, y_1(\mathbf{x}_n))^T$ and $\mathbf{y}_{n,2} = (y_2(\mathbf{x}_1), \dots, y_2(\mathbf{x}_n))^T$. Following (4), we have

$$Y_1(\mathbf{x})|\mathbf{y}_{n,1} \sim N(\hat{y}_1(\mathbf{x}), \hat{s}_1^2(\mathbf{x})) \quad \text{and} \quad Y_2(\mathbf{x})|\mathbf{y}_{n,2} \sim N(\hat{y}_2(\mathbf{x}), \hat{s}_2^2(\mathbf{x})), \tag{13}$$

where $\hat{y}_l(\mathbf{x})$ and $\hat{s}_l^2(\mathbf{x})$ for $l = 1, 2$ are the conditional means and variances, which can be calculated using (5) and (6). Under the assumption that $Y_1(\mathbf{x})$ and $Y_2(\mathbf{x})$ are independent, $Y_1(\mathbf{x})|\mathbf{y}_{n,1}$ and $Y_2(\mathbf{x})|\mathbf{y}_{n,2}$ are also independent. Note that the two Gaussian processes $Y_1(\mathbf{x})$ and $Y_2(\mathbf{x})$ can have different parameters and correlation functions. Given a reference point \mathbf{t} , EHI is defined by

$$\text{EHI}_n(\mathbf{x}; \mathcal{P}, \mathbf{t}) = \text{E}\{\text{HVI}((Y_1(\mathbf{x}), Y_2(\mathbf{x})), \mathcal{P}, \mathbf{t})\}, \tag{14}$$

where HVI is given by (12) with the new objective function evaluation $(Y_1(\mathbf{x}), Y_2(\mathbf{x}))$ and the expectation is taken with respect to the conditional distributions in (13). By maximizing $\text{EHI}_n(\mathbf{x}; \mathcal{P}, \mathbf{t})$ over $\mathbf{x} \in \mathcal{X}$, we find an input point for the new function evaluation. Given an approximation of the Pareto set \mathcal{P} by a set of nondominated points, EHI can be calculated empirically with a closed-form expression (Hupkens et al. 2015). We denote the nondominated points by $\mathbf{y}^j = (y_1^j, y_2^j) = (y_1(\mathbf{x}_{i_j}), y_2(\mathbf{x}_{i_j}))$ for $j = 1, \dots, p$ with $\{i_1, \dots, i_p\} \subset \{1, \dots, n\}$. Without loss of generality, we assume that those nondominated points are ordered based on the values of the first objective, i.e., $y_1(\mathbf{x}_{i_1}) < y_1(\mathbf{x}_{i_2}) < \dots < y_1(\mathbf{x}_{i_p})$. Then the approximation of the Pareto set \mathcal{P} is denoted by

$$\mathcal{P}^{(n)} = \{\mathbf{y}^0, \mathbf{y}^1, \dots, \mathbf{y}^p, \mathbf{y}^{p+1}\}, \tag{15}$$

where $\mathbf{y}^0 = (-\infty, t_2)$ and $\mathbf{y}^{p+1} = (t_1, -\infty)$ with t_1 and t_2 being the two coordinates of the reference point $\mathbf{t} = (t_1, t_2)$. Following Emmerich et al. (2011), EHI in (14) can be calculated empirically based on the conditional distributions in (13):

$$\begin{aligned} \text{EHI}_n(\mathbf{x}; \mathcal{P}^{(n)}, \mathbf{t}) &= \mathbb{E} \left\{ \text{HVI} \left((Y_1(\mathbf{x}), Y_2(\mathbf{x})), \mathcal{P}^{(n)}, \mathbf{t} \right) \right\} \\ &= \sum_{j=1}^{p+1} (y_1^{j-1} - y_1^j) \cdot \Phi \left(\frac{y_1^j - \hat{y}_1(\mathbf{x})}{\hat{s}_1(\mathbf{x})} \right) \cdot \Psi \left(y_2^j, y_2^j, \hat{y}_2(\mathbf{x}), \hat{s}_2(\mathbf{x}) \right) \\ &\quad + \sum_{j=1}^{p+1} \left(\Psi \left(y_1^{j-1}, y_1^{j-1}, \hat{y}_1(\mathbf{x}), \hat{s}_1(\mathbf{x}) \right) - \Psi \left(y_1^{j-1}, y_1^j, \hat{y}_1(\mathbf{x}), \hat{s}_1(\mathbf{x}) \right) \right) \\ &\quad \cdot \Psi \left(y_2^j, y_2^j, \hat{y}_2(\mathbf{x}), \hat{s}_2(\mathbf{x}) \right), \end{aligned} \tag{16}$$

where $\Psi(a, b, \mu, \sigma) = \sigma \phi \left(\frac{b-\mu}{\sigma} \right) + (a - \mu) \Phi \left(\frac{b-\mu}{\sigma} \right)$. It is worth noting that, the evaluation of this function is computationally non-trivial as demonstrated by Emmerich et al. (2016) and Yang et al. (2019). Since the key properties (such as concavity or convexity) of this function can not be justified, the number of function evaluations of EHI required to obtain a high-quality solution is not negligible for median to high dimensional decision spaces.

3 PROPOSED METHOD

We propose a bi-objective Bayesian optimization algorithm using transformed additive Gaussian processes (TAG) (Lin and Joseph 2020) as the surrogate model. TAG approximates each of the two objective functions by additive functions, each additive term associated with one-dimensional decision variable. Under this model assumption, maximizing the acquisition function EHI given by (14) can be simplified to maximize the EHI's for each decision variable. In this section, we introduce using TAG as a surrogate model for bi-objective problems, develop the EHI functions under the TAG model assumption, and summarize our algorithm implementation.

3.1 Bi-objective Surrogate with Transformed Additive Gaussian Processes

Transformed additive Gaussian processes aim to simplify a complex response to additive functions. Consider the bi-objective problem in (9). As noted in Section 2.3, we assume that $y_1(\mathbf{x})$ and $y_2(\mathbf{x})$ are the realizations of two independent transformed additive Gaussian processes $Y_1(\mathbf{x})$ and $Y_2(\mathbf{x})$ with $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^d$. Equivalently, following Lin and Joseph (2020), we have

$$g_{\lambda_l}(Y_l(\mathbf{x})) = \mu_l + \sum_{k=1}^d Z_{lk}(x_k) + \varepsilon_l, \quad \text{for } l = 1, 2. \tag{17}$$

We explain the notation in turn as follows. The function $g_{\lambda_l}(\cdot)$ is the Box-Cox transformation (Box and Cox 1964):

$$g_{\lambda_l}(y) = \begin{cases} \frac{y^{\lambda_l}-1}{\lambda_l} & \lambda_l \neq 0 \\ \log y & \lambda_l = 0, \end{cases} \quad \text{for } l = 1, 2, \quad (18)$$

where λ_l is the parameter of the transformation. The deterministic means of TAG are denoted by μ_l 's. For $k = 1, \dots, d$ and $l = 1, 2$, $Z_{lk}(x_k)$'s are mutually independent mean-zero Gaussian processes with variances τ_{lk}^2 's, and correlation functions $R_{lk}(\cdot)$'s. The additive noise $\varepsilon_l \sim N(0, \sigma_l^2)$ is independent of $Z_{lk}(x_k)$'s.

Under this model assumption, the correlation of $g_{\lambda_l}(Y_l(\mathbf{x}))$ and $g_{\lambda_l}(Y_l(\mathbf{x}'))$ for $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ is

$$R_l(\mathbf{x} - \mathbf{x}') = \sum_{k=1}^d \omega_{lk} R_{lk}(x_k - x'_k), \quad (19)$$

where $\tau_l^2 = \sum_{k=1}^d \tau_{lk}^2$ and $\omega_{lk} = \tau_{lk}^2 / \tau_l^2$. An example of the correlation function is given by (3), which is associated with some unknown correlation parameters. In Lin and Joseph (2020), empirical Bayes estimation and non-linear optimization toolboxes are used to estimate those unknown parameters, including τ_l , ω_{lk} , λ_l and correlation parameters, which are referred to as hyperparameters of TAG.

Assume that we collected outputs $\mathbf{y}_{n,1} = (y_1(\mathbf{x}_1) \dots, y_1(\mathbf{x}_n))^\top$ and $\mathbf{y}_{n,2} = (y_2(\mathbf{x}_1) \dots, y_2(\mathbf{x}_n))^\top$. Let $g_{\lambda_l}(\mathbf{y}_{n,l}) = (g_{\lambda_l}(y_l(\mathbf{x}_1)), \dots, g_{\lambda_l}(y_l(\mathbf{x}_n)))^\top$ for $l = 1, 2$. For $l = 1, 2$ and $k = 1, \dots, d$, we have that

$$\begin{pmatrix} g_{\lambda_l}(\mathbf{y}_{n,l}) - \mu_l \mathbf{1} \\ Z_{lk}(x_k) \end{pmatrix} \sim N_{n+1} \left[\begin{pmatrix} \mathbf{0}_n \\ 0 \end{pmatrix}, \begin{pmatrix} \tau_l^2 \mathbf{R}_l + \sigma_l^2 \mathbf{I} & \tau_{lk}^2 \mathbf{r}_{lk}(x_k) \\ \tau_{lk}^2 \mathbf{r}_{lk}(x_k)^\top & \tau_{lk}^2 \end{pmatrix} \right], \quad (20)$$

where $\mathbf{0}_n$ is the $n \times 1$ vector of zeros, \mathbf{I} is the $n \times n$ identify matrix, $\mathbf{r}_{lk}(x_k) = (R_{lk}(x_k - x_{1k}), \dots, R_{lk}(x_k - x_{nk}))^\top$ and \mathbf{R}_l is the $n \times n$ correlation matrix with the (i, j) -th element $R_l(\mathbf{x}_i - \mathbf{x}_j)$ given by (19). Using the conditional distribution for the normal distribution (e.g., Eaton (2007)), we can derive the conditional distribution of $Z_{lk}(x_k)$ given $\mathbf{y}_{n,l}$:

$$Z_{lk}(x_k) | \mathbf{y}_{n,l} \sim N(\hat{z}_{lk}(x_k), s_{lk}^2(x_k)) \quad \text{for } l = 1, 2, \quad (21)$$

with

$$\hat{z}_{lk}(x_k) = \omega_{lk} \mathbf{r}_{lk}(x_k)^\top (\mathbf{R}_l + \delta_l \mathbf{I})^{-1} (g_{\lambda_l}(\mathbf{y}_{n,l}) - \mu_l \mathbf{1}),$$

and

$$s_{lk}^2(x_k) = \tau_l^2 \omega_{lk} \left(1 - \mathbf{r}_{lk}(x_k)^\top (\mathbf{R}_l + \delta_l \mathbf{I})^{-1} \omega_{lk} \mathbf{r}_{lk}(x_k) \right),$$

where $\delta_l = \sigma_l^2 / \tau_l^2$ and μ_l can be replaced by $\hat{\mu}_l = \mathbf{1}^\top (\mathbf{R}_l + \delta_l \mathbf{I})^{-1} g_{\lambda_l}(\mathbf{y}_{n,l}) / \mathbf{1}^\top (\mathbf{R}_l + \delta_l \mathbf{I})^{-1} \mathbf{1}$ as a plug-in estimator. The conditional distribution of $Z_{lk}(x_k)$ will then be used to develop EHIs for each dimension of the decision variable.

3.2 Expected Hypervolume Improvement with Transformed Additive Gaussian Processes

Based on the conditional distributions in (21), we simplify the EHI in (14) to the expected hypervolume improvements for the bi-objective problems for each dimension of the decision variable. Provided by the approximation of TAG, we have

$$g_{\lambda_l}(y_l(\mathbf{x})) \approx \mu_l + \sum_{k=1}^d z_{lk}(x_k) \quad \text{for } l = 1, 2, \quad (22)$$

where $z_{lk}(x_k)$ is a realization of $Z_{lk}(x_k)$. Assume that the decision space $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_d$. Under this approximation, we obtain a surrogate optimization problem of the the original bi-objective optimization problem in (9)

$$\min_{\mathbf{x}_k \in \mathcal{X}_k} \mathbf{z}_k(x_k) = [z_{1k}(x_k), z_{2k}(x_k)] \quad \text{for } k = 1, \dots, d \quad (23)$$

Let \mathcal{P}_k be the Pareto set of the k -th problem in (23), i.e.,

$$\mathcal{P}_k = \{\mathbf{z}_k(x_k) \in \mathcal{R}^2 \mid \nexists \mathbf{z}_k(x'_k) \in \mathcal{R}^2 : \mathbf{z}_k(x'_k) \prec \mathbf{z}_k(x_k)\} \quad (24)$$

for $k = 1, \dots, d$. Under TAG, we can define the EHI for the problem in (23) by:

$$\text{EHI}_{n,k}(x_k; \mathcal{P}_k, \mathbf{t}_k) = \text{E}\{\text{HVI}((Z_{1k}(x_k), Z_{2k}(x_k)), \mathcal{P}_k, \mathbf{t}_k)\}, \quad (25)$$

where the expectation is taken with respect to the conditional distribution of $Z_{lk}(x_k)$ in (21). As a result, instead of maximizing $\text{EHI}_n(\mathbf{x}; \mathcal{P}, \mathbf{t})$ over $\mathbf{x} \in \mathcal{X}$, a new input point is given by maximizing $\text{EHI}_{n,k}(x_k; \mathcal{P}_k, \mathbf{t}_k)$ with one dimensional decision variable $x_k \in \mathcal{X}_k$ for $k = 1, \dots, d$. To be specific, the new design point is generated by selecting $\mathbf{x}_{new} = (x_1^*, \dots, x_d^*)$ with

$$x_k^* \in \max_{x \in \mathcal{X}_k} \text{EHI}_{n,k}(x_k; \mathcal{P}_k, \mathbf{t}_k) \quad \text{for } k = 1, \dots, d.$$

Given an approximation of the Pareto set \mathcal{P}_k , EHI can attain a close form expression as in (16) using the conditional distribution in (21). In the next subsection, we describe our detailed implementation of the proposed algorithm.

3.3 The Proposed Algorithm

We detail our implementation of the proposed bi-objective Bayesian optimization algorithm in this subsection. The surrogate model is a TAG model in (17) implemented by the R package TAG (Lin and Joseph 2021). The calculation of EHI depends on the approximation of the Pareto set in (15). We use the R package `ecr` (Bossek 2017) to approximate Pareto set. Given responses $\mathbf{y}_{m,1}$ and $\mathbf{y}_{m,2}$ of size m , the approximations of the Pareto sets of \mathcal{P} in (10) and \mathcal{P}_k in (24) are denoted by $\mathcal{P}^{(m)}$ and $\mathcal{P}_k^{(m)}$, respectively. We summarize our implementation in Algorithm 2. We first generate an initial dataset of size n (line 1 in Algorithm 2), and then add one more data point at each step until we exhaust the budget of function evaluation N (line 2-10 in Algorithm 2). In each step, we fit a TAG model for each objective function. For simplicity, the maximization of $\text{EHI}_{n,k}(x_k; \mathcal{P}_k, \mathbf{t}_k)$ is done by evaluating q random candidate points in \mathcal{X}_k (line 4-9 in Algorithm 2). In our implementation, we set $q = 10$ and generate a $q \times d$ random Latin hypercube design using the R package `lhs` (Carnell 2022) with k -th column being the candidate points in \mathcal{X}_k (line 4 in Algorithm 2). At each step, we compute the hypervolume indicator in (11) based on the approximated Pareto set $\mathcal{P}^{(m)}$ (line 10 in Algorithm 2).

The hyperparameters (such as λ_l 's and τ_l 's etc) in TAG are not updated in every step in our implementation. To save computation time and resources, we will refit those hyperparameters in every ten steps. The size of Latin hypercube design in line 4 of Algorithm 2 can be adjusted according to the computational requirement of the implementation of the algorithm. Also note that, for convenience, we use the enumeration approach to maximize univariate EHI functions in our implementation. In practice, alternative univariate optimization algorithms can be applied here to improve the efficiency and accuracy.

4 NUMERICAL RESULTS

In this section, we compare the proposed algorithm as described in Section 3.3 with the bi-objective Bayesian optimization procedure implemented with the function `GParetooptim` in the R package `GPareto` (Binois and Picheny 2019). This algorithm uses independent Gaussian processes as a surrogate model for the two objectives, and uses EHI in (16) as the acquisition function. To distinguish between the two approaches, our method is denoted by ‘‘Transformed Additive GP’’, and the method in `GPareto` is referred to as ‘‘Classical GP’’. The purpose of the numerical comparison is to explore the situations that the proposed algorithm shows competitive performances or advantages by incorporating the additive assumption. Therefore, ‘‘Classical GP’’ with original EHI is the key benchmark without the additive assumption.

Algorithm 2 Bi-objective Bayesian Optimization with TAG

- 1: Given an initial dataset with inputs $\mathbf{x}_1, \dots, \mathbf{x}_n$ and the outputs $\mathbf{y}_{n,1}$ and $\mathbf{y}_{n,2}$, compute an initial Pareto set approximation $\mathcal{P}^{(n)}$ and compute $H(\mathcal{P}^{(n)}, \mathbf{t})$.
 - 2: **for** $m = n, n + 1, \dots, N$ **do**
 - 3: Fit the TAG models based on $\mathbf{y}_{m,1}$ and $\mathbf{y}_{m,2}$ and \mathbf{x} , find $\mathcal{P}_k^{(m)}$ for $k = 1, \dots, d$.
 - 4: Generate a $q \times d$ matrix D as a random Latin Hypercube design with ij -th entries D_{ij} .
 - 5: **for** $k = 1, \dots, d$ **do**
 - 6: Evaluate $\text{EHI}_{n,k} \left(D_{ik}; \mathcal{P}_k^{(m)}, \mathbf{t}_k \right)$ for $i = 1, \dots, q$.
 - 7: Select point D_k^* that maximizes $\text{EHI}_{n,k} \left(D_{ik}; \mathcal{P}_k^{(m)}, \mathbf{t}_k \right)$ over $i = 1, \dots, q$.
 - 8: **end for**
 - 9: Set new input point $\mathbf{x}^{(m+1)} = (D_1^*, \dots, D_d^*)$ and obtain $\mathbf{y}(\mathbf{x}^{(m+1)})$.
 - 10: By including the new response, update $\mathbf{y}_{m,1} \rightarrow \mathbf{y}_{m+1,1}$, $\mathbf{y}_{m,2} \rightarrow \mathbf{y}_{m+1,2}$, and $\mathcal{P}^{(m)} \rightarrow \mathcal{P}^{(m+1)}$. Compute $H(\mathcal{P}^{(m+1)}, \mathbf{t})$.
 - 11: **end for**
- Return $H(\mathcal{P}^{(m)}, \mathbf{t})$ for $m = n, n + 1, \dots, N$.
-

First, we generate initial data using the R package `lhs` (Carnell 2022). Next, we fit an initial surrogate model to the data. For our method, the surrogate model is a TAG model as in (17). The classical GP approach will use the Gaussian processes as the surrogate models using the R package (Roustant et al. 2012) `DiceKriging`. The reference point \mathbf{t} is chosen as a point that bounds the outcome space of the bi-objective problem when $0 \leq x_k \leq 1$ for $k = 1, \dots, d$. In our numerical examples, we carefully select a common reference point \mathbf{t} for all \mathbf{t}_k 's. We sequentially add new points one by one using EHI for 20 steps. In each step, we compute $H(\mathcal{P}, \mathbf{t})$ in (11) based on approximations of \mathcal{P} following (Fonseca et al. 2006).

We replicate the whole procedure 30 times (with 30 different random initial datasets) for both methods. For the returned hypervolume values of each method in each step, we compute their average and use 10th and 90th quantiles to construct a confidence band. For our method, in the R package TAG, the default range of λ_l in (18) is -2 to 2 , we use this default range for the examples unless otherwise noted.

4.1 Example 1

We construct bi-objective functions as follows:

$$y_l(\mathbf{x}) = \exp \left(\sum_{k=1}^d (x_k - a_l)^2 + c_l x_1 x_2 \right) \quad \text{for } l = 1, 2, \tag{26}$$

which can be transformed to additive functions using the Box-Cox transformation (18) for $\lambda_l = 0$ and $c_l = 0$.

We first study the impact of the dimension of the decision space on the proposed method as compared to the classical GP approach. We consider the performance of the proposed algorithm under different sizes of d . We specify three cases with parameters of (26) in Table 1, where the dimension of input spaces changes from 4 to 8. To ensure there are a sufficient number of initial points for $d = 8$, we use 25 initial data points in this study. The results are depicted in Figure 1. As the dimension of the decision space increases, the advantage of the proposed method is greater in terms of both mean and variance of the approximated hypervolume indicators.

Next, we study the performance of the proposed method when the objective functions have an interaction term and are not strictly transformed to additive functions with one-dimensional input. We consider the performance of the proposed algorithm under different values of c_l . We specify three cases with parameters of (26) in Table 2, where the value of c_2 changes from 0 to 0.5. For this example, 15 initial data points are used. The results are depicted in Figure 2. The results show that when the interaction term is negligible as

Table 1: Table of parameter settings of (26) for the three cases in Figure 1.

Cases	(1)			(2)			(3)		
l	d	a_l	c_l	d	a_l	c_l	d	a_l	c_l
1	4	0.3	0	6	0.3	0	8	0.3	0
2	4	0.4	0	6	0.4	0	8	0.4	0

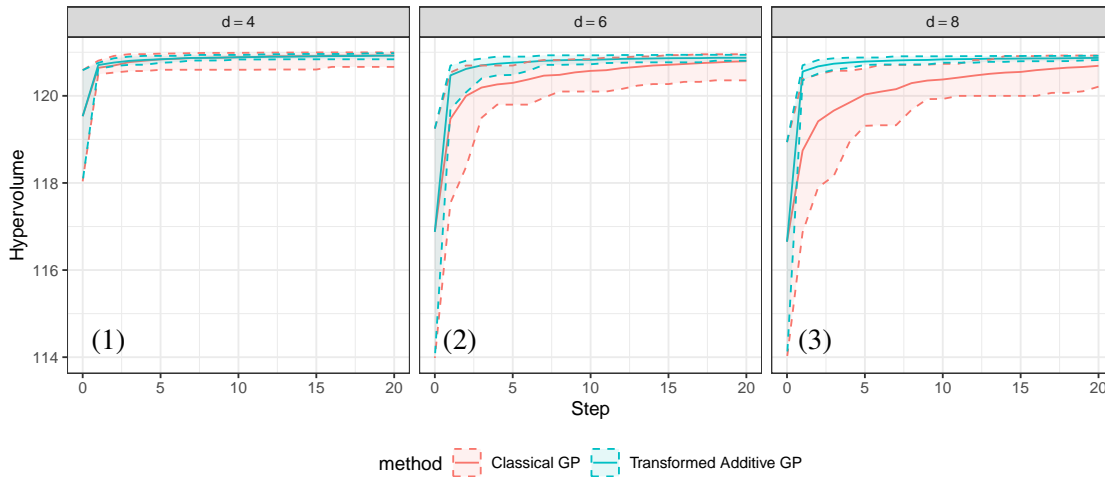


Figure 1: The mean, 10-th and 90-th quantiles of approximated hypervolume over 30 replications and 20 steps for optimizing the objective functions in (26) with the parameter settings in Table 1.

$c_2 = 0$ or 0.01 , our method is more robust than Classical GP in early stages. However, when the interaction term increases to $c_2 = 0.5$, classical GP slightly outperforms the proposed algorithm in terms of robustness.

Table 2: Table of parameter settings of (26) for the three cases in Figure 2

Cases	(1)			(2)			(3)		
l	d	a_l	c_1	d	a_l	c_1	d	a_l	c_1
1	4	0.3	0	4	0.3	0	4	0.3	0
2	4	0.4	0	4	0.4	0.01	4	0.4	0.5

4.2 Example 2

We use the example FES1 from the literature of bi-objective optimization (Fieldsend et al. 2003) to demonstrate the performance of the proposed method. The original problem of FES1 is given below:

$$\text{FES1} = \begin{cases} y_1(\mathbf{x}) = \sum_{i=1}^4 |x_i - \frac{\exp\{\frac{i}{4}\}}{3}|^{0.5}, \\ y_2(\mathbf{x}) = \sum_{i=1}^4 (x_i - 0.5 \cos(\frac{10\pi i}{4}) - 0.5)^2. \end{cases} \quad (27)$$

For a comparison purpose, we construct a modified version of FES1 by taking the exponential values of the original objectives:

$$\text{FES1-modified} = \begin{cases} y_1(\mathbf{x}) = \exp\{\sum_{i=1}^4 |x_i - \frac{\exp\{\frac{i}{4}\}}{3}|^{0.5}\} \\ y_2(\mathbf{x}) = \exp\{\sum_{i=1}^4 (x_i - 0.5 \cos(\frac{10\pi i}{4}) - 0.5)^2\} \end{cases} \quad (28)$$

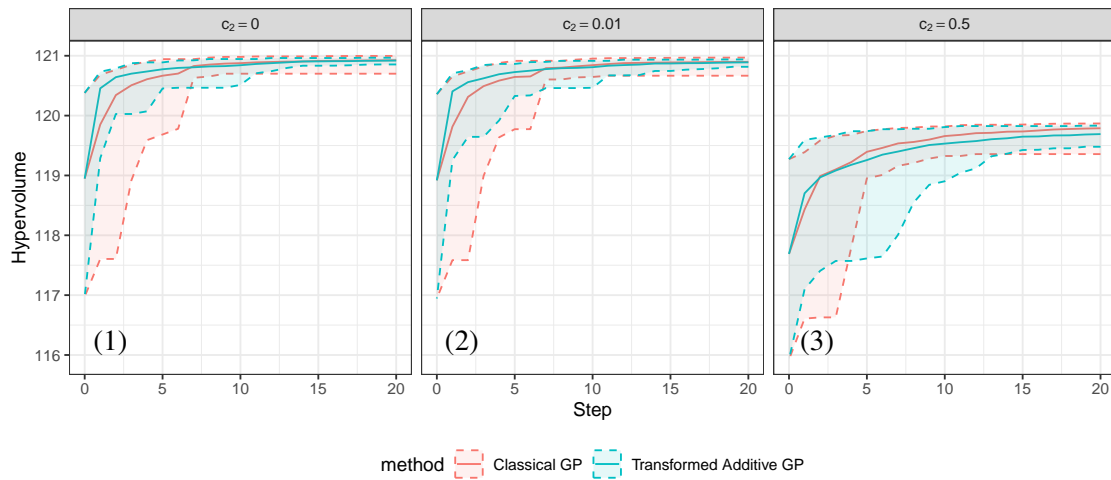


Figure 2: The mean, 10-th and 90-th quantiles of approximated hypervolume over 30 replications and 20 steps for optimizing the objective functions in (26) with the parameter settings in Table 2.

FES1 in (27) is an additive function, whereas FES1-modified in (28) can be transformed into an additive function. For FES1, we restrict λ_l for $l = 1, 2$ in TAG to be between 0.5 and 1.5 since the existing procedure does not estimate it properly. For this example, 20 initial data points are used. The results are depicted in Figure 3. Classical GP is more robust than our method when the function is originally additive without transformation. However, when the additive objective functions are not originally additive (for FES1-modified), our method outperforms the classical GP method in the early stages.

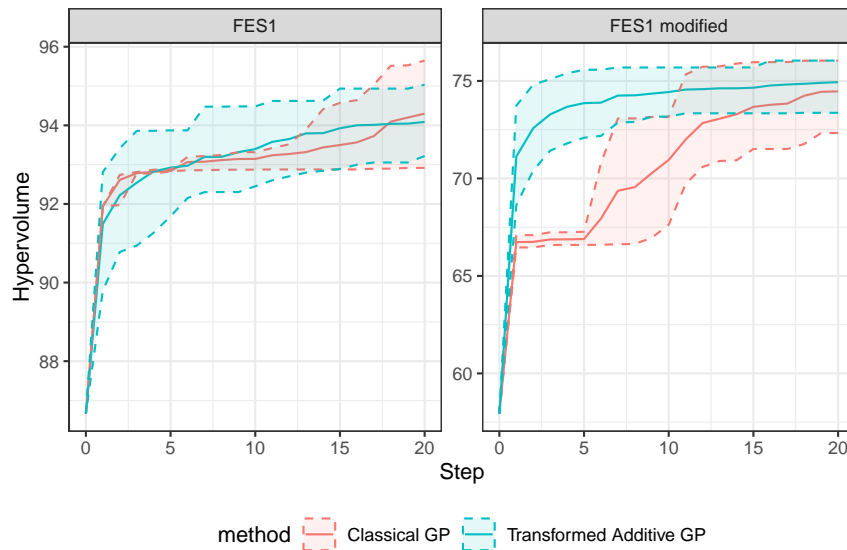


Figure 3: The mean, 10-th and 90-th quantiles of approximated hypervolume over 30 replications and 20 steps for the two problems in (27) and (28).

5 CONCLUSION

We have proposed a method for bi-objective Bayesian optimization using TAG as the surrogate model to simplify the optimization of EHI by decomposing the decision space into one-dimensional additive subspaces. Through a numerical study, we compare the performances of the proposed method with a classical bi-objective Bayesian optimization algorithm under different situations. The proposed algorithm can be further enhanced by improving the parameter tuning under the TAG framework, utilizing better algorithms in maximizing univariate EHI functions and integrating other existing approaches into the proposed algorithm to optimize EHI. Also, motivated by the promising computational results, it can be beneficial to develop theoretical justification of the proposed method. We aim to overcome the limitation of the current work in future and extend the method to the multi-objective case.

ACKNOWLEDGMENTS

This work was supported by Clemson University's Virtual Prototyping of Autonomy Enabled Ground Systems (VIPR-GS), under Cooperative Agreement W56HZV-21-2-0001 with the US Army DEVCOM Ground Vehicle Systems Center (GVSC).

DISCLAIMERS

All opinions, conclusions and findings wherein are those of the authors and may not be those of the affiliated institutions. DISTRIBUTION STATEMENT A. Approved for public release; distribution is unlimited. OPSEC # 8423.

REFERENCES

- Audet, C., J. Bignon, D. Cartier, and S. L. Digabel. 2021. "Performance Indicators in Multiobjective Optimization". *European Journal of Operational Research* 292:397–422.
- Binois, M. and V. Picheny. 2019. "GPareto: An R Package for Gaussian-Process-Based Multi-Objective Optimization and Analysis". *Journal of Statistical Software* 89(8):1–30.
- Bossek, J. 2017. "ecr 2.0: A Modular Framework for Evolutionary Computation in R". In *Proceedings of the Genetic and Evolutionary Computation Conference Companion, GECCO '17*, 1187–1193.
- Box, G. E. P. and D. R. Cox. 1964. "An Analysis of Transformations". *Journal of the Royal Statistical Society. Series B (Methodological)* 26(2):211–252.
- Carnell, R. 2022. *lhs: Latin Hypercube Samples*. R Package Version 1.1.6.
- Daulton, S., M. Balandat, and E. Bakshy. 2020. "Differentiable Expected Hypervolume Improvement for Parallel Multi-Objective Bayesian Optimization". *Advances in Neural Information Processing Systems* 33:9851–9864.
- de Castro, P., H. Stewart, C. Turner, M. Wiecek, G. Hartman, D. Rizzo, , *et al.* 2022. "Decomposition and Coordination to Support Tradespace Analysis for Ground Vehicle Systems". Technical report, SAE Technical Paper.
- Eaton, M. L. 2007. "Multivariate Statistics: A Vector Space Approach". *Lecture Notes-Monograph Series* 53:i–512.
- Ehrgott, M. 2005. *Multicriteria Optimization*, Volume 491. Springer Science & Business Media.
- Emmerich, M., K. Yang, A. Deutz, H. Wang and C. M. Fonseca. 2016. "A multicriteria generalization of Bayesian global optimization". *Advances in stochastic and deterministic global optimization*:229–242.
- Emmerich, M. T., A. H. Deutz, and J. W. Klinkenberg. 2011. "Hypervolume-Based Expected Improvement: Monotonicity Properties and Exact Computation". In *2011 IEEE Congress of Evolutionary Computation (CEC)*, 2147–2154. IEEE.
- Faulkenberg, S. L. and M. M. Wiecek. 2010. "On the Quality of Discrete Representations in Multiple Objective Programming". *Optimization and Engineering* 11:423–440.
- Fieldsend, J. E., R. M. Everson, and S. Singh. 2003. "Using Unconstrained Elite Archives for Multiobjective Optimization". *IEEE Transactions on Evolutionary Computation* 7(3):305–323.
- Fliege, J. and H. Xu. 2011. "Stochastic Multiobjective Optimization: Sample Average Approximation and Applications". *Journal of Optimization Theory and Applications* 151:135–162.
- Fonseca, C. M., L. Paquete, and M. López-Ibáñez. 2006. "An Improved Dimension-Sweep Algorithm for the Hypervolume Indicator". In *2006 IEEE International Conference on Evolutionary Computation*, 1157–1163. IEEE.
- Frazier, P. I. 2018. "A Tutorial on Bayesian Optimization". *arXiv preprint arXiv:1807.02811*.
- Guerreiro, A. P., C. M. Fonseca, and L. Paquete. 2021. "The Hypervolume Indicator: Computational Problems and Algorithms". *ACM Computing Surveys (CSUR)* 54(6):1–42.

- Herzel, A., S. Ruzika, and C. Thielen. 2021. "Approximation Methods for Multiobjective Optimization Problems: A Survey". *INFORMS Journal on Computing* 33(4):1284–1299.
- Hunter, S. R., E. A. Applegate, V. Arora, B. Chong, K. Cooper, O. Rincón-Guevara *et al.* 2019. "An Introduction to Multiobjective Simulation Optimization". *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 29(1):1–36.
- Hunter, S. R. and B. McClosky. 2016. "Maximizing Quantitative Traits in the Mating Design Problem via Simulation-Based Pareto Estimation". *IIE Transactions* 48(6):565–578.
- Hupkens, I., A. H. Deutz, K. Yang, and M. T. Emmerich. 2015. "Faster Exact Algorithms for Computing Expected Hypervolume Improvement." In *EMO (2)*, 65–79.
- Jones, D. R., M. Schonlau, and W. J. Welch. 1998. "Efficient Global Optimization of Expensive Black-Box Functions". *Journal of Global Optimization* 13:455–492.
- Knowles, J. D., D. W. Corne, and M. Fleischer. 2003. "Bounded Archiving Using the Lebesgue Measure". In *The 2003 Congress on Evolutionary Computation, 2003. CEC'03.*, Volume 4, 2490–2497. IEEE.
- Lee, L. H., C. U. Lee, and Y. P. Tan. 2007. "A Multi-Objective Genetic Algorithm for Robust Flight Scheduling Using Simulation". *European Journal of Operational Research* 177(3):1948–1968.
- Lin, L.-H. and V. R. Joseph. 2020. "Transformation and Additivity in Gaussian Processes". *Technometrics* 62(4):525–535.
- Lin, L.-H. and V. R. Joseph. 2021. *TAG: Transformed Additive Gaussian Processes*. R Package Version 0.5.1.
- Qin, C., D. Klabjan, and D. Russo. 2017. "Improving the Expected Improvement Algorithm". *Advances in Neural Information Processing Systems* 30.
- Roustant, O., D. Ginsbourger, and Y. Deville. 2012. "DiceKriging, DiceOptim: Two R Packages for the Analysis of Computer Experiments by Kriging-Based Metamodeling and Optimization". *Journal of Statistical Software* 51(1):1–55.
- Ruzika, S. and M. M. Wiecek. 2005. "Approximation Methods in Multiobjective Programming". *Journal of Optimization Theory and Applications* 126(3):473–501.
- Sacks, J., W. J. Welch, T. J. Mitchell, and H. P. Wynn. 1989. "Design and Analysis of Computer Experiments". *Statistical Science* 4(4):409 – 423.
- Yang, K., A. Deutz, Z. Yang, T. Back and M. Emmerich. 2016. "Truncated Expected Hypervolume Improvement: Exact Computation and Application". In *2016 IEEE Congress on Evolutionary Computation (CEC)*, 4350–4357. IEEE.
- Yang, K., M. Emmerich, A. Deutz, and T. Bäck. 2019. "Efficient computation of expected hypervolume improvement using box decomposition algorithms". *Journal of Global Optimization* 75:3–34.
- Yang, K., M. Emmerich, A. Deutz, and T. Bäck. 2019. "Multi-Objective Bayesian Global Optimization Using Expected Hypervolume Improvement Gradient". *Swarm and Evolutionary Computation* 44:945–956.
- Zhang, H. 2008. "Multi-Objective Simulation-Optimization for Earthmoving Operations". *Automation in Construction* 18(1):79–86.

AUTHOR BIOGRAPHIES

CAROLINE M. KERFONTA obtained her Ph.D in Mathematical Sciences at Clemson University in May 2024. Her research interests include Bayesian Optimization, experimental design, and statistical applications for manufacturing and engineering, particularly aerospace. Her email address is ckerfon@clemson.edu.

QIONG ZHANG is an Associate Professor of Statistics at Clemson University. Her research interests include the interface between information collection and statistical modeling, experimental design, modeling, and uncertainty quantification for computer experiments, statistical design for web experiments, and data analytics for engineering applications. Her email address is qiongz@clemson.edu.

MARGARET M. WIECEK is Professor of Mathematical Sciences in the School of Mathematical and Statistical Sciences and has held joint professorship with the Department of Mechanical Engineering at Clemson University. Her research area includes theory, methodology, and applications of mathematical programming with special interest in multiobjective optimization and decision-making. Part of her work is interdisciplinary since she has introduced new multiobjective optimization concepts and methods into engineering optimization to enrich the field of automotive and structural design. She has served as the President of the MCDM Section of INFORMS. In recognition of her lifetime contributions, she has been awarded the MCDM Gold Medal by the International Society on MCDM. Her email address is wmalgor@clemson.edu.