



HAL
open science

High Dimensional Bayesian Optimization Assisted by Principal Component Analysis

Elena Raponi, Hao Wang, Mariusz Bujny, Simonetta Boria, Carola Doerr

► **To cite this version:**

Elena Raponi, Hao Wang, Mariusz Bujny, Simonetta Boria, Carola Doerr. High Dimensional Bayesian Optimization Assisted by Principal Component Analysis. Parallel Problem Solving from Nature – PPSN XVI (PPSN 2020), Sep 2020, Leiden, Netherlands. pp.169-183, 10.1007/978-3-030-58112-1_12. hal-02935397

HAL Id: hal-02935397

<https://hal.sorbonne-universite.fr/hal-02935397v1>

Submitted on 10 Sep 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

High Dimensional Bayesian Optimization Assisted by Principal Component Analysis

Elena Raponi¹, Hao Wang², Mariusz Bujny³, Simonetta Boria¹, and Carola Doerr²

¹ University of Camerino, Via Madonna delle Carceri 9, 62032 Camerino, Italy elena.raponi@unicam.it

² Sorbonne Université, CNRS, LIP6, Paris, France

³ Honda Research Institute Europe GmbH, 63073 Offenbach am Main, Germany

Abstract. Bayesian Optimization (BO) is a surrogate-assisted global optimization technique that has been successfully applied in various fields, e.g., automated machine learning and design optimization. Built upon a so-called infill-criterion and Gaussian Process regression (GPR), the BO technique suffers from a substantial computational complexity and hampered convergence rate as the dimension of the search spaces increases. Scaling up BO for high-dimensional optimization problems remains a challenging task.

In this paper, we propose to tackle the scalability of BO by hybridizing it with a Principal Component Analysis (PCA), resulting in a novel PCA-assisted BO (PCA-BO) algorithm. Specifically, the PCA procedure learns a linear transformation from all the evaluated points during the run and selects dimensions in the transformed space according to the variability of evaluated points. We then construct the GPR model, and the infill-criterion in the space spanned by the selected dimensions.

We assess the performance of our PCA-BO in terms of the empirical convergence rate and CPU time on multi-modal problems from the COCO benchmark framework. The experimental results show that PCA-BO can effectively reduce the CPU time incurred on high-dimensional problems, and maintains the convergence rate on problems with an adequate global structure. PCA-BO therefore provides a satisfactory trade-off between the convergence rate and computational efficiency opening new ways to benefit from the strength of BO approaches in high dimensional numerical optimization.

Keywords: Bayesian optimization · Black-Box optimization · Principal component analysis · Dimensionality reduction

1 Introduction

Over the last few years, Gaussian Process Regression (GPR) [20] has been proven to be a very flexible and competitive tool in the modeling of functions that are expensive to evaluate or characterized by strong nonlinearities and measurement noises. A Gaussian Process (GP) is a stochastic process, i.e., a family of random variables, such that any finite collection of them have joint Gaussian distributions. GP is used in many application fields (engineering, geology, etc.) in order to evaluate datasets, predict unknown function values, and perform surrogate model-based optimization [7]. Based on a certain number of evaluations of the objective function (training data), surrogate models allow for the construction of computationally cheap-to-evaluate approximations and for replacing the direct optimization of the real objective with the model. As such, many more evaluations can be performed on the approximate model. GP uses a measure of the similarity between points – the kernel function – to predict the value for an unseen point from training data. GP is often chosen because it also provides a theoretical uncertainty quantification of the prediction error, among a variety of kernel-based methods, e.g., Radial Basis Functions [4] and Support Vector Regression [36,9]. Indeed, GP not only provides the estimate of the value of a function (the mean value), but also the variance at each domain point. This variance defines the uncertainty of the model while performing a prediction of the function value on an untested point of the search space. In particular, when GP is used in surrogate

model-based optimization, at each iteration of the optimization algorithm, the estimate of the error is a crucial information in the update phase of the approximation model.

A GP can be used as a prior probability distribution over functions in Bayesian Optimization (BO) [24,25], also known as Efficient Global Optimization (EGO) [16]. BO is a sequential design strategy targeting black-box global optimization problems. It generates an initial set of training points in the search space where the objective function is evaluated. Based on this training data, a GPR model is built to approximate the real objective function. Afterwards, the optimization procedure starts by iteratively choosing new candidate points for which the objective function is evaluated. These points are obtained through the optimization of an acquisition function. However, it has to be noted that the optimization of the acquisition function and the model construction become quite complicated and time-consuming as the dimension of the problem increases, due to the well-known *curse of dimensionality*. In particular, this is the case for most real-world optimization problems, where the number of design variables can range from a few (for simple parameter optimization problems) up to millions (for complex shape [12] or topology optimization [1] problems). In spite of this difficulty, BO techniques have been successfully applied in many fields of engineering, including materials science [35,17], aerospace engineering [15,6,18,21], turbomachinery design [14,2], or even fluid and structural topology optimization [41,27,28,29,22], thanks to a reduction of the problem dimensionality on the representation level. Nevertheless, since the efficiency of these methods decreases strongly with the rising number of design variables, the development of BO approaches able to address high-dimensional problems is crucial for their future applications in industry [37,26].

Different dimensionality reduction techniques using other tools than PCA have been introduced in the literature for the optimization of expensive problems through GPs. Huang et al. [13] proposed a scalable GP for regression by training a neural network with a stacked denoising auto-encoder and performing BO afterwards on the top layer of the trained network. In [40], Wang et al. proposed to embed a lower-dimensional subspace into the original search space, through the so-called random embedding technique. Moreover, Blanchet-Scalliet et al. [3] presented four algorithms to substitute classical kernels, which provide poor predictions of the response in high-dimensional problems. In [8], a kernel-based approach is devised to perform the dimensionality reduction in a feature space for the shape optimization problem in Computer Aided Design (CAD) systems. To the authors' knowledge, the only ones who used PCA to discover hidden features in GP regression are Vivarelli and Williams [38]. However, their strategy mainly allowed for estimating which are the relevant directions in a prescribed feature space, rather than for pursuing the intrinsic optimization process. Finally, Kapsoulis et al. [19] embedded Kernel and Linear PCA in optimization strategies using Evolutionary Algorithms (EAs). Nevertheless, they identify the principal component directions and map the design space to a new feature space only when a prescribed iteration is reached, while in the first iterations PCA is not applied. In this paper, we propose to tackle this problem by using the well-known Principal Component Analysis (PCA) procedure to identify a proper subspace, on which we deploy the BO algorithm. This treatment leads to a novel hybrid surrogate modeling method - the PCA-assisted Bayesian Optimization (PCA-BO), which starts with a Design of Experiments (DoE) [7] and adaptively learns a linear map to reduce the dimensionality by a weighted PCA procedure. Here, the weighting scheme is meant for taking the objective values into account. We assessed the empirical performance of PCA-BO by testing it on the well-known BBOB problem set [11]. Among these problems, we focus on the multi-modal ones, which are most representative of real-world optimization challenges.

2 Bayesian Optimization

In this paper, we consider the minimization of a real-valued objective function $f : S \subseteq \mathbb{R}^D \rightarrow \mathbb{R}$ with simple box constraints, i.e., $S = [\underline{\mathbf{x}}, \bar{\mathbf{x}}]$. Bayesian Optimization [16,25] (BO) starts with sampling an initial DoE of size n_0 : $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n_0}]^\top \subseteq S^{n_0}$. The DoE can be uniform random samples, or obtained through more sophisticated sampling methods [31]. We choose the so-called Optimal Latin Hypercube Sampling (OLHS) [5] in this paper. The corresponding objective function values are denoted as $\mathbf{y} = (f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_{n_0}))^\top$.

Conventionally, a centered Gaussian process prior is assumed on the objective function: $f \sim gp(0, k(\cdot, \cdot))$, where $k: S \times S \rightarrow \mathbb{R}$ is a positive definite function (a.k.a. *kernel*) that computes the autocovariance of the process. Often, a Gaussian likelihood is taken, leading to a conjugate posterior process [30], i.e., $f | \mathbf{y} \sim gp(\hat{f}(\cdot), k'(\cdot, \cdot))$, where \hat{f} and k' are the posterior mean and covariance function, respectively. On an unknown point \mathbf{x} , $\hat{f}(\mathbf{x})$ yields the maximum a posteriori (MAP) estimate of $f(\mathbf{x})$ whereas $\hat{s}^2(\mathbf{x}) := k'(\mathbf{x}, \mathbf{x})$ quantifies the uncertainty of this estimation. We shall refer to the posterior process as the Gaussian process regression (GPR) model. Based on the posterior process, we usually identify promising points via the so-called *infill-criterion* which balances \hat{f} with \hat{s}^2 (exploitation vs. exploration). A variety of infill-criteria has been proposed in the literature, e.g., Probability of Improvement [7,25], Expected Improvement [7], and the Moment-Generating Function of Improvement [39]. In this work, we adopt the Penalized Expected Improvement (PEI) criterion [28] to handle the box constraints,

$$\text{PEI}(\mathbf{x}) = \begin{cases} E\{\max\{0, \min \mathbf{y} - f(\mathbf{x})\} | \mathbf{y}\} & \text{if } \mathbf{x} \text{ is feasible,} \\ -P(\mathbf{x}) & \text{if } \mathbf{x} \text{ is infeasible,} \end{cases} \quad (1)$$

which essentially calculates the expected improvement when \mathbf{x} is feasible, and returns a negative penalty value otherwise. The penalty function $P(\mathbf{x})$ is proportional to the degree of infeasibility of the point \mathbf{x} . This will be further described in Sec. 3. BO then proceeds to select the next point by maximizing PEI, namely, $\mathbf{x}^* = \arg \max_{\mathbf{x} \in S} \text{PEI}(\mathbf{x})$. After evaluating \mathbf{x}^* , BO augments⁴ the data set, $\mathbf{X} \leftarrow \mathbf{X} \cup \{\mathbf{x}^*\}$, $\mathbf{y} \leftarrow \mathbf{y} \cup \{f(\mathbf{x}^*)\}$, on which the GPR model is retrained.

3 PCA-assisted Bayesian Optimization

In this section, we introduce the coupling of the principal component analysis procedure and Bayesian optimization method, called **PCA-BO**. Loosely speaking, this approach learns a linear transformation (from the initial design points of BO) that would, by design, identify directions (a.k.a. principal components) in \mathbb{R}^D along which the objective value changes rapidly. Also, this approach adapts this transformation through the optimization process. Intuitively, we could further drop the components to which objective function values are less sensitive, resulting in a lower-dimensional search space \mathbb{R}^r ($r < D$). Particularly for BO, the surrogate model shall be trained in this lower-dimensional search space, thus benefiting from the dimensionality reduction. Also, the infill-criterion optimization, which is typically an expensive sub-procedure in BO, becomes less costly because it now operates in the lower-dimensional space as well.

The PCA-BO algorithm. We first present a high-level overview of the proposed PCA-BO algorithm. A graphical representation is provided in Fig. 1, whereas a pseudo-code description is presented in Alg. 1. PCA-BO works as follows:

1. We perform an initial DoE in the original search space, generating as an outcome a set of evenly distributed points.
2. We design a weighting scheme to embed the information from objective function into the DoE points, where smaller weights are assigned to the points with worse function values.
3. We apply a PCA procedure to obtain a linear map from the original search space to a lower-dimensional space, using the weighted DoE points.
4. In the lower-dimensional space, we train a GPR model and maximize an infill-criterion to find a promising candidate point.

⁴ With an abuse of terminology, the operation $\mathbf{X} \cup \{\mathbf{x}^*\}$ is understood as appending \mathbf{x}^* at the bottom row of \mathbf{X} throughout this paper. $\mathbf{y} \cup \{f(\mathbf{x}^*)\}$ is defined similarly.

5. We map the candidate point back to the original search space, and evaluate it with the objective function.
6. We augment the data set with the candidate point and its objective value, and then proceed to step 2 for the next iteration.

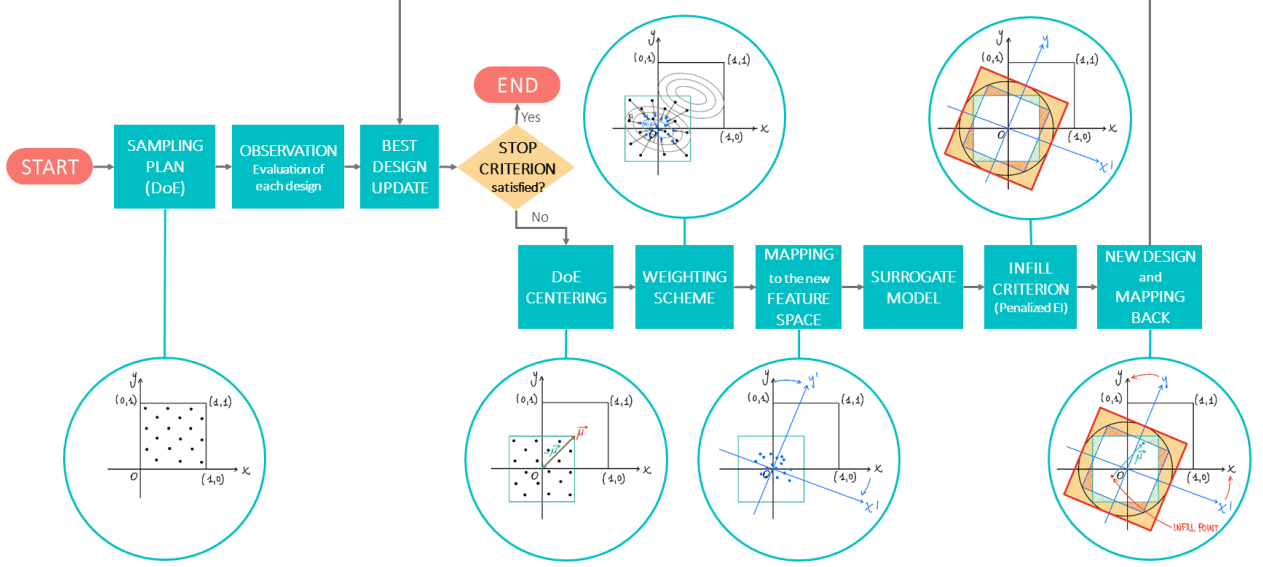


Fig. 1: Flowchart of the PCA-BO optimization algorithm in which PCA-related operations are depicted in circles.

Algorithm 1 PCA-assisted Bayesian Optimization

- 1: **procedure** PCA-BO(f, \mathcal{A}, S) $\triangleright f$: objective function, \mathcal{A} : infill-criterion, S : search space, S' : lower-dimensional search space
 - 2: Create $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n_0}]^\top \subset S^{n_0}$ with optimal Latin hypercube sampling
 - 3: $\mathbf{y} \leftarrow (f(\mathbf{x}_1), f(\mathbf{x}_2), \dots, f(\mathbf{x}_{n_0}))^\top$, $n \leftarrow n_0$
 - 4: **while** the stop criteria are not fulfilled **do**
 - 5: Centering: $\bar{\mathbf{X}} \leftarrow \mathbf{X} - \mathbf{1}_n \boldsymbol{\mu}$
 - 6: Rescaling: $\mathbf{X}' \leftarrow \bar{\mathbf{X}} \mathbf{W}$
 - 7: $\boldsymbol{\mu}', \mathbf{P}_r \leftarrow \text{PCA}(\mathbf{X}')$
 - 8: Mapping to \mathbb{R}^r : $\mathbf{Z}_r \leftarrow \mathbf{P}_r (\bar{\mathbf{X}} - \mathbf{1}_n \boldsymbol{\mu}')^\top$
 - 9: GPR training: $\hat{f}, \hat{s}^2 \leftarrow \text{GPR}(\mathbf{Z}_r, \mathbf{y})$
 - 10: $\mathbf{z}' \leftarrow \arg \max_{\mathbf{z} \in S'} \mathcal{A}(\mathbf{z}; \hat{f}, \hat{s}^2)$
 - 11: Mapping to \mathbb{R}^D : $\mathbf{x}' \leftarrow \mathbf{P}_r^\top \mathbf{z} + \boldsymbol{\mu}' + \boldsymbol{\mu}$
 - 12: $y' \leftarrow f(\mathbf{x}')$
 - 13: $\mathbf{X} \leftarrow \mathbf{X} \cup \{\mathbf{x}'\}$, $\mathbf{y} \leftarrow (\mathbf{y}^\top, y')^\top$, $n \leftarrow n + 1$
 - 14: **end while**
 - 15: **end procedure**
-

Rescaling data points for PCA. Because the goal of applying PCA, as stated above, is to spot directions in \mathbb{R}^D to which the function value is sensitive, it is then necessary to take into account the information on the objective function. We implement this consideration by rescaling the data points according to a weighting scheme, which depends on the corresponding function values. The weighting scheme aims to adjust the “importance” of data points when applying the PCA procedure, such that a point associated with a better (here lower) function value is assigned with a larger weight. We propose to use a rank-based weighting

scheme, in which the weight assigned to a point is solely determined by the rank thereof. Let $\{r_i\}_{i=1}^n$ be the ranking of the points in \mathbf{X} according to their function values \mathbf{y} . We calculate the rank-based (pre)-weights as follows:

$$\tilde{w}_i = \ln n - \ln r_i, \quad (2)$$

where n stands for the current number of data points. Afterwards, we normalize the pre-weights: $w_i = \tilde{w}_i / \sum \tilde{w}_i$. Notably, this weighting scheme will gradually decrease the weights of worse points when a better point is found and added to the data, hence leading to a self-adjusting discount factor on the data set.⁵ It is worth pointing out that this weighting scheme resembles the weight calculation in the well-known Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [10], although they are motivated differently. For ease of discussion, we will use a weight matrix $\mathbf{W} = \text{diag}(w_1, w_2, \dots, w_n)$ henceforth. We rescale the original data \mathbf{X} by removing its sample mean, i.e., $\bar{\mathbf{X}} = \mathbf{X} - \mathbf{1}_n \boldsymbol{\mu}$, where $\boldsymbol{\mu} = (\frac{1}{n} \sum_{i=1}^n X_{i1}, \dots, \frac{1}{n} \sum_{i=1}^n X_{id})$ and $\mathbf{1}_n = (1, \dots, 1)^\top$ (n 1's), and then adjust each point with the corresponding weight, namely, $\mathbf{X}' = \bar{\mathbf{X}} \mathbf{W}$.

Dimensionality reduction. The PCA procedure starts with computing the sample mean of data matrix \mathbf{X}' , i.e., $\boldsymbol{\mu}' = (\frac{1}{n} \sum_{i=1}^n X'_{i1}, \dots, \frac{1}{n} \sum_{i=1}^n X'_{id})$, after which it calculates the unbiased sample covariance matrix $\mathbf{C} = \frac{1}{n-1} \mathbf{X}''^\top \mathbf{X}''$, $\mathbf{X}'' = \mathbf{X}' - \mathbf{1}_n \boldsymbol{\mu}'$. The principle components (PCs) are computed via the eigendecomposition $\mathbf{C} = \mathbf{P} \mathbf{D} \mathbf{P}^{-1}$, where column vectors (a.k.a. eigenvectors) of \mathbf{P} identify the PCs, and the diagonal matrix $\mathbf{D} = \text{diag}(\sigma_1^2, \dots, \sigma_d^2)$ contains the variance of \mathbf{X}'' along each principal component, i.e., $\sigma_i^2 = \text{var}\{\mathbf{X}'' \mathbf{p}_i\}$ (\mathbf{p}_i is the i^{th} column of \mathbf{P}). For clarity of discussion, we consider the orthonormal matrix \mathbf{P} that defines a rotation operation, as a change of the standard basis in \mathbb{R}^D such that PCs are taken as the basis after the transformation. After changing the basis, we proceed to reduce the dimensionality of \mathbb{R}^D by sorting the PCs according to the decreasing order of their variances and keeping the first r components. Here r is chosen as the smallest integer such that the top- r variances sum up to at least α percent of the total variability of the data (we use $\alpha = 95\%$ in our experiments). By denoting the ordering of variances and PCs as $\sigma_{1:d}^2, \sigma_{2:d}^2, \dots, \sigma_{d:d}^2$ and $\mathbf{p}_{1:d}, \mathbf{p}_{2:d}, \dots, \mathbf{p}_{d:d}$, respectively, we formulate the selection of PCs as follows:

$$\mathbf{P}_r := [\mathbf{p}_{1:d}, \mathbf{p}_{2:d}, \dots, \mathbf{p}_{r:d}]^\top \in \mathbb{R}^{r \times d}, \quad r = \inf \left\{ k \in [1..d]: \sum_{i=1}^k \sigma_{i:d}^2 \geq \alpha \sum_{i=1}^d \sigma_i^2 \right\}.$$

Now, we could transform the data set \mathbf{X} into a lower-dimensional space, i.e., $\mathbf{Z}_r = \mathbf{P}_r (\bar{\mathbf{X}} - \mathbf{1}_n \boldsymbol{\mu}')^\top$. Note that, 1) in this transformation, the centered data set $\bar{\mathbf{X}}$ is not scaled by the weights because we only intend to incorporate the information on the objective values when determining \mathbf{P}_r . Using the scaled matrix $\bar{\mathbf{X}} \mathbf{W}$ will make it cumbersome to define the inverse mapping (see below). 2) It is crucial to subtract $\boldsymbol{\mu}'$ from $\bar{\mathbf{X}}$ since it estimates the center of ellipsoidal contours when f resembles a quadratic function globally. In this way, the principal axis of contours will be parallel to the PCs after applying \mathbf{P}_r . 3) Here, $\mathbf{P}_r: \mathbb{R}^D \rightarrow \mathbb{R}^r$ is not injective ($\dim(\ker \mathbf{P}_r) = D - r$). Therefore when transforming the centered data set to the lower-dimensional space, it seems that we might lose some data points in $\bar{\mathbf{X}}$. However, we argue that such an event $\{\mathbf{x} - \mathbf{x}' \in \ker \mathbf{P}_r: \mathbf{x}, \mathbf{x}' \in \mathbf{X}\}$ is a null set with respect to any non-singular measure on \mathbb{R}^D , and the probability distribution of \mathbf{X} , no matter which forms it takes, should not be singular (otherwise BO will only search in a subspace of \mathbb{R}^D). Thus, it is safe to ignore the possibility of losing data points. Now we are ready to train a GPR model on $(\mathbf{Z}_r, \mathbf{y})$, and propose candidate points by optimizing an infill-criterion in \mathbb{R}^r . Importantly, to evaluate a candidate point $\mathbf{z} \in \mathbb{R}^r$, we map it back to \mathbb{R}^D on which the objective function takes its domain, using the following linear map L :

$$L: \mathbb{R}^r \rightarrow \mathbb{R}^D, \quad \mathbf{z} \mapsto \mathbf{P}_r^\top \mathbf{z} + \boldsymbol{\mu}' + \boldsymbol{\mu}.$$

⁵ Alternatively, the weight can also be computed directly from the function value, e.g., through a parameterized hyperbolic function. However, we do not prefer this approach since it introduces extra parameters that require tuning, and does not possess the discount effect of the rank-based scheme since the weights remain static throughout the optimization.

Note that, 1) the linear map L is an injection; 2) it is necessary to include $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$ in this transformation since we subtract them from the design matrix \mathbf{X} when mapping it to the lower-dimensional space; 3) After evaluating point \mathbf{z} , i.e., $y = f(L(\mathbf{z}))$, we augment the data set in \mathbb{R}^D : $\mathbf{X} \cup \{L(\mathbf{z})\}$ and $\mathbf{y} \cup \{y\}$, from which the linear map \mathbf{P}_r is re-computed using the procedure described above; 4) After re-computing \mathbf{P}_r , we have to re-calculate \mathbf{Z}_r by mapping the augmented \mathbf{X} into \mathbb{R}^r : $\mathbf{Z}_r = \mathbf{P}_r(\bar{\mathbf{X}} - \mathbf{1}_n \boldsymbol{\mu}')^\top$, and retrain the GPR model on $(\mathbf{Z}_r, \mathbf{y})$.

Penalized infill-criterion. Once the regression model is fitted, we use it to construct and maximize the acquisition function that determines the next infill point. As previously stated, in this work we use a penalized version of the well-known EI, referred to as PEI. By definition, we search for the maximum of PEI through DE over a prescribed cubic domain C , which is orientated according to the axes of the transformed coordinate system and must contain the projection of the PCA-transform of the old domain S onto the lower-dimensional feature space. We hence define and penalize a bounding cube C in order to prevent 1) that the optimum is found in a region of C that falls outside S when mapped back, and 2) that we neglect some good regions for new query points, i.e., the regions that belong S but not to C . To achieve this, if ρ is the radius of the sphere tangent to S , we define $C = [C_1 - \rho, C_1 + \rho] \times \dots \times [C_r - \rho, C_r + \rho]$, where $C_i, i = 1, \dots, r$ is the i^{th} component of the center of the new feature space. Since, for simplicity reasons, C is wider than necessary, we introduce a penalty to be assigned to the points that would not fall into S when mapped back through the inverse PCA linear transformation – the *infeasible points*. In order to avoid a stagnation of the DE search, the penalty P assigned to infeasible points is defined as the additive inverse of the distance of their images through the PCA inverse transformation from the boundary of S . As a result, the infeasible points are automatically discarded in the maximization of PEI (Eq. (1)).

4 Experiments

Experimental Setup. We assess the performance of PCA-BO on ten multi-modal functions taken from the the BBOB problem set [11], and compare the experimental result to a standard BO. This choice of test functions aims to verify the applicability of PCA-BO to more difficult problems, which potentially resemble the feature of real-world applications. We test PCA-BO on three different dimensions $D \in \{10, 20, 40\}$ with the following budget of $10D + 50$ function evaluations, of which the initial DoE takes up to 20%. To allow for a statistically meaningful comparison, we performed 30 independent runs for each combination of algorithm, function, and dimension, on an Intel(R) Xeon(R) CPU E5620 @ 2.40GHz machine in single-threaded mode. Also, we recorded the best-so-far function value and measured CPU time taken by those algorithms, which are considered as performance metrics in this paper. We select a squared-exponential kernel function for the GPR model that underpins both PCA-BO and BO. A Differential Evolution (DE) [34] algorithm⁶ is adopted to maximize the PEI criterion in every iteration. We implement both algorithms in Python, where the PyDoE, PyOpt, and Scikit-learn open-source packages are used to handle the DoE, modeling, and PCA-transformation phases, respectively.

Results. In this section, we illustrate the results we obtained in the optimization of multi-modal benchmark functions from the BBOB suite by using PCA-BO and the state-of-the-art BO. In particular, each algorithm is tested on functions **F15-F24**, which can be categorized in two main groups: 1) multi-modal functions with adequate global structure: F15-19, and 2) multi-modal functions with weak global structure: F20-24.

⁶ We take the Scipy implementation of DE (https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.differential_evolution.html) with a population size of $20r$ and the “best1bin” strategy, which uses the binary crossover and calculates the differential vector based on the current best point. Here, we set the evaluation budget to $20020r^2$ to optimize the infill-criterion.

By keeping 95% of the variability of data points, our PCA-assisted approach, on average, reduces about 40% dimensions from the original search space across all problems and dimensions. Also, we compare the two algorithms in terms of the empirical convergence and CPU time. In Fig. 2, we report the mean target precision achieved on each test function, which is the difference between the best-so-far function value and the optimal function value averaged over 30 runs with the same initial DoE points. Also, we plot the mean target precision against the iteration number that counts the number of function evaluations after the DoE phase and illustrate the 95% confidence interval of the mean target precision.⁷ On each dimension, the plots clearly illustrate that PCA-BO is either superior or at least comparable to the standard BO when the test function exhibits an adequate global structure (**F15-19**). For each function in this group, PCA-BO reveals a steep convergence at the beginning of the optimization procedure. Notably, such an early and steep convergence is extremely valuable when we apply it in real-world scenarios, e.g., industrial optimizations, where we could only afford a very limited evaluation budget for expensive optimization problems. In contrast, on multi-modal functions with weak global structure (**F20-24**), PCA-BO has difficulties in maintaining the convergence speed, particularly on functions F21 and F22 whose global contour lines are composed of multiple local ellipsoids that have different orientations (see contour plots in [11]). We speculate that the poor performance of PCA-BO on those two functions is because we apply the PCA procedure on a global scale, which could be potentially confused by the landscape of F21 and F22. However, we argue that this issue could be alleviated effectively with a simple extension of applying the PCA procedure locally, which can be realized by either searching for the optimum with the trust region approach [42], or through clustering techniques [33]. We shall investigate this issue in future work. Moreover, we performed the well-known Wilcoxon rank-sum test on the target precision values reached by each algorithm upon termination in each dimension. Under the null hypothesis of this test, the target precision values has an equal distribution in BO and PCA-BO [23]. Based on the data, we reject the null hypothesis at the 5% significance level on all functions and dimensions except F23, which belongs to the category of multi-modal functions with weak global structure. Therefore, we argue that PCA-BO converges at least as fast as the standard BO algorithm except on problems **F20-F24**, where the problem landscape exhibits high local irregularities and lacks a global structure. This scenario will be addressed by the authors in future works. More importantly, the PCA-BO algorithm achieves this convergence speed with a significantly less computational power, which is supported by measured CPU times in Fig. 3. Here, we observe that the CPU time of PCA-BO, on average, is about 1.2209 times that of BO on 10D. Although this ratio does not imply any benefits from PCA-BO on 10D, it, however, decreases to 0.7519 for 20D, and further to 0.5910 for 40D. Therefore, PCA-BO can effectively reduce the CPU time incurred on higher-dimensional problems, leading to a 25% and a 41% reduction in the 20- and 40-variables tests, respectively.

5 Conclusions and future research

In this paper, we presented the novel PCA-assisted Bayesian Optimization (PCA-BO) method, which aims to efficiently address high-dimensional black-box optimization problems by means of surrogate modeling techniques. In fact, despite BO has been successfully applied in various fields, it demonstrated to be restricted to problems of moderate dimension, typically up to about 15. Therefore, since many real-world applications require the solution of very high-dimensional optimization problems, it becomes necessary to scale BO to more complex parameter spaces, e.g., through dimensionality reduction techniques. Essentially, PCA-BO allows for achieving this by performing iteratively an orthogonal linear transformation of the original data-set, in order to map it to a new feature space where the problem becomes *more separable* and features are selected in order to inherit the maximum possible variance from the original sample set. Once the transformation is applied, the model construction and the optimization of the acquisition function – i.e., the costly steps characterizing the BO sequential strategy – are performed in the lower-dimensional space so that we can

⁷ On the 10-dimensional **F20** problem, we observed that the standard deviation of BO over 30 runs gradually shrinks to zero after 50 iterations, making the confidence interval disappear in the corresponding subplot.

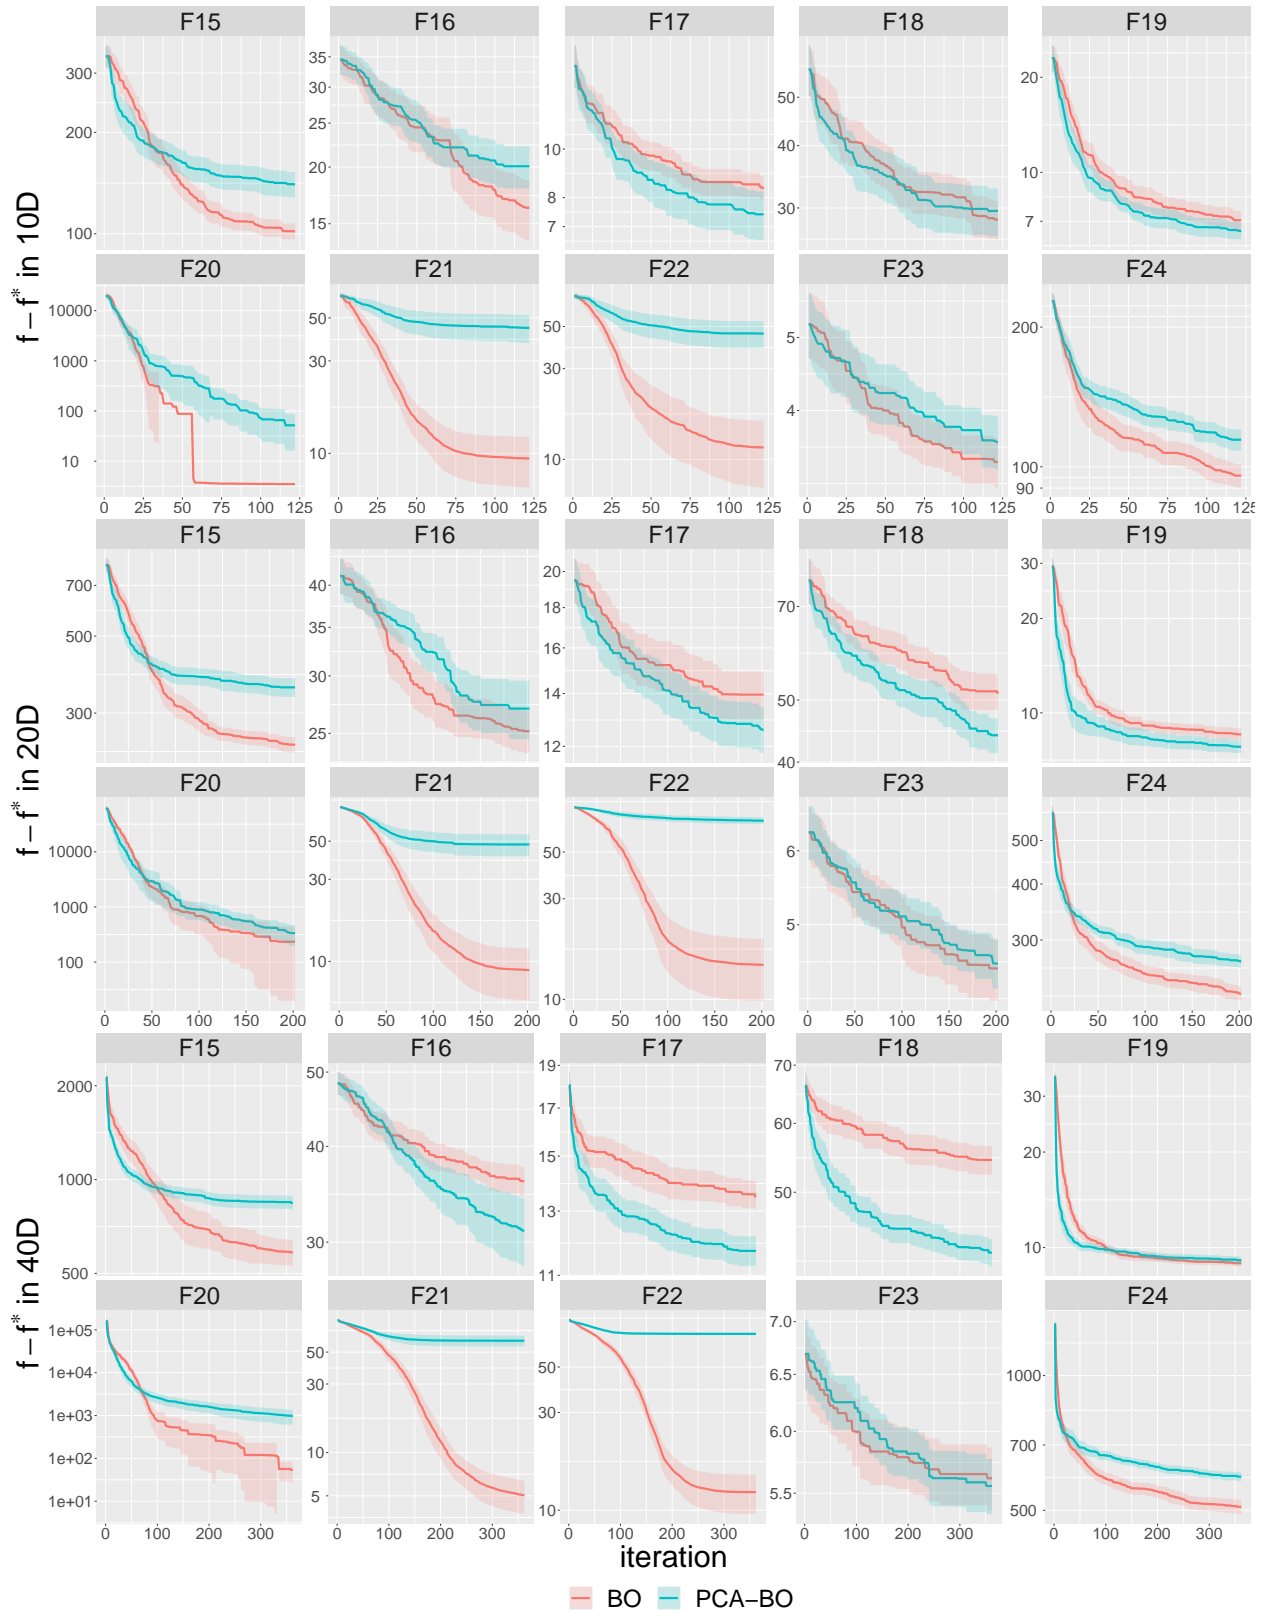


Fig. 2: Mean target precision to the optimal value against the iteration for BO (red) and PCA-BO (blue), on multi-modal functions F15-24 from the BBOB problem set, in three dimensionalities, 10D (top), 20D (middle), and 40D (bottom). 30 independent runs are conducted and the shade area indicates the 95% confidence interval of the mean target precision. The iteration counts the number of function evaluations after evaluating the initial design points.

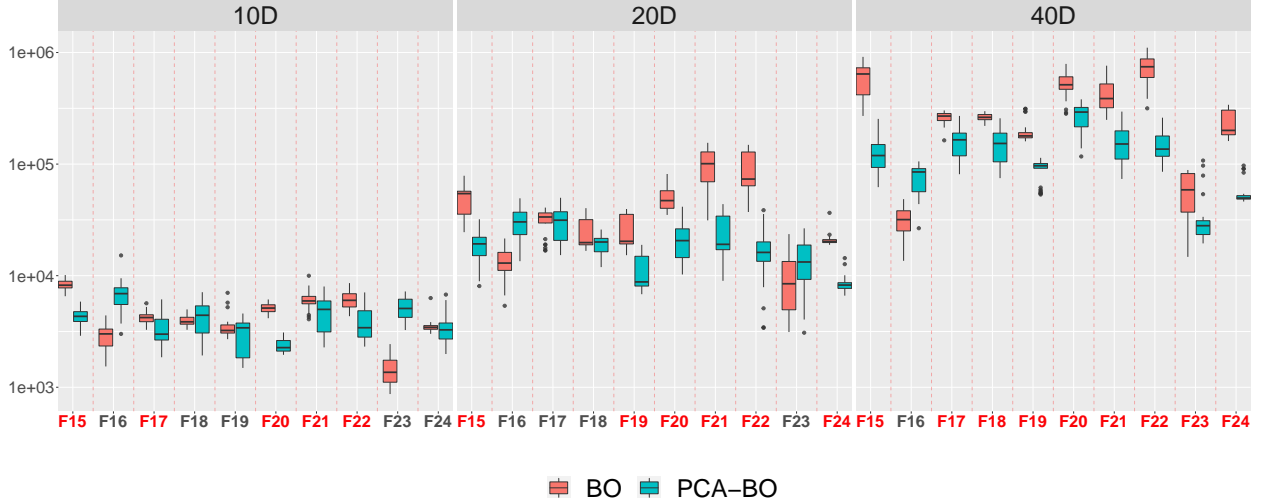


Fig. 3: CPU time (in seconds) taken by PCA-BO and BO to complete the evaluation budget for all functions and three dimensions. For each pair of problems and dimensions, we also applied the well-known Mann–Whitney U test on the CPU time with the Holm–Bonferroni correction of p -values and a significance level of 0.01. The red-colored test problems indicate that the CPU time of PCA-BO is significantly smaller than BO based on the test.

find a promising point by saving computational resources. Afterwards, the selected point is mapped back to the original feature space, it is evaluated, added to the sample set to increase the accuracy of the model, and compared to the previous points to eventually update the optimum.

We empirically evaluated PCA-BO by testing it on the well-known COCO/BBOB benchmark suite. In particular, we used the presented method to optimize two sets of multi-modal benchmarks, characterized by either an adequate or weak global structure. We also ran the standard BO for comparison purposes. It is clear from the results that PCA-BO achieves state-of-the-art BO performance for optimizing multi-modal benchmarks with an adequate global structure, while it finds some difficulties when dealing with multi-modal benchmarks with a weak global structure. We ascribe such inefficiency to the global nature of the PCA-BO algorithm. In fact, when learning the orthogonal linear transformation from all the evaluated points during the run, the proposed algorithm cannot detect a clear and well-definite trend of the isocontour within the function domain. Addressing this problem, by localizing the PCA strategy, is among our future goals. Moreover, in case of weak global structure benchmarks, PCA-BO might also benefit from a different scheme in the pre-tuning of the PCA algorithm: instead of weighting the sample set, a truncation scheme which only selects a prescribed percentage of samples – ordered according to their function values – to train the linear transformation parameters can be adopted. Within this study, we have assessed the performance of PCA-BO not only in terms of the empirical convergence rate, but also tracking the CPU time elapsed to complete the runs. Results showed that, except for the 10-variables test cases that can be efficiently addressed with standard BO, the average CPU time throughout the different tested functions can be consistently reduced by approaching the optimization problem with PCA-BO. In fact, training the GP model and optimizing the acquisition function is a much cheaper procedure in a search space of reduced dimensionality than in the full-dimensional space.

In the future, we will extend our study by applying a non-linear transformation when tuning the parameters for the PCA-mapping to the lower-dimensional feature space. One option is represented by the Kernel Principal Component Analysis (KPCA) [32], which is a non-linear dimensionality reduction through the use of kernels, more capable to capture the isocontour of the cost functions when variables are not linearly correlated. Moreover, we will extend PCA-BO to make it capable of addressing real-world constrained applications, such as the optimization of mechanical structures subjected to static/dynamic loads.

These problems are usually characterized by a great number of design variables and by costly evaluations, requiring optimization techniques of affordable computational cost.

Acknowledgments. Our work was supported by the Paris Ile-de-France Region and by COST Action CA15140 “Improving Applicability of Nature-Inspired Optimisation by Joining Theory and Practice (ImAppNIO)”.

References

1. Aage, N., Andreassen, E., Lazarov, B.S.: Topology optimization using PETSc: An easy-to-use, fully parallel, open source topology optimization framework. *Structural and Multidisciplinary Optimization* **51**(3), 565–572 (2015). <https://doi.org/10.1007/s00158-014-1157-0>
2. Arsenyev, I., Duddeck, F., Fischersworing-Bunk, A.: Adaptive surrogate-based multi-disciplinary optimization for vane clusters. In: *ASME Turbo Expo 2015: Turbine Technical Conference and Exposition*. American Society of Mechanical Engineers Digital Collection (2015). <https://doi.org/10.1115/GT2015-42164>
3. Blanchet-Scalliet, C., Helbert, C., Ribaud, M., Vial, C.: Four algorithms to construct a sparse kriging kernel for dimensionality reduction. *Computational Statistics* **34**(4), 1889–1909 (2019). <https://doi.org/10.1007/s00180-019-00874-2>
4. Broomhead, D.S.: Radial Basis Functions, Multi-variable Functional Interpolation and Adaptive Networks. *Royals Signals & Radar Establishment* (1988)
5. Fang, K.T., Li, R., Sudjianto, A.: Design and Modeling for Computer Experiments. CRC Press (2005)
6. Forrester, A.I., Bressloff, N.W., Keane, A.J.: Optimization using surrogate models and partially converged computational fluid dynamics simulations. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **462**(2071), 2177–2204 (2006). <https://doi.org/10.1098/rspa.2006.1679>
7. Forrester, A.I.J., Sóbester, A., Keane, A.J.: Engineering Design via Surrogate Modelling - A Practical Guide. John Wiley & Sons Ltd. (2008). <https://doi.org/10.1002/9780470770801>
8. Gaudrie, D., Le Riche, R., Picheny, V., Enaux, B., Herbert, V.: From CAD to Eigenshapes for Surrogate-based Optimization. In: *13th World Congress of Structural and Multidisciplinary Optimization*. Beijing, China (2019), <https://hal.archives-ouvertes.fr/hal-02142492>
9. Gunn, S.R.: Support Vector Machines for Classification and Regression (1998), Technical Report
10. Hansen, N.: The CMA Evolution Strategy: A Comparing Review, pp. 75–102. Springer Berlin Heidelberg, Berlin, Heidelberg (2006). https://doi.org/10.1007/3-540-32494-1_4
11. Hansen, N., Brockhoff, D., Mersmann, O., Tusar, T., Tusar, D., ElHara, O.A., Sampaio, P.R., Atamna, A., Varelas, K., Batu, U., Nguyen, D.M., Matzner, F., Auger, A.: COMparing Continuous Optimizers: numbbo/COCO on Github (2019). <https://doi.org/10.5281/zenodo.2594848>
12. Hojjat, M., Stavropoulou, E., Bletzinger, K.U.: The vertex morphing method for node-based shape optimization. *Computer Methods in Applied Mechanics and Engineering* **268**, 494–513 (2014). <https://doi.org/10.1016/j.cma.2013.10.015>
13. Huang, W., Zhao, D., Sun, F., Liu, H., Chang, E.: Scalable Gaussian process regression using deep neural networks. In: *Proceedings of the 24th International Conference on Artificial Intelligence*. pp. 3576–3582. IJCAI’15, AAAI Press, Buenos Aires, Argentina (2015)
14. Huang, Z., Wang, C., Chen, J., Tian, H.: Optimal design of aeroengine turbine disc based on kriging surrogate models. *Computers & structures* **89**(1-2), 27–37 (2011). <https://doi.org/10.1016/j.compstruc.2010.07.010>
15. Jeong, S., Murayama, M., Yamamoto, K.: Efficient optimization design method using kriging model. *Journal of aircraft* **42**(2), 413–420 (2005). <https://doi.org/10.1023/A:1008306431147>
16. Jones, D.R., Schonlau, M., Welch, W.J.: Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization* **13**(4), 455–492 (1998). <https://doi.org/10.1023/A:1008306431147>
17. Ju, S., Shiga, T., Feng, L., Hou, Z., Tsuda, K., Shiomi, J.: Designing nanostructures for phonon transport via bayesian optimization. *Physical Review X* **7**(2), 021024 (2017). <https://doi.org/10.1103/PhysRevX.7.021024>
18. Kanazaki, M., Takagi, H., Makino, Y.: Mixed-fidelity efficient global optimization applied to design of supersonic wing. *Procedia Engineering* **67**(1), 85–99 (2013). <https://doi.org/10.1016/j.proeng.2013.12.008>
19. Kapsoulis, D., Tsiakas, K., Asouti, V., Giannakoglou, K.: The use of Kernel PCA in evolutionary optimization for computationally demanding engineering applications. In: *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*. pp. 1–8 (2016). <https://doi.org/10.1109/SSCI.2016.7850203>
20. Kleijnen, J.P.C.: Kriging metamodeling in simulation: A review. *European Journal of Operational Research* **192**(3), 707–716 (2009). <https://doi.org/10.1016/j.ejor.2007.10.013>

21. Lam, R., Poloczek, M., Frazier, P., Willcox, K.E.: Advances in bayesian optimization with applications in aerospace engineering. In: 2018 AIAA Non-Deterministic Approaches Conference. p. 1656 (2018). <https://doi.org/10.2514/6.2018-1656>
22. Liu, K., Detwiler, D., Tovar, A.: Metamodel-Based Global Optimization of Vehicle Structures for Crashworthiness Supported by Clustering Methods. In: Advances in Structural and Multidisciplinary Optimization, pp. 1545–1557. Springer (2017). https://doi.org/10.1007/978-3-319-67988-4_116
23. MATLAB: version 9.5.0.944444 (R2018b). The MathWorks Inc., Natick, Massachusetts (2018)
24. Mockus, J.: On Bayesian Methods for Seeking the Extremum. In: Marchuk, G.I. (ed.) Optimization Techniques, IFIP Technical Conference, Novosibirsk, USSR, July 1-7, 1974. Lecture Notes in Computer Science, vol. 27, pp. 400–404. Springer (1974). https://doi.org/10.1007/3-540-07165-2_55
25. Mockus, J.: Bayesian approach to global optimization: theory and applications, vol. 37. Springer Science & Business Media (2012). <https://doi.org/10.1007/978-94-009-0909-0>
26. Palar, P.S., Liem, R.P., Zuhail, L.R., Shimoyama, K.: On the use of surrogate models in engineering design optimization and exploration: The key issues. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion. pp. 1592–1602 (2019). <https://doi.org/10.1145/3319619.3326813>
27. Raponi, E., Bujny, M., Olhofer, M., Aulig, N., Boria, S., Duddeck, F.: Kriging-guided level set method for crash topology optimization. In: 7th GACM Colloquium on Computational Mechanics for Young Scientists from Academia and Industry. GACM, Stuttgart, Germany (2017)
28. Raponi, E., Bujny, M., Olhofer, M., Aulig, N., Boria, S., Duddeck, F.: Kriging-assisted topology optimization of crash structures. *Computer Methods in Applied Mechanics and Engineering* **348**, 730–752 (2019). <https://doi.org/10.1016/j.cma.2019.02.002>
29. Raponi, E., Bujny, M., Olhofer, M., Boria, S., Duddeck, F.: Hybrid Kriging-assisted Level Set Method for Structural Topology Optimization. In: 11th International Conference on Evolutionary Computation Theory and Applications. Vienna, Austria (2019). <https://doi.org/10.5220/0008067800700081>
30. Rasmussen, C., Williams, C.: Gaussian Processes for Machine Learning. Adaptive computation and machine learning series, University Press Group Limited (2006)
31. Santner, T.J., Williams, B.J., Notz, W.I.: The Design and Analysis of Computer Experiments. Springer series in statistics, Springer (2003). <https://doi.org/10.1007/978-1-4757-3799-8>
32. Schölkopf, B., Smola, A., Müller, K.R.: Nonlinear Component Analysis as a Kernel Eigenvalue Problem. *Neural Computation* **10**(5), 1299–1319 (1998). <https://doi.org/10.1162/089976698300017467>
33. van Stein, B., Wang, H., Kowalczyk, W., Emmerich, M., Bäck, T.: Cluster-based kriging approximation algorithms for complexity reduction. *Appl. Intell.* **50**(3), 778–791 (2020). <https://doi.org/10.1007/s10489-019-01549-7>
34. Storn, R., Price, K.: Differential Evolution – A Simple and Efficient Heuristic for global Optimization over Continuous Spaces. *Journal of Global Optimization* **11**(4), 341–359 (1997). <https://doi.org/10.1023/A:1008202821328>
35. Ueno, T., Rhone, T.D., Hou, Z., Mizoguchi, T., Tsuda, K.: Combo: an efficient bayesian optimization library for materials science. *Materials discovery* **4**, 18–21 (2016). <https://doi.org/10.1016/j.md.2016.04.001>
36. Vapnik, V.N.: The Nature of Statistical Learning Theory. Springer-Verlag New York, Inc., New York, USA (1995). <https://doi.org/10.1007/978-1-4757-3264-1>
37. Viana, F.A., Simpson, T.W., Balabanov, V., Toropov, V.: Special section on multidisciplinary design optimization: metamodeling in multidisciplinary design optimization: how far have we really come? *AIAA journal* **52**(4), 670–690 (2014). <https://doi.org/10.2514/1.J052375>
38. Vivarelli, F., Williams, C.K.I.: Discovering Hidden Features with Gaussian Processes Regression. In: Kearns, M.J., Solla, S.A., Cohn, D.A. (eds.) *Advances in Neural Information Processing Systems 11*, pp. 613–619. MIT Press (1999)
39. Wang, H., van Stein, B., Emmerich, M., Bäck, T.: A new acquisition function for bayesian optimization based on the moment-generating function. In: 2017 IEEE International Conference on Systems, Man, and Cybernetics, SMC 2017, Banff, AB, Canada, October 5-8, 2017. pp. 507–512. IEEE (2017). <https://doi.org/10.1109/SMC.2017.8122656>
40. Wang, Z., Hutter, F., Zoghi, M., Matheson, D., De Freitas, N.: Bayesian optimization in a billion dimensions via random embeddings (2016)
41. Yoshimura, M., Shimoyama, K., Misaka, T., Obayashi, S.: Topology optimization of fluid problems using genetic algorithm assisted by the Kriging model. *International Journal for Numerical Methods in Engineering* **109**(4), 514–532 (2016). <https://doi.org/10.1002/nme.5295>
42. Yuan, Y.X.: A Review of Trust Region Algorithms for Optimization