



HAL
open science

High Dimensional Bayesian Optimization with Kernel Principal Component Analysis

Kirill Antonov, Elena Raponi, Hao Wang, Carola Doerr

► **To cite this version:**

Kirill Antonov, Elena Raponi, Hao Wang, Carola Doerr. High Dimensional Bayesian Optimization with Kernel Principal Component Analysis. 17th Proceedings of Parallel Problem Solving from Nature - (PPSN) 2022, Sep 2022, Dortmund, Germany. pp.118-131, 10.1007/978-3-031-14714-2_9. hal-03740759

HAL Id: hal-03740759

<https://hal.sorbonne-universite.fr/hal-03740759>

Submitted on 29 Jul 2022

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 with Kernel Principal Component Analysis

Kirill Antonov^{1,3*}[0000-0002-8757-8598], Elena Raponi^{2,4}[0000-0001-6841-7409],
Hao Wang³[0000-0002-4933-5181], and Carola Doerr⁴[0000-0002-4981-3227]

¹ ITMO University, Saint Petersburg, Russia,

² Technical University of Munich, TUM School of Engineering and Design, Munich, Germany

³ Leiden University, LIACS Department, Leiden, Netherlands

⁴ Sorbonne Université, CNRS, LIP6, Paris, France

Abstract. Bayesian Optimization (BO) is a surrogate-based global optimization strategy that relies on a Gaussian Process regression (GPR) model to approximate the objective function and an acquisition function to suggest candidate points. It is well-known that BO does not scale well for high-dimensional problems because the GPR model requires substantially more data points to achieve sufficient accuracy and acquisition optimization becomes computationally expensive in high dimensions. Several recent works aim at addressing these issues, e.g., methods that implement online variable selection or conduct the search on a lower-dimensional sub-manifold of the original search space. Advancing our previous work of PCA-BO that learns a linear sub-manifold, this paper proposes a novel kernel PCA-assisted BO (KPCA-BO) algorithm, which embeds a non-linear sub-manifold in the search space and performs BO on this sub-manifold. Intuitively, constructing the GPR model on a lower-dimensional sub-manifold helps improve the modeling accuracy without requiring much more data from the objective function. Also, our approach defines the acquisition function on the lower-dimensional sub-manifold, making the acquisition optimization more manageable.

We compare the performance of KPCA-BO to a vanilla BO and to PCA-BO on the multi-modal problems of the COCO/BBOB benchmark suite. Empirical results show that KPCA-BO outperforms BO in terms of convergence speed on most test problems, and this benefit becomes more significant when the dimensionality increases. For the 60D functions, KPCA-BO achieves better results than PCA-BO for many test cases. Compared to the vanilla BO, it efficiently reduces the CPU time required to train the GPR model and to optimize the acquisition function compared to the vanilla BO.

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

* Work done while visiting Sorbonne Université in Paris. Corresponding author, k.antonov@liacs.leidenuniv.nl

1 Introduction

Numerical black-box optimization problems are challenging to solve when the dimension of the problem’s domain becomes high [1]. The well-known *curse of dimensionality* implies that exponential growth of the data points is required to maintain a reasonable coverage of the search space. This is difficult to accommodate in numerical black-box optimization, which aims to seek a well-performing solution with a limited budget of function evaluations. Bayesian optimization (BO) [15, 18] suffers from high dimensionality more seriously compared to other search methods, e.g., evolutionary algorithms, since it employs a surrogate model of the objective function internally, which scales poorly with respect to the dimensionality (see Sec. 2 below). Also, BO proposes new candidate solutions by maximizing a so-called acquisition function (see Sec. 2), which assesses the potential of each search point for making progresses. The maximization task of the acquisition function is also hampered by high dimensionality. As such, BO is often taken only for small-scale problems (typically less than 20 search variables), and it remains an open challenge to scale it up for high-dimensional problems [3].

Recently, various methods have been proposed for enabling high-dimensional BO, which can be categorized into three classes: (1) variable selection methods that only execute BO on a subset of search variables [27], (2) methods that leverage the surrogate model to high dimensional spaces, e.g., via additive models [6, 7], and (3) conducting BO on a sub-manifold embedded in the original search space [14, 29]. Notably, in [9], a kernel-based approach is developed for parametric shape optimization in computer-aided design systems, which is not a generic approach since the kernel function is based on the representation of the parametric shape and is strongly tied to applications in mechanical design. In [21] we proposed PCA-BO, in which we conduct BO on a linear sub-manifold of the search space that is learned from the linear principal components analysis (PCA) procedure.

This paper advances the PCA-BO algorithm by considering the kernel PCA procedure [25], which is able to construct a non-linear sub-manifold of the original search space. The proposed algorithm - *Kernel PCA-assisted BO* (KPCA-BO) adaptively learns a nonlinear forward map from the original space to the lower-dimensional sub-manifold for reducing dimensionality and constructs a backward map that converts a candidate point found on the sub-manifold to the original search space for the function evaluation. We evaluate the empirical performance of KPCA-BO on the well-known BBOB problem set [12], focusing on the multi-modal problems.

This paper is organized as follows. In Sec. 2, we will briefly recap Bayesian optimization and some recent works on alleviating the issue of high dimension for BO. In Sec. 3, we describe the key components of KPCA-BO in detail. The experimental setting, results, and discussions are presented in Sec. 4, followed by the conclusion and future works in Sec. 5.

2 Related Work

Bayesian Optimization (BO) [15, 26] is a sequential model-based optimization algorithm which was originally proposed to solve single-objective black-box optimization problems that are expensive to evaluate. BO starts with sampling a small initial design of experiment (DoE, obtained with e.g., Latin Hypercube Sampling [24] or low-discrepancy sequences [20]) $X \subseteq S$. After evaluating $f(x)$ for all $x \in X$, it proceeds to construct a probabilistic model $\mathbb{P}(f | X, Y)$ (e.g., Gaussian process regression, please see the next paragraph). BO balances exploration and exploitation of the search by considering, for a decision point \mathbf{x} , two quantities: the predicted function value $\hat{f}(\mathbf{x})$ and the uncertainty of this prediction (e.g., the mean squared error $\mathbb{E}(f(\mathbf{x}) - \hat{f}(\mathbf{x}))^2$). Both of them are taken to form the acquisition function $\alpha: S \rightarrow \mathbb{R}$ used in this work, i.e., the expected improvement [15], which quantifies the potential of each point for making progresses. BO chooses the next point to evaluate by maximizing the acquisition function. After evaluating \mathbf{x}^* , we augment the data set with $(\mathbf{x}^*, f(\mathbf{x}^*))$ and proceed with the next iteration.

Gaussian Progress Regression (GPR) [22] models the objective function f as the realization of a Gaussian process $f \sim \text{gp}(0, c(\cdot, \cdot))$, where $c: S \times S \rightarrow \mathbb{R}$ is the covariance function, also known as kernel. That is, $\forall \mathbf{x}, \mathbf{x}' \in S$, it holds that $\text{Cov}\{f(\mathbf{x}), f(\mathbf{x}')\} = c(\mathbf{x}, \mathbf{x}')$. Given a set X of evaluated points and the corresponding function values Y , GPR learns a posterior Gaussian process to predict the function value at each point, i.e., $\forall \mathbf{x} \in S, f(\mathbf{x}) | X, Y, \mathbf{x} \sim \mathcal{N}(\hat{f}(\mathbf{x}), \hat{s}^2(\mathbf{x}))$, where \hat{f} and \hat{s}^2 are the posterior mean and variance functions, respectively. When equipped with a GPR and the expected improvement, BO has a convergence rate of $O(n^{-1/d})$ [4], which decreases quickly when the dimension increases.

High-dimensional Bayesian optimization. High dimensionality negatively affects the performance of BO in two aspects, the quality of the GPR model and the efficiency of acquisition optimization. For the quality of the GPR model, it is well-known [4] that many more data points are needed to maintain the modeling accuracy in higher dimensions. Moreover, acquisition optimization is a high-dimensional task requiring more surrogate evaluations to obtain a reasonable optimum. Notably, each surrogate evaluation takes $O(d)$ time to compute, making the acquisition optimization more time-consuming.

Depending on the expected structure of the high-dimensional problem, various strategies for dealing with this curse of dimensionality have been proposed in the literature, often falling into one of the following classes:

1. Variable selection or screening. It may be the case that a subset of parameters does not have any significant impact on solutions' quality, and it is convenient to identify and keep only the most influential ones. Different approaches may be considered: discarding variables uniformly [17], assigning weights to the variables based on the dependencies between them [27], identifying the most descriptive variables based on their length-scale value in the model [2], etc.
2. Additive models. They keep all the variables but limit their interaction as they are based on the idea of decomposing the problem into blocks. For

example, the model kernels can be seen as the sum of univariate ones [6, 7], the high-dimensional function can decompose as a sum of lower-dimensional functions on subsets of variables [23], or the additive model can be based on an ANOVA decomposition [10, 19].

3. Linear/nonlinear embeddings. They are based on the hypothesis that a large percentage of the variation of a high-dimensional function can be captured in a low-dimensional embedding of the original search space. The embedding can be either linear [21, 29] or nonlinear [9, 11].

We point the reader to [3] for a comprehensive overview of the state-of-the-art in high-dimensional BO.

3 Kernel-PCA assisted by Bayesian Optimization

In this paper, we deal with numerical black-box optimization problems $f: S \subseteq \mathbb{R}^d \rightarrow \mathbb{R}$, where the search domain is a hyperbox, i.e., $S = [l_1, u_1] \times [l_2, u_2] \times \dots \times [l_d, u_d]$. We reduce the dimensionality of the optimization problem *on-the-fly*, using a *Kernel Principal Component Analysis* (KPCA) [25] for learning, from the evaluated search points, a non-linear sub-manifold \mathcal{M} on which we optimize the objective function. Ideally, such a sub-manifold \mathcal{M} should capture important information of f for optimization. In other words, \mathcal{M} should “traverse” several basins of attractions of f . Loosely speaking, in contrast to a linear sub-manifold (e.g., our previous work [21]), the non-linear one would solve the issue that the correlation among search variables is non-linear (e.g., on multimodal functions), where it is challenging to identify a linear sub-manifold that passes through several local optima simultaneously. KPCA tackles this issue by first casting the search points to a high-dimensional Hilbert space \mathcal{H} (typically infinite-dimensional), where we learn a linear sub-manifold thereof. We consider a positive definite function $k: S \times S \rightarrow \mathbb{R}$, which induces a reproducing kernel Hilbert space (RKHS) \mathcal{H} constructed as the completion of $\text{span}\{k(\mathbf{x}, \cdot) : \mathbf{x} \in S\}$. The function $\phi(\mathbf{x}) := k(\mathbf{x}, \cdot)$ maps a point \mathbf{x} from the search space to \mathcal{H} , which we will refer as *the feature map*. An inner product on \mathcal{H} is defined with k , i.e., $\forall \mathbf{x}, \mathbf{x}' \in S, \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle_{\mathcal{H}} = k(\mathbf{x}, \mathbf{x}')$, known as the *kernel trick*.

The KPCA-BO algorithm. Fig. 1 provides an overview of the proposed KPCA-BO algorithm. We also present the pseudo-code of KPCA-BO in Alg. 1. Key differences to our previous work that employs the linear PCA method [21] are highlighted. Various building blocks of the algorithm will be described in the following paragraphs. Notably, the sub-routine KPCA indicates performing the standard kernel PCA algorithm, which returns a set of selected principal components, whereas, GPR represents the training of a Gaussian process regression model. In the following discussion, we shall denote by $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$ and $\mathbf{Y} = \{f(\mathbf{x}_i)\}_{i=1}^n$ the set of the evaluated points and their function values, respectively.

Rescaling data points. As an unsupervised learning method, KPCA disregards the distribution of the function values if applied directly, which contradicts our aim of capturing the objective function in the dimensionality reduc-




Fig. 1: Flowchart of the KPCA-BO optimization algorithm with detailed graphical representation of the KPCA subroutine.

tion. To mitigate this issue, we apply the weighting scheme proposed in [21], which scales the data points in S with respect to their objective values. In detail, we compute the rank-based weights for all points: w_i is proportional to $\ln n - \ln R_i$, $i = 1, \dots, n$, where R_1, R_2, \dots, R_n are the rankings of points with respect to Y in increasing order (minimization is assumed). Then we rescale each point with its weight, i.e., $\mathbf{x}_i = w_i(\mathbf{x}_i - n^{-1} \sum_{k=1}^n \mathbf{x}_k)$, $i = 1, \dots, n$. It is necessary to show that the feature map ϕ respects the rescaling operation performed in S . For any stationary and monotonic kernel (i.e., $k(\mathbf{x}, \mathbf{y}) = k(D_S(\mathbf{x}, \mathbf{y}))$ (D_S is a metric in S) and k decreases whenever $D_S(\mathbf{x}, \mathbf{y})$ increases), for all $\mathbf{x}, \mathbf{y}, \mathbf{c} \in X$ it holds that $D_S(\mathbf{x}, \mathbf{c}) \leq D_S(\mathbf{y}, \mathbf{c})$ implies that $D_{\mathcal{H}}(\phi(\mathbf{x}), \phi(\mathbf{c})) \leq D_{\mathcal{H}}(\phi(\mathbf{y}), \phi(\mathbf{c}))$. Consequently, the point pushed away from the center of the data in X will still have a large distance to the center of the data in \mathcal{H} after the feature map. The rescaling (or alternatives that incorporate the objective values into the distribution of data points in the domain) is an essential step in applying PCA to BO. On the one hand, DoE aims to span the search domain as evenly as possible and thereby the initial random sample from it has almost the same variability in all directions, which provides no information for PCA to learn. On the other hand, new candidates are obtained by the global optimization of the acquisition function in each iteration, which is likely to produce multiple clusters and/or isolated points that are not meaningful to the PCA procedure. This is in contrast to the direct application of PCA to evolutionary algorithms [16], where we apply PCA to the current population. Since the population is usually generated

Algorithm 1 KPCA-assisted Bayesian Optimization. Highlighted are those lines in which KPCA-BO differs from the linear PCA method [21]

```

1: procedure KPCA-BO( $f, S$ ) ▷  $f$ : objective function,  $S \subseteq \mathbb{R}^d$ : search space
2:   Create  $X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n_0}\} \subset S$  with Latin hypercube sampling
3:    $Y \leftarrow \{f(\mathbf{x}_1), \dots, f(\mathbf{x}_{n_0})\}$ ,  $n \leftarrow n_0$ 
4:   while the stop criteria are not fulfilled do
5:     if  $n = n_0$  or  $y^* \leq 20\%$ -percentile of  $Y$  then
6:        $R_1, R_2, \dots, R_n$  are the rankings of points in  $X$  w.r.t.  $Y$  (increasing order)
7:        $\mathbf{x}'_i \leftarrow \mathbf{x}_i - n^{-1} \sum_{k=1}^n \mathbf{x}_k$ ,  $i = 1, \dots, n$  ▷ centering
8:        $\mathbf{x}'_i \leftarrow w_i \mathbf{x}'_i$ ,  $w_i \propto \ln n - \ln R_i$ ,  $i = 1, \dots, n$  ▷ rescaling
9:        $\gamma^* \leftarrow \text{OPTIMIZE-RBF-KERNEL}(\{\mathbf{x}'_i\}_{i=1}^n)$  ▷ Eq. (2)
10:    end if
11:     $v_1, \dots, v_r \leftarrow \text{KPCA}(\{\mathbf{x}'_i\}_{i=1}^n, \gamma^*)$  ▷  $r < d \ll n$ 
12:    construct the forward map  $\mathcal{F}$  from  $\text{span}\{v_1, \dots, v_r\}$ . ▷ Eq. (1)
13:     $\mathbf{z}_i \leftarrow \mathcal{F}(\mathbf{x}_i)$ ,  $i = 1, \dots, n$  ▷ map the data to  $T := \text{span}\{v_1, \dots, v_r\}$ 
14:     $\hat{f}, \hat{s}^2 \leftarrow \text{GPR}(\{\mathbf{z}_i\}_{i=1}^n, Y)$  ▷ Gaussian process regression
15:     $\mathbf{z}^* \leftarrow \arg \max_{\mathbf{z} \in T} \text{EI}(\mathbf{z}; \hat{f}, \hat{s}^2)$ 
16:     $y^* \leftarrow f(\mathbf{x}^*)$ ,  $\mathbf{x}^* \leftarrow \mathcal{B}(\mathbf{z}^*)$  ▷ the backward map; Eq.(4)
17:     $X \leftarrow X \cup \{\mathbf{x}^*\}$ ,  $Y \leftarrow Y \cup \{y^*\}$ ,  $n \leftarrow n + 1$ 
18:  end while
19: end procedure

```

from a unimodal mutation distribution, it is well-suited for applying the PCA procedure.

Dimensionality reduction in Hilbert spaces. After the rescaling operation in X , we map the points to the feature space \mathcal{H} : $\phi(X) = \{\phi(\mathbf{x}_i)\}_i$, $i = 1, \dots, n$. After centering the feature points in \mathcal{H} , i.e., $\tilde{\phi}(\mathbf{x}_i) = \phi(\mathbf{x}_i) - n^{-1} \sum_{i=1}^n \phi(\mathbf{x}_i)$, we express the sample covariance⁵ of the feature points: $C = n^{-1} \sum_{i=1}^n \tilde{\phi}(\mathbf{x}_i) \tilde{\phi}(\mathbf{x}_i)^\top$. KPCA essentially computes the eigenvalues and eigenfunctions of C , namely $\forall i \in [1..n]$, $Cv_i = \lambda_i v_i$, $v_i \in \mathcal{H}$, $\|v_i\|_{\mathcal{H}} = 1$, and $\langle v_i, v_j \rangle_{\mathcal{H}} = 0$, if $i \neq j$. Note that (1) C is positive semi-definite; (2) since $\text{rank}(C) \leq \sum_{i=1}^n \text{rank}(\tilde{\phi}(\mathbf{x}_i) \tilde{\phi}(\mathbf{x}_i)^\top) = n$, there are maximally n nonzero eigenvalues and eigenfunctions; (3) the eigenfunction takes the following form $v_i = \sum_{j=1}^n a_j^{(i)} \tilde{\phi}(\mathbf{x}_j)$, $a_j^{(i)} \in \mathbb{R}$ and thereby all eigenfunctions can be represented by a matrix $\mathbf{V} = (a_j^{(i)})_{ij}$. Assume the eigenvalues are ordered in the decreasing manner (i.e., $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 0$. Eigenfunctions are sorted accordingly). It is not hard to show that the variance of $\phi(X)$ along v_i is exactly λ_i : $n^{-1} \sum_{k=1}^n \langle \tilde{\phi}(\mathbf{x}_k), v_i \rangle_{\mathcal{H}}^2 = \langle v_i, n^{-1} \sum_{k=1}^n [\tilde{\phi}(\mathbf{x}_k) \tilde{\phi}(\mathbf{x}_k)^\top] v_i \rangle_{\mathcal{H}} = \lambda_i$. Therefore, the eigenvalues can be used to select a linear subspace of \mathcal{H} which keeps the majority of the variability of $\phi(X)$. Specifically, we choose a subspace $T := \text{span}\{v_1, \dots, v_r\} \subset \mathcal{H}$ as the reduced search space of BO, where r is chosen as the smallest integer such that the first- r eigenvalues explain at least η percent of the total variability (we use $\eta = 90\%$ in our experiments). For

⁵ The outer product is a linear operator defined as $\forall h \in \mathcal{H}$, $[\phi(\mathbf{x})\phi(\mathbf{x})^\top](h) : h \mapsto \langle \phi(\mathbf{x}), h \rangle_{\mathcal{H}} \phi(\mathbf{x})$. Hence, the sample covariance is also a linear operator $C : \mathcal{H} \rightarrow \mathcal{H}$.

a point $\mathbf{x} \in S$, we can formulate a *forward map* that projects $\phi(\mathbf{x})$ onto the reduced space T : $\mathcal{F}: \mathbf{x} \mapsto \sum_{i=1}^r \langle \tilde{\phi}(\mathbf{x}), v_i \rangle_{\mathcal{H}} v_i$. Let $g_i(\mathbf{x}) = \langle \tilde{\phi}(\mathbf{x}), \tilde{\phi}(\mathbf{x}_i) \rangle_{\mathcal{H}} = k(\mathbf{x}, \mathbf{x}_i) - n^{-1} \sum_{j=1}^n k(\mathbf{x}, \mathbf{x}_j) - n^{-1} \sum_{j=1}^n k(\mathbf{x}_i, \mathbf{x}_j) + n^{-2} \sum_{i=1}^n \sum_{j=1}^n k(\mathbf{x}_i, \mathbf{x}_j)$, the forward map can be re-expressed as

$$\mathcal{F}: \mathbf{x} \mapsto \mathbf{V} \mathbf{g}(\mathbf{x}), \quad \mathbf{g}(\mathbf{x}) = (g_1(\mathbf{x}), \dots, g_n(\mathbf{x}))^\top. \quad (1)$$

Computationally, the eigenfunction representation \mathbf{V} can be calculated via eigen-decomposition of the Gram matrix ($\mathbf{G}_{ij} = \langle \tilde{\phi}(\mathbf{x}_i), \tilde{\phi}(\mathbf{x}_j) \rangle_{\mathcal{H}}$, i.e., $\mathbf{G} = \mathbf{V}^\top \mathbf{D} \mathbf{V}$, \mathbf{D} is a $n \times n$ diagonal matrix with the eigenvalues of C on its nonzero entries.

Learning the forward map. We use the *radial basis function* (RBF, a.k.a. Gaussian kernel) for KPCA in this paper. The RBF kernel $k(\mathbf{x}, \mathbf{x}') = \exp(-\gamma \|\mathbf{x} - \mathbf{x}'\|_2^2)$, contains a single length-scale hyperparameter $\gamma \in \mathbb{R}_{>0}$. To determine this length-scale, we minimize the number of eigenvalues/functions chosen to keep at least η percent of the variance, which effectively distributes more information of $\phi(X)$ on the first few eigenfunctions and hence allows for constructing a lower-dimensional space T . Also, we reward γ values which choose the same number of eigenfunctions and also yield a higher ratio of explained variance. In all, the cost function for tuning γ is:

$$\gamma^* = \arg \min_{\gamma \in (0, \infty)} r - \frac{\sum_{i=1}^r \lambda_i}{\sum_{i=1}^n \lambda_i}, \quad r = \inf \left\{ k \in [1..n]: \sum_{i=1}^k \lambda_i \geq \eta \sum_{i=1}^n \lambda_i \right\}. \quad (2)$$

This equation is solved numerically, using a quasi-Newton method (the L-BFGS-B algorithm [5]) with $\gamma \in [10^{-4}, 2]$ and maximally $200d$ iterations. It is worth noting that we do not consider anisotropic kernels (e.g., individual length-scales for each search variable) since such a kernel increases the number of hyperparameters to learn.

Also, note that the choice of the kernel can affect the smoothness of the manifold in the feature space \mathcal{H} , e.g., the Matérn 5/2 kernel induces a C^2 atlas for the manifold $\phi(S)$. We argue that the smoothness of $\phi(S)$ is less important to the dimensionality reduction task, comparing to the convexity and connectedness thereof. In this work, we do not aim to investigate the impact of the kernel on the topological properties of $\phi(S)$. Therefore, we use the RBF kernel for the construction of the forward map for its simplicity.

Learning the backward map. When performing the Bayesian optimization in the reduced space T , we need to determine a “pre-image” of a candidate point $\mathbf{z} \in T$ for the function evaluation. To implement such a backward map $\mathcal{B}: T \rightarrow S$, we base our construction on the approach proposed in [8], in which the pre-image of a point $\mathbf{z} \in T$ is a conical combination of some points in S : $\sum_{i=1}^d w_i \mathbf{p}_i$, $w_i \in \mathbb{R}_{>0}$. In this paper, the points $\{\mathbf{p}_i\}_{i=1}^d$ are taken as a random subset of the data points $\{\mathbf{x}_i\}_{i=1}^n$. The conical weights are determined by minimizing the distance between \mathbf{z} and the image of the conical combination under the forward map:

$$w_1^*, \dots, w_d^* = \arg \min_{\{w_i\}_{i=1}^d \subset \mathbb{R}_{>0}^d} \left\| \mathbf{z} - \mathcal{F} \left(\sum_{i=1}^d w_i \mathbf{p}_i \right) \right\|_2^2 + Q \left(\sum_{i=1}^d w_i \mathbf{p}_i \right), \quad (3)$$

$$Q(\mathbf{x}) = \exp \left(\sum_{i=1}^d \max(0, l_i - x_i) + \max(0, x_i - u_i) \right).$$

where the function Q penalizes the case that the pre-image is out of S . As with Eq. (2), the weights are optimized with the L-BFGS-B algorithm (starting from zero with $200d$ maximal iterations). Taking the optimal weights, we proceed to define the *backward map*: $\forall \mathbf{z} \in T$,

$$\mathcal{B}: \mathbf{z} \mapsto \text{CLIP} \left(\sum_{i=1}^d w_i^* \mathbf{p}_i \right), \quad (4)$$

where the function $\text{CLIP}(\mathbf{x})$ cuts off each component of \mathbf{x} at the lower and upper bounds of S , which ensures the pre-image is always feasible.

Remark. The event that $\{\mathbf{p}_i\}_{i=1}^d$ contains a co-linear relation is of measure zero and the conical form can procedure pre-images outside S , allowing for a complete coverage thereof. There exist multiple solutions to Eq. (3) (and hence multiple pre-images) since the forward map \mathcal{F} contains an orthogonal projection step, which is not injective. Those multiple pre-images can be obtained by randomly restarting the quasi-Newton method used to solve Eq. (3). However, since those pre-images do not distinguish from each other for our purpose, we simply take a random one in this work.

Bayesian optimization in the reduced space. Given the forward and backward maps, we are ready to perform the optimization task in the space T . Essentially, we first map the data set $X \subset S$ to T using the forward map (Eq. (1)): $\mathcal{F}(X) = \{\mathcal{F}(\mathbf{x}_i)\}_{i=1}^n$, which implicitly defines the counterpart $f' := f \circ \mathcal{B}$ of the objective function in T . Afterwards, we train a Gaussian process model with the data set $(\mathcal{F}(X), Y)$ to model f' , i.e., $\forall \mathbf{z} \in T, f'(\mathbf{z}) \mid \mathcal{F}(X), Y, \mathbf{z} \sim \mathcal{N}(\hat{f}(\mathbf{z}), \hat{s}^2(\mathbf{z}))$. The search domain in the reduced space T is determined as follows. Since the RBF kernel monotonically decreases w.r.t. the distance between its two input points, we can bound the set $\phi(X)$ by first identifying the point \mathbf{x}_{\max} with the largest distance to the center of data points \mathbf{c} and secondly computing the distance r between $\phi(\mathbf{x}_{\max})$ and $\phi(\mathbf{c})$ in the feature space \mathcal{H} . Since S is a hyperbox in \mathbb{R}^d , we simply take an arbitrary vertex of the hyperbox for \mathbf{x}_{\max} . Note that, as the orthogonal projection (from \mathcal{H} to T) does not increase the distance, the open ball $B := \{\mathbf{z} \in T: \|\mathbf{z}\|_2 < r\}$ always covers $\mathcal{F}(X)$. For the sake of optimization in T , we take the smallest hyperbox covering B as the search domain in T .

After the GPR model is created on the data set $(\mathcal{F}(X), Y)$, we maximize the expected improvement function $\text{EI}(\mathbf{z}; \hat{f}, \hat{s}) = \hat{s}(\mathbf{z})u \text{CDF}(u) + \hat{s}(\mathbf{z}) \text{PDF}(u), u = (\min Y - \hat{f}(\mathbf{z}))/\hat{s}(\mathbf{x})$ to pick a new candidate point \mathbf{z}^* , where CDF and PDF stand for the cumulative distribution and probability distribution functions of a standard normal random variable, respectively. Due to our construction of the search domain in T , it is possible that the global optimum \mathbf{z}^* of EI is associated with an infeasible pre-image in S . To mitigate this issue, we propose a multi-restart optimization strategy for maximizing EI (with different starting points in each restart), in which we only take the best outcome (w.r.t. its EI value) whose pre-image belongs to S . In our experiments, we used 10 random restarts

of the optimization. Also, it is unnecessary to optimize kernel’s hyperparameter γ in each iteration of BO since the new point proposed by EI would not make a significant impact on learning the feature map, if its quality is poor relative to the observed ones in Y (and hence assigned with a small weight). Therefore, it suffices to only re-optimize γ whenever we find a new point whose function value is at least as good as the 20% percentile of Y . Also, the 20% threshold is manually chosen to balance the convergence and computation time of KPCA-BO, after experimenting several different values on BBOB test problems.

4 Experiments

Experimental Setup. We evaluate the performance of KPCA-BO on ten multi-modal functions from the BBOB problem set [12] (F15 - F24), which should be sufficiently representative of the objective functions handled in real-world applications. We compare the experimental result of KPCA-BO to standard BO and the PCA-BO in our previous work [21] on three problem dimensions $d \in \{20, 40, 60\}$ with the evaluation budget in $\{100, 200, 300\}$, respectively. We choose a relatively large DoE size of $3d$, to ensure enough information for learning the first sub-manifold. On each function, we consider five problem instances (instance ID from 0 to 4) and conduct 10 independent runs of each algorithm. We select the Matérn 5/2 kernel for the GPR model. The L-BFGS-B algorithm [5] is employed to maximize the likelihood of GPR as well as the EI acquisition function at each iteration. We add to our comparison results for CMA-ES [13], obtained by executing the `pycma` package (<https://github.com/CMA-ES/pycma>) with 16 independent runs on each problem. The implementation of BO, PCA-BO, and KPCA-BO can be accessed at <https://github.com/wangronin/Bayesian-Optimization/tree/KPCA-BO>.

Results. All our data sets are available for interactive analysis and visualization in the IOHanalyzer [28] repository, under the *bbob-largescale* data sets of the *IOH* repository. In Fig. 2, we compare the convergence behavior of all four algorithms, where we show the evolution of the best-so-far target gap ($f_{\text{best}} - f^*$) with respect to the iteration for each function-dimension pair. In all dimensions, it is clear that both KPCA-BO and PCA-BO outperform BO substantially across functions F17 - F20, while both KPCA-BO and PCA-BO exhibit about the same convergence with BO on F23, and are surpassed by BO significantly on F16. The poor performance of PCA-BO and KPCA-BO on the Weierstrass function (F16) can be attributed to the nature of its landscape, which is highly rugged and moderately periodic with multiple global optima. Therefore, a basin of attraction is not clearly defined, which confuses both PCA variants. On functions F17, F19, and F24, we observe that KPCA-BO is outperformed by PCA-BO in 20D, while as the dimensionality increases, KPCA-BO starts to surpass the convergence rate of PCA-BO. For functions F21 and F22, KPCA-BO’s performance is indistinguishable from PCA-BO in 20D, and in higher dimensions, the advantage of KPCA-BO becomes prominent. For the remaining function-dimension pairs, KPCA-BO shows a comparable performance to PCA-BO. Compared to CMA-

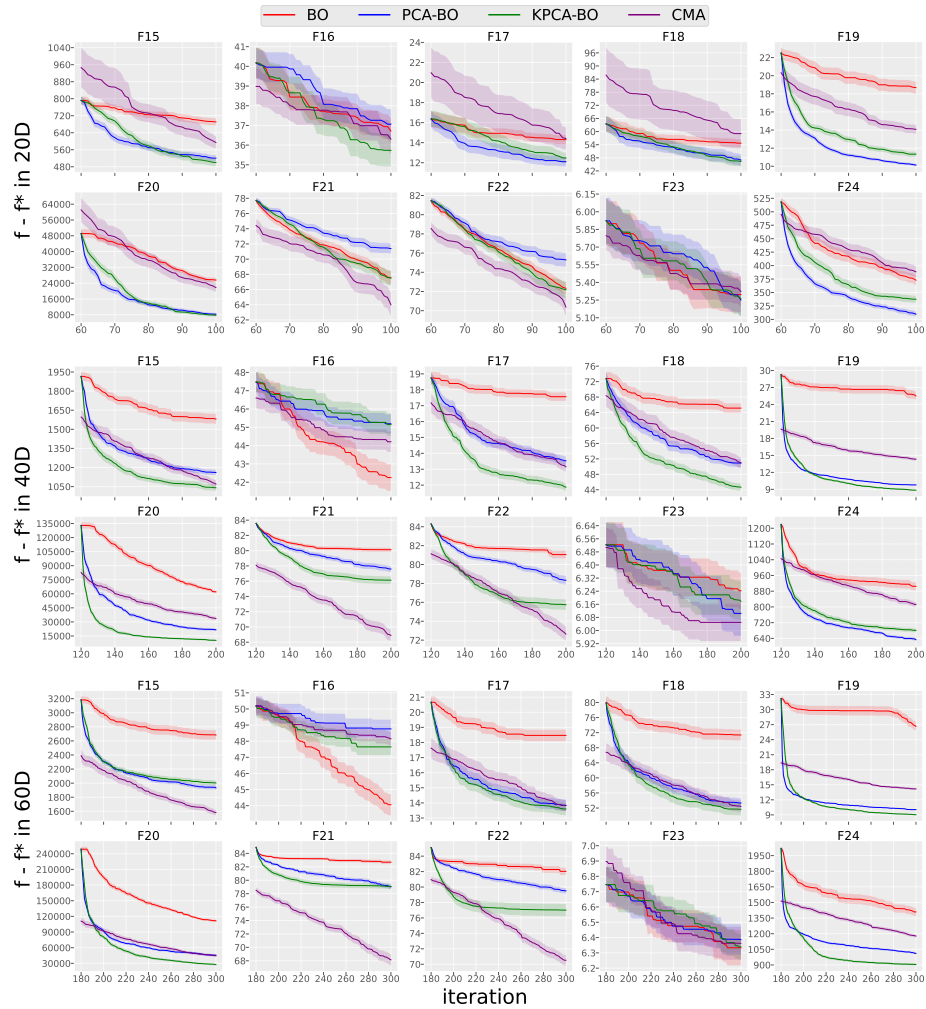


Fig. 2: The best-so-far target gap ($f_{\text{best}} - f^*$) against the iteration for CMA-ES (purple), BO (red), PCA-BO (blue), and KPCA-BO (green) is averaged over 50 independent runs on each test problem. We compare the algorithms in three dimensions, 20D (top), 40D (middle), and 50D (bottom). The shaded area indicates the standard error of the mean target gap. CMA-ES data is obtained from running the `pycma` package with the same evaluation budgets as BO.

ES, both KPCA-BO and PCA-BO either outperform CMA-ES or show roughly the same convergence except on F21 and F22 in 20D. In higher dimensions, although KPCA-BO still exhibits a steeper initial convergence rate (before about 200 function evaluations in 60D), CMA-ES finds a significantly better solution




Fig. 3: Mean CPU time taken by training the GPR model (dark cyan) and maximizing the EI acquisition function (red) in 60D for BO, PCA-BO, and KPCA-BO, respectively. In general, training the GPR model takes the majority of CPU time and both PCA-BO and KPCA-BO manages to reduce it significantly.

after the first $3d$ evaluations (the DoE phase of the BO variants), leading to better overall performance.

Also, we observe that KPCA-BO shows better relative performance when the dimensionality increases (e.g., on F18 and F22 across three dimensions), implying that the kernelized version is better suited for solving higher-dimensional problems. Interestingly, KPCA-BO shows an early faster convergence on F21 and F22 compared to PCA-BO and is gradually overtaken by PCA-BO, implying that KPCA-BO stagnates earlier than PCA-BO. We conjecture that the kernel function of KPCA-BO (and consequently the sub-manifold) stabilizes much faster than the linear subspace employed in PCA-BO, which might attribute to such a stagnation behavior. Therefore, KPCA-BO is more favorable than PCA-BO in higher dimensions (i.e., $d \geq 60$), while in lower dimensions ($d \approx 20$), it is competitive to PCA-BO when the budget is small for most test cases (it only loses a little to PCA-BO on F24).

In Fig. 3, we depict the CPU time (in seconds) taken to train the GPR model as well as maximize EI in 60D. As expected, the majority of the CPU time is taken by the training of the GPR model. PCA-BO and KPCA-BO achieve substantially smaller CPU time of GPR training than the vanilla BO with exception on F19, F20, and F24, where KPCA-BO actually takes more time. In all cases, the CPU time for maximizing EI is smaller for KPCA-BO and PCA-BO than for BO.

5 Conclusions

In this paper, we proposed a novel KPCA-assisted Bayesian optimization algorithm, the KPCA-BO. Our algorithm enables BO in high-dimensional numerical

optimization tasks by learning a nonlinear sub-manifold of the original search space from the evaluated data points and performing BO directly in this sub-manifold. Specifically, to capture the information about the objective function when performing the kernel PCA procedure, we rescale the data points in the original space using a weighting scheme based on the corresponding objective values of the data point. With the help of the KPCA procedure, the training of the Gaussian process regression model and the acquisition optimization – the most costly steps in BO – are performed in the lower-dimensional space. We also implement a backward map to convert the candidate point found in the lower-dimensional manifold to the original space.

We empirically evaluated KPCA-BO on the ten multimodal functions (F15-F24) from the BBOB benchmark suite. We also compare the performance of KPCA-BO with the vanilla BO, the PCA-BO algorithm from our previous work, and CMA-ES, a state-of-the-art evolutionary numerical optimizer. The results show that KPCA-BO performs better than PCA-BO in capturing the contour lines of the objective functions when the variables are not linearly correlated. The higher the dimensionality, the more significant this better capture becomes for the optimization of functions F20, F21, F22, and F24. Also, the mean CPU time measured in the experiments shows that after reducing the dimensionality, the CPU time needed to train the GPR model and maximize the acquisition is greatly reduced in most cases.

The learning of the lower-dimensional manifold is the crux of the proposed KPCA-BO algorithm. However, we observe that this manifold stabilizes too quickly for some functions, leading to unfavorable stagnation behavior. In further work, we plan to investigate the cause of this premature convergence and hope to identify mitigation methods. Since the manifold is learned from the data points evaluated so far, a viable approach might be to use a random subset of data points to learn the manifold, rather than taking the entire data set.

Another future direction is to improve the backward map proposed in this paper. Since an orthogonal projection is involved when mapping the data to the manifold, the point on the manifold has infinitely many pre-images. In our current approach, we do not favor one direction or another, whereas it might be preferable to bias the map using the information of the previously sampled points. For example, the term to minimize (to exploit) or maximize (to explore) the distance between the candidate pre-image point and the lower-dimensional manifold can be added to the cost function used in the backward map. These two approaches can be combined or switched in the process of optimization.

Acknowledgments. Our work is supported by the Paris Ile-de-France region (via the DIM RFSI project AlgoSelect), by the CNRS INS2I institute (via the RandSearch project), by the PRIME programme of the German Academic Exchange Service (DAAD) with funds from the German Federal Ministry of Education and Research (BMBF), and by RFBR and CNRS, project number 20-51-15009.

References

1. Bellman, R.: Dynamic programming. *Science* **153**(3731), 34–37 (1966). <https://doi.org/10.1126/science.153.3731.34>
2. Ben Salem, M., Bachoc, F., Roustant, O., Gamboa, F., Tomaso, L.: Sequential dimension reduction for learning features of expensive black-box functions (2019), <https://hal.archives-ouvertes.fr/hal-01688329>, preprint
3. Binois, M., Wycoff, N.: A survey on high-dimensional Gaussian process modeling with application to Bayesian optimization. arXiv:2111.05040 [math] (Nov 2021)
4. Bull, A.D.: Convergence Rates of Efficient Global Optimization Algorithms. *J. Mach. Learn. Res.* **12**, 2879–2904 (2011), <http://dl.acm.org/citation.cfm?id=2078198>
5. Byrd, R.H., Lu, P., Nocedal, J., Zhu, C.: A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.* **16**(5), 1190–1208 (1995). <https://doi.org/10.1137/0916069>
6. Delbridge, I., Bindel, D., Wilson, A.G.: Randomly Projected Additive Gaussian Processes for Regression. In: Proc. of the 37th International Conference on Machine Learning (ICML). pp. 2453–2463. PMLR (Nov 2020)
7. Duvenaud, D.K., Nickisch, H., Rasmussen, C.: Additive Gaussian Processes. In: Advances in Neural Information Processing Systems. vol. 24. Curran Associates, Inc. (2011)
8. García-González, A., Huerta, A., Zlotnik, S., Díez, P.: A kernel principal component analysis (kpca) digest with a new backward mapping (pre-image reconstruction) strategy. *CoRR abs/2001.01958* (2020)
9. Gaudrie, D., Riche, R.L., Picheny, V., Enaux, B., Herbert, V.: Modeling and Optimization with Gaussian Processes in Reduced Eigenbases – Extended Version. *Structural and Multidisciplinary Optimization* **61**(6), 2343–2361 (Jun 2020). <https://doi.org/10.1007/s00158-019-02458-6>
10. Ginsbourger, D., Roustant, O., Schuhmacher, D., Durrande, N., Lenz, N.: On ANOVA decompositions of kernels and Gaussian random field paths. arXiv:1409.6008 [math, stat] (Oct 2014)
11. Guhaniyogi, R., Dunson, D.B.: Compressed gaussian process for manifold regression. *Journal of Machine Learning Research* **17**(69), 1–26 (2016), <http://jmlr.org/papers/v17/14-230.html>
12. Hansen, N., Auger, A., Ros, R., Mersmann, O., Tušar, T., Brockhoff, D.: COCO: A platform for comparing continuous optimizers in a black-box setting. *Optimization Methods and Software* pp. 1–31 (2020)
13. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation* **9**(2), 159–195 (2001). <https://doi.org/10.1162/106365601750190398>, <https://doi.org/10.1162/106365601750190398>
14. Huang, W., Zhao, D., Sun, F., Liu, H., Chang, E.: Scalable Gaussian process regression using deep neural networks. In: Proc. of the 24th International Conference on Artificial Intelligence (IJCAI). pp. 3576–3582. AAAI Press (2015)
15. 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>
16. Kapsoulis, D., Tsiakas, K., Asouti, V., Giannakoglou, K.C.: The use of kernel PCA in evolutionary optimization for computationally demanding engineering applications. In: 2016 IEEE Symposium Series on Computational Intelligence, SSCI 2016, Athens, Greece, December 6–9, 2016. pp. 1–8. IEEE

- (2016). <https://doi.org/10.1109/SSCI.2016.7850203>, <https://doi.org/10.1109/SSCI.2016.7850203>
17. Li, C., Gupta, S., Rana, S., Nguyen, V., Venkatesh, S., Shilton, A.: High dimensional bayesian optimization using dropout. In: Proc. of the 26th International Joint Conference on Artificial Intelligence (IJCAI). p. 2096–2102. AAAI Press (2017)
 18. 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. LNCS, vol. 27, pp. 400–404. Springer (1974). https://doi.org/10.1007/3-540-07165-2_55
 19. Muehlenstaedt, T., Roustant, O., Carraro, L., Kuhnt, S.: Data-driven Kriging models based on FANOVA-decomposition. *Statistics and Computing* **22**(3), 723–738 (May 2012). <https://doi.org/10.1007/s11222-011-9259-7>
 20. Niederreiter, H.: Low-discrepancy and low-dispersion sequences. *Journal of number theory* **30**(1), 51–70 (1988)
 21. Raponi, E., Wang, H., Bujny, M., Boria, S., Doerr, C.: High Dimensional Bayesian Optimization Assisted by Principal Component Analysis. In: Proc. of Parallel Problem Solving from Nature (PPSN). LNCS, vol. 12269, pp. 169–183. Springer (2020). https://doi.org/10.1007/978-3-030-58112-1_12
 22. Rasmussen, C.E., Williams, C.K.I.: Gaussian processes for machine learning. Adaptive computation and machine learning, MIT Press (2006), <https://www.worldcat.org/oclc/61285753>
 23. Rolland, P., Scarlett, J., Bogunovic, I., Cevher, V.: High-Dimensional Bayesian Optimization via Additive Models with Overlapping Groups. In: Proc. of the Twenty-First International Conference on Artificial Intelligence and Statistics. pp. 298–307. PMLR (Mar 2018)
 24. Santner, T.J., Williams, B.J., Notz, W.I.: The Design and Analysis of Computer Experiments. Springer (2003). <https://doi.org/10.1007/978-1-4757-3799-8>
 25. 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>
 26. Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., de Freitas, N.: Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE* **104**(1), 148–175 (2016). <https://doi.org/10.1109/JPROC.2015.2494218>
 27. Ulmasov, D., Baroukh, C., Chachuat, B., Deisenroth, M., Misener, R.: Bayesian optimization with dimension scheduling: Application to biological systems. *Computer Aided Chemical Engineering* **38** (11 2015). <https://doi.org/10.1016/B978-0-444-63428-3.50180-6>
 28. Wang, H., Vermetten, D., Ye, F., Doerr, C., Bäck, T.: IOAnalyzer: Performance analysis for iterative optimization heuristic. *ACM Transactions on Evolutionary Learning and Optimization* (2022). <https://doi.org/10.1145/3510426>
 29. Wang, Z., Hutter, F., Zoghi, M., Matheson, D., De Freitas, N.: Bayesian optimization in a billion dimensions via random embeddings (2016)