
High-Dimensional Bayesian Optimization via Nested Riemannian Manifolds

Noémie Jaquier^{1,2}

¹Idiap Research Institute
1920 Martigny, Switzerland
noemie.jaquier@kit.edu

Leonel Rozo²

²Bosch Center for Artificial Intelligence
71272 Renningen, Germany
leonel.rozo@de.bosch.com

Abstract

Despite the recent success of Bayesian optimization (BO) in a variety of applications where sample efficiency is imperative, its performance may be seriously compromised in settings characterized by high-dimensional parameter spaces. A solution to preserve the sample efficiency of BO in such problems is to introduce domain knowledge into its formulation. In this paper, we propose to exploit the geometry of non-Euclidean search spaces, which often arise in a variety of domains, to learn structure-preserving mappings and optimize the acquisition function of BO in low-dimensional latent spaces. Our approach, built on Riemannian manifolds theory, features geometry-aware Gaussian processes that jointly learn a nested-manifold embedding and a representation of the objective function in the latent space. We test our approach in several benchmark artificial landscapes and report that it not only outperforms other high-dimensional BO approaches in several settings, but consistently optimizes the objective functions, as opposed to geometry-unaware BO methods.

1 Introduction

Bayesian optimization (BO) is considered as a powerful machine-learning based optimization method to globally maximize or minimize expensive black-box functions [50]. Thanks to its ability to model complex noisy cost functions in a data-efficient manner, BO has been successfully applied in a variety of applications ranging from hyperparameters tuning for machine learning algorithms [51] to the optimization of parametric policies in challenging robotic scenarios [11, 16, 41, 49]. However, BO performance degrades as the search space dimensionality increases, which recently opened the door to different approaches dealing with the curse of dimensionality.

A common assumption in high-dimensional BO approaches is that the objective function depends on a limited set of features, i.e. that it evolves along an underlying low-dimensional latent space. Following this hypothesis, various solutions based either on random embeddings [57, 43, 9] or on latent space learning [13, 23, 42, 59] have been proposed. Although these methods perform well on a variety of problems, they usually assume simple bound-constrained domains and may not be straightforwardly extended to complicatedly-constrained parameter spaces. Interestingly, several works proposed to further exploit the observed values of the objective function to determine or shape the latent space in a supervised manner [59, 42, 4]. However, the integration of *a priori* domain knowledge related to the parameter space is not considered in the learning process. Moreover, the aforementioned approaches may not comply easily to recover query points in a complex parameter space from those computed on the learned latent space.

Other relevant works in high-dimensional BO substitute or combine the low-dimensional assumption with an additive property, assuming that the objective function is decomposed as a sum of functions of low-dimensional sets of dimensions [33, 37, 21, 44, 24]. Therefore, each low-dimensional partition

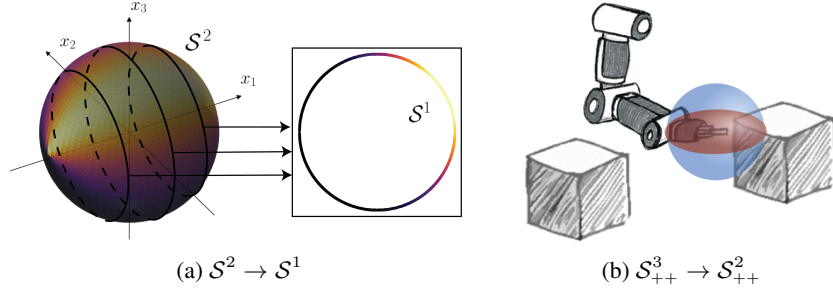


Figure 1: Illustration of the low-dimensional assumption on Riemannian manifolds. (a) The function on \mathcal{S}^2 is not influenced by the value of x_1 and may be represented more efficiently on the manifold \mathcal{S}^1 . (b) The stiffness matrix of a robot is optimized to push objects lying on a table. As the stiffness along the axis x_3 does not influence the pushing skill, the cost function may be better represented in a latent space \mathcal{S}^2_{++} . Note that the manifold's dimensionality is limited here due to the difficulty of visualizing high-dimensional parameter spaces. However, these examples are extensible to higher dimensions.

can be treated independently. In a similar line, inspired by the dropout algorithm in neural networks, other approaches proposed to deal with high-dimensional parameter spaces by optimizing only a random subset of the dimensions at each iteration [36]. Although the aforementioned strategies are well adapted for simple Euclidean parameter spaces, they may not generalize easily to complex domains. If the parameter space is not Euclidean or must satisfy complicated constraints, the problem of partitioning the space into subsets becomes difficult. Moreover, these subsets may not be easily and independently optimized as they must satisfy global constraints acting on the parameters domain.

Introducing domain knowledge into surrogate models and acquisition functions has recently shown to improve the performance and scalability of BO [11, 3, 45, 30, 14]. Following this research line, we hypothesize that building and exploiting geometry-aware latent spaces may improve the performance of BO in high dimensions by considering the intrinsic geometry of the parameter space. Fig. 1 illustrates this idea for two Riemannian manifolds widely used (see § 2 for a short background). The objective function on the sphere \mathcal{S}^2 (Fig. 1a) does not depend on the value x_1 and is therefore better represented on the low-dimensional latent space \mathcal{S}^1 . In Fig. 1b, the stiffness matrix $\mathbf{X} \in \mathcal{S}^3_{++}$ of a robot controller is optimized to push objects lying on a table, with \mathcal{S}^d_{++} the manifold of $d \times d$ symmetric positive definite (SPD) matrices. In this case, the stiffness along the vertical axis x_3 does not influence the robot's ability to push the objects. We may thus optimize the stiffness along the axes x_1 and x_2 , i.e., in the latent space \mathcal{S}^2_{++} . Therefore, similarly to high-dimensional BO frameworks where a Euclidean latent space of the Euclidean parameter space is exploited, the objective functions may be efficiently represented in a latent space that inherits the geometry of the original Riemannian manifold. In general, this latent space is unknown and may not be aligned with the coordinate axes.

Following these observations, this paper proposes a novel high-dimensional geometry-aware BO framework (hereinafter called HD-GaBO) for optimizing parameters lying on low-dimensional Riemannian manifolds embedded in high-dimensional spaces. Our approach is based on a geometry-aware surrogate model that learns both a mapping onto a latent space inheriting the geometry of the original space, and the representation of the objective in this latent space (see § 3). The next query point is then selected on the low-dimensional Riemannian manifold using geometry-aware optimization methods. We evaluate the performance of HD-GaBO on various benchmark functions and show that it efficiently and reliably optimizes high-dimensional objective functions that feature an intrinsic low dimensionality (see § 4). Potential applications of our approach are discussed in § 5.

2 Background

Riemannian Manifolds In machine learning, diverse types of data do not belong to a vector space and thus the use of classical Euclidean methods for treating and analyzing these variables is inadequate. A common example is unit-norm data, widely used to represent directions and orientations, that can be represented as points on the surface of a hypersphere. More generally, many data are normalized in a preprocessing step to discard superfluous scaling and hence are better explained through spherical representations [19]. Notably, spherical representations have been recently exploited to design

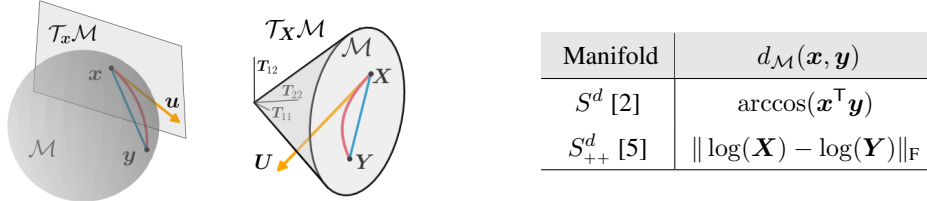


Figure 2: Illustrations of the manifolds S^2 (left) and S_{++}^2 (middle). *Left*: Points on the surface of the sphere, such as \mathbf{x} and \mathbf{y} belong to the manifold. *Middle*: One point corresponds to a matrix $\begin{pmatrix} T_{11} & T_{12} \\ T_{12} & T_{22} \end{pmatrix} \in \text{Sym}^2$ in which the manifold is embedded. For both graphs, the shortest path between \mathbf{x} and \mathbf{y} is the geodesic represented as a red curve, which differs from the Euclidean path depicted in blue. \mathbf{u} lies on the tangent space of \mathbf{x} . The *right* table describes the distance operations on S^d and S_{++}^d .

variational autoencoders [58, 12]. SPD matrices are also extensively used: They coincide with the covariance matrices of multivariate distributions and are employed as descriptors in many applications, such as computer vision [56] and brain-computer interface classification [8]. SPD matrices are also widely used in robotics in the form of stiffness and inertia matrices, controller gains, manipulability ellipsoids, among others.

Both the sphere and the space of SPD matrices can be endowed with a Riemannian metric to form Riemannian manifolds. Intuitively, a Riemannian manifold \mathcal{M} is a mathematical space for which each point locally resembles a Euclidean space. For each point $\mathbf{x} \in \mathcal{M}$, there exists a tangent space $\mathcal{T}_{\mathbf{x}}\mathcal{M}$ equipped with a smoothly-varying positive definite inner product called a Riemannian metric. This metric permits us to define curve lengths on the manifold. These curves, called geodesics, are the generalization of straight lines on the Euclidean space to Riemannian manifolds, as they represent the minimum length curves between two points in \mathcal{M} . Fig. 2 illustrates the two manifolds considered in this paper and details the corresponding distance operations. The unit sphere S^d is a d -dimensional manifold embedded in \mathbb{R}^{d+1} . The tangent space $\mathcal{T}_{\mathbf{x}}S^d$ is the hyperplane tangent to the sphere at \mathbf{x} . The manifold of $d \times d$ SPD matrices S_{++}^d , endowed here with the Log-Euclidean metric [5], can be represented as the interior of a convex cone embedded in its tangent space Sym^d . Supplementary manifold operations used to optimize acquisition functions in HD-GaBO are detailed in Appendix A.

Geometry-aware Bayesian Optimization The geometry-aware BO (GaBO) framework [30] aims at finding a global maximizer (or minimizer) of an unknown objective function f , so that $\mathbf{x}^* = \arg\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$, where the design space of parameters \mathcal{X} is a Riemannian manifold or a subspace of a Riemannian manifold, i.e. $\mathcal{X} \subseteq \mathcal{M}$. With GaBO, geometry-awareness is first brought into BO by modeling the unknown objective function f with a GP adapted to manifold-valued data. This is achieved by defining geometry-aware kernels measuring the similarity of the parameters on \mathcal{M} . In particular, the geodesic generalization of the SE kernel is given by $k(\mathbf{x}_i, \mathbf{x}_j) = \theta \exp(-\beta d_{\mathcal{M}}(\mathbf{x}_i, \mathbf{x}_j)^2)$, where $d_{\mathcal{M}}(\cdot, \cdot)$ denotes the Riemannian distance between two observations and the parameters β and θ control the horizontal and vertical scale of the function [31]. For manifolds that are not isometric to a Euclidean space, this kernel is valid, i.e. positive definite, only for parameters values $\beta > \beta_{\min}$ [18], where β_{\min} can be determined experimentally [17, 30]. Other types of kernels are available for specific manifolds and may also be used in BO (see e.g., [45, 18, 25]).

Secondly, the selection of the next query point \mathbf{x}_{n+1} is achieved by optimizing the acquisition function on the manifold \mathcal{M} . To do so, optimization algorithms on Riemannian manifolds are exploited [2]. These geometry-aware algorithms reformulate constrained problems as an unconstrained optimization on manifolds and consider the intrinsic structure of the space of interest. Also, they tend to show lower computational complexity and better numerical properties [29].

3 High-Dimensional Geometry-aware Bayesian Optimization

In this section, we present the high-dimensional geometry-aware BO (HD-GaBO) framework that naturally handles the case where the design space of parameters \mathcal{X} is (a subspace of) a high-dimensional Riemannian manifold, i.e. $\mathcal{X} \subseteq \mathcal{M}^D$. We assume here that the objective function satisfies the low-dimensional assumption (i.e., some dimensions of the original parameter space

do not influence its value) and thus only varies within a low-dimensional latent space. Moreover, we assume that this latent space can be identified as a low-dimensional Riemannian manifold \mathcal{M}^d inheriting the geometry of the original manifold \mathcal{M}^D , with $d \ll D$. Notice that the same assumption is generally made by Euclidean high-dimensional BO frameworks, as the objective function is represented in a latent space \mathbb{R}^d of \mathbb{R}^D . In particular, we model the objective function $f : \mathcal{M}^D \rightarrow \mathbb{R}$ as a composition of a structure-preserving mapping $m : \mathcal{M}^D \rightarrow \mathcal{M}^d$ and a function $g : \mathcal{M}^d \rightarrow \mathbb{R}$, so that $f = g \circ m$. A model of the objective function is thus available in the latent space \mathcal{M}^d , which is considered as the optimization domain to maximize the acquisition function. As the objective function can be evaluated only in the original space \mathcal{M}^D , the query point $z \in \mathcal{Z}$, with $\mathcal{Z} \subseteq \mathcal{M}^d$, obtained by the acquisition function is projected back into the high-dimensional manifold with the right-inverse projection mapping $m^\dagger : \mathcal{M}^d \rightarrow \mathcal{M}^D$.

In HD-GaBO, the latent spaces are obtained via nested approaches on Riemannian manifolds featuring parametric structure-preserving mappings $m : \mathcal{M}^D \rightarrow \mathcal{M}^d$. Moreover, the parameters Θ_m and Θ_g of the mapping m and function g are determined jointly in a supervised manner using a geometry-aware GP model, as detailed in § 3.1. Therefore, the observed values of the objective function are exploited not only to design the BO surrogate model, but also to drive the dimensionality reduction process towards expressive latent spaces for a data-efficient high-dimensional BO. Considering nested approaches also allows us to build a mapping m^\dagger that can be viewed as the pseudo-inverse of the mapping m . As explained in § 3.3, the corresponding set of parameters Θ_{m^\dagger} includes the projection mapping parameters Θ_m and a set of reconstruction parameters Θ_r , so $\Theta_{m^\dagger} = \{\Theta_m, \Theta_r\}$. Therefore, the parameters Θ_r are determined as to minimize the reconstruction error, as detailed in § 3.2. Similarly to GaBO [30], geometry-aware kernel functions are used in HD-GaBO (see § 3.1), and the acquisition function is optimized using techniques on Riemannian manifolds, although the optimization is carried out on the latent Riemannian manifold in HD-GaBO. The proposed HD-GaBO framework is summarized in Algorithm 1.

Algorithm 1: HD-GaBO

Input: Initial observations $\mathcal{D}_0 = \{(\mathbf{x}_i, y_i)\}_{i=1}^{N_0}$, $\mathbf{x}_i \in \mathcal{M}^D$, $y_i \in \mathbb{R}$

Output: Final recommendation \mathbf{x}_N

```

1 for  $n = 0, 1, \dots, N$  do
2   Update the hyperparameters  $\{\Theta_m, \Theta_g\}$  of the geometry-aware mGP model ;
3   Project the observed data into the latent space, so that  $\mathbf{z}_i = m(\mathbf{x}_i)$  ;
4   Select the next query point  $\mathbf{z}_{n+1} \in \mathcal{M}^d$  by optimizing the acquisition function in the latent space, i.e.,
       $\mathbf{z}_{n+1} = \operatorname{argmax}_{\mathbf{z} \in \mathcal{Z}} \gamma_n(\mathbf{z}; \{(\mathbf{z}_i, y_i)\})$  ;
5   Update the hyperparameters  $\Theta_{m^\dagger}$  of the pseudo-inverse projection ;
6   Obtain the new query point  $\mathbf{x}_{n+1} = m^\dagger(\mathbf{z}_{n+1})$  in the original space ;
7   Query the objective function to obtain  $y_{n+1}$  ;
8   Augment the set of observed data  $\mathcal{D}_{n+1} = \{\mathcal{D}_n, (\mathbf{x}_{n+1}, y_{n+1})\}$  ;
9 end

```

3.1 HD-GaBO Surrogate Model

The choice of latent spaces is crucial for the efficiency of HD-GaBO as it determines the search space for the selection of the next query point \mathbf{x}_{n+1} . In this context, it is desirable to base the latent-space learning process not only on the distribution of the observed parameters \mathbf{x}_n in the original space, but also on the quality of the corresponding values y_n of the objective function. Therefore, we propose (i) to supervisedly learn a structure-preserving mapping onto a low-dimensional latent space, and (ii) to learn the representation of the objective function in this latent space along with the corresponding mapping. To do so, we exploit the so-called manifold Gaussian process (mGP) model introduced in [10]. It is important to notice that the term *manifold* denotes here a latent space, whose parameters are learned by the mGP, which does not generally correspond to a Riemannian manifold.

In a mGP, the regression process is considered as a composition $g \circ m$ of a parametric projection m onto a latent space and a function g . Specifically, a mGP is defined as a GP so that $f \sim \mathcal{GP}(\mu_m, k_m)$ with mean function $\mu_m : \mathcal{X} \rightarrow \mathbb{R}$ and positive-definite covariance function $k_m : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ defined as $\mu_m(\mathbf{x}) = \mu(m(\mathbf{x}))$ and $k_m(\mathbf{x}_i, \mathbf{x}_j) = k(m(\mathbf{x}_i), m(\mathbf{x}_j))$, with $\mu : \mathcal{Z} \rightarrow \mathbb{R}$ and $k : \mathcal{Z} \times \mathcal{Z} \rightarrow \mathbb{R}$ a kernel function. The mGP parameters are estimated by maximizing the marginal likelihood of the model, so that $\{\Theta_m^*, \Theta_g^*\} = \operatorname{argmax}_{\Theta_m, \Theta_g} p(\mathbf{y}|\mathbf{X}, \Theta_m, \Theta_g)$.

In mGP [10], the original and latent spaces are subspaces of Euclidean spaces, so that $\mathcal{X} \subseteq \mathbb{R}^D$ and $\mathcal{Z} \subseteq \mathbb{R}^d$, respectively. Note that the idea of jointly learning a projection mapping and a representation of the objective function with a mGP was also exploited in the context of high-dimensional BO in [42]. In [10, 42], the mapping $m : \mathbb{R}^D \rightarrow \mathbb{R}^d$ was represented by a neural network. However, in the HD-GaBO framework, the design parameter space $\mathcal{X} \subseteq \mathcal{M}^D$ is a high-dimensional Riemannian manifold and we aim at learning a geometry-aware latent space $\mathcal{Z} \subseteq \mathcal{M}^d$ that inherits the geometry of \mathcal{X} . Thus, we define a structure-preserving mapping $m : \mathcal{M}^D \rightarrow \mathcal{M}^d$ as a nested projection from a high- to a low-dimensional Riemannian manifold of the same type, as described in § 3.3. Moreover, as in GaBO, we use a geometry-aware kernel function k that allows the GP to properly measure the similarity between parameters $z = m(x)$ lying on the Riemannian manifold \mathcal{M}^d . Therefore, the surrogate model of HD-GaBO is a geometry-aware mGP, that leads to a geometry-aware representation of the objective function in a locally optimal low-dimensional Riemannian manifold \mathcal{M}^d .

Importantly, the predictive distribution for the mGP $f \sim \mathcal{GP}(\mu_m, k_m)$ at test input \tilde{x} is equivalent to the predictive distribution of the GP $g \sim \mathcal{GP}(\mu, k)$ at test input $\tilde{z} = m(\tilde{x})$. Therefore, the predictive distribution can be straightforwardly computed in the latent space. This allows the optimization function to be defined and optimized in the low-dimensional Riemannian manifold \mathcal{M}^d instead of the original high-dimensional parameter space \mathcal{M}^D . Then, the selected next query point z_{n+1} in the latent space needs to be projected back onto \mathcal{M}^D in order to evaluate the objective function.

3.2 Input Reconstruction from the Latent Embedding to the Original Space

After optimizing the acquisition function, the selected query point z_{n+1} in the latent space needs to be projected back onto the manifold \mathcal{M}^D in order to evaluate the objective function. For solving this problem in the Euclidean case, Moriconi et al. [42] proposed to learn a reconstruction mapping $r : \mathbb{R}^d \rightarrow \mathbb{R}^D$ based on multi-output GPs. In contrast, we propose here to further exploit the nested structure-preserving mappings in order to project the selected query point back onto the original manifold. As shown in § 3.3, a right-inverse parametric projection $m^\dagger : \mathcal{M}^d \rightarrow \mathcal{M}^D$ can be built from the nested Riemannian manifold approaches. This pseudo-inverse mapping depends on a set of parameters $\Theta_{m^\dagger} = \{\Theta_m, \Theta_r\}$. Note that the parameters Θ_m are learned with the mGP surrogate model, but we still need to determine the reconstruction parameters Θ_r . While the projection mapping m aimed at finding an optimal representation of the objective function, the corresponding pseudo-inverse mapping m^\dagger should (ideally) project the data z lying on the latent space \mathcal{M}^d onto their original representation x in the original space \mathcal{M}^D . Therefore, the parameters Θ_r are obtained by minimizing the sum of the squared residuals on the manifold \mathcal{M}^D , so that

$$\Theta_r^* = \underset{\Theta_r}{\operatorname{argmin}} \sum_{i=1}^n d_{\mathcal{M}^D}^2(x_i, m^\dagger(z_i; \Theta_m, \Theta_r)). \quad (1)$$

3.3 Nested Manifolds Mappings

As mentioned previously, the surrogate model of HD-GaBO learns to represent the objective function in a latent space \mathcal{M}^d inheriting the geometry of the original space \mathcal{M}^D . To do so, the latent space is obtained via nested approaches, which map a high-dimensional Riemannian manifold to a low-dimensional latent space inheriting the geometry of the original Riemannian manifold. While various other dimensionality reduction techniques have been proposed on Riemannian manifolds [20, 52, 53, 28, 46], the resulting latent space is usually formed by curves on the high-dimensional manifold \mathcal{M}^D . This would still require to optimize the acquisition function on \mathcal{M}^D with complex constraints, which may not be handled efficiently by optimization algorithms. In contrast, nested manifold mappings reduce the dimension of the search space in a systematic and structure-preserving manner, so that the acquisition function can be efficiently optimized on a low-dimensional Riemannian manifold with optimization techniques on Riemannian manifolds. Moreover, intrinsic latent spaces may naturally be encoded with nested manifold mappings in various applications (see Fig. 1). Nested mappings for the sphere and SPD manifolds are presented in the following.

Sphere manifold The concept of nested spheres, introduced in [32], is illustrated in Fig. 3. Given an axis $v \in \mathcal{S}^D$, the sphere is first rotated so that v aligns with the origin, typically defined as the north pole $(0, \dots, 0, 1)^\top$. Then, the data $x \in \mathcal{S}^D$ (in purple) are projected onto the subsphere \mathcal{A}^{D-1} defined as $\mathcal{A}^{D-1}(v, r) = \{w \in \mathcal{S}^D : d_{\mathcal{S}^D}(v, w) = r\}$, where $r \in (0, \pi/2]$, so that

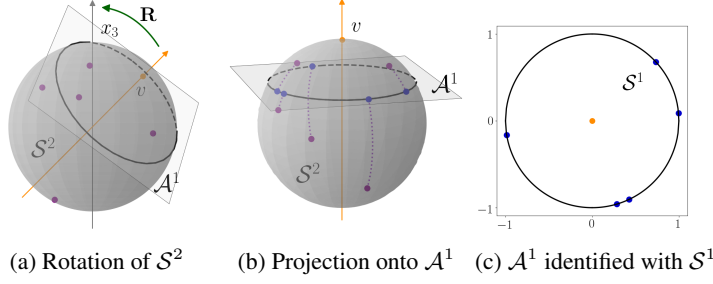


Figure 3: Illustration of the nested sphere projection mapping. Data on the sphere \mathcal{S}^2 , depicted by purple dots, are projected onto the subsphere \mathcal{A}^1 , which is then identified with the sphere \mathcal{S}^1 .

$x_D = \cos(r)$. The last coordinate of \mathbf{x} is then discarded and the data $\mathbf{z} \in \mathcal{S}^{D-1}$ (in blue) are obtained by identifying the subsphere \mathcal{A}^{D-1} of radius $\sin(r)$ with the nested unit sphere \mathcal{S}^{D-1} via a scaling operation. Specifically, given an axis $\mathbf{v}_D \in \mathcal{S}^D$ and a distance $r_D \in (0, \pi/2]$, the projection mapping $m_D : \mathcal{S}^D \rightarrow \mathcal{S}^{D-1}$ is computed as

$$\mathbf{z} = m_D(\mathbf{x}) = \underbrace{\frac{1}{\sin(r_D)}}_{\text{scaling}} \underbrace{\mathbf{R}_{\text{trunc}}}_{\text{rotation + dim. red.}} \underbrace{\left(\frac{\sin(r_D)\mathbf{x} + \sin(d_{\mathcal{S}^D}(\mathbf{v}_D, \mathbf{x}) - r_D)\mathbf{v}_D}{\sin(d_{\mathcal{S}^D}(\mathbf{v}_D, \mathbf{x}))} \right)}_{\text{projection onto } \mathcal{A}^{D-1}}, \quad (2)$$

with $d_{\mathcal{S}^D}$ defined as in the table of Fig. 2, $\mathbf{R} \in \text{SO}(D)$ is the rotation matrix that moves \mathbf{v} to the origin on the manifold and $\mathbf{R}_{\text{trunc}}$ the matrix composed of the $D-1$ first rows of \mathbf{R} . Notice also that the order of the projection and rotation operations is interchangeable. In (2), the data are simultaneously rotated and reduced after being projected onto \mathcal{A}^{D-1} . However, the same result may be obtained by projecting the rotated data $\mathbf{R}\mathbf{x}$ onto \mathcal{A}^{D-1} using the rotated axis $\mathbf{R}\mathbf{v}$ and multiplying the obtained vector by the truncated identity matrix $\mathbf{I}_{\text{trunc}} \in \mathbb{R}^{(D-1) \times D}$. This fact will be later exploited to define the SPD nested mapping. Then, the full projection mapping $m : \mathcal{S}^D \rightarrow \mathcal{S}^d$ is defined via successive mappings (2), so that $m = m_{d+1} \circ \dots \circ m_{D-1} \circ m_D$, with parameters $\{\mathbf{v}_D, \dots, \mathbf{v}_{d+1}, r_D, \dots, r_{d+1}\}$ such that $\mathbf{v}_k \in \mathcal{S}^k$ and $r_k \in (0, \pi/2]$. Importantly, notice that the distance $d_{\mathcal{S}^d}(m(\mathbf{x}_i), m(\mathbf{x}_j))$ between two points $\mathbf{x}_i, \mathbf{x}_j \in \mathcal{S}^D$ projected onto \mathcal{S}^d is invariant w.r.t the distance parameters $\{r_D, \dots, r_{d+1}\}$ (see Appendix B for a proof). Therefore, when using distance-based kernels, the parameters set of the mGP projection mapping corresponds to $\Theta_m = \{\mathbf{v}_D, \dots, \mathbf{v}_{d+1}\}$. The mGP parameters optimization is thus carried out with techniques on Riemannian manifolds on the domain $\mathcal{S}^D \times \dots \times \mathcal{S}^{d+1} \times \mathcal{M}_g$, where \mathcal{M}_g is the space of GP parameters Θ_g (usually $\mathcal{M}_g \sim \mathbb{R} \times \dots \times \mathbb{R}$).

As shown in [32], an inverse transformation $m_D^{-1} : \mathcal{S}^{D-1} \rightarrow \mathcal{S}^D$ can be computed as

$$\mathbf{x} = m_D^{-1}(\mathbf{z}) = \mathbf{R}^\top \begin{pmatrix} \sin(r_{d+1})\mathbf{z} \\ \cos(r_{d+1}) \end{pmatrix}. \quad (3)$$

Therefore, the query point selected by the acquisition function in the latent space can be projected back onto the original space with the inverse projection mapping $m^\dagger : \mathcal{S}^d \rightarrow \mathcal{S}^D$ given by $m^\dagger = m_D^\dagger \circ \dots \circ m_{d+1}^\dagger$. As the axes parameters are determined within the mGP model, the set of reconstruction parameters is given by $\Theta_r = \{r_D, \dots, r_{d+1}\}$.

SPD manifold Although not explicitly named as such, the dimensionality reduction technique for the SPD manifold introduced in [26, 27] can be understood as a nested manifold mapping. Specifically, Harandi et al. [26, 27] proposed a projection mapping $m : \mathcal{S}_{++}^D \rightarrow \mathcal{S}_{++}^d$, so that

$$\mathbf{Z} = m(\mathbf{X}) = \mathbf{W}^\top \mathbf{X} \mathbf{W}, \quad (4)$$

with $\mathbf{W} \in \mathbb{R}^{D \times d}$. Note that the matrix $\mathbf{Z} \in \mathcal{S}_{++}^d$ is guaranteed to be positive definite if \mathbf{W} has a full rank. As proposed in [26, 27], this can be achieved, without loss of generality, by imposing orthogonality constraint on \mathbf{W} such that $\mathbf{W} \in \mathcal{G}_{D,d}$, i.e., $\mathbf{W}^\top \mathbf{W} = \mathbf{I}$, where $\mathcal{G}_{D,d}$ denotes the Grassmann manifold corresponding to the space of d -dimensional subspaces of \mathbb{R}^D [15]. Therefore, in the case of the SPD manifold, the projection mapping parameter set is $\Theta_m = \{\mathbf{W}\}$. Specifically, the mGP parameters are optimized on the product of Riemannian manifolds $\mathcal{G}_{D,d} \times \mathcal{M}_g$. Also, the optimization of the mGP on the SPD manifold can be simplified as shown in Appendix C.

In order to project the query point $\mathbf{Z} \in \mathcal{S}_{++}^d$ back onto the original space \mathcal{S}_{++}^D , we propose to build an inverse projection mapping based on m . It can be easily observed that using the pseudo-inverse \mathbf{W} so

that $\mathbf{X} = \mathbf{W}^\top \mathbf{Z} \mathbf{W}^\dagger$ does not guarantee the recovered matrix \mathbf{X} to be positive definite. Therefore, we propose a novel inverse mapping inspired by the nested sphere projections. To do so, we observe that an analogy can be drawn between the mappings (2) and (4). Namely, the mapping (4) first consists of a rotation $\mathbf{R}^\top \mathbf{X} \mathbf{R}$ of the data $\mathbf{X} \in \mathcal{S}_{++}^D$ with \mathbf{R} a rotation matrix whose D first columns equal \mathbf{W} , i.e., $\mathbf{R} = (\mathbf{W} \ \mathbf{V})$, where \mathbf{W} can be understood as $\mathbf{R}_{\text{trunc}}$ in Eq. (2). Similarly to the nested sphere case, the rotated data can be projected onto a subspace of the manifold \mathcal{S}_{++}^D by fixing their last coordinates. Therefore, the subspace is composed of matrices $\begin{pmatrix} \mathbf{W}^\top \mathbf{X} \mathbf{W} & \mathbf{C} \\ \mathbf{C}^\top & \mathbf{B} \end{pmatrix}$, where $\mathbf{B} \in \mathcal{S}_{++}^{D-d}$ is a constant matrix. Finally, this subspace may be identified with \mathcal{S}_{++}^d by multiplying the projected matrix $\begin{pmatrix} \mathbf{W}^\top \mathbf{X} \mathbf{W} & \mathbf{C} \\ \mathbf{C}^\top & \mathbf{B} \end{pmatrix}$ with a truncated identity matrix $\mathbf{I}_{\text{trunc}} \in \mathbb{R}^{D \times d}$. Therefore, the mapping (4) is equivalently expressed as $\mathbf{Z} = m(\mathbf{X}) = \mathbf{I}_{\text{trunc}}^\top \begin{pmatrix} \mathbf{W}^\top \mathbf{X} \mathbf{W} & \mathbf{C} \\ \mathbf{C}^\top & \mathbf{B} \end{pmatrix} \mathbf{I}_{\text{trunc}} = \mathbf{W}^\top \mathbf{X} \mathbf{W}$.

From the properties of block matrices with positive block-diagonal elements, the projection is positive definite if and only if $\mathbf{W}^\top \mathbf{X} \mathbf{W} \geq \mathbf{C} \mathbf{B} \mathbf{C}^\top$ [6]. This corresponds to defining the side matrix as $\mathbf{C} = (\mathbf{W}^\top \mathbf{X} \mathbf{W})^{\frac{1}{2}} \mathbf{K} \mathbf{B}^{\frac{1}{2}}$, where $\mathbf{K} \in \mathbb{R}^{d \times D-d}$ is a contraction matrix, so that $\|\mathbf{K}\| \leq 1$ [6]. Based on the aforementioned equivalence, the inverse mapping $m^\dagger : \mathcal{S}_{++}^d \rightarrow \mathcal{S}_{++}^D$ is given by

$$\mathbf{X} = m^\dagger(\mathbf{Z}) = \mathbf{R} \begin{pmatrix} \mathbf{Z} & \mathbf{Z}^{\frac{1}{2}} \mathbf{K} \mathbf{B}^{\frac{1}{2}} \\ \mathbf{B}^{\frac{1}{2}} \mathbf{K}^\top \mathbf{Z}^{\frac{1}{2}} & \mathbf{B} \end{pmatrix} \mathbf{R}^\top, \quad (5)$$

with reconstruction parameters $\Theta_r = \{\mathbf{V}, \mathbf{K}, \mathbf{B}\}$. The optimization (1) is thus carried out on the product of manifolds $\mathcal{G}_{D-d,d} \times \mathbb{R}^{d,D-d} \times \mathcal{S}_{++}^{D-d}$ subject to $\|\mathbf{K}\| \leq 1$ and $\mathbf{W}^\top \mathbf{V} = \mathbf{0}$. The latter condition is necessary for \mathbf{R} to be a valid rotation matrix. We solve this optimization problem with the augmented Lagrangian method on Riemannian manifolds [38].

4 Experiments

In this section, we evaluate the proposed HD-GaBO framework to optimize high-dimensional functions that lie on an intrinsic low-dimensional space. We consider benchmark test functions defined on a low-dimensional manifold \mathcal{M}^d embedded in a high-dimensional manifold \mathcal{M}^D . Therefore, the test functions are defined as $f : \mathcal{M}^D \rightarrow \mathbb{R}$, so that $y = f(m(\mathbf{x}))$ with $m : \mathcal{M}^D \rightarrow \mathcal{M}^d$ being the nested projection mapping, as defined in Section 3.3. The projection mapping parameters are randomly set for each trial. The search space corresponds to the complete manifold for \mathcal{S}^D and to SPD matrices with eigenvalues $\lambda \in [0.001, 5]$ for \mathcal{S}_{++}^D . We carry out the optimization by running 30 trials with random initialization. Both GaBO and HD-GaBO use the geodesic generalization of the SE kernel and their acquisition functions are optimized using trust region on Riemannian manifolds [1] (see Appendix D). The other state-of-the-art approaches use the classical SE kernel and the constrained acquisition functions are optimized using sequential least squares programming [34]. All the tested methods use EI as acquisition function and are initialized with 5 random samples. The GP parameters are estimated using MLE. All the implementations employ GPyTorch [22], BoTorch [7] and Pymanopt [55]. Source code is available at <https://github.com/NoemieJaquier/GaBOtorch>. Supplementary results are presented in Appendix F.

In the case of the sphere manifold \mathcal{S}^D , we compare HD-GaBO against GaBO, the Euclidean BO and three high-dimensional BO approaches, namely dropout BO [36], SIR-BO [59], and REMBO [57], which carry out all the operations in the Euclidean space. The optimization of the acquisition function of each Euclidean BO method was adapted to fulfill the constraint $\|\mathbf{x}\| = 1$. Other approaches, such as the MGPC-BO of [42], are not considered here due to the difficulty of adapting them when the parameters lie on Riemannian manifolds. We minimize the Rosenbrock, Ackley, and product-of-sines functions (see also Appendix E) defined on the low-dimensional manifold \mathcal{S}^5 embedded in \mathcal{S}^{50} . Fig. 4a- 4c display the median of the logarithm of the simple regret along 300 BO iterations and the distribution of the logarithm of the BO recommendation \mathbf{x}_N for the three functions. We observe that HD-GaBO generally converges fast and provides good optimizers for all the test cases. Moreover, it outperforms all the other BO methods for the product-of-sines function: it provides fast convergence and better optimizer with low variance. In contrast, SIR-BO, which leads to the best optimizer for the Rosenbrock function, performs poorly to optimize the product-of-sines function. Similarly, dropout achieves a similar performance as HD-GaBO for the Ackley function, but it is outperformed by HD-GaBO in the two other test cases. Moreover, it is worth noticing that GaBO converges faster to

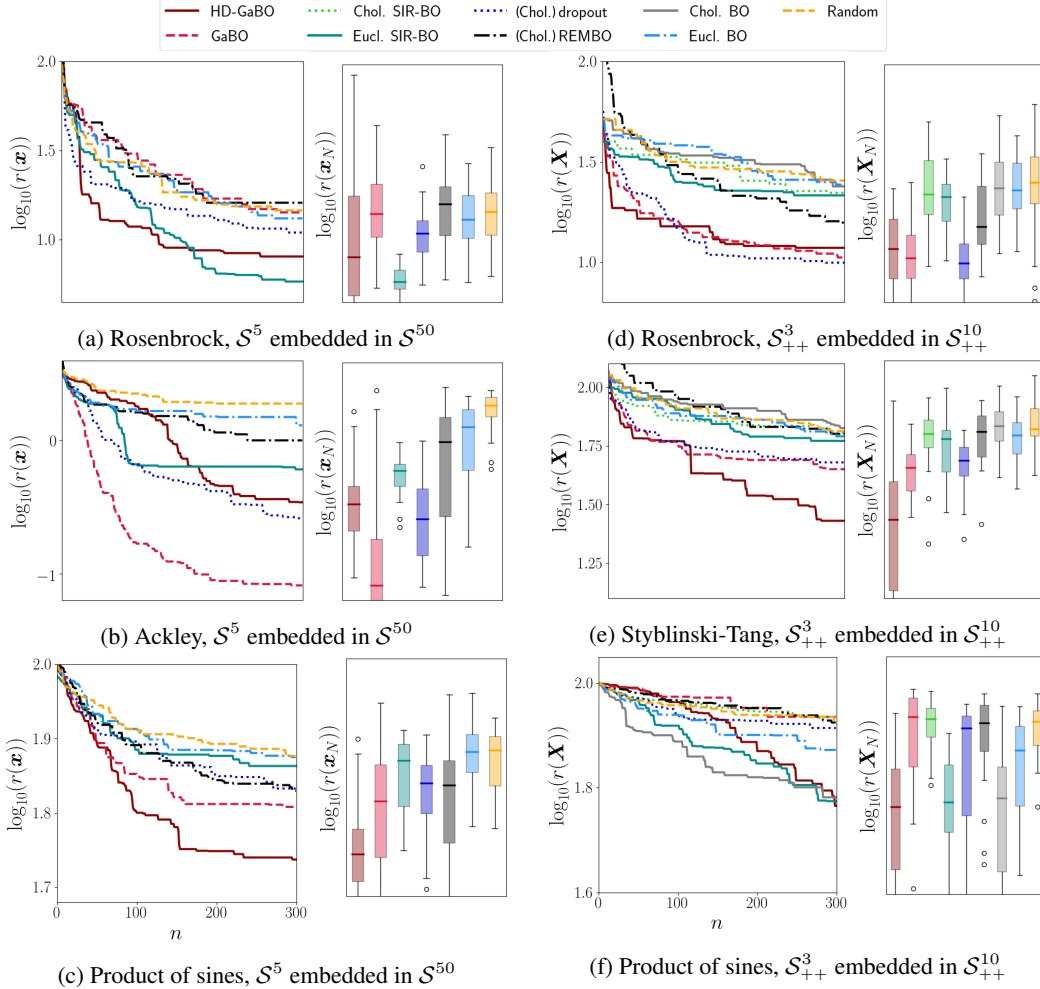


Figure 4: Logarithm of the simple regret for benchmark test functions over 30 trials. The *left* graphs show the evolution of the median for the BO approaches and the random search baseline. The *right* graphs display the distribution of the logarithm of the simple regret of the BO recommendation \mathbf{x}_N after 300 iterations. The boxes extend from the first to the third quartiles and the median is represented by a horizontal line. Supplementary results are provided in Appendix F.

the best optimizer than the other approaches for the Ackley function and performs better than all the geometry-unaware approaches for the product-of-sines function. This highlights the importance of using geometry-aware approaches for optimizing objective functions lying on Riemannian manifolds.

Regarding the SPD manifold \mathcal{S}_{++}^D , we compare HD-GaBO against GaBO, the Euclidean BO and SIR-BO (augmented with the constraint $\lambda_{\min} > 0$). Moreover, we consider alternative implementations of BO, dropout, SIR-BO and REMBO that exploit the Cholesky decomposition of an SPD matrix $\mathbf{A} = \mathbf{L}\mathbf{L}^T$, so that the resulting parameter is the vectorization of the lower triangular matrix \mathbf{L} (hereinafter denoted as Cholesky-methods). Note that we do not consider here the Euclidean version of the dropout and REMBO methods due to the difficulty of optimizing the acquisition function in the latent space while satisfying the constraint $\lambda_{\min} > 0$ for the query point in the high-dimensional manifold. We minimize the Rosenbrock, Styblinski-Tang, and product-of-sines functions defined on the low-dimensional manifold \mathcal{S}_{++}^3 embedded in \mathcal{S}_{++}^{10} . The corresponding results are displayed in Fig. 4d-4f (in logarithm scale). We observe that HD-GaBO consistently converges fast and provides good optimizers for all the test cases. Moreover, it outperforms all the other approaches for the Styblinski-Tang function. Similarly to the sphere cases, some methods are still competitive with respect to HD-GaBO for some of the test functions but perform poorly in other cases. Interestingly, GaBO performs well for both Rosenbrock and Styblinski-Tang functions. Moreover, the Euclidean

BO methods generally perform poorly compared to their Cholesky equivalences, suggesting that, although they do not account for the manifold geometry, Cholesky-based approaches provide a better representation of the SPD parameter space than the Euclidean methods.

5 Potential Applications

After evaluating the performance of HD-GaBO in various benchmark artificial landscapes, we discuss potential real-world applications of the proposed approach. First, HD-GaBO may be exploited for the optimization of controller parameters in robotics. Of particular interest is the optimization of the error gain matrix $\mathbf{Q}_t \in \mathcal{S}_{++}^{D_x}$ and control gain matrix $\mathbf{R}_t \in \mathcal{S}_{++}^{D_u}$ in linear quadratic regulators (LQR), where D_x and D_u are the dimensionality of the system state and control input, respectively. The system state may consist of the linear and angular position and velocity of the robot end-effector, so that $D_x = 13$, and D_u corresponds to Cartesian accelerations or wrench commands. Along some parts of the robot trajectory, the error w.r.t. some dimensions of the state space may not influence the execution of the task, i.e., affect negligibly the LQR cost function. Therefore, the matrix \mathbf{Q}_t for this trajectory segment may be efficiently optimized in a latent space $\mathcal{S}_{++}^{d_x}$ with $d_x < D_x$. A similar analysis applies for \mathbf{R} . Notice that, although BO has been applied to optimize LQR parameters [40, 41], the problem was greatly simplified as only diagonal matrices \mathbf{Q} and \mathbf{R} were considered in the optimization, resulting in a loss of flexibility in the controller. From a broader point of view, the low-dimensional assumption may also apply in the optimization of gain matrices for other types of controllers.

Another interesting application is the identification of dynamic model parameters of (highly-) redundant robots. These parameters typically include the inertia matrix $\mathbf{M} \in \mathcal{S}_{++}^D$ with D being the number of robot joints. As discussed in [60], a low-dimensional representation of the parameter space and state-action space may be sufficient to determine the system dynamics. Therefore, the inertia matrix may be more efficiently represented and identified in a lower-dimensional SPD latent space.

In the context of directional statistics [54, 47], HD-GaBO may be used to adapt mixtures of von Mises-Fisher distributions, whose mean directions belong to \mathcal{S}^D . On a different topic, object shape spaces are typically characterized on high-dimensional unit spheres \mathcal{S}^D . Several works have shown that the main features of the shapes are efficiently represented in a low-dimensional latent space \mathcal{S}^d inheriting the geometry of the original manifold (see e.g., [32]). Therefore, such latent spaces may be exploited for shape representation optimization. Along a similar line, skeletal models, which seek at capturing the interior of objects, lie on a Cartesian product of manifolds that involves the unit hypersphere [48]. The relevant data structure is efficiently expressed in a product of low-dimensional manifolds of the same types, so that HD-GaBO may be exploited to optimize skeletal models.

6 Conclusion

In this paper, we proposed HD-GaBO, a high-dimensional geometry-aware Bayesian optimization framework that exploited geometric prior knowledge on the parameter space to optimize high-dimensional functions lying on low-dimensional latent spaces. To do so, we used a geometry-aware GP that jointly learned a nested structure-preserving mapping and a representation of the objective function in the latent space. We also considered the geometry of the latent space while optimizing the acquisition function and took advantage of the nested mappings to express the next query point in the high-dimensional parameter space. We showed that HD-GaBO not only outperformed other BO approaches in several settings, but also consistently performed well while optimizing various objective functions, unlike geometry-unaware state-of-the-art methods.

An open question, shared across various high-dimensional BO approaches, concerns the model dimensionality mismatch. In order to avoid suboptimal solutions where the optimum of the function may not be included in the estimated latent space, we hypothesize that the dimension d should be selected slightly higher in case of uncertainty on its value [35]. A limitation of HD-GaBO is that it depends on nested mappings that are specific to each Riemannian manifold. Therefore, such mappings may not be available for all kinds of manifolds. Also, the inverse map does not necessarily exist if the manifold contains self-intersection. In this case, a non-parametric reconstruction mapping may be learned (e.g., based on wrapped GP [39]). However, most of the Riemannian manifolds encountered in machine learning and robotics applications do not self-intersect, so that this problem is avoided. Future work will investigate the aforementioned aspects.

Broader Impact

The HD-GaBO formulation presented in this paper makes a step towards more explainable and interpretable BO approaches. Indeed, in addition to the benefits in terms of performance, the inclusion of domain knowledge via Riemannian manifolds into the BO framework permits to treat the space parameters in a principled way. This can notably be contrasted with approaches based on random features, that generally remain hard to interpret for humans. As often, the gains in terms of explainability and interpretability come at the expense of the low computational cost that characterizes random-based approaches. However, the carbon footprint of the proposed approach remains low compared to many deep approaches used nowadays in machine learning applications.

Acknowledgments and Disclosure of Funding

This work was mainly developed during a PhD sabbatical at the Bosch Center for Artificial Intelligence (Renningen, Germany). This work was also partially supported by the FNS/DFG project TACT-HAND, as part of the PhD thesis of the first author, carried out at the Idiap Research Institute (Martigny, Switzerland), while also affiliated to the Ecole Polytechnique Fédérale de Lausanne (Lausanne, Switzerland). Noémie Jaquier is now affiliated with the Karlsruhe Institute of Technology (Karlsruhe, Germany).

References

- [1] P.-A. Absil, C. G. Baker, and K. A. Gallivan. Trust-region methods on Riemannian manifolds. *Foundations of Computational Mathematics*, 7:303–330, 2007.
- [2] P. A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2007.
- [3] R. Antonova, A. Rai, and C. Atkeson. Deep kernels for optimizing locomotion controllers. In *Conference on Robot Learning (CoRL)*, pages 47–56, 2017.
- [4] R. Antonova, A. Rai, T. Li, and D. Kragic. Bayesian optimization in variational latent spaces with dynamic compression. In *Conference on Robot Learning (CoRL)*, 2019.
- [5] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 56(2):411–421, 2006.
- [6] Rajendra B. *Positive Definite Matrices*. Princeton University Press, 2007.
- [7] M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy. BoTorch: Programmable bayesian optimization in PyTorch. *arXiv preprint 1910.06403*, 2019.
- [8] A. Barachant, S. Bonnet, M. Congedo, and C. Jutten. Multiclass brain-computer interface classification by Riemannian geometry. *IEEE Trans. on Biomedical Engineering*, 59(4):920–928, 2012.
- [9] M. Binois, D. Ginsbourger, and O. Roustant. On the choice of the low-dimensional domain for global optimization via random embeddings. *Journal of Global Optimization*, 76(1):69–90, 2020.
- [10] R. Calandra, J. Peters, C. E. Rasmussen, and M. P. Deisenroth. Manifold Gaussian processes for regression. In *Proc. IEEE Intl Joint Conf. on Neural Networks (IJCNN)*, 2016.
- [11] A. Cully, J. Clune, D. Tarapore, and J. B. Mouret. Robots that can adapt like animals. *Nature*, 521:503–507, 2015.
- [12] T. R. Davidson, L. Falorsi, N. De Cao, T. Kipf, and J. M. Tomczak. Hyperspherical variational auto-encoders. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2018.
- [13] J. Djolonga, A. Krause, and V. Cevher. High-dimensional Gaussian process bandits. In *Neural Information Processing Systems (NeurIPS)*, 2013.
- [14] D. K. Duvenaud. *Automatic Model Construction with Gaussian Processes*. PhD thesis, University of Cambridge, 2014.
- [15] A. Edelman, T. A. Arias, and S. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal of Matrix Analysis and Applications*, 20(2):303–351, 1998.

- [16] Peter Englert and Marc Toussaint. Combined optimization and reinforcement learning for manipulations skills. In *Robotics: Science and Systems (R:SS)*, 2016.
- [17] A. Feragen and S. Hauberg. Open problem: Kernel methods on manifolds and metric spaces. what is the probability of a positive definite geodesic exponential kernel? In *29th Annual Conference on Learning Theory*, pages 1647–1650, 2016.
- [18] A. Feragen, F. Lauze, and S. Hauberg. Geodesic exponential kernels: When curvature and linearity conflict. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [19] N. I. Fisher, T. Lewis, and B. J. J. Embleton. *Statistical analysis of spherical data*. Cambridge University Press, 1987.
- [20] P. T. Fletcher and S. C. Joshi. Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. In *In Proc. of CVAMIA and MMBIA Workshops*, pages 87–98, 2004.
- [21] J. R. Gardner, C. Guo, K. Q. Weinberger, R. Garnett, and R. Grosse. Discovering and exploiting additive structure for Bayesian optimization. In *Proc. of the Intl Conf. on Artificial Intelligence and Statistics (AISTATS)*, pages 1311–1319, 2017.
- [22] J. R. Gardner, G. Pleiss, D. Bindel, K. Q. Weinberger, and A. G. Wilson. GPyTorch: Blackbox matrix-matrix gaussian process inference with GPU acceleration. In *Neural Information Processing Systems (NeurIPS)*, 2018.
- [23] R. Garnett, M. A. Osborne, and P. Hennig. Active learning of linear embeddings for Gaussian processes. In *Conference of Uncertainty in Artificial Intelligence (UAI)*, pages 230–239, 2014.
- [24] D. Gaudrie, R. Le Riche, V. Picheny, B. Enaux, and V. Herbert. Modeling and optimization with Gaussian processes in reduced eigenbases. *Structural and Multidisciplinary Optimization*, 61(6):2343–2361, 2020.
- [25] B. Gong, Y. Shi, F. Sha, and K. Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2066–2073, 2012.
- [26] M. Harandi, M. Salzmann, and R. Hartley. From manifold to manifold: Geometry-aware dimensionality reduction for spd matrices. In *Proc. European Conf. on Computer Vision (ECCV)*, 2014.
- [27] M. Harandi, M. Salzmann, and R. Hartley. Dimensionality reduction on spd manifolds: The emergence of geometry-aware methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(1):48–62, 2018.
- [28] S. Hauberg. Principal curves on Riemannian manifolds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(9):1915–1921, 2016.
- [29] J. Hu, X. Liu, Z. Wen, and Y. Yuan. A brief introduction to manifold optimization. *arXiv preprint 1906.05450*, 2019.
- [30] N. Jaquier, L. Rozo, S. Calinon, and M. Bürger. Bayesian optimization meets Riemannian manifolds in robot learning. In *Conference on Robot Learning (CoRL)*, 2019.
- [31] S. Jayasumana, R. Hartley, M. Salzmann, H. Li, and M. Harandi. Kernel methods on Riemannian manifolds with Gaussian RBF kernels. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 37(12):2464–2477, 2015.
- [32] S. Jung, I. L. Dryden, and J. S. Marron. Analysis of principal nested spheres. *Biometrika*, 99(3): 551–568, 2012.
- [33] K. Kandasamy, J. Schneider, and B. Póczos. High dimensional Bayesian optimisation and bandits via additive models. In *Intl. Conf. on Machine Learning (ICML)*, 2015.
- [34] D. Kraft. A software package for sequential quadratic programming. Technical report, Technical Report DFVLR-FB 88-28, Institut für Dynamik der Flugsysteme, Oberpfaffenhofen, 1988.
- [35] B. Letham, R. Calandra, A. Rai, and E. Bakshy. Re-examining linear embeddings for high-dimensional Bayesian optimization. In *Neural Information Processing Systems (NeurIPS)*, 2020.
- [36] C. Li, S. Gupta, S. Rana, V. Nguyen, S. Venkatesh, and A. Shilton. High dimensional Bayesian optimization using dropout. In *Intl. Joint Conf. on Artificial Intelligence (IJCAI)*, pages 2096–2102, 2017.

- [37] C.-L. Li, K. Kandasamy, B. Póczos, and J. Schneider. High dimensional Bayesian optimization via restricted projection pursuit models. In *Proc. of the Intl Conf. on Artificial Intelligence and Statistics (AISTATS)*, 2016.
- [38] C. Liu and N. Boumal. Simple algorithms for optimization on Riemannian manifolds with constraints. *Applied Mathematics & Optimization*, pages 1–33, 2019.
- [39] A. Mallasto and A. Feragen. Wrapped Gaussian process regression on Riemannian manifolds. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 5580–5588, 2018.
- [40] A. Marco, P. Hennig, J. Bohg, S. Schaal, and S. Trimpe. Automatic LQR tuning based on Gaussian process global optimization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 270–277, 2016.
- [41] A. Marco, P. Hennig, S. Schaal, and S. Trimpe. On the design of LQR kernels for efficient controller learning. In *IEEE Conference on Decision and Control (CDC)*, pages 5193–5200, 2017.
- [42] R. Moriconi, M. P. Deisenroth, and K. S. Sesh Kumar. High-dimensional Bayesian optimization using low-dimensional feature spaces. *Machine Learning*, 109:1925–1943, 2020.
- [43] A. Munteanu, A. Nayebi, and M. Poloczek. A framework for Bayesian optimization in embedded subspaces. In *Intl. Conf. on Machine Learning (ICML)*, volume 97, pages 4752–4761, 2019.
- [44] M. Mutný and A. Krause. Efficient high dimensional Bayesian optimization with additivity and quadrature fourier features. In *Neural Information Processing Systems (NeurIPS)*, 2018.
- [45] C. Oh, E. Gavves, and M. Welling. BOCK: Bayesian optimization with cylindrical kernels. In *Intl. Conf. on Machine Learning (ICML)*, pages 3868–3877, 2018.
- [46] X. Pennec. Barycentric subspace analysis on manifolds. *Annals of Statistics*, 46(6A):2711–2746, 2018.
- [47] Arthur Pewsey and Eduardo García-Portugués. Recent advances in directional statistics. *arXiv preprint 2005.06889*, 2020.
- [48] S. M. Pizer, S. Jung, D. Goswami, J. Vicory, X. Zhao, R. Chaudhuri, J. N. Damon, S. Huckemann, and J. S. Marron. Nested sphere statistics of skeletal models. *Innovations for Shape Analysis*, pages 93–115, 2012.
- [49] A. Rai, R. Antonova, S. Song, W. Martin, H. Geyer, and C. Atkeson. Bayesian optimization using domain knowledge on the ATRIAS biped. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1771–1778, 2018.
- [50] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- [51] J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Neural Information Processing Systems (NeurIPS)*, page 2951–2959, 2012.
- [52] S. Sommer, F. Lauze, S. Hauberg, and M. Nielsen. Manifold valued statistics, exact principal geodesic analysis and the effect of linear approximations. In *European Conf. On Computer Vision*, pages 43–56, 2010.
- [53] S. Sommer, F. Lauze, and M. Nielsen. Optimization over geodesics for exact principal geodesic analysis. *Advances in Computational Mathematics*, 40(2):283–313, 2014.
- [54] S. Sra. Directional statistics in machine learning: a brief review. In C. Ley and T. Verdebout, editors, *Applied Directional Statistics*, Chapman & Hall/CRC Interdisciplinary Statistics Series, pages 259–276. CRC Press, Boca Raton, 2018.
- [55] J. Townsend, N. Koep, and S. Weichwald. Pymanopt: A python toolbox for optimization on manifolds using automatic differentiation. *Journal of Machine Learning Research*, 17(137): 1–5, 2016.
- [56] O. Tuzel, F. Porikli, and P. Meer. Region covariance: A fast descriptor for detection and classification. In *European Conference on Computer Vision (ECCV)*, pages 589–600, 2006.
- [57] Z. Wang, M. Zoghiy, F. Hutterz, D. Matheson, and N. De Freitas. Bayesian optimization in high dimensions via random embeddings. In *Intl. Joint Conf. on Artificial Intelligence (IJCAI)*, pages 1778–1784, 2013.

- [58] J. Xu and G. Durrett. Spherical latent spaces for stable variational autoencoders. In *In Proc. of Conf. on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- [59] M. Zhang, H. Li, and S. Su. High dimensional Bayesian optimization via supervised dimension reduction. In *Proc. of Intl Joint Conf. on Artificial Intelligence (IJCAI)*, 2019.
- [60] S. Zhu, D. Surovik, K. Bekris, and A. Boularias. Efficient model identification for tensegrity locomotion. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 2985–2990, 2018.