
High-Dimensional Gaussian Process Bandits

Josip Djolonga
ETH Zürich
josipd@ethz.ch

Andreas Krause
ETH Zürich
krausea@ethz.ch

Volkan Cevher
EPFL
volkan.cevher@epfl.ch

Abstract

Many applications in machine learning require optimizing unknown functions defined over a high-dimensional space from noisy samples that are expensive to obtain. We address this notoriously hard challenge, under the assumptions that the function varies only along some low-dimensional subspace and is smooth (i.e., it has a low norm in a Reproducible Kernel Hilbert Space). In particular, we present the SI-BO algorithm, which leverages recent low-rank matrix recovery techniques to learn the underlying subspace of the unknown function and applies Gaussian Process Upper Confidence sampling for optimization of the function. We carefully calibrate the exploration–exploitation tradeoff by allocating the sampling budget to subspace estimation and function optimization, and obtain the first subexponential cumulative regret bounds and convergence rates for Bayesian optimization in high-dimensions under noisy observations. Numerical results demonstrate the effectiveness of our approach in difficult scenarios.

1 Introduction

The optimization of non-linear functions whose evaluation may be noisy and expensive is a challenge that has important applications in sciences and engineering. One approach to this notoriously hard problem takes a Bayesian perspective, which uses the predictive uncertainty in order to trade exploration (gathering data for reducing model uncertainty) and exploitation (focusing sampling near likely optima), and is often called Bayesian Optimization (BO). Modern BO algorithms are quite successful, surpassing even human experts in learning tasks: e.g., gait control for the SONY AIBO, convolutional neural networks, structural SVMs, and Latent Dirichlet Allocation [1, 2, 3].

Unfortunately, the theoretical efficiency of these methods depends exponentially on the—often high—dimension of the domain over which the function is defined. A way to circumvent this “curse of dimensionality” is to make the assumption that only a small number of the dimensions actually matter. For example, the cost function of neural networks effectively varies only along a few dimensions [2]. This idea has been also at the root of nonparametric regression approaches [4, 5, 6, 7].

To this end, we propose an algorithm that learns a low dimensional, *not necessarily axis-aligned*, subspace and then applies Bayesian optimization on this estimated subspace. In particular, our SI-BO approach combines low-rank matrix recovery with Gaussian Process Upper Confidence Bound sampling in a carefully calibrated manner. We theoretically analyze its performance, and prove bounds on its cumulative regret. *To the best of our knowledge, we prove the first subexponential bounds for Bayesian optimization in high dimensions under noisy observations.* In contrast to existing approaches, which have an exponential dependence on the ambient dimension, our bounds have in fact polynomial dependence on the dimension. Moreover, our performance guarantees depend explicitly on what we could have achieved if we had known the subspace in advance.

Previous work. Exploration–exploitation tradeoffs were originally studied in the context of finite multi-armed bandits [8]. Since then, results have been obtained for continuous domains, starting with the linear [9] and Lipschitz-continuous cases [10, 11]. A more recent algorithm that enjoys theoretical bounds for functions sampled from a Gaussian Process (GP), or belong to some Repro-

ducible Kernel Hilbert Space (RKHS) is GP-UCB [12]. The use of GPs to negotiate exploration–exploitation tradeoffs originated in the areas of response surface and Bayesian optimization, for which there are a number of approaches (cf., [13]), perhaps most notably the Expected Improvement [14] approach, which has recently received theoretical justification [15], albeit only in the noise-free setting.

Bandit algorithms that exploit low-dimensional structure of the function appeared first for the linear setting, where under sparsity assumptions one can obtain bounds, which depend only weakly on the ambient dimension [16, 17]. In [18] the more general case of functions sampled from a GP under the same sparsity assumptions was considered. The idea of random projections to BO has been recently introduced [19]. They provide bounds on the simple regret under noiseless observations, while we also analyze the cumulative regret and allow noisy observations. Also, unless the low-dimensional space is of dimension 1, our bounds on the simple regret improve on theirs. In [7] the authors approximate functions that live on low-dimensional subspaces using low-rank recovery and analysis techniques. While providing uniform approximation guarantees, their approach is not tailored towards exploration–exploitation tradeoffs, and does not achieve sublinear cumulative regret. In [20] the stochastic and adversarial cases for axis-aligned Hölder continuous functions are considered.

Our specific contributions in this paper can be summarized as follows:

- We introduce the SI-BO algorithm for Bayesian bandit optimization in high dimensions, admitting a large family of kernel functions. Our algorithm is a natural but non-trivial fusion of modern low-rank subspace approximation tools with GP optimization methods.
- We derive theoretical guarantees on SI-BO’s cumulative and simple regret in high-dimensions *with noise*. To the best of our knowledge, these are the first theoretical results on the sample complexity and regret rates that are subexponential in the ambient dimension.
- We provide experimental results on synthetic data and classical benchmarks.

2 Background and Problem Statement

Goal. In plain words, we wish to sequentially optimize a bounded function over a compact, convex subset $D \subset \mathbb{R}^d$. Without loss of generality, we denote the function by $f : D \rightarrow [0, 1]$ and let \mathbf{x}^* be a maximizer. The algorithm proceeds in a total of T rounds. In each round t , it asks an oracle for the function value at some point \mathbf{x}_t and it receives back the value $f(\mathbf{x}_t)$, possibly corrupted by noise. Our goal is to choose points such that their values are close to the optimum $f(\mathbf{x}^*)$.

As performance metric, we consider the regret, which tells us how much better we could have done in round t had we known \mathbf{x}^* , or formally $r_t = f(\mathbf{x}^*) - f(\mathbf{x}_t)$. In many applications, such as recommender systems, robotic control, etc., we care about the quality of the points chosen at every time step t . Hence, a natural quantity to consider is the *cumulative regret* defined as $R_T = \sum_{t=1}^T r_t$. One can also consider the *simple regret*, defined as $S_T = \min_{t=1}^T r_t$, measuring the quality of the best solution found so far. We will give bounds on the more challenging notion of cumulative regret, which also bounds the simple regret via $S_T \leq R_T/T$.

Low-dimensional functions in high-dimensions. Unfortunately, our problem cannot be tractably solved without further assumptions on the properties of the function f . What is worse is that the usual compact support and smoothness assumptions cannot achieve much: the minimax lower bound on the sample complexity is exponential in d [21, 6, 7]. We hence assume that the function effectively varies only along a small number of true active dimensions: i.e., the function lives on a $k \ll d$ -dimensional subspace. Typically, k or an upper bound on k is assumed known [4, 5, 7, 6].

Formally, we suppose that there exists some function $g : \mathbb{R}^k \rightarrow [0, 1]$ and a matrix $A \in \mathbb{R}^{k \times d}$ with orthogonal rows so that $f(\mathbf{x}) = g(A\mathbf{x})$. We will additionally assume that $g \in \mathcal{C}^2$, which is necessary to bound the errors from the linear approximation that we will make. Further, w.l.o.g., we assume that $D = \mathbb{B}^d(1 + \bar{\varepsilon})$ for some $\bar{\varepsilon} > 0$, where we define $\mathbb{B}^d(r)$ to be the closed ball around 0 of radius r in \mathbb{R}^d .¹ To be able to recover the subspace we also need the condition that g has Lipschitz continuous second order derivatives and a full rank Hessian at $\mathbf{0}$, which is satisfied for many functions [7].

Smooth, low-complexity functions. In addition to the low-dimensional subspace assumption, we also assume that g is smooth. One way to encode our prior is to assume that the function g resides in

¹Our method can be extended to any convex compact set, see Section 5.2 in [22].

Algorithm 1 The SI-BO algorithm

Require: $m_X, m_\Phi, \lambda, \varepsilon, k$, oracle for f , kernel κ

$\mathcal{C} \leftarrow m_X$ samples uniformly from \mathbb{S}^{d-1}

for $i \leftarrow 1$ to m_X **do**

$\Phi_i \leftarrow m_\Phi$ samples uniformly from $\{\pm 1/\sqrt{m_\Phi}\}^k$

$\mathbf{y} \leftarrow$ compute using Equation 1

$\hat{X}_{DS} \leftarrow$ Dantzig Selector using \mathbf{y} , see Equation 2 and compute the SVD $\hat{X}_{DS} = \hat{U}\hat{\Sigma}\hat{V}^T$

$\hat{A} \leftarrow \hat{U}^{(k)}$ // Principal k vectors of \hat{U} , $\mathcal{D} \leftarrow$ all $(\hat{A}\mathbf{x}, y)$ pairs queried so far

Use GP inference to obtain $\mu_1(\cdot), \sigma_1(\cdot)$.

for $t \leftarrow 1$ to $T - m_X(m_\Phi + 1)$ **do**

$\mathbf{z}_t \leftarrow \arg \max_{\mathbf{z}} \mu_t(\mathbf{z}) + \beta_t^{1/2} \sigma_t(\mathbf{z}), y_t \leftarrow f(\hat{A}^T \mathbf{z}_t) + \text{noise}, \mathcal{D}.add(\mathbf{z}_t, y_t)$

a Reproducing Kernel Hilbert Space (RKHS; cf., [23]), which allows us to quantify g 's complexity via its norm $\|g\|_{\mathcal{H}_\kappa}$. The RKHS for some positive semidefinite kernel $\kappa(\cdot, \cdot)$ can be constructed by completing the set of functions $\sum_{i=1}^n \alpha_i \kappa(\mathbf{x}_i, \cdot)$ under a suitable inner product. In this work, we use isotropic kernels, i.e., those that depend only on the distance between points, since the problem is rotation invariant and we can only recover A up to some rotation.

Here is a final summary of our problem and its underlying assumptions:

1. We wish to maximize $f : \mathbb{B}^d(1 + \varepsilon) \rightarrow [0, 1]$, where $f(\mathbf{x}) = g(A\mathbf{x})$ for some matrix $A \in \mathbb{R}^{k \times d}$ with orthogonal rows and g belongs to some RKHS \mathcal{H}_κ .
2. The kernel κ is isotropic $\kappa(\mathbf{x}, \mathbf{x}') = \kappa'(\|\mathbf{x} - \mathbf{x}'\|_2) = \kappa''(\|\mathbf{x} - \mathbf{x}'\|_2)$ and κ' is continuous, integrable and with a Fourier transform $\mathcal{F}\kappa'$ that is isotropic and radially non-increasing.²
3. The function g has Lipschitz continuous 2nd-order derivatives and a full rank Hessian at 0.
4. The function g is \mathcal{C}^2 on a compact support and $\max_{|\beta| \leq 2} \|D^\beta g\|_\infty \leq C_2$ for some $C_2 > 0$.
5. The oracle noise is Gaussian with zero mean with a known variance σ^2 .

3 The SI-BO Algorithm

The SI-BO algorithm performs two separate exploration and exploitation stages: (1) subspace identification (SI), i.e. estimating the subspace on which the function is supported, and then (2) Bayesian optimization (BO), in order to optimize the function on the learned subspace. A key challenge here is to carefully allocate samples between these phases.

We first give a detailed outline for SI-BO in Alg. 1, deferring its theoretical analysis to Section 4. Given the (noisy) oracle for f , we first evaluate the function at several suitably chosen points and then use a low-rank recovery algorithm to compute a matrix \hat{A} that spans a subspace well aligned with the one generated by the true matrix A . Once we have computed \hat{A} , similarly to [22, 7], we define the function which we optimize as $\hat{g}(\mathbf{z}) = f(\hat{A}^T \mathbf{z}) = g(A\hat{A}^T \mathbf{z})$. Thus, we effectively work with an approximation \hat{f} to f given by $\hat{f}(\mathbf{x}) = \hat{g}(\hat{A}\mathbf{x}) = g(A\hat{A}^T \hat{A}\mathbf{x})$. With the approximation at hand, we apply BO, in particular the GP-UCB algorithm, on \hat{g} for the remaining steps.

Subspace Learning. We learn A using the approach from [7], which reduces the learning problem to that of low rank matrix recovery. We construct a set of m_X points $\mathcal{C} = [\xi_1, \dots, \xi_{m_X}]$, which we call sampling centers, and consider the matrix X of gradients at those points $X = [\nabla f(\xi_1), \dots, \nabla f(\xi_{m_X})]$. Using the chain rule, we have $X = A^T [\nabla g(A\xi_1), \dots, \nabla g(A\xi_{m_X})]$. Because A is a matrix of size $k \times d$ it follows that the rank of X is at most k . This suggests that using low-rank approximation techniques, one may be able to (up to rotation) infer A from X .

Given that we have no access to the gradients of f directly, we approximate X using a linearization of f . Consider a fixed sampling center ξ . If we make a linear approximation with step size ε to the directional derivative at center ξ in direction φ then, by Taylor's theorem, for a suitable $\zeta(\mathbf{x}, \varepsilon, \varphi)$:

$$\langle \varphi, A^T \nabla g(A\xi) \rangle = \frac{1}{\varepsilon} (f(\xi + \varepsilon\varphi) - f(\xi)) - \underbrace{\frac{\varepsilon}{2} \varphi^T \nabla^2 f(\zeta) \varphi}_{E(\xi, \varepsilon, \varphi)}.$$

²This is the same assumption as in [15]. Radially non-increasing means that if $\|w\| \leq \|w'\|$ then $\mathcal{F}\kappa'(w) \geq \mathcal{F}\kappa'(w')$. Note that this is satisfied by the RBF and Matérn kernels.

Thus, sampling the finite difference $f(\xi + \varepsilon\varphi) - f(\xi)$ provides (up to the curvature error $E(\xi, \varepsilon, \varphi)$, and sampling noise) information about the one-dimensional subspace spanned by $A^T \nabla g(A\xi)$. To estimate it accurately, we must observe multiple directions φ . Further, to infer the full k -dimensional subspace A , we need to consider at least $m_X \geq k$ centers. Consequently, for each center ξ_i , we define a set of m_Φ directions and arrange them in a total of m_Φ matrices $\Phi_i = [\varphi_{i,1}, \varphi_{i,2}, \dots, \varphi_{i,m_X}] \in R^{d \times m_X}$. We can now define the following linear system:

$$\mathbf{y} = \mathcal{A}(X) + \mathbf{e} + \mathbf{z}, \quad y_i = \frac{1}{\varepsilon} \sum_{j=1}^{m_X} (f(\xi_j + \varepsilon\varphi_{i,j}) - f(\xi_j)), \quad (1)$$

where the linear operator \mathcal{A} is defined as $\mathcal{A}(X)_i = \text{tr}(\Phi_i^T X)$, the curvature errors have been accumulated in \mathbf{e} and the noise has been put in the vector \mathbf{z} which is distributed as $z_i \sim \mathcal{N}(0, 2m_X \sigma^2 / \varepsilon)$.

Given the structure of the problem, we can make use of several low-rank recovery algorithms. For concreteness, we choose the Dantzig Selector (DS, [24]), which recovers low rank matrices via

$$\underset{M}{\text{minimize}} \quad \|M\|_* \quad \text{subject to} \quad \|\underbrace{\mathcal{A}^*(\mathbf{y} - \mathcal{A}(M))}_{\text{residual}}\| \leq \lambda, \quad (2)$$

where $\|\cdot\|_*$ is the nuclear norm and $\|\cdot\|$ is the spectral norm. The DS will successfully recover a matrix \hat{X} close to the true solution in the Frobenius norm and moreover this distance decreases linearly with λ . As shown in [7], choosing the centers \mathcal{C} uniformly at random from the unit sphere \mathbb{S}^{d-1} , choosing each direction vector uniformly at random from $\{\pm 1/\sqrt{m_\Phi}\}^k$, and—in the case of noisy observations, resampling f repeatedly—suffices to obtain an accurate \hat{X} w.h.p., as long as m_Φ and m_X are sufficiently large. The precise choices of these quantities are analyzed in Section 4.

Finally, we extract the matrix \hat{A} from the SVD of \hat{X} , by taking its top k left singular vectors. Because the DS will find a matrix \hat{X} close to X , due to a result by Wedin [25] we know that the learned subspace will be close to the true one.

Optimizing \hat{g} . Once we have an approximate \hat{A} , we optimize the function $\hat{g}(\mathbf{z}) = f(\hat{A}^T \mathbf{z})$ on the low-dimensional domain $\mathcal{Z} = \mathbb{B}^k(1+\bar{\varepsilon})$. Concretely, we use GP-UCB [12], because it exhibits state of the art empirical performance, and enjoys strong theoretical bounds for the cumulative regret. It requires that \hat{g} belongs to the RKHS and the noise when conditioned on the history is zero-mean and almost surely bounded by some $\hat{\sigma}$. Section 4 shows that this is indeed true with high probability.

In order to trade exploration and exploitation, the GP-UCB algorithm computes, for each point \mathbf{z} , a score that combines the predictive mean that we have inferred for that point with its variance, which quantifies the uncertainty in our estimate. They are combined linearly with a time-dependent weighting factor β_t in the following surrogate function

$$\text{ucb}(\mathbf{z}) = \mu_t(\mathbf{z}) + \beta_t^{1/2} \sigma_t(\mathbf{z}) \quad (3)$$

for a suitably chosen $\beta_t = 2B + 300\gamma_t \log^3(t/\delta)$. Here, B is an upper bound on the squared RKHS norm of the function that we optimize, δ is an upper bound on the failure probability and γ_t depends on the kernel [12]: cf., Section 4.³ The algorithm then greedily maximizes the ucb score above.

Note that finding the maximum of this non-convex and in general multi-modal function, while considered to be cheaper than evaluating f at a new point, is by itself a hard problem and it is usually approached by either sampling on a grid in the domain, or using some global Lipschitz optimizer [13]. Hence, by reducing the dimension of the domain \mathcal{Z} over which we have to optimize, our algorithm has the additional benefit that this process can be performed more efficiently.

Handling the noise. The last ingredient that we need is theory on how to pick $\hat{\sigma}$ so that it bounds the noise during the execution of GP-UCB w.h.p., and how to select λ in (2) so that the true matrix X is feasible in the DS. Due to the fast decay of the tails of the Gaussian distribution we can pick $\hat{\sigma} = \left(2 \log \frac{1}{\delta} + 2 \log T + \log \frac{1}{2\pi}\right)^{1/2} \sigma$, where T is the number of GP-UCB iterations and σ^2 is the variance of the noise. Then the noise will be trapped in $[-\hat{\sigma}, \hat{\sigma}]$ with probability at least $1 - \delta$.

³If the bound B is not known beforehand then one can use a doubling trick.




Figure 1: A 2-dimensional function $f(x, y)$ varying along a 1-dimensional subspace and its projections on different subspaces. The numbers are the respective cosine distances.

The analysis on λ comes from [7]. They bound $\|\mathcal{A}^*(\mathbf{e} + \mathbf{z})\|$ using the assumption that the second order derivatives are bounded and, as shown in [24], because \mathbf{z} has a Gaussian distribution,

$$\|\mathcal{A}^*(\mathbf{e} + \mathbf{z})\| \leq 1.2 \left(\frac{C_2 \varepsilon d m_X k^2}{2\sqrt{m_\Phi}} + \frac{5\sqrt{m_X m_\Phi} \sigma}{\varepsilon} \right) \quad (4)$$

If there is no noise it still holds by setting $\sigma = 0$. This bound, intuitively, relates the approximation quality λ of the subspace to the quantities m_Φ , m_X as well as the step size ε .

4 Theoretical Analysis

Overview. A crucial choice in our algorithm is how to allocate samples (by choosing m_Φ and m_X appropriately) to the tasks of subspace learning and function optimization. We now analyze both phases, and determine how to split the queries in order to optimize the cumulative regret bounds.

Let us first consider the regret incurred in the second phase, in the ideal (but unrealistic) case that the subspace is estimated exactly (i.e., $\hat{A} = A$). This question was answered recently in [12], where it is proven that it is bounded by $\mathcal{O}^*(\sqrt{T}(B\sqrt{\gamma_t} + \gamma_t))^4$. Hereby, the quantity γ_T is defined as

$$\gamma_T = \max_{S \subseteq D, |S|=T} H(\mathbf{y}_S) - H(\mathbf{y}_S | f),$$

where \mathbf{y}_S are the values of f at the points in S , corrupted by Gaussian noise, and $H(\cdot)$ is the entropy. It quantifies the maximal gain in information that we can obtain about f by picking a set of T points. In [12] sublinear bounds for γ_T have been computed for several popular kernels. For example, for the RBF kernel in k dimensions, $\gamma_T = \mathcal{O}((\log T)^{k+1})$. Further, B is a bound on the squared norm $\|g\|_{\mathcal{H}_\kappa}^2$ of g w.r.t. kernel κ . Note that generally γ_T grows exponentially with k , rendering the application of GP-UCB directly to the high-dimensional problem intractable.

What happens if the subspace \hat{A} is estimated incorrectly? Fortunately, w.h.p. the estimated function \hat{g} still remains in the RKHS associated with kernel κ . However, the norm $\|\hat{g}\|_{\mathcal{H}_\kappa}$ may increase, and consequently may the regret. Moreover, the considered \hat{f} disagrees with the true f , and consequently additional regret per sample may be incurred by $\eta = \|\hat{f} - f\|_\infty$. As an illustration of the effect of misestimated subspaces see Figure 1. We can observe that subspaces far from the true one stretch the function more, thus increasing its RKHS norm.

We now state a general result that formalizes these insights by bounding the cumulative regret in terms of the samples allocated to subspace learning, and the subspace approximation quality.

Lemma 1 *Assume that we spend $0 < n \leq T$ samples to learn the subspace such that $\|f - \hat{f}\|_\infty \leq \eta$, $\|\hat{g}\| \leq B$ and the error is bounded by $\hat{\sigma}$, each w.p. at least $1 - \delta/4$. If we run the GP-UCB algorithm for the remaining $T - n$ steps with the suggested $\hat{\sigma}$ and $\delta/4$, then the following bound on the cumulative regret holds w.p. at least $1 - \delta$*

$$R_T \leq n + \underbrace{\eta T}_{\text{approx. error}} + \underbrace{\mathcal{O}^*(\sqrt{T}(B\sqrt{\gamma_t} + \gamma_t))}_{R_{UCB}(T, \hat{g}, \kappa)}$$

⁴We have used the notation $\mathcal{O}^*(f) = \mathcal{O}(f \log f)$ to suppress the log factors. $\Omega^*(\cdot)$ is analogously defined.




Figure 2: Approximations \hat{g} resulting from differently aligned subspaces. Note that inaccurate estimation (the middle two cases) can wildly distort the objective.

where $R_{UCB}(T, \hat{g}, \kappa)$ is the regret of GP-UCB when run for T steps using \hat{g} and kernel κ ⁵.

Lemma 1 breaks down the regret in terms of the approximation error incurred by subspace-misestimation, and the optimization error incurred by the resulting increased complexity $\|\hat{g}\|_{\mathcal{H}_\kappa}^2 \leq B$. We now analyze these effects, and then prove our main regret bounds.

Effects of Subspace Alignment. A notion that will prove to be very helpful for analyzing both, the approximation precision η and the norm of \hat{g} , is the set of angles between the subspaces that are defined by A and \hat{A} . The following definition [26] makes this notion precise.

Definition 2 Let $A, \hat{A} \in \mathbb{R}^{k \times d}$ be two matrices with orthogonal rows so that $AA^T = \hat{A}\hat{A}^T = I$. We define the vector of cosines between the spanned subspaces $\cos \Theta(A, \hat{A})$ to be equal to the singular values of $A\hat{A}^T$. Analogously $\sin \Theta(A, \hat{A})_i = (1 - \cos \Theta(A, \hat{A})_i^2)^{1/2}$.

Let us see how \hat{A} affects \hat{g} . Because $\hat{g}(\mathbf{z}) = g(A\hat{A}^T\mathbf{z})$, the matrix $M = A\hat{A}^T$, which converts any point from its coordinates determined by \hat{A} to the coordinates defined by A , will be of crucial importance. First, note that its singular values are cosines and are between -1 and 1 . This means that it can only shrink the vectors that we apply it to (possibly by different amounts in different directions). The effect on \hat{g} is that it might only “see” a small part of the whole space, and its shape might be distorted, which in turn will increase its RKHS complexity (see Figure 2 for an illustration).

Lemma 3 If $g \in \mathcal{H}_\kappa$ for a kernel that is isotropic with a radially non-increasing Fourier transform and $\hat{g}(x) = g(A\hat{A}^T x)$ for some A, \hat{A} with orthogonal rows, then for $C = C_2\sqrt{2k}(1 + \varepsilon)$,

$$\|f - \hat{f}\|_\infty \leq C\|\sin \Theta(A, \hat{A})\|_2 \quad \text{and} \quad \|\hat{g}\|_{\mathcal{H}_\kappa}^2 \leq |\text{prod} \cos \Theta(A, \hat{A})|^{-1} \|g\|_{\mathcal{H}_\kappa}^2. \quad (5)$$

Here, we use the notation $\text{prod} \mathbf{x} = \prod_{i=1}^d x_i$ to denote the product of the elements of a vector. By decreasing the angles we tackle both issues: the approximation error $\eta = \|f - \hat{f}\|_\infty$ is reduced and the norm of \hat{g} gets closer to the one of g . There is one nice interpretation of the product of the cosines. It is equal to the determinant of the matrix M . Hence, \hat{g} will not be in the RKHS only if M is rank deficient as dimensions are collapsed.

Regret Bounds. We now present our main bounds on the cumulative regret. In order to achieve sublinear regret, we need a way of controlling η and $\|\hat{g}\|_{\mathcal{H}_\kappa}$. In the following, we show how this goal can be achieved. As it turns out, subspace learning is substantially harder in the case of noisy observations. Therefore, we focus on the easier, noise-free setting first.

Noiseless Observations. We should note that the theory behind GP-UCB still holds in the deterministic case, as it only requires the noise to be bounded a.s. by $\hat{\sigma}$. The following theorem guarantees that in this setting for non-linear kernels we have a regret dominated by GP-UCB, which is of order $\Omega^*(\sqrt{T}\gamma_T)$, as it is usually exponential in k .

Theorem 4 If the observations are noiseless we can pick $m_x = \mathcal{O}(kd \log 1/\delta)$, $\varepsilon = \frac{1}{k^{2.25}d^{3/2}T^{1/2}}$ and $m_\varphi = \mathcal{O}(k^2d \log 1/\delta)$ so that with probability at least $1 - \delta$ we have the following

$$R_T \leq \mathcal{O}(k^3d^2 \log^2(1/\delta)) + 2 R_{UCB}(T, g, \kappa).$$

⁵ Because the noise parameter $\hat{\sigma}$ depends on T , we have to slightly change the bounds from [12] as we have a term of order $\mathcal{O}(\sqrt{\log T + \log(1/\delta)})$; c.f. supplementary material.

Noisy Observations. Equation 4 hints that the noise can have a dramatic effect in learning efficiency. As already mentioned, the DS gets better results as we decrease λ . In the noiseless case, it suffices to increase the number of directions m_Φ and decrease the step size ε in estimating the finite differences. However, the second term in λ can only be reduced by decreasing the variance σ^2 .

As a result, each point that we evaluate is sampled n times and we take as its value the average. Moreover, note that because the standard deviation decreases as $1/\sqrt{n}$, we have to resample at least ε^{-2} times and this significantly increases the number of samples that we need. Nevertheless, we are able to obtain cumulative regret bounds (and thereby the first convergence guarantees and rates) for this setting, which *only polynomially depend on d* . Unfortunately, the dependence on T is now weaker than those in the noiseless setting (Theorem 4), and the regret due to the subspace learning might dominate that of GP-UCB.

Theorem 5 *If the observations are noisy, we can pick $\varepsilon = \frac{1}{k^{2.25}d^{1.5}T^{1/5}}$ and all other parameters as in the previous theorem. Moreover, we have to resample each point $\mathcal{O}(\sigma^2 k^2 d T^{2/5} m_\Phi / \varepsilon^2)$ times. Then, with probability at least $1 - \delta$*

$$R_T \leq \mathcal{O}\left(\sigma^2 k^{11.5} d^7 T^{4/5} \log^3(1/\delta)\right) + 2 R_{UCB}(T, g, \kappa).$$

Mismatch on the effective dimension k . All models are imperfect in some sense and the structure of a general f is impossible to identify unless we have further scientific evidence beyond the data. In our case, the assumption $f(\mathbf{x}) = g(A\mathbf{x})$ for some k more or less takes the weakest form for indicating our hope that BO can succeed from a sub-exponential sample size. In general, we must tune k in a degree to reflect the anticipated complexity in the learning problem. Fortunately, all the guarantees are preserved if we assume a $k > k_{\text{true}}$, for some true synthetic model, where $f(\mathbf{x}) = g(A\mathbf{x})$ holds. Underfitting k leads to additional errors that are well-controlled in low-rank subspace estimation [24]. The impact of under fitting in our setting is left for future work.

5 Experiments

The main intent of our experiments is to provide a proof of concept, confirming that SI-BO not just in theory provides the first subexponential regret bounds, but also empirically obtains low average regret for Bayesian optimization in high dimensions.

Baselines. We compare SI-BO against the following baseline approaches:

- RandomS-UCB, which runs GP-UCB on a random subspace.
- RandomH-UCB, which runs GP-UCB on the high-dimensional space. At each iteration we pick 1000 points at random and choose the one with highest UCB score.
- Exact-UCB, which runs GP-UCB on the exact (but in practice unknown) subspace.

The β_t parameter in the GP-UCB score was set as recommended in [12] for finite sets. To optimize the UCB score we sampled on a grid on the low dimensional subspace. For all of the measurements we have added Gaussian zero-mean noise with $\sigma = 0.01$.

Data sets. We carry out experiments in the following settings:

- *GP Samples.* We generate random 2-dimensional samples from a GP with Matérn kernel with smoothness parameter $\nu = 5/2$, length scale $\ell = 1/2$ and signal variance $\sigma_f^2 = 1$. The samples are “hidden” in a random 2-dimensional subspace in 100 dimensions.
- *Gabór Filters.* The second data set is inspired by experimental design in neuroscience [27]. The goal is to determine visual stimuli that maximally excite some neuron, which reacts to edges in the images. We consider the function $f(\mathbf{x}) = \exp(-(\theta^T \mathbf{x} - 1)^2)$, where θ is a Gabór filter of size 17×17 and the set of admissible signals is $[0, 1]^d$.

In the appendix we also include results for the Branin function, a classical optimization benchmark.

Results. The results are presented in Figure 3. We show the averages of 20 runs (10 runs for *GP-Posterior*) and the shaded areas represent the standard error around the mean. We show both the average regret and simple regret (i.e., suboptimality of the best solution found so far). We find that although SI-BO spends a total of $m_X(m_\Phi + 1)$ samples to learn the subspace and thus incurs




Figure 3: Performance comparison on different datasets. Our SI-BO approach outperforms the natural benchmarks in terms of cumulative regret, and competes well with the unrealistic Exact-UCB approach that knows the true subspace A .

much regret during this phase, learning the subspace pays off, both for average and simple regret, and SI-BO ultimately outperforms the baseline methods on both data sets. This demonstrates the value of accurate subspace estimation for Bayesian optimization in high dimensions.

Mis-specified k . What happens if we do not know the dimensionality k of the low dimensional subspace? To test this, we experimented with the stability of SI-BO w.r.t. k . We sampled 2-dimensional GP-Posterior functions and ran SI-BO with k set to 1, 2 and 3. From the figure above we can see that in this scenario SI-BO is relatively stable to this parameter mis-specification.

6 Conclusion

We have addressed the problem of optimizing high dimensional functions from noisy and expensive samples. We presented the SI-BO algorithm, which tackles this challenge under the assumption that the objective varies only along a low dimensional subspace, and has low norm in a suitable RKHS. By fusing modern techniques for low rank matrix recovery and Bayesian bandit optimization in a carefully calibrated manner, it addresses the exploration-exploitation dilemma, and enjoys cumulative regret bounds, which only polynomially depend on the ambient dimension. Our results hold for a wide family of RKHS's, including the popular RBF and Matérn kernels. Our experiments on different data sets demonstrate that our approach outperforms natural benchmarks.

Acknowledgments. A. Krause acknowledges SNF 200021-137971, DARPA MSEE FA8650-11-1-7156, ERC StG 307036 and a Microsoft Faculty Fellowship. V. Cevher acknowledges MIRG-268398, ERC Future Proof, SNF 200021-132548, SNF 200021-146750, and SNF CRSII2-147633.

References

- [1] D. Lizotte, T. Wang, M. Bowling, and D. Schuurmans. Automatic gait optimization with gaussian process regression. In *Proc. of IJCAI*, pages 944–949, 2007.
- [2] J. Bergstra and Y. Bengio. Random search for hyper-parameter optimization. *The Journal of Machine Learning Research*, 13:281–305, 2012.
- [3] J. Snoek, H. Larochelle, and R. P. Adams. Practical bayesian optimization of machine learning algorithms. In *Neural Information Processing Systems*, 2012.
- [4] Ker-Chau Li. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327, 1991.

- [5] G. Raskutti, M.J. Wainwright, and B. Yu. Minimax rates of estimation for high-dimensional linear regression over ℓ_q -balls. *Information Theory, IEEE Transactions on*, 57(10):6976–6994, 2011.
- [6] S. Mukherjee, Q. Wu, and D. Zhou. Learning gradients on manifolds. *Bernoulli*, 16(1):181–207, 2010.
- [7] H. Tyagi and V. Cevher. Learning non-parametric basis independent models from point queries via low-rank methods. *Applied and Computational Harmonic Analysis*, (0):–, 2014.
- [8] H. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.
- [9] P. Auer. Using confidence bounds for exploitation-exploration trade-offs. *The Journal of Machine Learning Research*, 3:397–422, 2003.
- [10] R. Kleinberg, A. Slivkins, and E. Upfal. Multi-armed bandits in metric spaces. In *STOC*, pages 681–690, 2008.
- [11] S. Bubeck, R. Munos, G. Stoltz, and C. Szepesvári. Online optimization in X-armed bandits. In *NIPS*, 2008.
- [12] N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, May 2012.
- [13] E. Brochu, V.M. Cora, and N. De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.
- [14] J. Močkus. On bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference Novosibirsk, July 1–7, 1974*, pages 400–404. Springer, 1975.
- [15] A.D. Bull. Convergence rates of efficient global optimization algorithms. *The Journal of Machine Learning Research*, 12:2879–2904, 2011.
- [16] A. Carpentier and R. Munos. Bandit theory meets compressed sensing for high dimensional stochastic linear bandit. *Journal of Machine Learning Research - Proceedings Track*, 22:190–198, 2012.
- [17] Y. Abbasi-Yadkori, D. Pal, and C. Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Conference on Artificial Intelligence and Statistics (AISTATS)*, 2012.
- [18] B. Chen, R. Castro, and A. Krause. Joint optimization and variable selection of high-dimensional gaussian processes. In *Proc. International Conference on Machine Learning (ICML)*, 2012.
- [19] Z. Wang, M. Zoghi, F. Hutter, D. Matheson, and N. de Freitas. Bayesian optimization in high dimensions via random embeddings. In *In Proc. IJCAI*, 2013.
- [20] H. Tyagi and B. Gärtner. Continuum armed bandit problem of few variables in high dimensions. *CoRR*, abs/1304.5793, 2013.
- [21] R.A. DeVore and G.G. Lorentz. *Constructive approximation*, volume 303. Springer Verlag, 1993.
- [22] M. Fornasier, K. Schnass, and J. Vybiral. Learning functions of few arbitrary linear parameters in high dimensions. *Foundations of Computational Mathematics*, pages 1–34, 2012.
- [23] B. Schölkopf and A.J. Smola. *Learning with kernels: Support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- [24] E.J. Candes and Y. Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *Information Theory, IEEE Transactions on*, 57(4):2342–2359, 2011.
- [25] P. A. Wedin. Perturbation bounds in connection with singular value decomposition. *BIT Numerical Mathematics*, 12(1):99–111, 1972.
- [26] G.W. Stewart and J. Sun. *Matrix Perturbation Theory*, volume 175. Academic Press New York, 1990.
- [27] J. G. Daugman. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Optical Society of America, Journal, A: Optics and Image Science*, 2:1160–1169, 1985.
- [28] J.A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.

APPENDIX : SUPPLEMENTARY MATERIAL

A Effects of Subspace Alignment

There is an alternative view of looking at RKHS's using the Fourier Transform (FT) of the kernel. The following two lemmas can be found in [15][§3.1] where proofs and additional references can also be found. Note that because of the conditions on the kernel its FT is integrable and non-negative.

Lemma 6 (Lemma 1 in [15]) $\mathcal{H}_\kappa(\mathbb{R}^d)$ is the space of real continuous functions $f \in L^2(\mathbb{R}^d)$ whose norm

$$\|f\|_{\mathcal{H}_\kappa(\mathbb{R}^d)}^2 = \int_{\mathbb{R}^d} \frac{|\mathcal{F}f(\omega)|^2}{\mathcal{F}\kappa(\omega)} d\omega$$

is finite, taking $0/0 = 0$.

Lemma 7 (Lemma 2 in [15]) $\mathcal{H}_\kappa(S)$ is the space of functions $f = g|_S$ for some $g \in \mathcal{H}_\kappa(\mathbb{R}^d)$ with norm

$$\|f\|_{\mathcal{H}_\kappa(S)} = \inf_{g|_S=f} \|g\|_{\mathcal{H}_\kappa(\mathbb{R}^d)}$$

and there is a unique g minimizing this expression.

We will need the following well known result, which can be interpreted as the multivariate version of the popular stretch lemma for the one dimensional case.

Lemma 8 Let $g \in L^2(\mathbb{R}^n)$ be a function with Fourier transform $\mathcal{F}g$. For any function g' that has the form

$$g'(\mathbf{z}) = g(M\mathbf{z})$$

for some non-singular matrix M we have that

$$\mathcal{F}g'(\omega) = |\det M|^{-1} g(M^{-T}\omega)$$

Proof

$$\begin{aligned} \mathcal{F}g'(\omega) &= \int_{\mathbb{R}^n} \exp(-2\pi i \mathbf{z}^T \omega) g'(\mathbf{z}) d\mathbf{z} \\ &= \int_{\mathbb{R}^n} \exp(-2\pi i \mathbf{z}^T \omega) g(M\mathbf{z}) d\mathbf{z} \\ &= |\det M|^{-1} \int_{\mathbb{R}^n} \exp(-2\pi i \mathbf{z}^T M^{-T}\omega) g(\mathbf{z}) d\mathbf{z} \\ &= |\det M|^{-1} \mathcal{F}g(M^{-T}\omega) \end{aligned}$$

The third step followed from integration by substitution. □

We can now show how the RKHS norm of the learned function relates to that of the true low-dimensional function.

Lemma 9 If $g \in \mathcal{H}(\mathcal{Z})$ and $\hat{g} : \hat{\mathcal{Z}} \rightarrow \mathbb{R}$ is defined as $\hat{g}(\mathbf{z}) = g(M\mathbf{z})$ for some non-singular M that satisfies $0 < \|M\| \leq 1$ for some matrix norm compatible with the Euclidian norm ℓ_2 , then the following holds

$$\|\hat{g}\|_{\mathcal{H}_\kappa(\hat{\mathcal{Z}})}^2 \leq \frac{1}{|\det M|} \|g\|_{\mathcal{H}_\kappa(\mathcal{Z})}^2$$

Proof The proof is very simple and is similar to that of [15][Lem. 4]. If $\mathcal{Z} = \mathbb{R}^n$, then

$$\begin{aligned}
\|\hat{g}\|_{\mathcal{H}_k(\mathbb{R}^n)}^2 &\stackrel{(1)}{=} \int \frac{|\mathcal{F}\hat{g}(\omega)|^2}{\mathcal{F}k(\omega)} d\omega \\
&\stackrel{(2)}{=} |\det M|^{-2} \int \frac{|\mathcal{F}g(M^{-T}\omega)|^2}{\mathcal{F}k(\omega)} d\omega \\
&\stackrel{(3)}{=} |\det M|^{-1} \int \frac{|\mathcal{F}g(\omega)|^2}{\mathcal{F}k(M^T\omega)} d\omega \\
&\stackrel{(4)}{\leq} |\det M|^{-1} \int \frac{|\mathcal{F}g(\omega)|^2}{\mathcal{F}k(\omega)} d\omega \\
&\stackrel{(5)}{=} \frac{1}{|\det M|} \|g\|_{\mathcal{H}_k(\mathbb{R}^n)}^2
\end{aligned}$$

The second equality follows from Lemma 8 and the third one from integration by substitution and the fact that $\det M^T = \det M$. Noticing that $0 < \|M\| = \|M^T\| \leq 1$ we can use the compatibility of $\|\cdot\|$ with the ℓ_2 norm to conclude that $\|Mw\| \leq \|M\|\|w\|_2 \leq \|w\|_2$. Then, we apply the fact that $\mathcal{F}k$ is radially non-increasing to obtain (4). Finally, we use Lemma 6. If $\mathcal{Z} \subset \mathbb{R}^n$, let $g^H \in \mathcal{H}_k(\mathbb{R}^n)$ be the unique minimum norm extension of g (Lemma 7). Then $\hat{g}^H(\mathbf{z}) = g^H(M\mathbf{z})$ agrees with \hat{g} on $\hat{\mathcal{Z}}$ and we can use the above argument. \square

We can now finally prove the central result in this section.

Proof (Lemma 3) The first part of the lemma is already known, e.g. see Appendix F in [7]. The second conclusion follows from Lemma 9. \square

B Regret Bounds

B.1 Almost Surely Bounded Error

Assume that the noise is $\mathcal{N}(0, \sigma^2)$ and let e_i denote the noise at iteration i . We denote by z a random variable distributed as $\mathcal{N}(0, 1)$. The probability of having an error bigger than $\hat{\sigma} = t\sigma$ is equal to

$$\begin{aligned}
\mathbb{P}\left\{\max_{i=1}^T |\varepsilon_i| > t\sigma\right\} &\leq T\mathbb{P}\{|\varepsilon_1| > t\sigma\} \\
&= T\mathbb{P}\{|z| > t\} \\
&\leq \frac{T}{\sqrt{2\pi}} \exp(-t^2/2)
\end{aligned}$$

where we have used a standard Gaussian tail inequality. Hence if we pick

$$\hat{\sigma} = (2 \log T + 2 \log \frac{1}{\delta} + \log \frac{1}{2\pi})^{1/2} \sigma$$

the probability of observing noise outside $[-\hat{\sigma}, \hat{\sigma}]$ is at most δ .

B.2 Proof of Lemma 1

Proof The first n samples are spent on learning the subspace and a trivial bound on the regret of these first steps is n . Because by assumption the maxima of g and \hat{g} differ by at most η and $\text{GP-UCB}(T - n, \hat{g}, k)$ is the regret w.r.t. \hat{g} , at each step we will accrue at most η additional regret. Hence, with probability at least $1 - \delta$ the total regret is bounded by

$$\underbrace{n}_{\text{subs. learning}} + \underbrace{(T - n)\eta}_{\text{approx.}} + \text{GP-UCB}(T - n, \hat{g}, k)$$

We now apply Theorem 3 from [12] to $\text{GP-UCB}(T, \hat{g}, k)$ and this completes the proof. \square

Because for non-linear functions $\text{GP-UCB}(\cdot) = \Omega^*(\sqrt{T})$ if we can assure that $\eta = \mathcal{O}(1/\sqrt{T})$ and $n = \mathcal{O}^*(\sqrt{T})$, then the regret will be dominated by GP-UCB.

Note that our $\hat{\sigma}$ is a function of T and the bounds in [12] depend on $\hat{\sigma}$, which is treated as a constant. More precisely, in Theorem 3 in [12] we have to change

$$\frac{1}{\log(1 + \frac{1}{\sigma^2})} \quad \text{to} \quad \frac{1}{\log(1 + \frac{1}{(2 \log T + 2 \log \frac{1}{\delta} + \log \frac{1}{2\pi})\sigma^2})}.$$

Note that the function $1/\log(1+1/x)$ is bounded by $x+1$, so in the worst case we will be multiplying the cumulative regret bounds by a term of order $\mathcal{O}(\sqrt{\log T + \log(1/\delta)})$.

B.3 Notation

Because there are plenty of variables and constants that we use in the proofs of the regret bounds, we provide their definitions in the table below.

Symbol	Meaning
X	Contains the derivatives at the centers
σ_k	k -th largest singular value of X
\hat{X}	The matrix obtained from the DS
$\hat{X}^{(k)}$	Closest rank k matrix to \hat{X} in $\ \cdot\ _F$
d	Ambient dimension
k	The dimension of the subspace
m_X	Number of centers
m_Φ	Number of directions
ε	The step in the numerical gradient
C_2	Bound on second derivatives of g
α	The k -th singular value of H^f
κ	Bound on the RIP constant
η	Bound on $\ f - \hat{f}\ _\infty$
δ	Bound on the failure probability
λ	Controls the feasibility region in the DS

B.4 Background

Before showing the regret bounds results we must first set the stage by stating some results on subspace learning. Many of these results have been already known from [7] and [22], but we extend them if necessary for our setting.

As we have already mentioned, the error of the DS decreases with λ and this is stated precisely in the next lemma. The probability is taken over the m_Φ sampling directions.

Lemma 10 (Corollary 1 in [7]) *If X is feasible in the DS then the matrix $\hat{X}^{(k)}$ satisfies*

$$\|X - \hat{X}^{(k)}\|_F^2 \leq \tau^2 \leq 4k\lambda^2$$

w.p. at least

$$1 - \underbrace{2 \exp(-m_\Phi q(\kappa) + 4k(d + m_X + 1)u(\kappa))}_{p_1(m_X, m_\Phi)}$$

for some $0 \leq \kappa < \sqrt{2} - 1$ and $u(\kappa) = \log(\frac{36\sqrt{2}}{\kappa})$, $q(\kappa) = \frac{1}{144}(\kappa^2 - \frac{\kappa^3}{9})$.

Due to a result by Wedin [25] we can show that if two matrices are close in the Frobenius norm, then the subspaces spanned from their left singular vectors are aligned.

Lemma 11 *If $\|X - \hat{X}^{(k)}\|_F \leq \tau$ and we extract \hat{A} from the top k left singular vectors of $\hat{X}^{(k)}$ then*

$$\|\sin \Theta(A, \hat{A})\|_2 \leq \frac{\sqrt{2}}{\sigma_k - \tau} \tau$$

Proof Follows from the proof of Lemma 2 in [7] and Theorem I.5.5 in [26]. \square

We will need the following technical lemma for one of the central results that we will show.

Lemma 12 *Let*

$$X = \{\mathbf{x} \in (0, 1]^n \mid \sum_{i=1}^n x_i = n - 1 + \mu\}$$

for some $\mu \in (0, 1]$. Then $\min_{\mathbf{x} \in X} \prod_{i=1}^n x_i = \mu$ and the minimum is achieved at points \mathbf{x} which have one coordinate equal to μ and all other to 1.

Proof First I will show that at the minimum there must exist some i such that $x_i = \mu$. Assume the contrary. Then all of the elements are greater than μ (if at least one of them is smaller than μ then the equality can not be satisfied). Moreover we assume w.l.o.g. that they are ordered so that $\mu < \underbrace{\mu + m_1}_{x_1} \leq \dots \leq \underbrace{\mu + m_n}_{x_n}$ and all m_i are positive. We consider two cases

(i) Assume $\mu + m_1 + m_2 \leq 1$. Consider the modified vector $\hat{\mathbf{x}}$ where we set

$$\hat{x}_1 = \mu, \hat{x}_2 = \mu + m_1 + m_2, \hat{x}_i = x_i \text{ for } i > 2$$

The constraints are again satisfied, so this is a valid assignment. The value of this assignment is $A' = C\mu(\mu + m_1 + m_2)$ while the old one had a value of $A = C(\mu + m_1)(\mu + m_2)$. It is easy to check that $A' < A$ and this is a contradiction.

(ii) Assume $\mu + m_1 + m_2 = 1 + \beta$ for some $\beta > 0$. Note that $x_1 + x_2 = \mu + \mu + m_1 + m_2 = 1 + \beta + \mu \leq 2$. Hence $\beta + \mu \leq 1$ and we can create a new vector $\hat{\mathbf{x}}$ as

$$\hat{x}_1 = \beta + \mu, \hat{x}_2 = 1, \hat{x}_i = x_i \text{ for } i > 2$$

The new value is $A' = (2\mu + m_1 + m_2 - 1)C$ while the old one was $A = (\mu + m_1)(\mu + m_2)C$. By assumption $A' \geq A$ which is possible only if the following quadratic is not positive

$$q(\mu) = \mu^2 + \mu(m_1 + m_2 - 2) + (1 + m_1 m_2 - m_2 - m_1)$$

By finding the zeros of q we see that $A' \geq A$ only if $\mu \in [1 - m_2, 1 - m_1]$. However if $\mu \geq 1 - m_2$, then $x_2 = \mu + m_2 \geq 1$. This is impossible, because it would imply that $x_2 = x_3 = \dots = x_n = 1$. Then $\sum_{i=1}^n x_n = x_1 + n - 1 > \mu + n - 1$, which is a contradiction.

Hence we can conclude that at least one of the elements in the optimal solution has to be μ . Then the set X over which we optimize consists of points which have one element equal to μ and all other to 1. \square

The following lemma will tell us how precise our solutions from the DS has to be in order to have a good approximation accuracy and have the norm of the \hat{g} under control.

Lemma 13 *If for some $F, \eta \in (0, 1)$ the recovered matrix $\hat{X}^{(k)}$ satisfies*

$$\underbrace{\|X - \hat{X}^{(k)}\|_F}_{\tau} \leq \underbrace{\min\left\{\frac{1}{1 + \sqrt{\frac{2}{1-F}}}, \frac{\eta}{\eta + 2C_2\sqrt{k}(1 + \bar{\varepsilon})}\right\}}_{\Xi} \sigma_k$$

then $\|f - \hat{f}\|_\infty \leq \eta$ and $\|\hat{g}\|_{\mathcal{H}_k}^2 \leq \frac{1}{F}\|g\|_{\mathcal{H}_k}^2$.

Proof The first claim follows directly from Lemmas 3 and 11. For brevity let us denote $\sin \Theta = (\lambda_1, \dots, \lambda_k)$ and $\cos \Theta = (\gamma_1, \dots, \gamma_k)$. From Lemma 11 it follows that

$$\begin{aligned} \sum_{i=1}^k \lambda_i^2 &= k - \sum_{i=1}^k \gamma_i^2 \leq 2\left(\frac{\tau}{\sigma_k - \tau}\right)^2 \\ &\Downarrow \\ \sum_{i=1}^k |\gamma_i| &\stackrel{|\gamma_i| \leq 1}{\geq} \sum_{i=1}^k \gamma_i^2 \geq k - 2\left(\frac{\tau}{\sigma_k - \tau}\right)^2 \end{aligned}$$

Assume that we want to assure that the minimal cosine is at least C for some $C \in (0, 1)$. This is obviously satisfied if

$$\sum_{i=1}^k |\gamma_i| \geq k - 2\left(\frac{\tau}{\sigma_k - \tau}\right)^2 > k - 1 + C$$

because then

$$\min_i |\gamma_i| \geq 1 - 2\left(\frac{\tau}{\sigma_k - \tau}\right)^2 > C$$

Which will be satisfied if

$$\frac{\tau}{\sigma_k - \tau} < \sqrt{\frac{1-C}{2}} \implies \tau < \frac{1}{1 + \sqrt{\frac{2}{1-C}}} \sigma_k$$

We can now apply Lemma 12 and Lemma 3 to obtain the final result. \square

The final ingredient necessary before showing the regret bounds is a lower bound on σ_k which can be obtained from a concentration inequality on the sum of positive semi-definite matrices [28]. The bound that we use depends on the quantity α , which is equal to the k -th largest singular value of the following ‘‘Hessian’’ matrix

$$H^f = \int_{\mathbb{S}^{d-1}} \nabla f(\mathbf{x}) \nabla f(\mathbf{x})^T d\mu,$$

where μ is the uniform measure on \mathbb{S}^{d-1} .

Lemma 14 (Lemma 4.5 in [22]) *For any $\rho \in (0, 1)$ the k -th singular value σ_k of the matrix X that we try to retrieve with the DS satisfies*

$$\sigma_k \geq \sqrt{(1-\rho)m_X \alpha}$$

w.p at least $1 - \underbrace{k \exp\left(-\frac{m_X \alpha \rho^2}{2kC_2^2}\right)}_{p_2(m_X)}$.

Moreover, it has been proven in [7] that under the assumptions that we made for the Hessian of f and the Lipschitz continuity of the second derivatives of f , α behaves as $\alpha = \Theta(1/d)$. We will use this fact when proving the cumulative regret bounds.

B.5 Bounding The Failure Probability

Other than the noise failing to be bounded a.s., there is another possible source of failure – the subspace learning algorithm can fail to learn the subspace within the required accuracy. The following lemma precisely bounds the probability of that happening.

Lemma 15 *The residual is bounded by λ , i.e. Equation 4 holds, with probability at least*

$$1 - \underbrace{\exp(-cm_X)}_{p_3(m_X)} \text{ for some } c > 0$$

Proof First appeared in [7] based on Lemma 1.1 in [24]. \square

Lemma 16 *We can pick $m_X = \mathcal{O}\left(\frac{k}{\alpha} \log(1/\delta)\right)$ and $m_\Phi = \mathcal{O}\left(\frac{kd}{\alpha} \log(1/\delta)\right)$ so that the failure probability is at most δ . Moreover, if $\alpha = \Theta(1/d)$ then we have a better dependence on d and we can use $m_\Phi = \mathcal{O}(k^2 d \log(1/\delta))$.*

Proof The failure probability for the algorithm is at most $p_1(m_X, m_\Phi) + p_2(m_X) + p_3(m_X)$. From the previous lemma we can obviously achieve $p_3 = \delta/3$ if we have $m_X = \mathcal{O}(\log(1/\delta))$. Because only m_X appears in p_2 if we pick (assuming $\delta < 3/k$)

$$m_X \geq \frac{4kC_2^2}{\alpha\rho^2} \log(3/\delta) \geq \frac{2kC_2^2}{\alpha\rho^2} (\log(3/\delta) + \log k)$$

then $p_2 \leq \delta/3$. Similarly, by taking a look at $p_1(m_X, m_\Phi)$ we see that by choosing m_Φ as

$$m_\Phi \geq \frac{1}{q(\kappa)} [4k(d+1)u(\kappa) + \log(6/\delta) + 4ku(\kappa)] \frac{4kC_2^2}{\alpha\rho^2} \log(3/\delta)$$

we also have $p_1 \leq \delta/3$. The special case for $\alpha = \Theta(1/d)$ follows easily as none of the terms above contain both d and α . \square

B.6 Regret Bounds

The following lemma combines many of the results that we have stated so far and will be very useful in showing the regret bounds.

Lemma 17 *If*

$$4k \underbrace{1.2 \left(\frac{C_2 \varepsilon d m_X k^2}{2\sqrt{m_\Phi}} + \frac{5\sqrt{m_X m_\Phi} \sigma}{\varepsilon} \right)^2}_{\lambda^2} \leq \Xi^2 (1 - \rho) m_X \alpha$$

then

$$\tau \leq \Xi \sigma_k \text{ w.p. at least } 1 - p_1(m_X, m_\Phi) - p_2(m_X) - p_3(m_\Phi)$$

Proof We have that with at least the claimed probability Lemmas 15, 14 and 10 hold, in which case

$$\begin{aligned} \tau &\stackrel{\text{Lemma 10}}{\leq} 2\sqrt{k}\lambda \leq \Xi \sqrt{(1 - \rho) m_X \alpha} \\ &\stackrel{\text{Lemma 14}}{\leq} \Xi \sigma_k \end{aligned}$$

where the second inequality follows from the assumption. \square

We can now state the regret bounds by combining this result with Lemma 13.

B.6.1 Noiseless Case

Proof (Theorem 4) If we set $C = 1/2$ in Lemma 13 we see that to have precision η and a guarantee that $\|\hat{g}\|_{\mathcal{H}_k}^2 \leq 2\|g\|_{\mathcal{H}_k}^2$ we will need the following to hold

$$\underbrace{\|X - \hat{X}^{(k)}\|_F}_{\tau} \leq \underbrace{\min\left\{1/2, \frac{\eta}{\eta + 2C_2\sqrt{k}(1 + \bar{\varepsilon})}\right\}}_{\Xi} \sigma_k$$

The result will follow from picking m_X, m_Φ and ε in Lemma 17 so that the above holds and the sampling complexity is $m_X(m_\Phi + 1) = \mathcal{O}^*(\sqrt{T})$. We will pick m_X as in proof of Lemma 16. Now, the inequality above will hold if

$$\begin{aligned} m_\Phi &\geq \Xi^{-2} \frac{1.44 C_2^3 k^5 \varepsilon^2 d^2}{(1 - \rho) \alpha} m_X \\ &\iff \\ m_\Phi &\geq \Xi^{-2} \frac{1.44 C_2^3 k^6 \varepsilon^2 d^2}{(1 - \rho) \alpha^2 \rho^2} \log(3/\delta) \end{aligned}$$

If $\Xi = 1/2$ then the claim can be easily proven as the dependence on T is very weak. Otherwise the above inequality transforms to (setting $\eta = T^{-1/2}$)

$$m_\Phi \geq (1 + 2\sqrt{kT}C_2(1 + \bar{\varepsilon}))^2 \frac{1.44 C_2^3 k^6 \varepsilon^2 d^2}{(1 - \rho) \alpha^2 \rho^2} \log(3/\delta)$$

Which is satisfied by setting for some constant C depending on ρ, C_2 and $\bar{\varepsilon}$

$$m_\Phi \geq C \frac{k^{6.5} \varepsilon^2 d^2}{\alpha^2} T \log(3/\delta)$$

We can now apply the fact that $\alpha = \Theta(1/d)$. Hence, the claimed guarantees will hold if we pick $m_\Phi = \mathcal{O}(k^2 d \log(1/\delta))$ and $\varepsilon = \frac{1}{k^{2.25} d^{3/2} T^{1/2}}$. We finally apply Lemma 1. \square

B.6.2 Noisy Case

The proofs are similar to the previous case, we just have an additional term that we have to bound in Lemma 17. We do this by resampling which will reduce σ .

Proof (Theorem 5) Unfortunately for this case we can not achieve cumulative regret of $\mathcal{O}^*(\sqrt{T})$. Let us try to see how many samples do we need to have $\eta = \mathcal{O}(T^{-\beta})$ for some $\beta \in (0, 1)$. We will pick m_X the same way as in Lemma 16. We will average out each point $m_\Phi n / \varepsilon^2$ times in which case in order the condition in Lemma 17 to hold we need

$$\sqrt{4.8k} \left(\frac{C_2 \varepsilon d m_X k^2}{2\sqrt{m_\Phi}} + \frac{6\sqrt{m_X} \sigma}{\sqrt{n}} \right) \leq \Xi \sqrt{(1-\rho)m_X \alpha}$$

Following the proof from the noiseless case we can make

$$\sqrt{4.8k} \frac{C_2 \varepsilon d m_X k^2}{2\sqrt{m_\Phi}} \leq \frac{1}{2} \Xi \sqrt{(1-\rho)m_X \alpha} \quad (6)$$

by choosing $m_\Phi = \mathcal{O}(k^2 d \log(1/\delta))$ and $\varepsilon = \frac{1}{k^{2.25} d^{3/2} T^\beta}$. We will now assure that

$$\sqrt{4.8k} \frac{6\sqrt{m_X} \sigma}{\sqrt{n}} \leq \frac{1}{2} \Xi \sqrt{(1-\rho)m_X \alpha} \quad (7)$$

which is satisfied if for some C depending on ρ



$$n \geq \Xi^{-2} \frac{C k \sigma^2}{\alpha}$$

Again the case $\Xi = 1/2$ easily yields the claimed bounds as then Ξ is independent of T . Otherwise we need

$$\begin{aligned} n &\geq (1 + 2\sqrt{k} T^\beta C_2 (1 + \bar{\varepsilon}))^2 \frac{C k \sigma^2}{\alpha} \\ &\iff \text{For some } C' \text{ depending on } C_2, \rho \\ n &\geq C' \frac{k^2 \sigma^2}{\alpha} T^{2\beta} \end{aligned}$$

Hence, as $\alpha = \Theta(1/d)$ we have to resample $\mathcal{O}(k^2 d \sigma^2 T^{2\beta} m_\Phi / \varepsilon^2)$ times. By combining Equations 6 and 7 and applying Lemma 17 we can achieve $\eta = \mathcal{O}(T^{-\beta})$ and also have $\|\hat{g}\|_{\mathcal{H}_k} \leq 2\|\hat{g}\|_{\mathcal{H}_k}$ using a total of $m_X m_\Phi^2 n / \varepsilon^2 = \mathcal{O}(k^{11.5} d^7 \sigma^2 T^{4\beta} / \varepsilon^2 \log(1/\delta)^3)$ samples. To bound the regret we use Lemma 1 and we will get both $T^{1-\beta}$ and $T^{4\beta}$ as terms in the bound. If we optimize for β we get $\beta = \arg \min_\beta \{1 - \beta, 4\beta\} = 1/5$ and this proves the claim. \square

C Optimizing the Branin Function



The Branin function is a classical benchmark for global optimization. The function is two-dimensional and we hide it in a random two-dimensional subspace in 100 dimensions. It is flat near the origin and peaks very sharply on two opposite corners of the square. The heights of these peaks differ, which makes it even more complicated to optimize. To encourage more exploration we have increased β_t by a factor of 15. This objective is very challenging because it is not well modeled with a GP prior and it is relatively flat near the origin, which means that it does not satisfy our assumptions. For example we see that even when we know the subspace exactly GP-UCB does not perform well. Moreover, its maximum is concentrated at the boundary of the region – hence we would need an extremely good approximation to the subspace in order to be able to achieve low regret.