

Exploiting Variance Reduction Potential in Local Gaussian Process Search

Chih-Li Sung^a, Robert B. Gramacy^b, Benjamin Haaland^a

^aGeorgia Institute of Technology

^bVirginia Tech

Abstract

Gaussian process models are commonly used as emulators for computer experiments. However, developing a Gaussian process emulator can be computationally prohibitive when the number of experimental samples is even moderately large. Local Gaussian process approximation (Gramacy and Apley (2015)) was proposed as an accurate and computationally feasible emulation alternative. Constructing local sub-designs specific to predictions at a particular location of interest remains a substantial computational bottleneck to the technique. In this paper, two computationally efficient neighborhood search limiting techniques are proposed, a maximum distance method and a feature approximation method. Two examples demonstrate that the proposed methods indeed save substantial computation while retaining emulation accuracy.

1 Introduction

Due to continual advances in computational capabilities, researchers across fields increasingly rely on computer simulations in lieu of prohibitively costly or infeasible physical experiments. One example is Eckstein (2013), who use computer simulations to investigate the interaction of energetic particles with solids. Physical effects such as elastic energy loss when a particle penetrates a solid, particle transmission through solids, and radiation damage are explored. These processes can be approximated by simulating the trajectories of all moving particles in a solid based on mathematical models. An example in linguistics is the study of language evolution (Cangelosi and Parisi (2012)), which is made challenging by the unobserved nature of language origin. Modeling techniques such as genetic algorithms can be used to simulate the process of natural selection and make it possible to explore a virtual evolution. While computer simulations provide a feasible alternative to many physical experiments, simulating from mathematical models is often itself expensive, in terms of both time and computation, and many researchers seek inexpensive approximations to their computationally demanding computer models—so-called *emulators*.

Gaussian process (GP) models (Sacks et al. (1989)) play an important role as emulators for computationally expensive computer experiments. They provide an accurate approximation to the relationship between simulation output and untried inputs at a reduced computational cost, and provide appropriate (statistical) measures of predictive uncertainty. A major challenge in building a GP emulator for a large-scale computer experiment is that it necessitates decomposing a large ($N \times N$) correlation matrix. For dense matrices, this requires around $O(N^3)$ time, where N is the number of experimental runs. Inference for unknown parameters can demand hundreds of such decompositions to evaluate the likelihood, and its derivatives,

under different parameter settings for even the simplest Newton-based maximization schemes. This means that for a computer experiment with as few as $N = 10^4$ input-output pairs, accurate GP emulators cannot be constructed without specialized computing resources.

There are several recent approaches to emulating large-scale computer experiments, most of which focus on approximation of the GP emulator due to the infeasibility of actual GP emulation. Examples include covariance tapering which replaces the dense correlation matrix with a sparse version (Furrer et al. (2006)), multi-step interpolation which successively models global, then more and more local behavior while controlling the number of non-zero entries in the correlation matrix at each stage (Haaland et al. (2011)), and multiresolution modeling with Wendland’s compactly supported basis functions (Nychka et al. (2015)). Alternatively, Paciorek et al. (2015) developed an R package called `bigGP` that combines symmetric-multiprocessors and GPU facilities to handle N as large as 67,275 without approximation. Nevertheless, computer model emulation is meant to avoid expensive computer simulation, not be a major consumer of it. Another approach, proposed by Plumlee (2014), is to sample input-output pairs according to a specific design structure, which leads to substantial savings in building a GP emulator. That method, however, can be limited in practice due to the restriction to sparse grid designs.

In this paper, Gramacy and Apley (2015)’s local GP approach is considered. The approach is modern, scalable, and easy to implement with limited resources. The essential idea focuses on approximating the GP emulator at a particular location of interest via a relatively small subset of the original design, thus requiring computation on only a modest subset of the rows and columns of the large $(N \times N)$ covariance matrix. This process is then repeated across predictive locations of interest, ideally largely in parallel. The determination of this

local subset for each location of interest is crucial since it greatly impacts the accuracy of the corresponding local GP emulator. Gramacy and Apley (2015) proposed a greedy search to sequentially augment the subset according to an appropriate criteria and that approach yields reasonably accurate GP emulators. More details are presented in Section 2.

A bottleneck in this approach is that a complete iterative search for the augmenting point requires looping over $O(N)$ data points at each iteration. In Section 3, motivated by the intuition that there is little potential benefit in including a data point far from the prediction location, two new neighborhood search limiting techniques are proposed, the *maximum distance method* and the *feature approximation method*. Two examples in Section 4 show that the proposed methods substantially speed up the local GP approach while retaining its accuracy. A brief discussion follows in Section 5. Mathematical proofs are provided in the online supplement.

2 Preliminaries

2.1 Gaussian Process Model

A Gaussian process (GP) is a stochastic process whose finite dimensional distributions are defined via a mean function $\mu(x)$ and a covariance function $\Sigma(x, x')$, for d -dimensional inputs x and x' . In particular, for N input x -values, say X_N , which define the N -vector $\mu(X_N)$ and $N \times N$ matrix $\Sigma(X_N, X_N)$, and a corresponding N -vector of responses Y_N , the responses have distribution $Y_N \sim \mathcal{N}(\mu(X_N), \Sigma(X_N, X_N))$. The scale $\sigma^2 > 0$ is commonly separated from the process correlation function, $Y_N \sim \mathcal{N}(\mu(X_N), \sigma^2 \Phi(X_N, X_N))$, where the $N \times N$ matrix $\Phi(X_N, X_N) = (\Phi(x_i, x_j))$ is defined in terms of a correlation function $\Phi(\cdot, \cdot)$, with

$\Phi(x, x) = 1$. As an example, consider the often-used separable *Gaussian correlation function*

$$\Phi_{\Theta}(x, x') = \exp \left\{ - \sum_{j=1}^d (x_j - x'_j)^2 / \theta_j \right\}, \text{ where } \Theta = (\theta_1, \dots, \theta_d), \theta_j > 0, j = 1, \dots, d. \quad (1)$$

As this correlation decays exponentially fast in the squared distance between x_j and x'_j at rate θ_j , the sample paths are infinitely differentiable and the resulting predictor is an interpolator.

The GP model is popular because inference for $\mu(\cdot), \sigma^2$, and Θ is easy and prediction is highly accurate. A popular inferential choice is maximum likelihood, with corresponding log likelihood (up to an additive constant)

$$\begin{aligned} \ell(\mu, \sigma^2, \Theta) = & - \frac{1}{2} \{ n \log(\sigma^2) + \log(\det(\Phi_{\Theta}(X_N, X_N))) + \\ & (Y_N - \mu(X_N))^T \Phi_{\Theta}(X_N, X_N)^{-1} (Y_N - \mu(X_N)) / \sigma^2 \} \end{aligned}$$

and the MLEs of $\mu(\cdot), \sigma^2$, and Θ are

$$(\hat{\mu}(\cdot), \hat{\sigma}^2, \hat{\Theta}) = \arg \max_{\mu, \sigma^2, \Theta} \ell(\mu, \sigma^2, \Theta). \quad (2)$$

Here, $\mu(\cdot)$ and its estimate are described somewhat vaguely. Common choices are $\mu(\cdot) \equiv 0$, $\mu(\cdot) = \mu$, or $\mu(\cdot) = h(\cdot)^T \beta$, for a vector of relatively simple basis functions $h(\cdot)$. More details on inference can be found in Fang et al. (2005) or Santner et al. (2013). Importantly, the predictive distribution of $Y(x)$ at a new setting x can be derived for fixed parameters by properties of the conditional multivariate normal distribution. In particular, it can be shown that $Y(x)|X_N, Y_N \sim \mathcal{N}(\mu_N(x), V_N(x))$, where

$$\mu_N(x) = \mu(x) + \Phi_{\Theta}(x, X_N) \Phi_{\Theta}(X_N, X_N)^{-1} (Y_N - \mu(X_N)), \quad (3)$$

$$V_N(x) = \sigma^2(\Phi_\Theta(x, x) - \Phi_\Theta(x, X_N)\Phi_\Theta(X_N, X_N)^{-1}\Phi_\Theta(X_N, x)). \quad (4)$$

In a practical context, the parameters $\mu(\cdot)$, σ^2 , and Θ can be replaced by their estimates (2) and it can be argued that the corresponding predictive distribution is better approximated by a t -distribution than normal (see 4.1.3 in Santner et al. (2013)). Either way, $\hat{\mu}_N(x)$ is commonly taken as the *emulator*, and $V_N(x)$ captures uncertainty.

2.2 Local Gaussian Process Approximation

A major difficulty in computing the emulator (3) and its predictive variance (4) is solving the linear system $\Phi_{\hat{\Theta}}(X_N, X_N)y = \Phi_{\hat{\Theta}}(X_N, x)$, since it requires $O(N^2)$ storage and around $O(N^3)$ computation for dense matrices. A promising approach is to search small sub-designs that approximate GP prediction and inference from the original design (Gramacy and Apley (2015)). The idea of the method is to focus on prediction at a location, x , using a subset of the full data $X_n(x) \subseteq X_N$. Intuitively, the sub-design $X_n(x)$ may be expected to be comprised of X_N close to x . For typical choices of $\Phi_\Theta(x, x')$, correlation between elements x, x' in the input space decays quickly for x' far from x , and x' 's that are far from x have vanishingly small influence on prediction. Ignoring them in order to work with much smaller, $n \times n$ matrices brings big computational savings, ideally with little impact on accuracy. Figure 1 displays a smaller sub-design ($n = 7$) near location $x = 0.5$ extracted from the original design ($N = 21$). Although the emulator (dotted line) performs very poorly from 0 to 0.3 and from 0.6 to 1.0, the sub-design provides accurate and robust prediction at $x = 0.5$.

For an accurate and robust emulator, a smaller predictive variance (4) for each x is desirable. We seek a small sub-design $X_n(x) \subseteq X_N$ for each location of interest x , which minimizes the predictive variance (4) corresponding to the sub-design $X_n(x)$. This procedure is




Figure 1: An example sub-design $X_7(x)$ for a one dimensional input. Dots represent the full design, X_{21} , the triangle represents the point of interest $x = 0.5$ and the diamonds represent the sub-design, $X_7(x)$. Based on the sub-design $X_7(x)$, the emulator is represented as the dotted line, with the shaded region providing a pointwise 95% confidence band.

then repeated for each location of interest x . The identification of sub-designs and subsequent prediction at each such x can be parallelized immediately, providing a substantial leap in computational scalability. However, searching for the optimal sub-design, which involves choosing n from N input sites, is a combinatorially huge undertaking. A sensible idea is to build up $X_n(x)$ by n nearest neighbors (NNs) close to x , and the result is a valid probability model for $Y(x)|X_n(x), Y(X_n(x))$ (Datta et al. (2016)). Gramacy and Apley (2015) proposed a greedy, iterative search for the sub-design, starting from a small NN set X_{n_0} and sequentially choosing the x_{j+1} which provides the greatest reduction in predictive variance to augment $X_j(x)$, for $j = n_0, n_0 + 1, \dots, n$. That is,

$$x_{j+1} = \arg \min_{\substack{u \in X_N \setminus X_j(x), \\ X_{j+1} = X_j(x) \cup u}} V_{j+1}(x) \quad (5)$$

and $X_{j+1}(x) = X_j(x) \cup x_{j+1}$. Both the greedy and NN schemes can be shown to have computational order $O(n^3)$ (for fixed N) when the scheme is efficiently deployed for each

update $j \rightarrow j + 1$. Specifically, the matrix inverse $\Phi_{\Theta}(X_{j+1}, X_{j+1})^{-1}$ in $V_{j+1}(x)$ can be updated efficiently using partitioned inverse equations (Harville (1997)). Before the greedy subsample selection proceeds, correlation parameters can be initialized to reasonable fixed values to be used throughout the sub-design search iterations. After a sub-design has been selected for a particular location, a local MLE can be constructed. Thus, only $O(n^3)$ cost is incurred for building the local subset and subsequent local parameter estimation. For details and implementation, see the `1aGP` package for R (Gramacy (2016)). An initial *overall* estimate of the correlation parameters can be obtained using the Latin hypercube design-based block bootstrap subsampling scheme proposed by Liu and Hung (2015), which has been shown to consistently estimate *overall* lengthscale θ_j -values in a computationally tractable way, even with large N .

The greedy scheme, searching for the next design point in $X_N \setminus X_j(x)$ to minimize the predictive variance (5), is still computationally expensive, especially when the design size N is large. For example, the new x_{j+1} based on (5) involves searching over $N - j$ candidates. In that case, the greedy search method still contains a serious computational bottleneck in spite of its improvements relative to solving the linear system in (3) for GP prediction and inference. Gramacy et al. (2014) recognized this issue and accelerated the search by exporting computation to graphical processing units (GPUs). They showed that the GPU scheme with local GP approximation and massive parallelization can lead to an accurate GP emulator for a one-million-run full design, with the GPUs providing approximately an order of magnitude speed increase. Gramacy and Haaland (2016) noticed that the progression of $x_{j+1}, j = 1, 2, \dots$ *qualitatively* takes on a ribbon and ring pattern in the input space and suggested a computationally efficient heuristic based on one dimensional searches along rays

emanating from the predictive location of interest x .

In Section 3, two computationally efficient and accuracy preserving neighborhood search methods are proposed. Both neighborhood searches reduce computation by decreasing the number of candidate design points examined. It is shown that only locations within a particular distance of either the prediction location x or the current sub-design, or locations in particular regions within a *feature space*, can have *substantial* influence on prediction. Using these techniques, it is possible to search a *much* smaller candidate set at each stage, leading to huge reductions in computation and increases in scalability.

3 Reduced Search in Local Gaussian Process

For prediction at location x , there is intuitively little benefit to considering input locations that are distant from x (relative to the correlation decay) as the response value at these locations is nearly independent of the response at x . In Section 3.1, a *maximum distance* bound and corresponding algorithm are provided, and in Section 3.2, a *feature approximation* bound and corresponding modification to the algorithm are provided. The algorithms furnish a dramatically reduced set of potential design locations which need to be examined, in a computationally efficient and scalable manner. Notably, for the algorithms presented below, Θ is fixed. Updating of Θ could follow Gramacy and Apley (2015), where an overall estimate of Θ is generated initially, then the local design formed, then a *local* Θ estimated. The below algorithms, and subsequent complexity comparison, focus on updating the local design, while the updating of Θ is considered as an offline procedure.

3.1 Maximum Distance Method

Here x is the particular location of interest, in terms of emulation/prediction, and $X_j(x)$ is the greedy sub-design at stage j . To augment the sub-design $X_j(x)$, we downplay locations distant from x with little loss.

Assume that the underlying correlation function is radially decreasing after appropriate linear transformation of the inputs: there is a strictly decreasing function ϕ so that $\Phi_\Theta(x, x') = \phi(\|\Theta(x - x')\|_2)$ for some Θ . In practice, Θ can be estimated using the local MLE as discussed in Section 2.2, using as a starting value the *overall*, consistent estimate from the sub-design search iterations. Now, consider a candidate input location x_{j+1} at stage $j + 1$ of the greedy sub-design search for an input location to add to the design and take $d_{\min}(x_{j+1})$ to be the minimum (Mahalanobis-like) distance between the candidate point x_{j+1} and the current design and location of interest,

$$d_{\min}(x_{j+1}) = \min\{\|\Theta(x - x_{j+1})\|_2, \|\Theta(x_1 - x_{j+1})\|_2, \|\Theta(x_2 - x_{j+1})\|_2, \dots, \|\Theta(x_j - x_{j+1})\|_2\}. \quad (6)$$

We use the term Mahalanobis-*like* distance to emphasize that the rescaling and rotation of the inputs induced by Θ is not related to the variance-covariance matrix of the input locations. For example, consider the sub-design $X_j(x)$ with two-dimensional inputs shown in Figure 2 for $j = 8$.

Based on the local design scheme introduced in Section 2 and (5), the sub-design $X_n(x)$ is built up through the choices of x_{j+1} to sequentially augment $X_j(x)$, at each stage aiming to minimize predictive variance. Proposition 1 provides an alternate formula for this variance,




Figure 2: An example sub-design $X_8(x)$ with two-dimensional inputs. The circled \times represents the location of interest. With $\Theta = \text{diag}(1/\sqrt{3}, 1/\sqrt{3})$, the dots \bullet represent current design points $X_8(x)$, the dot \bullet represents the new input location x_9 , and the shaded region represents the candidate points x_* with $d_{\min}(x_*) < 3.07$.

which is used to greatly reduce the number of candidates in the minimization problem.

Proposition 1. The predictive variance $V_j(x)$ in (4) can be represented via the recurrence

$$V_{j+1}(x) = V_j(x) - \sigma^2 R(x_{j+1}). \quad (7)$$

Here, $R(x_{j+1})$ is the (scaled) reduction in variance,

$$R(x_{j+1}) = \frac{(\Phi_{\Theta}(x, x_{j+1}) - \Phi_{\Theta}(x_{j+1}, X_j)\Phi_{\Theta}(X_j, X_j)^{-1}\Phi_{\Theta}(X_j, x))^2}{\Phi_{\Theta}(x_{j+1}, x_{j+1}) - \Phi_{\Theta}(x_{j+1}, X_j)\Phi_{\Theta}(X_j, X_j)^{-1}\Phi_{\Theta}(X_j, x_{j+1})}. \quad (8)$$

The recurrence relation (7) is useful for searching candidates to entertain. Further, minimizing variance after adding the new input location x_{j+1} is equivalent to maximizing reduction in variance $R(x_{j+1})$.

Theorem 1. *Suppose $\Phi : \Omega \times \Omega \rightarrow \mathbb{R}$ is a symmetric positive-definite kernel on a compact set $\Omega \subseteq \mathbb{R}^d$ and there exists a strictly decreasing function $\phi : \mathbb{R}^+ \rightarrow \mathbb{R}$ such that $\Phi_{\Theta}(x, y) = \phi(\|\Theta(x - y)\|_2)$ for some Θ . Then, for $\delta > 0$, $R(x_{j+1}) \leq \delta$ if*

$$d_{\min}(x_{j+1}) \geq \phi^{-1} \left(\sqrt{\frac{\delta}{(1 + \sqrt{j}\|\Phi_{\Theta}(X_j, X_j)^{-1}\Phi_{\Theta}(X_j, x)\|_2)^2 + j\delta/\lambda_{\min}}} \right), \quad (9)$$

where λ_{\min} is the minimum eigenvalue of $\Phi_{\Theta}(X_j, X_j)$.

This result indicates that candidate locations that are sufficiently distant from the location of interest and the current sub-design do not have potential to reduce the variance more than δ . Importantly, if the full set of design locations X_N is stored in a data structure such as a *k-d tree* (Bentley (1975)), then the set of candidate locations that do not satisfy (9) can be identified in $O(\log N)$ time, with constant depending on δ , dimension of the input space, and stage j ; This provides a computationally efficient and readily scalable technique for reducing the set of potential candidate locations.

Our Algorithm 1 is a starting point for efficiently selecting sub-designs for prediction at location x . In the algorithm, a larger value of δ is desirable since this leads to fewer candidate design locations to search. One way to obtain a relatively large value of δ is to examine the variance reductions on the set of k nearest neighbors which are not yet in the sub-design, as is shown in Step 2. The number of nearest neighbors k is a tuning parameter. A larger value of k provides a larger variance reduction and therefore excludes more candidate design locations, albeit at an additional computational expense since the variance reduction must be checked at each of these locations. Intuitively, larger-scale and higher-dimensional problems might be expected to benefit from a larger k —larger-scale problems because the cost of computing the

variance reductions across k might be relatively modest compared to the potential cost of computing the variance reductions across a large candidate set, higher-dimensional problems because more points are needed to explore the surface of the hyper-sphere of points near the location of interest. A large value of δ can be obtained by applying the heuristic proposed in Gramacy and Haaland (2016). From the result of Theorem 1, $T(X_j)$ in Step 3, which indicates the region such that

$$d_{\min}(x_{j+1}) \leq \phi^{-1} \left(\sqrt{\frac{\delta}{(1 + \sqrt{j} \|\Phi_{\Theta}(X_j, X_j)^{-1} \Phi_{\Theta}(X_j, x)\|_2)^2 + j\delta/\lambda_{\min}}} \right), \quad (10)$$

gives the subset of candidate locations that have potential to reduce the variance more than δ .

For each update $j \rightarrow j + 1$, the algorithm *ideally* involves $O(j^2 + j \log N)$ computations in Step 3: $O(j \log N)$ for eliminating search locations, and $O(j^2)$ for computing the right-hand side of (10). In particular, the matrix inverse $\Phi_{\Theta}(X_j, X_j)^{-1}$ can be updated via the partitioned inverse equations (Harville (1997)) with $O(j^2)$ cost at each iteration. Analysis of the computational complexity of obtaining (an approximation to) the minimum eigenvalue of $\Phi_{\Theta}(X_j, X_j)$ is more challenging. It is convenient to work with the reciprocal of the maximum eigenvalue of $\Phi_{\Theta}(X_j, X_j)^{-1}$, for which relatively efficient algorithms, such as the power or Lanczos method, exist (Golub and Van Loan (1996)). If the starting vector is not orthogonal to the target eigenvector, then convergence of the (less efficient, but easier to analyze) power method is geometric with rate depending on the ratio between the two largest eigenvalues of $\Phi_{\Theta}(X_j, X_j)^{-1}$ (see equation 9.1.5 in Golub and Van Loan (1996)). While this rate and the constants in front are not fixed across j , they can be bounded, with the exception of the

influence of the starting vector, across all subsets of the full dataset. The starting vector might be expected to be increasingly collinear with the target eigenvector as j increases, thereby improving the rate bound. All together this implies an approximately constant number of iterations, each costing $O(j^2)$, is required to approximate λ_{\min} for each j . Another perspective chooses a random starting vector, for which Kuczynski and Wozniakowski (1992) provide respective average and probabilistic bounds of $O(j^2 \log j)$ for the power method and $O(j^2 \log^2 j)$ for the Lanczos method. If standard eigen-decomposition routines that return all the eigenvalues are used, then the j^2 term in the computational complexity is j^3 . The inverse function $\phi^{-1} : \mathbb{R} \rightarrow \mathbb{R}$ can be computed in roughly constant time by a root-finding algorithm or even computed exactly for many choices of Φ . For example, for the power correlation function, $\Phi_{\Theta}(x, y) = \exp\{-\|\Theta(x - y)\|_2^p\}$, the ϕ can be formed as $\phi(u) = \exp\{-u^p\}$, so $\phi^{-1}(v) = (-\log v)^{1/p}$. When a large n is required, computation of λ_{\min} might be numerically unstable. A remedy in that case may be to stop the search when λ_{\min} falls below a prespecified threshold, or perhaps to introduce a penalty inversely proportional to λ_{\min} .

3.2 Feature Approximation Method

In addition to the maximum distance method and associated algorithm, an approximation via eigen-decomposition can be applied to reduce the potential locations in a computationally efficient manner. Suppose that Φ is a symmetric positive-definite kernel on a compact set $\Omega \subseteq \mathbb{R}^d$ and $P : L_2(\Omega) \rightarrow L_2(\Omega)$ is the integral operator

$$Pv(x) := \int_{\Omega} \Phi(x, y)v(y)dy, \quad v \in L_2(\Omega), \quad x \in \Omega. \quad (13)$$

Algorithm 1 Maximum distance search method in local Gaussian process.

- 1: Set $j = 1$ and x_1 as the point closest to the predictive location x . Throughout, let $X_j(x) \equiv X_j = \{x_1, x_2, \dots, x_j\}$, dropping the explicit (x) argument.
- 2: Let $N_{jk}(x)$ denote the k nearest neighbors to x in $X_N \setminus X_j$, the candidate locations not currently in the sub-design. Set δ_{j+1} equal to the maximum variance reduction from $N_{jk}(x)$,

$$\delta_{j+1} = \max_{u \in N_{jk}(x)} R(u), \quad (11)$$

where $R(\cdot)$ is shown in (8).

- 3: Set $y = \phi^{-1} \left(\sqrt{\frac{\delta_{j+1}}{(1+\sqrt{j}\|\Phi_{\Theta}(X_j, X_j)^{-1}\Phi_{\Theta}(X_j, x)\|_2)^2 + j\delta_{j+1}/\lambda_{\min}}} \right)$, where $\Phi_{\Theta}(x, x') = \phi(\|\Theta(x-x')\|_2)$ and λ_{\min} is the minimum eigenvalue of $\Phi_{\Theta}(X_j, X_j)$. For

$$T(X_j) = \{u \in X_N \setminus X_j : \|\Theta(u - v)\|_2 \leq y \text{ for some } v \in \{x, X_j\}\}, \quad (12)$$

take

$$x_{j+1} = \arg \max_{u \in T(X_j)} R(u).$$

- 4: Set $j = j + 1$ and repeat 2 and 3 until either the reduction in variance $R(x_{j+1})$ falls below a prespecified threshold, or the local design budget is met.
-

Mercer's theorem guarantees the existence of a countable set of positive eigenvalues $\{\lambda_j\}_{j=1}^{\infty}$ and an orthonormal set $\{\varphi_j\}_{j=1}^{\infty}$ in $L_2(\Omega)$ consisting of the corresponding eigenfunctions of P , $P\varphi_j = \lambda_j\varphi_j$ (Wendland (2004)). The eigenfunctions φ 's here are continuous on Ω and Φ has the absolutely and uniformly convergent representation

$$\Phi(x, y) = \sum_{j=1}^{\infty} \lambda_j \varphi_j(x) \varphi_j(y).$$

In particular, Φ can be approximated uniformly over inputs in terms of a finite set of eigenfunctions

$$\Phi(x, y) \approx \sum_{j=1}^D \lambda_j \varphi_j(x) \varphi_j(y) \quad (14)$$

for some moderately large integer D . For some kernel functions, closed form expressions exist.

For example, the Gaussian correlation function (1) (on \mathbb{R}^d , with weighted integral operator)

has eigenfunctions given by products of Gaussian correlations and Hermite polynomials (Zhu et al. (1997)). More generally, Williams and Seeger (2001) show high-quality approximations to these eigen-decompositions can be obtained via Nyström’s method.

Theorem 2. *Assume $\Phi : \Omega \times \Omega \rightarrow \mathbb{R}$ is a symmetric positive-definite kernel on a compact set $\Omega \subseteq \mathbb{R}^d$ which can be approximated via D eigenfunctions (14). Then, the reduction in variance (8) has approximate representation*

$$R(x_{j+1}) \approx \|C_{X_j}(x)\|_2^2 \cos^2(\vartheta), \quad (15)$$

where ϑ is the angle between $C_{X_j}(x)$ and $C_{X_j}(x_{j+1})$,

$$\begin{aligned} C_{X_j}(t) &= [I - U(X_j)[U^T(X_j)U(X_j)]^{-1}U^T(X_j)]U(t), \\ U(t) &= \left(\sqrt{\lambda_1}\varphi_1(t), \dots, \sqrt{\lambda_D}\varphi_D(t)\right)^T, \quad \text{and} \\ U(X_j) &= [U(x_1), \dots, U(x_j)], \end{aligned}$$


for eigenfunctions $\varphi_i(t)$, $t \in \Omega$ and corresponding ordered eigenvalues $\lambda_1 \geq \dots \geq \lambda_D$.

According to this approximation, the candidate set can be reduced by transforming the inputs into a feature space. A modified algorithm is suggested as follows. The variance reduction threshold in (11) now places a restriction on the *angle* between $C_{X_j}(x)$ and $C_{X_j}(x_{j+1})$, where we would like to exclude points *outside* the cones

$$\cos^2(\vartheta) \leq \frac{\delta_{j+1}}{\|C_{X_j}(x)\|_2^2}. \quad (16)$$

A feature approximation is shown in Algorithm 2. To reduce the computational burden

in checking (16), the values of the first D eigenfunctions at the full dataset X_N , $U(X_N)$, could be computed in advance and stored based on a locality-sensitive hashing (LSH) scheme (Indyk and Motwani (1998)). LSH is a method for answering approximate similarity-search queries in high-dimensional spaces. The basic idea is to use special locality-sensitive functions to *hash* points into “buckets” such that “nearby” points map to the same bucket with high probability. Many similarity measures have corresponding LSH functions that achieve this property. For instance, the hashing functions for cosine-similarity are the normal vectors of random hyperplanes through the origin, denoted for example as v_1, \dots, v_k . Depending on its side of these random hyperplanes, a point p is placed in bucket $h_1(p), \dots, h_k(p)$, where $h_i(p) = \text{sign}(v_i^T p)$. A simple example, following Van Durme and Lall (2010), is provided in Figure 3. Figure 3a illustrates the hashing process for a point p , where the point p is hashed into the bucket $(h_1(p), \dots, h_6(p)) = (-1, -1, 1, 1, 1, 1)$ by the definition $h_i(p) = \text{sign}(v_i^T p)$, $i = 1, \dots, 6$ (when the point p is above the hyperplane, the inner product is negative, otherwise the inner product is positive). Similarly, other points are placed in their corresponding buckets. In the search process, shown in Figure 3b, the query point q is mapped to the bucket $(h_1(q), \dots, h_6(q)) = (-1, 1, 1, 1, 1, 1)$, which matches the bucket of point p' . Thus, the hashing and search processes retrieve p' as the *most similar* neighbor of q . Also, since the one different label in the buckets of p and q implies that the angular difference is close to $\pi/6$ (six hyperplanes), p is retrieved when querying the points whose angular difference from q is less than $\pi/6$. Many more than six hyperplanes are needed to ensure that the returned angle similarity is approximately correct. In a standard LSH scheme, the hashing process is performed several times by different sets of random hyperplanes, and the search procedure iterates over these random sets of hyperplanes. More details and examples



(a) Hashing.

(b) Search.

Figure 3: Illustration of locality-sensitive hashing (LSH) scheme. Lines — are random hyperplanes through origin, and v_1, \dots, v_6 (arrows \rightarrow) are the corresponding normal vectors. Dots \bullet present stored data points, and the dot \bullet with circle presents the query data point.

can be seen in Indyk and Motwani (1998), Van Durme and Lall (2010), and Leskovec et al. (2014).

Apart from cosine-similarity, Jain et al. (2008) showed for the pairwise similarity

$$\frac{y_k^T A_j y_h}{\|G_j y_k\|_2 \|G_j y_h\|_2},$$

where $y_k, y_h \in \mathbb{R}^d$, $G_j^T G_j = A_j$ and A_j is a $d \times d$ positive-definite matrix that is updated for each iteration j , the hash function can be defined as:

$$h_{A_j}(y) = \begin{cases} 1 & r^T G_j y \geq 0 \\ 0 & \text{otherwise} \end{cases}, \quad (17)$$

where the vector r is chosen at random from a d -dimensional Gaussian distribution. Let $G_j = I - U(X_j)[U^T(X_j)U(X_j)]^{-1}U^T(X_j)$ and $A_j = G_j$ (G_j is symmetric and idempotent),

then $\cos(\vartheta)$ in (15) can be represented as

$$\cos(\vartheta) = \frac{U(x_{j+1})^T A_j U(x)}{\|G_j U(x_{j+1})\|_2 \|G_j U(x)\|_2}.$$

Thus, in the feature approximation method, an LSH scheme can be employed by storing $U(X_N)$ in advance and updating the hash function (17) at each iteration, where y is replaced by $U(y)$. At query time, similar points are hashed to the same bucket with the query $U(x)$ and the results are guaranteed to have a similarity within a small error after repeating the procedure several times. In particular, for each update $j \rightarrow j + 1$, given that the LSH method guarantees retrieval of points within the radius $(1 + \epsilon)M$ from the query point $U(x)$, where M is the distance of the true nearest neighbor from $U(x)$, the method requires $O(D^2 + jDN^{1/(1+\epsilon)})$ computational cost, $O(D^2)$ for updating matrix G_j (via the partitioned inverse equations (Harville (1997))) and computing the hash function $h_{A_j}(y)$ (via the implicit update in Jain et al. (2008)), and $O(jDN^{1/(1+\epsilon)})$ for identifying the hashed query (Jain et al. (2008)), where D is the number of eigenfunctions in Theorem 2. In Section 4, two examples show the benefit from the LSH approach in the feature approximation method.

Algorithm 2 Feature approximation modification to Algorithm 1.

In Step 3 of Algorithm 1, replace $T(X_j)$ with $T^*(X_j)$, where

$$T^*(X_j) = \{u \in X_N \setminus X_j : \|\Theta(u - v)\|_2 \leq y \text{ and } \cos^2(\vartheta) \geq \delta_{j+1} / \|C_{X_j}(x)\|_2^2 \\ \text{for some } v \in \{x, X_j\}\},$$

and ϑ , $C_{X_j}(x)$ are defined in Theorem 2. Then,

$$x_{j+1} = \arg \max_{u \in T^*(X_j)} R(u).$$

As an illustration of how cones in feature space relate to the design space, consider a

full design X_N consisting of 2500 Unif(0, 1) data points, plotted in dots in the *left* panel of Figure 4. The correlation function is $\Phi(x, x') = \exp\{-\|(x - x')/10\|_2^2\}$ and the predictive location of interest is $x = (0.5, 0.5)$, shown as a triangle in the *left* panel. The first 7 design points are chosen greedily and indicated with numbers. The *right* panel shows the first two components of the feature space (the first two eigenfunctions evaluated at the design points), labeled correspondingly. The vector $C_{X_7}(x)$ is denoted as the middle dotted line in the *right* panel, with $|\vartheta| \leq \pi/20$ shown as the outer dotted lines. Design points falling within these cones are shown in shaded region in both panels. The design points in the *left* panel which fall in the stripe have the most potential to reduce predictive variance.




Figure 4: Dots \bullet and \bullet represent design points in the original space (left) and a $D = 2$ dimensional feature space approximation (right). Location of interest and current design are annotated with triangle and numbers, respectively. Vector $C_{X_7}(x)$ and cones $|\vartheta| \leq \pi/20$ shown with dotted lines. Design points falling within these cones are shown in shaded region in both panels.

The storage requirements and computational complexity of each update $j \rightarrow j + 1$ at each predictive location of interest for the proposed algorithms are summarized in Table 1. Notably, if n local data points are needed, then the updating costs are incurred for $j = 1, \dots, n$. Costs across predictive locations are ideally incurred largely in parallel. Here, the original greedy approach proposed in Gramacy and Apley (2015) is referred to as exhaustive search. Recall that $T(X_j)$ and $T^*(X_j)$ are the candidate sets from maximum distance method and feature approximation method, respectively. Let $|\cdot|$ denote the cardinality of a set. Since $|T(X_j)|$

and $|T^*(X_j)|$ are expected to be much smaller than N , the computational cost of the two proposed algorithms can be substantially reduced at each stage j relative to the original greedy search. Overhead cost, for computing benchmarks and eliminating search locations, is required for both methods. Also, with a k-d tree or LSH search method, the specially adapted data structure indeed improves overhead computational costs ($O(j^2 + jN) \rightarrow O(j^2 + j \log(N))$ and $O(j^2 + D^2N) \rightarrow O(j^2 + D^2 + jDN^{1/(1+\epsilon)})$, respectively). Considering the two proposed methods, $|T^*(X_j)|$ might be expected to be much smaller than $|T(X_j)|$ if the correlation function is well approximated by the finite set of eigenfunctions and eigenvalues, and the dimension of input is not too large, since distance becomes a very powerful exclusion criteria in even moderately high-dimensional space. The maximum distance method has smaller storage and overhead requirements. Section 4 presents two examples implementing the two proposed methods and shows the comparison.

	Exhaustive Search	Maximum Distance Method		Feature Approximation Method with D Features*	
		w/o k-d tree	w/ k-d tree	w/o LSH	w/ LSH
Storage	N	N	N	ND	ND
Overhead		$O(j^2 + jN)$	$O(j^2 + j \log(N))$	$O(j^2 + D^2N)$	$O(j^2 + D^2 + jDN^{1/(1+\epsilon)})$
Search	$O(j^2N)$	$O(j^2 T(X_j))$	$O(j^2 T(X_j))$	$O(j^2 T^*(X_j))$	$O(j^2 T^*(X_j))$

Table 1: Complexity comparison between exhaustive search and two proposed methods for each update $j \rightarrow j + 1$. The notation $|\cdot|$ denotes the cardinality of a set, and ϵ is a pre-specified value for the LSH method. *The complexity of pre-computation for feature approximation method is $O(D^3)$.

4 Examples

Two examples are discussed in this section: a two-dimensional example that demonstrates the algorithm and visually illustrates the reduction of candidates; a larger-scale, higher-dimensional example. Both examples show the proposed methods considerably outperforming

the original search method with respect to computation time. All numerical studies were conducted using R (R Core Team (2015)) on a laptop with 2.4 GHz CPU and 8GB of RAM. The k-d tree and LSH were implemented via R package RANN (Arya et al. (2015)) and modifications to the source code of the Python package `scikit-learn` (Pedregosa et al. (2011); Bawa et al. (2005)), and accessed in R through the `rPython` package (Bellosta (2015)). Notably, the Lanczos method offered little consistent advantage for the scale of local datasets ($j \leq 30$) entertained below, and relatively typical of practical situations. As such, reported timings correspond to computing $\lambda_{\max}(\Phi(X_j, X_j)^{-1})$ via the full eigen-decomposition using `eigen`.

4.1 Two-dimensional problem of size $N = 50^2$

Consider a computer experiment with full set of design locations X_N consisting of a regular 50×50 grid on $[-10, 10]^2$ (2500 design points, light small dots in Figure 5), and take the predictive location of interest x to be $(0.216, 0.303)$ (circled \times in Figure 5). Set $\sigma^2 = 1$, and consider the Gaussian correlation function

$$\Phi_{\Theta}(x, y) = \exp \left\{ - \left(\frac{(x_1 - y_1)^2}{\theta_1} + \frac{(x_2 - y_2)^2}{\theta_2} \right) \right\},$$

with $\theta_1 = \theta_2 = 3$. This correlation function implies the ϕ in Algorithm 1 is $\phi(u) = \exp\{-u^2\}$ and $\Theta = \text{diag}(1/\sqrt{\theta_1}, 1/\sqrt{\theta_2})$. Then we have $\phi^{-1}(v) = \sqrt{-\log v}$.

Figure 5 illustrates the sub-design selection procedure shown in Algorithm 1, in which $k = 8$ nearest neighbors (from the candidate set) are used to generate the threshold in Step 2. In Figure 5, the dots represent the current design $X_j(x)$ and the optimal augmenting point

x_{j+1} , and the points that are excluded from the search for that location are those which fall outside the shaded region. The panels in the figure correspond to greedy search steps $j \in \{3, 16, 29\}$. Notably, the optimal additional design points illustrated in Figure 5 are not always the nearest neighbors to the location of interest. In this example, only 7.40% (185/2500) of candidates need to be searched in the beginning. Even after choosing thirty data points, there is no need to search much more than half of the full data (56.92%=1423/2500).

Continuing the same example, Figure 5 shows substantial improvement from the feature approximation method. In the example, a $D = 500$ -dimensional feature space approximation is pre-computed using Nyström’s method (Williams and Seeger (2001)). The points annotated with +s are the points that are not excluded from the search. In fact, the number of candidates which need to be searched is usually reduced at least 10-fold and in many cases 50- or 100-fold, or more.




Figure 5: *Left, middle, and right panels respectively illustrate selection at $j = 3, 16,$ and 29 . The circled \times is the location of interest, $(0.216, 0.303)$. Dots \bullet are the current design points; dots \bullet are the optimal x_{j+1} ; points which are excluded from the search based on maximum distance method are those which fall outside the shaded region. Points which are excluded from the search based on feature approximation method are those which are not annotated with a $+$.*

While the maximum distance method and original greedy approach proposed in Gramacy and Apley (2015) produce the same sub-designs and, in turn, the same predictive variances, the feature approximation method is *approximate* and can produce different sub-designs and

slightly different predictive variances (not *necessarily* inflated due to greedy nature of search). Table 2 shows relative differences in predictive variance resulting from feature approximation method with $D = 10, 200,$ and 500 features as compared to take maximum distance method (or equivalently the original greedy approach). The number of search candidates is listed in parentheses. The relative difference in predictive variance is defined as

$$\frac{V_{j,FA}(x) - V_{j,MD}(x)}{V_{j,MD}(x)},$$

where $V_{j,FA}(x)$ and $V_{j,MD}(x)$ denote the predictive variance of the emulator at location x at stage j using the feature approximation method and maximum distance method, respectively. As might be expected, a larger number of features, $D = 500$, reduces the search candidates without any loss in variance reduction. For $D = 200$, although there are small differences in predictive variance, the discrepancies may be small enough to be of little practical consequence. At stage 15, 20, and 25, the predictive variance for $D = 200$ is even smaller than that of the maximum distance method, due to the greedy nature of the searches. Notably, if a small number of features, say $D = 10$, is chosen, feature approximation search offers little improvement over the maximum distance method in terms of search reduction, even though the predictive variances are similar to maximum distance method. In this case, $D = 200$ features might be a reasonable choice, balancing ease of computation and small predictive variance.

To further compare the performance of the proposed methods with the original greedy approach (exhaustive search), a Sobol's quasi-random sequence (Bratley and Fox (1988)) of 100 predictive locations was generated. Table 3 shows the average computation time and

Relative Difference (# of searching candidates)	$D = 10$		$D = 200$		$D = 500$		Variance by Maximum Distance Method
Stage 10	0	(842)	0.178	(4)	0	(47)	1.95×10^{-6} (844)
Stage 15	0.006	(1057)	-0.76	(7)	0	(69)	9.35×10^{-7} (1040)
Stage 20	0.018	(1149)	-0.722	(30)	0	(12)	6.12×10^{-7} (1168)
Stage 25	-0.155	(1332)	-0.091	(4)	0	(116)	1.66×10^{-7} (1295)
Stage 30	0.009	(1459)	0.024	(20)	0	(2)	1.28×10^{-8} (1423)

Table 2: *The relative difference in variance of the emulator at location (0.216, 0.303) between maximum distance search as a baseline and feature approximation search with number of features D : 10, 200 and 500. Baseline variance by maximum distance search is shown in the last column. The value in parentheses is the number of search candidates.*

proportion of search candidates for the proposed methods and exhaustive search over the 100 predictive locations. The proportion of search candidates for the maximum distance method ran from 22.78% to 39.15%. The method also marginally speeds up computation time from 21 to 18 seconds with a k-d tree data structure. Although the feature approximation method needed 6 seconds for computing the features in advance, the proportion of search candidates for feature approximation method with $D = 200$ was reduced to 5.88% at stage 30. The computation time, on an ordinary laptop, was less than 15 seconds for 30 stages of iteration in the $N = 50^2$ experiment. Relative average predictive variance increases due to using the feature approximation method, both with and without LSH, are shown in Table 4. At stage 30 the average predictive variance increases due to using the feature approximation method are, with and without LSH, 4.6% and 1.4%, respectively, potentially small enough to be disregarded in a practical context. The LSH data structure also marginally reduces search time from 13 to 11 seconds. Recall that the feature approximation method with LSH approximates both the covariance function and the cosine similarity measure, so the candidate set is slightly different from the one without LSH. While in this moderately-sized problem the k-d tree and LSH data structures do not greatly improve the computational cost

(at stage 30, k-d tree: $21 \rightarrow 18$, LSH: $13 \rightarrow 11$), in a larger-scale problem the improvements due to incorporating a k-d tree or LSH data structure can be relatively substantial.

Seconds (Candidates %)	Exhaustive Search	Maximum Distance Method		*Feature Approximation Method with $D = 200$	
		w/o KD-tree	w/ KD-tree	w/o LSH	w/ LSH
Stage 10	11	3 (22.78%)	2 (22.78%)	3 (2.69%)	2 (1.98%)
Stage 15	19	5 (28.04%)	4 (28.04%)	5 (3.31%)	4 (2.90%)
Stage 20	30	9 (32.68%)	7 (32.68%)	8 (7.57%)	6 (5.56%)
Stage 25	44	14 (36.10%)	12 (36.10%)	10 (6.41%)	8 (5.66%)
Stage 30	61	21 (39.15%)	18 (39.15%)	13 (5.88%)	11 (6.65%)

Table 3: Average time (seconds) comparison between exhaustive search and two proposed methods in two-dimensional setting with $N = 50^2$ over 100 Sobol predictive locations. The values in parentheses are the average percentage searched of full design. *Pre-computation time for feature approximation method is 6 seconds.

Relative Difference	Feature Approximation Method with $D = 200$		Average Variance by Maximum Distance Method
	w/o LSH	w/ LSH	
Stage 10	0.192	0.168	4.15×10^{-6}
Stage 15	0.348	0.140	1.08×10^{-6}
Stage 20	0.171	0.066	4.34×10^{-7}
Stage 25	0.011	-0.124	2.26×10^{-7}
Stage 30	0.046	0.014	1.62×10^{-7}

Table 4: The relative difference in average predictive variance of the emulator between maximum distance search as a baseline and feature approximation search with number of features $D = 200$ over 100 Sobol predictive locations in 2-dimensional setting.

4.2 6-dimensional problem of size $N = 5 \times 10^4$

Even more substantial reductions in the number of search candidates are seen for both methods in a larger-scale, higher-dimensional setting. In this example, we generated a 6-dimensional Sobol’s quasi-random sequence of size $N = 5 \times 10^4$ in a $[-1, 1]^6$ for the design space and the predictive locations were chosen from a Sobol’s quasi-random sequence of

size 20. We took $\sigma^2 = 1$ and tuning parameter $k = 30$, and used the correlation function $\Phi_{\Theta}(x, y) = \exp\{-\sum_{i=1}^6 \frac{(x_i - y_i)^2}{\theta_i}\}$ with $\theta_i = 1.5, i = 1, \dots, 6$.

Table 5 shows the comparison between exhaustive search and the two proposed methods. The two methods outperform exhaustive search in terms of computation time. Further, the number of candidates searched for both methods are less than 10% ($= 5000/50000$) across all 30 stages. While the exhaustive search took 3423 seconds (≈ 1 hours) for 30 stage iterations, 240 seconds (4 minutes) were required for maximum distance method. Incorporating a k-d tree data structure, the computation time decreased to 193 seconds (≈ 3.2 minutes). Compared to the 2-dimensional example in Section 4.1, incorporating a k-d tree data structure has moderately more computational benefit in this larger-scale setting.

The feature approximation method, as expected, has a smaller-sized candidate set than the maximum distance method. Using $D = 300$ features, less than 2% average predictive variance increases at stage 30 are observed due to approximation, as shown in Table 6. On the other hand, due to the moderately expensive computation in Algorithm 2 using $D = 300$ features, feature approximation search without LSH is more time-consuming than the maximum distance method. As shown in Table 1, the computation of more design points incurs higher computational costs in order of D^2 for feature approximation search without LSH (complexity $O(j^2 + D^2N)$). With an LSH approximate similarity-search method, computation time is reduced by 189 seconds (≈ 3 minutes) across all 30 stages. While the feature approximation approach outperforms exhaustive search, it appears to be most useful when the maximum distance approach is very conservative, such as in the two-dimensional case in Section 4.1.

Seconds (Candidates %)	Exhaustive Search	Maximum Distance Method		*Feature Approximation Method with $D = 300$	
		w/o KD-tree	w/ KD-tree	w/o LSH	w/ LSH
Stage 10	488	24 (2.77%)	10 (2.77%)	74 (1.71%)	26 (1.7%)
Stage 15	953	50 (4.27%)	28 (4.27%)	126 (3.34%)	45 (3.68%)
Stage 20	1601	93 (5.84%)	62 (5.84%)	199 (4.77%)	76 (5.16%)
Stage 25	2423	154 (7.34%)	115 (7.34%)	296 (5.28%)	121 (5.21%)
Stage 30	3423	240 (8.62%)	193 (8.62%)	435 (6.70%)	189 (6.38%)

Table 5: Time (seconds) comparison between exhaustive search and two proposed methods in 6-dimensional setting with $N = 5 \times 10^4$ over 20 Sobol predictive locations. The values in parentheses shows the percentage searched of full design. *Pre-computation time for feature approximation method was 26 seconds.

Relative Difference	Feature Approximation Method with $D = 300$		Average Variance by Maximum Distance Method
	w/o LSH	w/ LSH	
Stage 10	0.049	0.047	0.2328
Stage 15	0.030	0.032	0.2120
Stage 20	0.023	0.022	0.1997
Stage 25	0.017	0.017	0.1913
Stage 30	0.016	0.016	0.1850

Table 6: The relative difference in average predictive variance of the emulator between maximum distance search as a baseline and feature approximation search with number of features $D = 300$ over 20 Sobol predictive locations in 6-dimensional setting.

5 Conclusion and Discussion

The two methods considered here can be extended for selecting more than one point in each stage j in a straight-forward manner. For example, suppose two points are to be selected in each stage. Let $j' = 2j$, $X_{j'}$ be the current sub-design at stage j , and $x_{j'+1}$ and $x_{j'+2}$ be two points selected at the stage $j + 1$. Proposition 1 can be extended to $V_{j+1}(x) = V_j(x) - \sigma^2 R^*(x_{j'+1}, x_{j'+2})$ for a function R^* , Theorem 1 can narrow the window of potential *pairs* of candidate locations, to say $T'(X_j)$, and Algorithm 1 can be updated accordingly. On the other hand, retaining good computational properties in a batch-sequential

framework is not straight-forward. For example, searching for the optimal candidates, $(x_{j'+1}, x_{j'+2}) = \arg \max_{(u_1, u_2) \in T'(X_j)} R^*(u_1, u_2)$, can be very expensive, say $O(|T'(X_j)|^2)$, compared to searching for one point in each stage. Efficiently augmenting multiple points at each stage, for example by alternating maximizations on $x_{j'+1}$ and $x_{j'+2}$, might be worth exploring in future work.

The essential ideas of the proposed approaches have potential for application in search space reduction in global optimization. Consider the following example. Lam and Notz (2008) modified the *maximum entropy design* (Shewry and Wynn (1987)) for use as a sequential algorithm to efficiently construct a space-filling design in computer experiments. They showed that the algorithm can be simplified to selecting a new point that maximizes the so-called *sequential maximum entropy criterion*

$$x_{j+1} = \arg \min_{u \in \mathcal{D} \setminus X_j} \Phi_{\Theta}(u, X_j) \Phi_{\Theta}(X_j, X_j)^{-1} \Phi_{\Theta}(X_j, u),$$

where \mathcal{D} is a *discrete* design space. For this global optimization problem, let

$$d_{\max}(x_{j+1}) = \max\{\|\Theta(x_1 - x_{j+1})\|_2, \|\Theta(x_2 - x_{j+1})\|_2, \dots, \|\Theta(x_j - x_{j+1})\|_2\}$$

and $\delta > 0$. It can be shown that if $d_{\max}(x_{j+1}) \leq \phi^{-1}(\sqrt{\lambda_{\max}\delta})$, where λ_{\max} is the maximum eigenvalue of $\Phi_{\Theta}(X_j, X_j)$, then the objective function $\Phi_{\Theta}(x_{j+1}, X_j) \Phi_{\Theta}(X_j, X_j)^{-1} \Phi_{\Theta}(X_j, x_{j+1}) > \delta$. Thus, similar to Algorithm 1, a maximum distance approach could be used to eliminate search candidates. For other specific global optimization problems, detailed examination is needed.

An implicit disadvantage of these methods is the impact of the correlation parameters.

Take the example in Figure 2, where $d_{\min}(x_9) < 3.07$. From the definition (6) of $d_{\min}(x_{j+1})$, if $\Theta = (1/\sqrt{\theta}, 1/\sqrt{\theta})$, then the larger θ is, the bigger the search area, the shaded region in Figure 2. When θ is large, the correlation is close to one and the data points tend to be highly correlated, implying that every data point in the full design carries important information for each predictive location. Thus, the algorithm requires more computation for “easier” problems—i.e., with a “flatter” surfaces. On the other hand “flatter” surfaces do not require large sub-designs to achieve small predictive variance.

An improvement worth exploring is how to determine of the number of features D in the feature approximation method. Cross-validation to minimize predictive variance of an emulator may present an attractive option. An examination of the choice between the maximum distance and feature approximation methods might be desirable. Although using them both in concert guarantees a smaller candidate set in the feature approximation method, pre-computation of the features constitutes a moderately expensive sunk cost in terms of computation and storage. In the example in Section 4.2, a $500 \times 50,000$ matrix needed to be computed and stored in advance. In either case, the two methods outperform exhaustive search, as shown in Table 5.

Acknowledgments

The authors gratefully acknowledge funding from the National Science Foundation, awards DMS-1564438, DMS-1621722, and DMS-1621746.

Supplementary Materials

The proofs of Proposition 1 and Theorems 1 and 2 are given in the online supplement.

References

- Arya, S., Mount, D., Kemp, S. E., and Jefferis, G. (2015). *RANN: Fast Nearest Neighbour Search (Wraps Arya and Mount's ANN Library)*. R package version 2.5.
- Bawa, M., Condie, T., and Ganesan, P. (2005). Lsh forest: self-tuning indexes for similarity search. In *Proceedings of the 14th International Conference on World Wide Web*, 651–660. ACM.
- Bellosta, C. J. G. (2015). *rPython: Package Allowing R to Call Python*. R package version 0.0-6.
- Bentley, J. L. (1975). Multidimensional binary search trees used for associative searching. *Communications of the ACM*, **18**, 509–517.
- Bratley, P. and Fox, B. L. (1988). Algorithm 659: Implementing sobol's quasirandom sequence generator. *ACM Transactions on Mathematical Software (TOMS)*, **14**, 88–100.
- Cangelosi, A. and Parisi, D. (2012). *Simulating the Evolution of Language*. Springer Science & Business Media.
- Datta, A., Banerjee, S., Finley, A. O., and Gelfand, A. E. (2016). Hierarchical nearest-neighbor gaussian process models for large geostatistical datasets. *Journal of the American Statistical Association*, to appear. arXiv:1406.7343.
- Eckstein, W. (2013). *Computer Simulation of Ion-Solid Interactions*, volume 10. Springer Science & Business Media.
- Fang, K.-T., Li, R., and Sudjianto, A. (2005). *Design and Modeling for Computer Experiments*. CRC Press.
- Furrer, R., Genton, M. G., and Nychka, D. (2006). Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*, **15**, 502–523.

- Golub, G. H. and Van Loan, C. F. (1996). *Matrix Computations*. JHU Press.
- Gramacy, R. B. (2016). 1aGP: Large-scale spatial modeling via local approximate Gaussian processes in R. *Journal of Statistical Software*, **72**, 1–46.
- Gramacy, R. B. and Apley, D. W. (2015). Local gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics*, **24**, 561–578.
- Gramacy, R. B. and Haaland, B. (2016). Speeding up neighborhood search in local gaussian process prediction. *Technometrics*, **58**, 294–303.
- Gramacy, R. B., Niemi, J., and Weiss, R. M. (2014). Massively parallel approximate gaussian process regression. *SIAM/ASA Journal on Uncertainty Quantification*, **2**, 564–584.
- Haaland, B., Qian, P. Z., et al. (2011). Accurate emulators for large-scale computer experiments. *The Annals of Statistics*, **39**, 2974–3002.
- Harville, D. A. (1997). *Matrix Algebra from a Statistician’s Perspective*, volume 157. Springer.
- Indyk, P. and Motwani, R. (1998). Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the Thirtieth annual ACM symposium on Theory of computing*, 604–613. ACM.
- Jain, P., Kulis, B., and Grauman, K. (2008). Fast image search for learned metrics. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, 1–8. IEEE.
- Kuczynski, J. and Wozniakowski, H. (1992). Estimating the largest eigenvalue by the power and lanczos algorithms with a random start. *SIAM Journal on Matrix Analysis and Applications*, **13**, 1094–1122.
- Lam, C. Q. and Notz, W. I. (2008). *Sequential adaptive designs in computer experiments for response surface model fit*. PhD thesis, The Ohio State University.

- Leskovec, J., Rajaraman, A., and Ullman, J. D. (2014). *Mining of Massive Datasets*. Cambridge University Press.
- Liu, Y. and Hung, Y. (2015). Latin hypercube design-based block bootstrap for computer experiment modeling. Technical report, Rutgers.
- Nychka, D., Bandyopadhyay, S., Hammerling, D., Lindgren, F., and Sain, S. (2015). A multi-resolution gaussian process model for the analysis of large spatial data sets. *Journal of Computational and Graphical Statistics*, **24**, 579–599.
- Paciorek, C. J., Lipshitz, B., Zhuo, W., Kaufman, C. G., Thomas, R. C., et al. (2015). Parallelizing gaussian process calculations in r. *Journal of Statistical Software*, **63**, 1–23.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, **12**, 2825–2830.
- Plumlee, M. (2014). Fast prediction of deterministic functions using sparse grid experimental designs. *Journal of the American Statistical Association*, **109**, 1581–1591.
- R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989). Design and analysis of computer experiments. *Statistical Science*, **4**, 409–423.
- Santner, T. J., Williams, B. J., and Notz, W. I. (2013). *The Design and Analysis of Computer Experiments*. Springer Science & Business Media.
- Shewry, M. C. and Wynn, H. P. (1987). Maximum entropy sampling. *Journal of Applied Statistics*, **14**, 165–170.

- Van Durme, B. and Lall, A. (2010). Online generation of locality sensitive hash signatures. In *Proceedings of the ACL 2010 Conference Short Papers*, 231–235. Association for Computational Linguistics.
- Wendland, H. (2004). *Scattered Data Approximation*, volume 17. Cambridge University Press.
- Williams, C. and Seeger, M. (2001). Using the nyström method to speed up kernel machines. In *Proceedings of the 14th Annual Conference on Neural Information Processing Systems*, number EPFL-CONF-161322, 682–688.
- Zhu, H., Williams, C. K., Rohwer, R., and Morciniec, M. (1997). Gaussian regression and optimal finite dimensional linear models. *Aston University*.