



HAL
open science

Small ensembles of kriging models for optimization

Hossein Mohammadi, Rodolphe Le Riche, Eric Touboul

► **To cite this version:**

Hossein Mohammadi, Rodolphe Le Riche, Eric Touboul. Small ensembles of kriging models for optimization. [Research Report] LIMOS; Ecole National Supérieure des Mines de Saint-Etienne. 2016. hal-01285225

HAL Id: hal-01285225

<https://hal.science/hal-01285225v1>

Submitted on 8 Mar 2016

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

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

Small ensembles of kriging models for optimization

Hossein Mohammadi^{1, 2}, Rodolphe Le Riche ^{*2, 1}, and Eric Touboul^{1, 2}

¹Ecole des Mines de Saint Etienne, H. Fayol Institute

²CNRS LIMOS, UMR 5168

Abstract

The Efficient Global Optimization (EGO) algorithm uses a conditional Gaussian Process (GP) to approximate an objective function known at a finite number of observation points and sequentially adds new points which maximize the Expected Improvement criterion according to the GP. The important factor that controls the efficiency of EGO is the GP covariance function (or kernel) which should be chosen according to the objective function. Traditionally, a parameterized family of covariance functions is considered whose parameters are learned through statistical procedures such as maximum likelihood or cross-validation. However, it may be questioned whether statistical procedures for learning covariance functions are the most efficient for optimization as they target a global agreement between the GP and the observations which is not the ultimate goal of optimization. Furthermore, statistical learning procedures are computationally expensive. The main alternative to the statistical learning of the GP is self-adaptation, where the algorithm tunes the kernel parameters based on their contribution to objective function improvement. After questioning the possibility of self-adaptation for kriging based optimizers, this paper proposes a novel approach for tuning the length-scale of the GP in EGO: At each iteration, a small ensemble of kriging models structured by their length-scales is created. All of the models contribute to an iterate in an EGO-like fashion. Then, the set of models is densified around the model whose length-scale yielded the best iterate and further points are produced. Numerical experiments are provided which motivate the use of many length-scales. The tested implementation does not perform better than the classical EGO algorithm in a sequential context but show the potential of the approach for parallel implementations.

Keywords: Continuous Global Optimization; EGO; Gaussian processes; Kernel parameters; optimization based on surrogate ensembles.

*Corresponding author: Ecole Nationale Supérieure des Mines de Saint Etienne, Institut H. Fayol, 158, Cours Fauriel, 42023 Saint-Etienne cedex 2 - France
Tel : +33477420023
Email: leriche@emse.fr

1 Introduction

The EGO optimization algorithm uses a kriging model, which is a conditional Gaussian process (GP) [19], for predicting objective function values and quantifying the prediction uncertainty. The shapes of sample paths of a GP such as its smoothness, periodicity, etc. are controlled by the covariance function of the process, also known as its kernel. Traditionally, a parameterized family of covariance functions is considered whose parameters are estimated.

The kernel parameters are often estimated by statistical approaches like maximum likelihood (ML)[24] or cross validation (CV) [25]. ML and CV are compared in [2] when the covariance structure of a GP is misspecified. It is recommended in [16] to use a penalized likelihood for the kriging models when the sample size is small. However, the efficiency of such statistical approaches, which aims at learning the objective function globally, remains questionable in the context of optimization. For example, in the EGO algorithm if the design points do not carry enough information about the true function, the parameters are not estimated correctly. These parameters are then plugged into the expected improvement (EI) criterion that may lead to disappointing results [13, 4].

Not surprisingly, several methods alternative to ML and CV have been proposed to tune the kernel parameters. For instance, in [7] the kernel parameters are estimated with a log normal prior density assumption over them. A fully Bayesian approach is used in [4, 20]. In [14, 6], the process of estimating parameters and searching for the optimum are combined together through a likelihood which encompasses a targeted objective. In [23], the bounds on the length-scales values are changing with the iterations following an a priori schedule.

Another drawback of statistical learning procedures such as ML and CV in the context of moderately expensive functions¹ is their computational complexity as they involve the repeated inversion of an $n \times n$ covariance matrix (where n is the number of available observations) where each inversion needs of the order of n^3 operations.

This paper considers isotropic kernels and investigates an alternative approach to tuning the length-scale parameter. In this approach, a small set of length scales (hence GP models) is first tested as alternative ways to consider the objective function, independently of their statistical relevance. The set is completed based on the direct contribution of the best model to the optimization. The method is based on ensembles of surrogates. It can also be seen as weakly self-adaptive in the sense of self-adaptive algorithms [3, 9] where no statistical measure intervenes in the building of the representation which the optimization algorithm has of the objective function.

Ensembles of surrogates have attracted a lot of attention from the machine learning community for prediction [10], but fewer contributions seem to address surrogate ensembles for optimizing. Several approaches have been proposed that aggregate the metamodels of the ensemble into a hopefully better

¹We call “moderately expensive” functions that take between 10 seconds and an hour to be evaluated at one point.

metamodel either by model selection or by mixing the models. This better metamodel is then used by the optimization algorithm [1, 5, 8].

On the opposite, other previous optimization methods take advantage of all the metamodels in the set as a diversity preserving mechanism (in addition to, of course, a way to reduce the number of calls to the objective function), in the context of evolutionary computation [12, 17] or more generally [21]. The algorithm studied in this text belongs to this category.

Another classification can be made with respect to the homogeneity (all metamodels are of the same type) or heterogeneity of the ensemble. There has been recent contributions to optimization algorithms that rely on a homogeneous set of kriging models: in [15] the ensembles are built by bootstrap on the data and serve as a way to estimate model uncertainty for later use in optimization; in [22], the metamodels are the trajectories of a GP and their contributions are aggregated through an uncertainty reduction criterion (on the entropy of the global optima of the trajectories). The optimization algorithm investigated here also relies on an homogeneous ensemble of GP models.

2 EGO algorithm overview

EGO is a sequential model-based optimization algorithm. It starts with an initial design of experiments (DoE). At each iteration, one point which maximizes the Expected Improvement (EI) according to the current kriging model is added to the DoE. Then, the kernel parameters are re-estimated and the kriging model is updated. The location of \mathbf{x}^{n+1} , where $\mathbf{x}^{n+1} = \arg \max_{x \in \mathcal{S}} EI(\mathbf{x})$, depends

Algorithm 1 Efficient Global Optimization Algorithm (EGO)

Create an initial design: $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^n]^T$.
 Evaluate the functions at \mathbf{X} , $\mathbf{y} = f(\mathbf{X})$.
 Fit a kriging model to the data points $(\mathbf{X}, \mathbf{y}) =$ estimate θ, μ, σ^2
while not stop do
 $\mathbf{x}^{n+1} \leftarrow \arg \max_{\mathbf{x} \in \mathcal{S}} EI(\mathbf{x})$ and add \mathbf{x}^{n+1} to \mathbf{X} .
 $y^{n+1} \leftarrow f(\mathbf{x}^{n+1})$ and add y^{n+1} to \mathbf{y} .
 Re-estimate the parameters (θ, μ, σ^2) and update the kriging model.
end while

on the current DoE, \mathbf{X} , \mathbf{y} , the kriging trend, μ , and the kernel parameters: the length-scale, θ , and the process variance, σ^2 . We use $\mathbf{x}^{n+1} = g(\mathbf{X}, \mu, \theta, \sigma^2)$ to denote that \mathbf{x}^{n+1} is a function of the above-mentioned parameters. Figure 1 illustrates how the DoE and the magnitude of length-scale affect the EI.

Among the parameters of the EI criterion, \mathbf{X} and θ play a prominent role because once \mathbf{X} and θ are fixed, the ML estimations of μ and σ^2 have a closed-form expression [19]:

$$\hat{\mu} = \frac{\mathbf{1}^\top \mathbf{R}^{-1}(\theta) \mathbf{y}}{\mathbf{1}^\top \mathbf{R}^{-1}(\theta) \mathbf{1}}, \quad (1)$$

$$\hat{\sigma}^2 = \frac{(\mathbf{y} - \hat{\mu} \mathbf{1})^\top \mathbf{R}^{-1}(\theta) (\mathbf{y} - \hat{\mu} \mathbf{1})}{n}. \quad (2)$$




Figure 1: Effect of DoE and length-scale on EI function. The function to be optimized is the Sphere whose global minimum is located at 2.5. The blue and magenta curves represent the EI of kriging models with length-scales equal to 5 and 0.2, respectively. The crosses indicate the location of design points. The other parameters are fixed. The location of the third sample point changes from 2 to 1.5 in the right picture.

Accordingly, \mathbf{x}^{n+1} can be expressed as a function of \mathbf{X} and θ . For example, Figure 2 shows all plausible next infill sample points by changing the length-scale for a given DoE.




Figure 2: Illustration of all possible next infill sample points with $\mathbf{X} = \{-5, -2, 2, 5\}$ as the DoE. The true functions are Sphere (left, as in Figure 1) and Ackley (right) in dimension 1. For θ values larger than, say $\theta \geq 8$, the location of x^{n+1} is quite stable and close to 2.5, the location of the global minimum. While large θ 's lead to the global optimum of the Sphere for any \mathbf{X} , it is a coincidence for Ackley's function.

3 Tuning the length-scale from an optimization point of view: a study on self-adaptation

When the kernel parameters are estimated by ML, the selected kriging model has statistical best agreement with the observed data. However, the goal of using EGO, like other optimization algorithms, is to solve an optimization problem with the least number of function evaluations. In other words, the main goal is the fast convergence of EGO even if the kriging model does not represent well the true function. This idea is similar to the notion of “self-adaptation” in evolutionary optimization [3, 9].

To investigate the potential of tuning the length-scale θ in an optimization oriented, greedy, self-adaptive way, we first tested a theoretical algorithm that tries all possible values of θ in the range $[0.01, 20]$. The true objective functions of the points that maximize the expected improvement for each of these length-scale θ value is calculated, $\mathbf{x}^{n+1}(\theta|\mathbf{X}) = \arg \max_{x \in \mathcal{S}} EI(\mathbf{x}; \theta)$. This makes this algorithm not practical in the context of expensive problems. Then, the iterate associated to the best objective function, $\mathbf{x}^{\text{sel}} = \arg \min_{\mathbf{x}^{n+1}} f(\mathbf{x}^{n+1}(\theta|\mathbf{X}))$, is added to the Design of Experiment \mathbf{X} , the kriging model is updated, and the algorithm loops. This algorithm is sketched in the flow chart 2.

From a one step ahead optimization point of view, the “best” length-scale, denoted by θ^* , is the one that yields the next infill sample with the lowest objective function value, $\theta^* = \arg \min_{\theta} f(\mathbf{x}^{n+1}(\theta|\mathbf{X}))$. In the examples provided in Figure 3, the best length-scales are shown for the two test functions (Ackley and Sphere). In this example, the best length-scales are different from the length-scales estimated by ML, see the caption of Figure 3.

Algorithm 2 Toy EGO with greedy θ tuning

```

Create an initial design:  $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^n]^T$ 
Evaluate the functions at  $\mathbf{X}$ ,  $\mathbf{y} = f(\mathbf{X})$ 
Select  $\mathbf{x}^j$ ,  $1 \leq j \leq n$ , for which  $f(\mathbf{x}^j) = \max(\mathbf{y})$ 
while not stop (typically a limit on budget) do
  Set  $\mathbf{x}^{\text{sel}} = \mathbf{x}^j$ 
  for  $\theta_i \in [\theta_{\min}, \dots, \theta_{\max}]$  do
     $\mathbf{x}^{n+1}(\theta_i|\mathbf{X}) = \arg \max_{x \in \mathcal{S}} EI(\mathbf{x}; \theta_i)$ 
    if  $f(\mathbf{x}^{n+1}(\theta_i|\mathbf{X})) < f(\mathbf{x}^{\text{sel}})$  then
       $\mathbf{x}^{\text{sel}} \leftarrow \mathbf{x}^{n+1}(\theta_i|\mathbf{X})$ 
    end if
  end for
   $\mathbf{X} \leftarrow \mathbf{X} \cup \mathbf{x}^{\text{sel}}$ 
end while

```




Figure 3: Function values of x^{n+1} already shown in Figure 2. The asterisk indicate the correlation length-scale, θ^* , which causes the maximum improvement in the objective function. In this example, θ^* is different from $\hat{\theta}_{ML}$, estimated by ML,: $\theta^* = 0.61271$ and $\hat{\theta}_{ML} = 5.34$ (Sphere; left), $\theta^* = 12.7674$ and $\hat{\theta}_{ML} = 0.01$ (Ackley; right). Both functions have their global minimum at 2.5 and the DoE is $\mathbf{X} = \{-5, -2, 2, 5\}$.

We now analyze this approach in more details by providing some examples in $2D$. Figure 4 illustrates the first and the second iterations of this algorithm again on the Sphere and Ackley functions. In this Figure, the location of the points that maximize the expected improvement for different length-scale values is plotted on the top of the true function contour lines. In total, 64 length-scales, started from 0.01, are used. The length-scales are divided into eight groups. Each group consists of eight length-scales in ascending order. The i th group is denoted by $\theta^{(i:8)}, i = 1, \dots, 8$ and is defined as $[0.01 + 8(i - 1) \times \alpha_{\text{increment}}, 0.01 + 8i \times \alpha_{\text{increment}})$ where $\alpha_{\text{increment}} \approx 0.1$. The infill sample points obtained by the length-scales of a particular group have identical color, see the legend of Figure 4.

The first remark that can be done, and which motivates this study, is that the points visited as θ changes make a one dimensional manifold (obviously since it is parameterized by the scalar θ), continuous by parts and, most interestingly, often curved towards the global optimum of the function. The discontinuities of the trajectory are associated to changes of basin of attraction during the maximization of the expected improvement. This simple observation, even though only based on a few cases, is a hint that the volume search of global optimization algorithms might be iteratively transformed into a one dimensional search in θ , with potentials for containing the ‘‘curse of dimensionality’’ (the geometric increase in search space volume as the number of dimensions increases). The difficulties of the associated problem and a possible implementation will be discussed in the next section.

In Figure 4, it can be seen that the magnitude of the ‘‘best’’ length-scale in the first iteration is between 2 and 3, i.e., $\theta^* \in \theta^{(3:8)}$ or $\theta^{(4:8)}$. While EGO

with a small length-scale samples near the best observed point (cf. the black points), EGO with large length-scale is more explorative (see yellow and grey points) [18]. The search points and the length-scales obtained by the algorithm after 15 iterations are given in Figure 5. It can be observed that, after the first iterations where the “best” length-scale magnitude, θ^* , is of order 1, θ^* oscillates at usually small values. Because θ^* oscillates, self-adaptive strategies and Bayesian strategies based on assuming a prior density over the length-scale may not be a good strategy for optimization (at least if θ^* makes an efficient strategy).




Figure 4: First (top row) and second (bottom row) iterations of EGO in which $\mathbf{x}^{n+1}(\theta^*|\mathbf{X}) = \arg \max_{\mathbf{x} \in \mathcal{S}} EI(\mathbf{x}|\theta^*)$ is added to the existing DoE, the crosses, on the Sphere (left) and the Ackley (right) functions. 64 equally distant length-scales are grouped into eight equal sized intervals, $\theta^{(i:8)}$, $i = 1, \dots, 8$. The infill sample points obtained by the length-scales of a particular group have identical color.




Figure 5: DoEs created by the toy greedy algorithm 2 after 15 iterations on the Sphere (left) and the Ackley (middle) functions. Right: plot of “best” length-scale, θ^* . θ^* oscillates during optimization iterations and usually has a small magnitude after the first iterations. The y-axis is in logarithmic scale.

In order to investigate the effect of initial DoE on the algorithm performance, the above experiments are repeated with another initial DoE. Figure 6 shows the results which are similar to the previous experiments. For example, the length-scales tend to be small especially in the case of highly multimodal Ackley function. The algorithm’s behavior, typical of small θ ’s (as explained in details in [18]) is greedy, that of a local search algorithm: local convergences can be seen in Figure 8 where the function to be optimized is Rastrigin with several local minima.

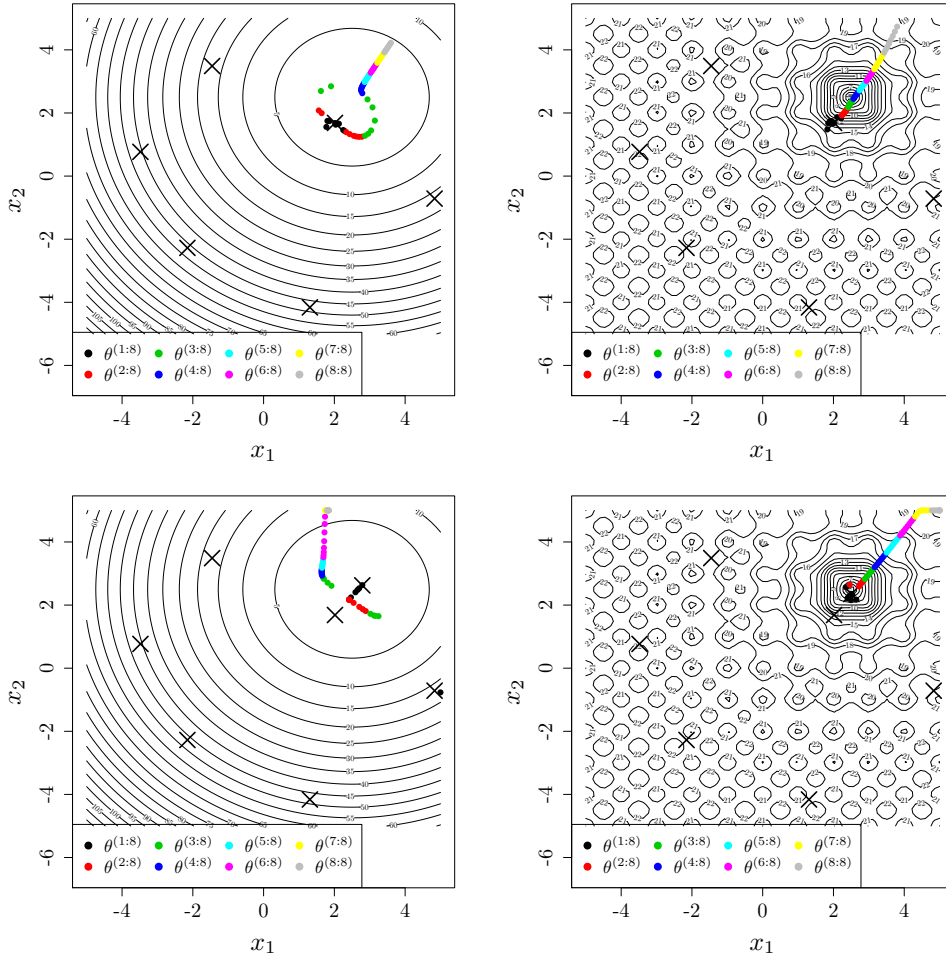


Figure 6: First (top row) and second (bottom row) iteration of the toy greedy algorithm 2 on the Sphere (left) and the Ackley functions(right). The initial DoE is different from the one shown in Figure 4. For more information see the caption of Figure 4.




Figure 7: DoEs created by the toy greedy algorithm 2 after 15 iterations on the Sphere (left) and the Ackley (middle) functions. Right: plot of “best” length-scale, θ^* . The initial DoE is different from the one shown in Figure 5.




Figure 8: DoEs created by the toy greedy algorithm 2 after 15 iterations on the Rastrigin function with two DoEs (left and middle). Right: plot of “best” length-scale, θ^* . The global minimum is located at (2.5, 2.5).

4 An EGO algorithm with a small ensemble of kriging models

4.1 Description of the algorithm

EGO is used for the optimization of computationally intensive functions. So, it is practically impossible to calculate $f(\mathbf{x}^{n+1}(\theta|\mathbf{X}))$ for many length-scales in order to obtain θ^* . Herein, we propose an approach that works with a limited number of kriging models. The ensemble of kriging models is structured by the length-scales. The pseudo-code is given below (Algorithm 3) followed by a detailed explanation of the approach.

Let (\mathbf{X}, \mathbf{y}) be the initial design of experiments. The covariance function we use here is the isotropic Matérn 5/2 kernel [19]. Thus, there exists only one length-scale to be tuned. The first reason for using an isotropic kernel is simplicity and clarity in the analysis. By taking isotropic functions and kernels, a difficult aspect of the algorithm (anisotropy, which is related to variables sensitivity) is neutralized to focus on other (also quite complex) phenomena. By taking isotropic kernels, the results of the numerical experiments are more stable. The second reason is that isotropic kernels have been found to perform well for EGO in high-dimension in the context of expensive-to-evaluate functions [11].

At each iteration, five length-scales are generated. They are sampled on a basis 10 logarithmic scale from $[-2, 1]$ based on a Latin Hypercube Sampling (LHS) plan (that is θ ranges from 10^{-2} to 10^1). Then, they are sorted and scaled back, $\theta_i = 10^{\log \theta_i}$, $1 \leq i \leq 5$; $\theta_1 < \theta_2 < \dots < \theta_5$. Corresponding to each length-scale θ_i , a kriging model is created which gives a new infill sample: $\mathbf{x}^{n+1}(\theta_i|\mathbf{X}) = \arg \max_{\mathbf{x} \in \mathcal{S}} EI(\mathbf{x}; \theta_i)$.

In the next step, the $\mathbf{x}^{n+1}(\theta_i|\mathbf{X})$, $1 \leq i \leq 5$, that are not close to the design points are selected and the function is evaluated there. The notion of closeness is expressed by defining a neighborhood of radius $R^{(t)}$ around design points, see Figure 9. It is important to prevent the points from converging around early good performers, otherwise such greedy algorithm where decisions are taken solely on the account of objective function values would not

Algorithm 3 EGO based on a small ensemble of kriging models

Create an initial design: $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^n]^\top$.
Evaluate function at \mathbf{X} and set $\mathbf{y} = f(\mathbf{X})$.
Set the maximum number of evaluations, t_{\max} .
for $t \leftarrow n+1$ **to** t_{\max} **do**
 Define a neighborhood of radius $R^{(t)}$ around the current sample points.
 Set $\mathbf{X}^{(n+1)} = \emptyset$ and $\mathbf{X}^{\text{sel}} = \emptyset$.
 Generate q length-scales, $\theta_1, \dots, \theta_q$.
 for $i \leftarrow 1$ **to** q **do**
 $\mathbf{x}^{n+1} \leftarrow \arg \max_{\mathbf{x} \in \mathcal{S}} EI(\mathbf{x}; \theta_i)$.
 $\mathbf{X}^{(n+1)} \leftarrow \mathbf{X}^{(n+1)} \cup \mathbf{x}^{n+1}$.
 if \mathbf{x}^{n+1} is not inside the defined neighborhoods **then**
 $\mathbf{X}^{\text{sel}} \leftarrow \mathbf{X}^{\text{sel}} \cup \mathbf{x}^{n+1}$.
 end if
 end for
 if $\mathbf{X}^{\text{sel}} = \emptyset$ **then**
 $\mathbf{X}^{\text{sel}} \leftarrow \arg \max (\min_{\mathbf{x} \in \mathbf{X}^{(n+1)}} \text{dist}(\mathbf{x}, \mathbf{X}))$
 end if
 Evaluate function at \mathbf{X}^{sel} and set $\mathbf{y}^{\text{sel}} = f(\mathbf{X}^{\text{sel}})$.
 Select θ^* , for which $f(\arg \max_{\mathbf{x} \in \mathcal{S}} EI(\mathbf{x}; \theta^*)) = \min(\mathbf{y}^{\text{sel}})$.
 Generate two length-scales close to θ^* . This yields two new infill samples
 by EI maximization, $\mathbf{X}^{\text{new}} = [\mathbf{x}^{\text{new}1}, \mathbf{x}^{\text{new}2}]^\top$.
 Evaluate function at \mathbf{X}^{new} and set $\mathbf{y}^{\text{new}} = f(\mathbf{X}^{\text{new}})$.
 Update the DoE: $\mathbf{X} \leftarrow \mathbf{X} \cup \mathbf{X}^{\text{sel}} \cup \mathbf{X}^{\text{new}}$, $\mathbf{y} \leftarrow \mathbf{y} \cup \mathbf{y}^{\text{sel}} \cup \mathbf{y}^{\text{new}}$.
end for

be sufficiently explorative for global optimization. Further explanations about the neighborhood definition are provided in the next paragraph. The eligible $\mathbf{x}^{n+1}(\theta_i|\mathbf{X})$, $1 \leq i \leq 5$, are selected and stored in the matrix \mathbf{X}^{sel} . \mathbf{y}^{sel} contains the function values at \mathbf{X}^{sel} .




Figure 9: DoE and neighborhoods as balls around the design points (blue circles). The infill samples occurring inside any neighborhood are not considered by the optimizer.

The neighborhood defined around every design point is a ball with radius $R^{(t)}$ where the index t is the iteration. As the optimization progresses, the radius shrinks according to the following linear scheme:

$$R^{(t)} = \begin{cases} R^{(1)} - \frac{R^{(1)}}{t_{\text{threshold}}} \times (t - 1) & \text{if } t \leq t_{\text{threshold}} \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

in which $t_{\text{threshold}}$ is 70% of total number of iterations, t_{max} . The initial radius $R^{(1)}$, is half of the distance between the best initial DoE (based on its f value) and the closest design point to it. Again, defining such neighborhoods prevents the algorithm from focusing around good points too early.

Now, among the five generated length-scales, the best one is selected and is denoted by θ^* . Recall that the best length-scale is the one that yields $f(\mathbf{x}^{n+1}(\theta_i|\mathbf{X})) = \min(\mathbf{y}^{\text{sel}})$. Then, two length-scales, θ_-^* and θ_+^* , close to θ^* are generated. They are defined as:

- If $\theta^* = \theta_i$, $2 \leq i \leq 4$, $\theta_-^* = \theta^* - \frac{1}{3}(\theta^* - \theta_{i-1})$ and $\theta_+^* = \theta^* + \frac{1}{3}(\theta_{i+1} - \theta^*)$.
- If $\theta^* = \theta_1$, $\theta_-^* = 0.01$ and $\theta_+^* = \theta^* + \frac{1}{3}(\theta_2 - \theta^*)$.
- If $\theta^* = \theta_5$, $\theta_-^* = \theta^* - \frac{1}{3}(\theta^* - \theta_4)$ and $\theta_+^* = 10$.

The two new infill samples obtained with the kriging models with length-scales θ_-^* and θ_+^* are stored in the \mathbf{X}^{new} matrix,

$$\mathbf{X}^{\text{new}} = [\mathbf{x}^{n+1}(\theta_-^*|\mathbf{X}), \mathbf{x}^{n+1}(\theta_+^*|\mathbf{X})]^\top. \quad (4)$$

Finally, the current DoE (\mathbf{X}, \mathbf{y}) is updated by adding \mathbf{X}^{new} and \mathbf{X}^{sel} to \mathbf{X} and \mathbf{y}^{new} and \mathbf{y}^{sel} to \mathbf{y} . This procedure continuous until the budget is exhausted.

4.2 Tests of the algorithm

The performance of this EGO method that is based on a small ensemble of kriging models (5+2 models) is tested on three isotropic functions, Sphere, Ackley and Rastrigin. The functions are defined in $\mathcal{S} = [-5, 5]^d$ where $d = 5$. The total number of iterations is $15 \times d$. Each optimization run is repeated eight times (thin black lines). Figure 10 shows the results and the performance of the standard EGO method (thin blue lines) which is repeated five times with a budget equals to $70 \times d$. The plots show the best objective functions observed so far. The initial DoE is fixed for both algorithms and has a size equal to $3 \times d$. The thick lines are the median of the runs.

The small ensemble version of EGO is slightly better on the sphere function because it benefits from its greedy choice of points that are never misleading. On Rastrigin and Ackley, the small ensemble EGO is slower early in the search, which might be due to the schedule of $R^{(t)}$. Later on, still on Rastrigin and Ackley, EGO with a small ensemble shows both the worst and best performances, therefore illustrating a tendency to get trapped in local optima. In terms of median performance, after 250 evaluations of the objective function (at the time when the neighborhood control ceases), the small ensemble EGO is equivalent to EGO on Rastrigin and worse on Ackley.




Figure 10: Best objective function vs. number of calls of EGO with the ensemble of kriging models (thin black lines) and standard EGO (thin blue lines) on Sphere(top left), Ackley (top right) and Rastrigin (bottom) functions. The thick lines show the median of the runs.

5 Conclusions

We have investigated a variant of the EGO optimization algorithm where, instead of using at each iteration a kriging model learned through a statistical estimation procedure such as maximum likelihood, a small set of models with a fixed length-scale is employed. The motivations are threefolds. Firstly, it has been noticed in two-dimensions that the manifolds of the points that maximize expected improvement for various length-scales approach rapidly the global optimum. Secondly, ensemble methods have a lower computational complexity since the number of kriging covariance matrices inversions is limited to the number of elements in the ensemble, seven in the current work. On the contrary, maximum likelihood or cross-validation approaches require the inversion of the covariance matrix at each of their internal iteration. Thirdly, ensemble

methods may more easily lead to parallel versions of EGO as the maximization of expected improvement can be distributed on several computing nodes, one for each kriging model.

Our first investigations have led to the following conclusions: tuning the length-scale to achieve an immediate improvement in the objective function may not be as efficient a strategy as two-dimensional plots of the manifold seem to indicate; the greediness of the method is a source of premature convergence to good performing points; optimal values of the length scale (in the sense of short term improvement) change a lot from one iteration to the next as the design of experiments evolves, rendering self-adaptive and Bayesian strategies not efficient for this purpose.

Nevertheless, we believe that the idea of searching in the space of length-scales as a proxy for searching in the space of optimization variables deserves further investigations because of its potential for tackling the curse of dimensionality. In particular, the schedule of the neighborhood radius, an iteration-smoothing learning procedure for the length-scales, and alternative strategies for making the ensemble of kriging should be studied.

References

- [1] Erdem Acar and Masoud Rais-Rohani. Ensemble of metamodels with optimized weight factors. *Structural and Multidisciplinary Optimization*, 37(3):279–294, 2009.
- [2] Francois Bachoc. Cross validation and maximum likelihood estimations of hyper-parameters of Gaussian processes with model misspecification. *Computational Statistics & Data Analysis*, 66:55–69, 2013.
- [3] Thomas Bäck. *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*. Oxford University Press, Oxford, UK, 1996.
- [4] Romain Benassi, Julien Bect, and Emmanuel Vazquez. Robust gaussian process-based global optimization using a fully bayesian expected improvement criterion. In Carlos A. Coello Coello, editor, *Learning and Intelligent Optimization*, volume 6683 of *Lecture Notes in Computer Science*, pages 176–190. Springer Berlin Heidelberg, 2011.
- [5] Anirban Chaudhuri, Rodolphe Le Riche, and Mickael Meunier. Estimating Feasibility Using Multiple Surrogates and ROC Curves. In *54th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Boston, France, April 2013.
- [6] Alexander I.J. Forrester and Donald R. Jones. Global optimization of deceptive functions with sparse sampling. In *12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*. American Institute of Aeronautics and Astronautics, 2008.

- [7] Marcus R. Frean and Phillip Boyle. Using Gaussian processes to optimize expensive functions. In *AI 2008: Advances in Artificial Intelligence, 21st Australasian Joint Conference on Artificial Intelligence, Auckland, New Zealand, December 1-5, 2008. Proceedings*, pages 258–267, 2008.
- [8] Tushar Goel, RaphaelT. Haftka, Wei Shyy, and NestorV. Queipo. Ensemble of surrogates. *Structural and Multidisciplinary Optimization*, 33(3):199–216, 2007.
- [9] Nicolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195, 2001.
- [10] Stefan Hess, Tobias Wagner, and Bernd Bischl. *Learning and Intelligent Optimization: 7th International Conference, LION 7, Catania, Italy, January 7-11, 2013, Revised Selected Papers*, chapter PROGRESS: Progressive Reinforcement-Learning-Based Surrogate Selection, pages 110–124. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [11] Frank Hutter, Holger Hoos, and Kevin Leyton-Brown. An evaluation of sequential model-based optimization for expensive blackbox functions. In *Proceedings of the 15th Annual Conference Companion on Genetic and Evolutionary Computation, GECCO '13 Companion*, pages 1209–1216, New York, NY, USA, 2013. ACM.
- [12] Yaochu Jin and Bernhard Sendhoff. *Genetic and Evolutionary Computation – GECCO 2004: Genetic and Evolutionary Computation Conference, Seattle, WA, USA, June 26-30, 2004. Proceedings, Part I*, chapter Reducing Fitness Evaluations Using Clustering Techniques and Neural Network Ensembles, pages 688–699. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004.
- [13] Donald R. Jones. A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization*, 21:345–383, 2001.
- [14] Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, 1998.
- [15] Jack P. C. Kleijnen. Simulation-optimization via kriging and bootstrapping: a survey. *J. Simulation*, 8(4):241–250, 2014.
- [16] Runze Li and Agus Sudjianto. Analysis of computer experiments using penalized likelihood in Gaussian kriging models. *Technometrics*, 47(2), 2005.
- [17] Jianfeng Lu, Bin Li, and Yaochu Jin. An evolution strategy assisted by an ensemble of local gaussian process models. In *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation, GECCO '13*, pages 447–454, New York, NY, USA, 2013. ACM.

- [18] Hossein Mohammadi, Rodolphe Le Riche, and Eric Touboul. A detailed analysis of kernel parameters in Gaussian process-based optimization. Technical report, Ecole Nationale Supérieure des Mines ; LIMOS, 2015.
- [19] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning. The MIT Press, 2005.
- [20] Sam Davanloo Tajbakhsh, Enrique Castillo, and James L Rosenberger. A Bayesian approach to sequential optimization based on computer experiments. *Quality and Reliability Engineering International*, 31(6):1001–1012, 2015.
- [21] Felipe A.C. Viana, Raphael T. Haftka, and Layne T. Watson. Efficient global optimization algorithm assisted by multiple surrogate techniques. *Journal of Global Optimization*, 56(2):669–689, 2013.
- [22] Julien Villemonteix, Emmanuel Vazquez, and Eric Walter. An informational approach to the global optimization of expensive-to-evaluate functions. *Journal of Global Optimization*, 44(4):509–534, 2008.
- [23] Ziyu Wang, Masrour Zoghi, Frank Hutter, David Matheson, and Nando de Freitas. Bayesian optimization in high dimensions via random embeddings. In *International Joint Conferences on Artificial Intelligence (IJCAI)*, 2013.
- [24] Zhiliang Ying. Asymptotic properties of a maximum likelihood estimator with data from a Gaussian process. *Journal of Multivariate Analysis*, 36(2):280 – 296, 1991.
- [25] Hao Zhang and Yong Wang. Kriging and cross-validation for massive spatial data. *Environmetrics*, 21(3-4):290–304, 2010.