Innovative Applications of O.R.

Comparison of Kriging-based algorithms for simulation optimization with heterogeneous noise

Hamed Jalali^{a,*}, Inneke Van Nieuwenhuyse^a, Victor Picheny^b^a Research Center for Operations Management, Department of Decision Sciences and Information Management, KU Leuven, Naamsestraat 69, Leuven 3000, Belgium^b INRA-French National Institute for Agricultural Research, 31326 Castanet-Tolosan, France

ARTICLE INFO

Article history:

Received 12 November 2015

Accepted 24 January 2017

Available online 27 January 2017

Keywords:

Simulation

Stochastic Kriging

Heterogeneous noise

Ranking and selection

Optimization via simulation

ABSTRACT

In this article we investigate the unconstrained optimization (minimization) of the performance of a system that is modeled through a discrete-event simulation. In recent years, several algorithms have been proposed which extend the traditional Kriging-based simulation optimization algorithms (assuming deterministic outputs) to problems with noise. Our objective in this paper is to compare the relative performance of a number of these algorithms on a set of well-known analytical test functions, assuming different patterns of heterogeneous noise. We also apply the algorithms to a popular inventory test problem. The conclusions and insights obtained may serve as a useful guideline for researchers aiming to apply Kriging-based algorithms to solve engineering and/or business problems, and may be useful in the development of future algorithms.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Assume that we would like to find the solution that minimizes some goal function $f(\mathbf{x})$:

$$\min_{\mathbf{x} \in \Theta} f(\mathbf{x}), \quad (1)$$

where $f: \Theta \rightarrow \mathbb{R}$ and $\mathbf{x} = (x_1, x_2, \dots, x_d)$ (with d the dimension of the solution space). This goal function cannot be directly observed, and must be estimated through a stochastic simulation model: we thus only have access to noisy observations $\tilde{f}_j(\mathbf{x}^i) = f(\mathbf{x}^i) + \varepsilon_j(\mathbf{x}^i)$, where $\tilde{f}_j(\mathbf{x}^i)$ represents the observed goal value in the j th simulation replication at point $\mathbf{x}^i$. The noise is additive; the distribution of $\varepsilon_j(\mathbf{x}^i)$ has mean zero, and its variance depends on $\mathbf{x}^i$ (denoted by $\tau^2(\mathbf{x}^i)$, we thus have heterogeneous noise). We usually estimate the value of $f(\mathbf{x}^i)$ by performing n^i simulation replications: $\bar{f}(\mathbf{x}^i) = \sum_{j=1}^{n^i} \tilde{f}_j(\mathbf{x}^i) / n^i$.

This type of problem is very common in engineering and business applications since in many practical problems, the objective function is analytically intractable (Law, 2015). In inventory management problems, for instance, we often need to optimize a

system modeled through a discrete event simulation to obtain the optimal stocking quantities (Jalali & Van Nieuwenhuyse, 2015). Examples in other fields include supply chain design (Saif & Elhedhli, 2015), pump scheduling (Naoum-Sawaya, Ghaddar, Arandia, & Eck, 2015), and ambulance fleet allocation (McCormack & Coates, 2015).

In this paper, we focus on problems where Θ is continuous and is bounded by box constraints; for reasons explained in Section 3.1, we discretize the function domain to obtain a finite set of points (as common in the literature; see, e.g., Frazier, Powell, and Dayanik, 2009; Kleijnen, van Beers, and Van Nieuwenhuyse, 2012, and Sun, Hong, & Hu, 2014). Many methods are available for solving problem (1) with a finite set of points such as the ranking and selection (R&S) method (see Hong, Nelson, and Xu, 2015, for a review). Most of them are, however, unsuitable when (1) the number of feasible solutions is large, or (2) simulation replications are time-consuming (Xu, 2012). Kriging-based or Bayesian optimization is among the few techniques that can handle this problem with low problem dimensionality ($d \leq 20$, Brochu, Cora, & De Freitas, 2010; Preuss, Wagner, & Ginsbourger, 2012). Based on the observed $\bar{f}(\mathbf{x}^i)$ for several $\mathbf{x}^i$, Kriging provides an approximation or *metamodel* for $f(\mathbf{x})$. The traditional *Efficient global optimization* (EGO) approach of Jones, Schonlau, and Welch (1998) is one of the most popular Kriging-based algorithms for optimizing *noiseless* simulation; in this case, the fitted metamodel is a deterministic Kriging model (see Kleijnen, 2015, for references on the use of EGO in constrained and multi-objective optimization problems). In stochastic simulation (e.g., discrete event simulation), however, EGO might not be

* Corresponding author.

E-mail addresses: hamed.jalali@kuleuven.be, hamedsi2000@gmail.com (H. Jalali),inneke.vannieuwenhuyse@kuleuven.be (I. Van Nieuwenhuyse), Victor.Picheny@toulouse.inra.fr (V. Picheny).

Fig. 1. Typical steps in Kriging-based optimization.

very appropriate since it ignores the noise in the observations, assuming they were sampled with infinite precision (Quan, Yin, Ng, & Lee, 2013). Being aware of this deficiency, researchers have tried to extend EGO for stochastic simulation. Most of the approaches assume *homogeneous* simulation noise, meaning that the variance of the noise does not depend on $\mathbf{x}$ (Picheny, Ginsbourger, Richet, & Caplin, 2013a); Picheny, Wagner, and Ginsbourger (2013b) compare several Kriging-based algorithms for optimizing functions with this type of noise.

In practice, however, the noise is heterogeneous (Kim & Nelson, 2006; Kleijnen & Van Beers, 2005; Yin, Ng, & Ng, 2011). In recent years, a number of Kriging-based optimization algorithms have been proposed that can handle heterogeneous noise, based on stochastic Kriging models (see Ankenman, Nelson, and Staum, 2010; Cressie, 1993, and Yin et al., 2011). To the best of our knowledge, the performance of these algorithms has not yet been compared. Our objective in this paper is to evaluate the effectiveness and the relative performance of these algorithms for simulation optimization problems with heterogeneous noise. More specifically, we evaluate the performance of six Kriging-based algorithms (details will be explained in Section 2.2; see the overview in Table 2):

- MQ: the minimum quantile criterion (Picheny et al., 2013b), which is similar to the Gaussian process upper confidence bound (GP-UCB) method of Auer, Cesa-Bianchi, and Fischer (2002) and Jones (2001);
- SKO: an adapted version of the sequential Kriging optimization approach (Huang, Allen, Notz, & Zeng, 2006);
- CKG: correlated knowledge-gradient (Frazier et al., 2009);
- EQI: expected quantile improvement (Picheny et al., 2013a);
- TSSO: two-stage stochastic optimization (Quan et al., 2013);
- eTSSO: extended two-stage stochastic optimization (Liu, Pedrielli, & Ng, 2014).

To this end, we apply these algorithms to minimize three well-known analytical test functions (Rescaled Branin, Six-hump Camelback, and Hartmann-6), perturbed with heterogeneous Gaussian noise, over a finite set of solutions. We also apply the algorithms to

optimize the (s, S) inventory system of Fu and Healy (1997), which has been a popular test problem in the simulation optimization literature (Jalali & Van Nieuwenhuyse, 2015).

Though our work is closely related to Picheny et al. (2013b), it clearly differs in the following respects: (1) we consider heterogeneous noise, (2) we add two recent Kriging-based algorithms (TSSO and eTSSO), (3) we adapt the SKO algorithm (Huang et al., 2006) to settings with heterogeneous noise, and (4) we explicitly consider the *replication strategy* used by the different algorithms, i.e., the algorithms need to smartly allocate the limited replication budget to obtain reasonably good estimates. In Picheny et al. (2013b), the noise magnitude is known and the replication strategy is limited to simple revisits (i.e., each time the algorithms visit $\mathbf{x}^i$, $n^i = 1$ replication is performed). While Picheny et al. (2013b) compare the algorithms only in terms of how they perform the search for interesting alternatives in Θ , we also study the effectiveness and importance of the replication strategies and their influence on the choice of the final solution (i.e., the identification step, see Fig. 1, last block, and Section 2.1, last paragraph).

The complexity and computational requirements vary widely across the algorithms. As the MQ algorithm is the most straightforward, we consider it as a benchmark in our experiments. Some algorithms (i.e., CKG, SKO, and EQI, see Section 2) require information on the noise variance function $\tau^2(\mathbf{x})$ at any $\mathbf{x} \in \Theta$. This may limit the applicability of these algorithms in practice, as an unknown noise variance function is “a fact of life in system simulation problems” (Kim & Nelson, 2006), so it needs to be estimated. In our experiments, we test whether these “informed” algorithms outperform the ones that don’t need this information.

The following conclusions may serve as a useful guideline for researchers who want to apply Kriging-based algorithms to solve engineering or business problems, and may be useful in the development of future algorithms:

1. Not only the magnitude but also the structure (i.e., shape) of noise can have a large impact on the performance of the algorithms.

Table 1
Overview of notations.

Notation	Description
$f(\mathbf{x}^i)$	The true objective value at point $\mathbf{x}^i$, which is assumed to be unknown.
$\tilde{f}(\mathbf{x}^i)$	Estimate of the objective value at point $\mathbf{x}^i$: $\tilde{f}(\mathbf{x}^i) = \sum_{j=1}^n \tilde{f}_j(\mathbf{x}^i)/n^i$.
$\hat{f}(\mathbf{x}^i)$	Kriging estimate: Kriging prediction of the objective value at point $\mathbf{x}^i$.
$s^2(\mathbf{x}^i)$	Estimate of Kriging prediction variance at point $\mathbf{x}^i$.

2. Among the algorithms that require an estimate of $\tau^2(\mathbf{x})$, SKO and CKG usually outperform MQ while EQI performs worse.
3. Two-stage stochastic optimization (TSSO) simulates the most promising alternatives but usually fails to identify them at the end. If the identification is improved, TSSO will outperform MQ and will provide very competitive results, making it preferable over SKO or CKG since it does not require an estimate of the noise variance function. Also, eTSSO does not seem to provide better results than TSSO.
4. None of the existing algorithms have an appropriate replication strategy. While the focus of the literature has been primarily on designing better search methods, our results show that there is a need for smarter replication strategies (e.g., by adding an appropriate ranking and selection procedure). This can likely lead to significant improvements in the performance of all algorithms.

The remainder of this article is organized as follows. Section 2 provides a brief explanation of the Kriging-based algorithms studied, Section 3 details the experiments, and Section 4 contains the main results (additional results are provided in an appendix). We present conclusions in Section 5.

2. Overview of Kriging-based algorithms

Section 2.1 provides a brief introduction to Kriging-based simulation optimization, and explains the general steps involved in this process. Section 2.2 highlights the similarities and differences between the algorithms studied in this paper.

2.1. Kriging-based simulation optimization

As noted in Section 1, we need a stochastic simulation model to estimate the objective value at a point within Θ (see Eq. (1)). As the simulation output is noisy, we use the sample mean of the simulation output across several replications to obtain an estimate. This estimate is denoted by $\tilde{f}(\mathbf{x}^i)$ with estimated variance $\widehat{\text{Var}}[\tilde{f}(\mathbf{x}^i)]$ at any arbitrary point $\mathbf{x}^i$.

As shown in Fig. 1, the Kriging-based algorithms studied in this article are *sequential*. In the first step, they fit an initial Kriging metamodel using n_0 initial points. Space filling designs (e.g., Latin hypercube sampling) are often used for sampling these initial points (Kleijnen, 2009); simulation replications are used to estimate the objective value at these points. Since our observations are noisy, we use stochastic Kriging (Ankenman et al., 2010) to fit an initial Kriging metamodel. This metamodel provides us with (1) $\hat{f}(\mathbf{x})$: a prediction for the objective function value at any $\mathbf{x} \in \Theta$ and (2) $s^2(\mathbf{x})$: an estimate of the variance of $\hat{f}(\mathbf{x})$, which provides a measure for the Kriging prediction uncertainty or error (Kleijnen, 2014); see Appendix A for related expressions. Table 1 provides an overview of the most important notations in this paper. We refer to Quan et al. (2013), Ankenman et al. (2010), Yin et al. (2011), and Cressie (1993) for further details on stochastic Kriging.

In the *search step*, the algorithm iteratively chooses a point and simulates it for some number of replications to obtain $\tilde{f}(\mathbf{x})$ and $\widehat{\text{Var}}[\tilde{f}(\mathbf{x})]$; it then refits the Kriging metamodel, re-estimating the Kriging parameters with the new point included. The algorithm uses an *infill criterion* to choose these points (known as *infill*

points): the purpose of the infill criterion is to discover the good alternatives in Θ , and guide the search to regions with promising objective values (Forrester, Sobester, & Keane, 2008).

Optionally, the *replication step* allows us to perform extra replications at some of the points already sampled (either initial design points or infill points), to increase the precision of both $\tilde{f}(\mathbf{x})$ and the estimated variance in those points. Such a replication step is present in TSSO, eTSSO, and EQI (see Section 2.2), but the actual choice of points and the number of replications allocated to this step differ between these algorithms.

Each time we execute the search and replication steps, we consume part of our finite replication budget. Hence, after some iterations, the replication budget will end and we enter the *identification step*: the algorithm looks back at all simulated or tested (T) alternatives (Θ^T) and returns one of them as the best solution to the optimization problem. In noiseless simulation, this is trivial since $f(\mathbf{x})$ is available for all the simulated points: $\mathbf{x}^r = \arg \min_{\mathbf{x} \in \Theta^T} f(\mathbf{x})$, where $\mathbf{x}^r$ denotes the proposed solution. In the presence of noise, only $\tilde{f}(\mathbf{x})$ is available, along with $s^2(\mathbf{x})$ and the observed $\tilde{f}(\mathbf{x})$ at each alternative $\mathbf{x} \in \Theta^T$. Consequently, the success of an algorithm is impacted by the infill criterion (only alternatives that have been simulated are considered in the identification step), by the criterion used to choose $\mathbf{x}^r$ in the identification step (i.e., the *identification criterion*), and by the replication strategy (more precise estimates can help to distinguish superior alternatives in the identification step).

2.2. Description of the algorithms

Due to the limited budget, all algorithms must make a trade-off between the total number of infill points and their precision (Picheny et al., 2013a). If the replication strategy performs only a few replications at each point, we can sample many infill points but the estimates $\tilde{f}(\mathbf{x})$ may remain very noisy. This may result in inaccurate Kriging models, which in turn affects the search step (i.e., we might sample points with bad objective values). Additionally, it may cause problems in the identification step, as it becomes more difficult to differentiate between good and inferior solutions. Allocating more replications to the points mitigates these issues, but decreases the number of infill points, implying that we might not get the chance to discover (some) interesting alternatives.

Table 2 summarizes the main characteristics of the algorithms (infill criterion, replication strategy, identification criterion). As evident from this table, the identification criterion varies between the algorithms: yet, most of them rely on $\hat{f}(\mathbf{x})$ and $s^2(\mathbf{x})$ for proposing the final solution. The remainder of this section briefly explains the search and the replication strategy for each algorithm.

2.2.1. MQ algorithm

MQ is straightforward and acts as a benchmark for other algorithms: it simply chooses the point with minimum Kriging quantile, $q(\mathbf{x}) = \hat{f}(\mathbf{x}) + \Phi^{-1}(\beta)s(\mathbf{x})$ with $\beta \in (0, 0.5]$, as the infill point in each iteration k : $\mathbf{x}^k = \arg \min_{\mathbf{x} \in \Theta} q(\mathbf{x})$.

Obviously, MQ does not require information on $\tau^2(\mathbf{x})$. The algorithm allocates a fixed number B of replications per iteration: the

Table 2
Kriging-based optimization algorithms.

Algorithm	Source	Infill criterion	Replication strategy	Identification criterion	Info on $\tau^2(\mathbf{x})$ required?
1) MQ	Picheny et al. (2013b)	Min quantile	Revisits	$\min_{\mathbf{x} \in \Theta^k} q(\mathbf{x})$	No
2) SKO	Huang et al. (2006)	AEI	Revisits	$\min_{\mathbf{x} \in \Theta^k} q(\mathbf{x})$	Yes
3) CKG	Frazier et al. (2009)	CKG	Revisits	$\min_{\mathbf{x} \in \Theta^k} \hat{f}(\mathbf{x})$	Yes
4) EQI	Picheny et al. (2013a)	EQI	Replication step & Revisits	$\min_{\mathbf{x} \in \Theta^k} q(\mathbf{x})$	Yes
5) TSSO	Quan et al. (2013)	MEI	Replication step	$\min_{\mathbf{x} \in \Theta^k} \tilde{f}(\mathbf{x})$	No
6) eTSSO	Liu et al. (2014)	MEI	Replication step	$\min_{\mathbf{x} \in \Theta^k} \tilde{f}(\mathbf{x})$	No

Note: $q(\mathbf{x})$ is the Kriging quantile, $q(\mathbf{x}) = \hat{f}(\mathbf{x}) + \Phi^{-1}(\beta)s(\mathbf{x})$ where $\Phi^{-1}(\cdot)$ is the inverse CDF of standard normal distribution and $\beta \in (0, 0.5]$ in MQ and $\beta \in [0.5, 1]$ in SKO and EQI (Picheny & Ginsbourger, 2014).

analyst needs to decide on B prior to running the algorithm. However, MQ allows for revisits: at any iteration k , we might sample a point that has been simulated previously, and add B additional replications. The sampled points can thus receive a different total number of replications depending on how often they are revisited.

2.2.2. Adapted sequential Kriging optimization (SKO)

The SKO algorithm of Huang et al. (2006) chooses the alternative with maximum *augmented expected improvement* (AEI) as the next infill point:

$$AEI(\mathbf{x}) = E \left[\max \left(\hat{f}(\mathbf{x}^{**}) - \hat{f}(\mathbf{x}), 0 \right) \right] \left(1 - \frac{\tau(\mathbf{x})}{\sqrt{s^2(\mathbf{x}) + \tau^2(\mathbf{x})}} \right), \quad (2)$$

where $\hat{f}(\mathbf{x}^{**})$ is the Kriging prediction at the current *effective best solution* $\mathbf{x}^{**}$: i.e., the point with minimum $q(\mathbf{x})$ among the simulated points, with $\beta \in [0.5, 1]$. Thus, the first term of the above equation represents how much we expect the cost value at $\mathbf{x}$ to be lower than $\hat{f}(\mathbf{x}^{**})$, see Huang et al. (2006) for details. This algorithm was originally designed for homogeneous noise: in the second term of Eq. (2), we introduce $\tau^2(\mathbf{x})$ instead of τ^2 to reflect the presence of heterogeneous noise. SKO also uses a fixed number of replications per iteration (B) and allows for revisits. Note that the second term in Eq. (2) prevents the algorithm from getting trapped in a given point $\mathbf{x}$, as continued resampling causes $s^2(\mathbf{x})$ to gradually approach zero, which in turn pushes $AEI(\mathbf{x})$ towards zero (Huang et al., 2006).

2.2.3. Correlated knowledge-gradient (CKG)

The idea behind the correlated knowledge-gradient (CKG) infill criterion is that in noisy environments, the Kriging prediction $\hat{f}(\mathbf{x})$ may be closer to $f(\mathbf{x})$ than the sample mean $\bar{f}(\mathbf{x})$; therefore, points are selected based on their effect on the Kriging prediction (Frazier et al., 2009; Picheny & Ginsbourger, 2014). More specifically, at iteration $k+1$, the improvement that would result from sampling alternative $\mathbf{x}^{k+1}$ is defined as:

$$I(\mathbf{x}^{k+1}) = \min_{\mathbf{x} \in \Theta^k \cup \mathbf{x}^{k+1}} \hat{f}_k(\mathbf{x}) - \min_{\mathbf{x} \in \Theta^k \cup \mathbf{x}^{k+1}} \hat{f}_{k+1}(\mathbf{x}), \quad (3)$$

where Θ^k is the set of simulated points at the end of iteration k , $\hat{f}_k(\mathbf{x})$ is the current Kriging prediction and $\hat{f}_{k+1}(\mathbf{x})$ is the Kriging prediction after refitting the Kriging metamodel with $\mathbf{x}^{k+1}$ included. We choose the point with highest $CKG(\mathbf{x}) = E[I(\mathbf{x})]$ as the next point. As with SKO, we need to know $\tau^2(\mathbf{x})$ (or its estimate) for calculating this expectation (see Frazier et al., 2009, for further details on this calculation). Analogous to MQ and SKO, the CKG algorithm allows for revisits, and uses a fixed number of replications B per iteration.

2.2.4. Expected quantile improvement (EQI)

In EQI, at iteration $k+1$, simulating an alternative $\mathbf{x}^{k+1}$ is beneficial if its quantile using the updated Kriging model is lower than the current minimum quantile:

$$I(\mathbf{x}^{k+1}) = \left(\min_{\mathbf{x} \in \Theta^k} q_k(\mathbf{x}) - q_{k+1}(\mathbf{x}^{k+1}) \right)^+, \quad (4)$$

where $q_{k+1}(\mathbf{x}^{k+1})$ is the quantile of the updated Kriging metamodel at $\mathbf{x}^{k+1}$, with $\beta \in [0.5, 1]$. The next point is the point with maximum $EQI(\mathbf{x}) = E[I(\mathbf{x})]$; Picheny et al. (2013a) explain the calculation of this expectation. Again, this calculation requires an estimate of $\tau^2(\mathbf{x})$.

EQI allows for revisits, and additionally includes a replication step after each search step (see Fig. 1). The number of replications per iteration is no longer fixed to a constant B . In the search step of iteration $k+1$, EQI samples the infill point $\mathbf{x}^{k+1}$ with n_{inc} replications and refits the Kriging metamodel. In the replication step, it checks whether performing another n_{inc} replications at $\mathbf{x}^{k+1}$ is beneficial, by looking at the evolution of $EQI(\mathbf{x}^{k+1})$ (see Picheny et al., 2013a, for details). If so, the algorithm adds another n_{inc} replications, updates $\bar{f}(\mathbf{x}^{k+1})$ and $\widehat{\text{Var}}[\bar{f}(\mathbf{x}^{k+1})]$, and refits the Kriging metamodel. This is repeated until adding extra replications is no longer beneficial: at this moment EQI leaves the replication step.

2.2.5. Two-stage sequential optimization (TSSO)

Analogous to EQI, TSSO has a search and a replication step (called *allocation stage* in Quan et al., 2013). In the search step of iteration $k+1$, TSSO looks at all *unvisited* alternatives ($\mathbf{x} \in \Theta \setminus \Theta^k$) and simulates the one with maximum modified expected improvement (MEI):

$$MEI(\mathbf{x}) = E \left[\max \left(\hat{f}(\mathbf{x}_{min}) - f_p^*(\mathbf{x}), 0 \right) \right], \quad (5)$$

where $\hat{f}(\mathbf{x}_{min})$ is the stochastic Kriging prediction at $\mathbf{x}_{min} = \arg \min_{\mathbf{x} \in \Theta^k} \bar{f}(\mathbf{x})$ (i.e., the alternative with the lowest sample mean among the already simulated points) and $f_p^*(\mathbf{x})$ is a normal random variable: $f_p^*(\mathbf{x}) \sim N(\hat{f}(\mathbf{x}), s_D^2(\mathbf{x}))$ where the mean $\hat{f}(\mathbf{x})$ is the stochastic Kriging prediction at solution $\mathbf{x}$, and $s_D^2(\mathbf{x})$ is the estimate of *deterministic* Kriging prediction error. TSSO does not require $\tau^2(\mathbf{x})$ for calculating $MEI(\mathbf{x})$, see Quan et al. (2013) for details. Note that, as TSSO considers only unvisited alternatives in the search step, revisits are not allowed.

The number of replications per iteration is fixed to a predetermined value B . This number is shared between the search and replication steps according to a heuristic (Quan et al., 2013). In the replication step, the algorithm uses the optimal computing budget allocation (OCBA) which is a Bayesian ranking and selection procedure (Chen, Lin, Yücesan, & Chick, 2000). OCBA allocates most of the replication step budget to points with low $\bar{f}(\mathbf{x})$ and high $\widehat{\text{Var}}[\bar{f}(\mathbf{x})]$ (Quan et al., 2013), in view of maximizing the probability of selecting the best point among the sampled ones.

Table 3
Analytical test functions.

Function	Description
Six-hump Camelback	$f(x_1, x_2) = 4x_1^2 - 2.1x_1^4 + x_1^6/3 + x_1x_2 - 4x_2^2 + 4x_2^4$ $-2 \leq x_1 \leq 2$ and $-1 \leq x_2 \leq 1$ $x^* = (0.0977, -0.6973)$ $f(x^*) = -1.0294, R_f \approx 7.3$
Rescaled Branin	$f(x_1, x_2) = \frac{1}{51.95} \left[\left(\tilde{x}_2 - \frac{5.1\tilde{x}_1^2}{4\pi^2} + \frac{5\tilde{x}_1}{\pi} - 6 \right)^2 + \left(10 - \frac{10}{8\pi} \right) \cos(\tilde{x}_1) - 44.81 \right]$, $\tilde{x}_1 = 15x_1 - 5, \tilde{x}_2 = 15x_2$ $0 \leq x_1 \leq 1$ and $0 \leq x_2 \leq 1$ $x^* = (0.541, 0.1348)$ $f(x^*) = -1.0459, R_f \approx 6$
Hartmann-6	$f(x_1, \dots, x_6) = -\sum_{i=1}^6 c_i \exp \left[-\sum_{j=1}^6 \alpha_{ij} (x_j - p_{ij})^2 \right]$ c_i, α_{ij}, p_{ij} are parameters, see Dixon and Szegö (1978) $0 \leq x_i \leq 1$ for $i = 1, \dots, 6$ $x^* = (0.2382, 0.1391, 0.3665, 0.3286, 0.3519, 0.7018)$ $f(x^*) = -3.02, R_f \approx 3.3$

Note: R_f is the range of the response ($\max_{x \in \Theta} f(x) - \min_{x \in \Theta} f(x)$) in the region of interest.

2.2.6. Extended two-stage sequential optimization (eTSSO)

The eTSSO algorithm differs from the TSSO approach in the number of replications per iteration; while this was fixed to a given value B in TSSO, it now differs for each iteration k :

$$B_k = B_{k-1} \left(1 + \frac{\widehat{\text{Var}}[\tilde{f}(\mathbf{x}_{OCBA})]}{\widehat{\text{Var}}[\tilde{f}(\mathbf{x}_{OCBA})] + s_D^2(\mathbf{x}^k)} \right), \quad (6)$$

where $\widehat{\text{Var}}[\tilde{f}(\mathbf{x}_{OCBA})]$ is the estimate of the variance of the sample mean at the point where OCBA allocates the highest number of replications ($\mathbf{x}_{OCBA}$), and $s_D^2(\mathbf{x}^k)$ is the estimate of the deterministic Kriging prediction error at the point that maximizes MEI at iteration k . In eTSSO, the search step budget of each iteration is fixed to a value (n_s) prior to running the algorithm: at iteration k , the replication step budget of eTSSO is thus $B_k - n_s$.

3. Designing numerical tests

In this section, we explain our experiments in detail. We apply the Kriging-based algorithms to three well-known analytical test functions; [Section 3.1](#) provides the details. We also test our algorithms on the inventory problem of [Fu and Healy \(1997\)](#), see [Section 3.2](#). [Section 3.3](#) explains the method used for sampling the initial design and [Section 3.4](#) discusses the different scenarios tested in our experiments. We discuss the implementation of the algorithms in [Section 3.5](#). Finally, [Section 3.6](#) explains the performance measure used to compare the algorithms.

3.1. Analytical test functions and candidate points

The algorithms are applied to minimize three well-known analytical test functions (six-hump Camelback, rescaled Branin, and Hartmann-6; [Dixon & Szegö, 1978](#)), perturbed with heterogeneous Gaussian noise. Hence, we obtain noisy observations $\tilde{f}_j(\mathbf{x}^i) = f(\mathbf{x}^i) + \varepsilon_j(\mathbf{x}^i)$, with $\varepsilon_j(\mathbf{x}^i) \sim N(0, \tau(\mathbf{x}^i))$. [Table 3](#) gives further details on these analytical test functions. All test functions are continuous; however, as discussed in the introduction, we discretize the domain to obtain a large but finite set of solutions. In this way, we avoid the challenging problem of maximizing the expected improvement criteria in continuous space; also, some of the criteria might be more difficult to maximize than others. We discretize the domains using Faure low discrepancy sequences (a quasi-Monte Carlo sampling method with good space-filling properties, see chapter 5 of [Lemieux, 2009](#)) and we take 1000 points for the Camelback and Branin functions ($d = 2$), and 10000 points for the Hartmann-6 function ($d = 6$). These points represent our set of candidate points Θ ; our objective is to find the global minimum

among these alternatives (referred to as $\mathbf{x}^*$, having function value $f(\mathbf{x}^*)$) using our Kriging-based algorithms.

[Figs. 2 and 3](#) show an illustration of the rescaled Branin and the Camelback functions, along with their respective candidate points. As evident from the figures, the global minimum of Camelback lies within a small valley, while Branin has a large flat valley. Although the exact shape of Hartmann-6 is unknown, it is well-known that it is a multimodal function.

3.2. The (s, S) inventory problem and candidate points

In the inventory control problem of [Fu and Healy \(1997\)](#), a firm employs a periodic review (s, S) policy to manage the inventory of a single product. At the end of each period (e.g., each week), the firm checks the inventory. If the inventory position (i.e., inventory on-hand + on-order - backorder) is above the level s , the firm does not order. If it is below s , the firm orders the difference between the order-up-to level S and the inventory position. Demand is stochastic and continuous and is i.i.d across the periods. The goal is to determine the optimal s and S in view of minimizing the long-run expected cost per period. The costs include the fixed ordering (K), per-unit ordering (c), holding (h), and backorder (b) cost. Zero replenishment lead time and full backlogging are assumed.

To simulate this problem, we use a simulation runlength of 1000 periods per replication with a warm-up length of 100 periods. The inventory parameters (s, S) are continuous; we again use the Faure sequence to take 1000 space filling solutions ([Fig. 4](#)). The expected cost function is convex and with exponential demand, it has a closed form, such that the optimal s and S can be calculated analytically; we can thus evaluate the performance of our Kriging-based algorithms (see [Table 4](#) for details about the parameters used in our experiment; based on [Table 1](#) in [Fu & Healy, 1997](#)).

3.3. Initial design and Kriging metamodeling

For the initial design, we follow the suggestion of [Jones et al. \(1998\)](#): $n_0 = 10d$ (so 20 initial points for the Camelback and rescaled Branin functions, as well as for the (s, S) inventory problem; 60 initial design points for the Hartmann-6 function). A maximin Latin hypercube sample (LHS) is used to obtain these points (as is common in Kriging metamodeling, see [Picheny et al., 2013b](#)); we determine $\tilde{f}(\mathbf{x}^i)$ for each $i = 1, \dots, n_0$ using 55 replications (in view of controlling the noise, see [Section 3.4](#)). As prevalent in the literature, we consider only a constant trend (i.e., an intercept) in the Kriging metamodel ([Kleijnen, 2009](#)). We also don't use common random numbers (CRN): as shown by [Chen, Ankenman, and Nelson \(2012\)](#), CRN increases the variance of the Kriging predictor

Fig. 2. Rescaled Branin function and candidate points; an arrow points to $\mathbf{x}^*$.

Fig. 3. Six-hump Camelback function and candidate points; an arrow points to $\mathbf{x}^*$.

Fig. 4. Expected cost function of the (s, S) inventory problem with exponential(λ) demand; an arrow points to $\mathbf{x}^*$.

Table 4
 (s, S) inventory problem with exponential(λ) demand.

Description	
Periodic review (s, S) problem	$f(s, S) = \frac{c}{\lambda} + [K + h(s - \frac{1}{\lambda} + 0.5\lambda(S^2 - s^2)) + \frac{h+b}{\lambda}e^{-\lambda s}]/[1 + \lambda(S - s)]$ $K = b = 100, c = h = 1, \lambda = 0.0002$ $10000 \leq s \leq 22500$ and $22600 \leq S \leq 35000$ $\mathbf{x}^* = (22084.9609, 23060.1563)$ $f(\mathbf{x}^*) = 28165.0049, R_f \approx 8584$

Note: R_f is the range of the response $(\max_{\mathbf{x} \in \Theta} f(\mathbf{x}) - \min_{\mathbf{x} \in \Theta} f(\mathbf{x}))$ in the region of interest.

Table 5
External factors considered.

Factors	Levels
Replication budget	Low: 550 replications High: 2750 replications
Number of replications per iteration (for SKO, CKG, TSSO, and MQ)	$B = 55$
Noise magnitude for the 3 analytical test functions	Light noise $\min_{\mathbf{x} \in \Theta} \tau(\mathbf{x}) = 0.15R_f, \max_{\mathbf{x} \in \Theta} \tau(\mathbf{x}) = 0.6R_f$ Heavy noise $\min_{\mathbf{x} \in \Theta} \tau(\mathbf{x}) = 1.5R_f, \max_{\mathbf{x} \in \Theta} \tau(\mathbf{x}) = 6R_f$
Noise structure for the 3 analytical test functions	Best case $\min_{\mathbf{x} \in \Theta} \tau(\mathbf{x})$ at $\min_{\mathbf{x} \in \Theta} f(\mathbf{x})$ $\max_{\mathbf{x} \in \Theta} \tau(\mathbf{x})$ at $\max_{\mathbf{x} \in \Theta} f(\mathbf{x})$ Worst case $\min_{\mathbf{x} \in \Theta} \tau(\mathbf{x})$ at $\max_{\mathbf{x} \in \Theta} f(\mathbf{x})$ $\max_{\mathbf{x} \in \Theta} \tau(\mathbf{x})$ at $\min_{\mathbf{x} \in \Theta} f(\mathbf{x})$

Note: R_f is the range of the objective value in the region of interest.

and variance of the constant trend estimator. We fit the Kriging metamodel by means of maximum likelihood; Matern 2.5 is used as the covariance function (see Rasmussen and Williams, 2006, chapter 4).

3.4. Algorithmic factors

There are many factors that can influence the performance of the algorithms: some factors are *internal* or algorithm-specific parameters (e.g., β in SKO), while others are *external* (such as the replication budget, the noise magnitude and noise structure). Table 5 provides an overview which we now detail.

We test our algorithms using low and high replication budgets; this refers to the replication budget available *after* simulating the n_0 initial points, for performing the search and replication steps. We focus on problems for which a very small number of solution evaluations are possible, e.g., because the simulation is very time consuming. The number of replications per iteration (B) is equal to 55 for SKO, CKG, TSSO, and MQ, implying that with low (resp. high) budget we simulate 10 (resp. 50) infill points for these algorithms. In TSSO, the number of replications per infill point is usually less than 55, since 55 replications are shared between the search and replication steps (see Section 2.2.5). In EQI and eTSSO, B is not fixed; the number of infill points is thus not known prior to running these algorithms.

The maximum and minimum values of $\tau(\mathbf{x})$ are linked to R_f , i.e., the range of the objective value in the region of interest (as in Huang et al., 2006 and Picheny & Ginsbourger, 2014). With light noise, $\tau(\mathbf{x})$ varies between 15% and 60% of R_f . With heavy noise, these numbers increase to 150% and 600% of R_f . We sample the initial designs with 55 replications: for these points the noise on $\hat{f}(\mathbf{x})$ varies between $1.5/\sqrt{55}R_f = 0.2R_f$ and $6/\sqrt{55}R_f = 0.8R_f$ with heavy noise.

In real-life problems, the noise can follow any type of structure. In the (s,S) inventory problem this structure is unknown. For the 3 analytical test functions, we assume that the standard deviation of the noise (i.e., $\tau(\mathbf{x})$) varies linearly in terms of the objective value: $\tau(\mathbf{x}) = af(\mathbf{x}) + a \times b$ (see Appendix B for details). Here, we focus on two cases:

1. *Best case*: the noise standard deviation decreases linearly as the objective value decreases; therefore we have minimum noise at the global minimum.
2. *Worst case*: the noise standard deviation increases linearly as the objective value decreases; we then have maximum noise at the global minimum.

As mentioned earlier, CKG, EQI, and SKO need an estimate of $\tau(\mathbf{x})$ at all the candidate points. In the analytical test functions,

Table 6
Parameter settings for different algorithms.

Algorithm	Factors
MQ	$\beta = 0.1$
SKO	$\beta = 0.84$
EQI	$\beta = 0.5, n_{inc} = 10, \gamma = 0.5$
TSSO	$n_{min} = 2$
eTSSO	$n_s = 10$

Note: CKG does not require any parameters.

$\tau(\mathbf{x})$ is known: as mentioned above, we assume a linear relation between $\tau(\mathbf{x})$ and the function value. We then estimate $\tau(\mathbf{x})$ at $\mathbf{x} \in \Theta$ by using the Kriging prediction $\hat{f}(\mathbf{x})$ instead of the unknown true function value: $\hat{\tau}(\mathbf{x}) = a\hat{f}(\mathbf{x}) + a \times b$.

In practical implementations, $\tau(\mathbf{x})$ will be unknown. Therefore, for the analytical test functions, the noise estimates offered to CKG, EQI, and SKO are much better than the estimates that one would likely obtain in practice, where no prior information on $\tau(\mathbf{x})$ is available. This is the case in the inventory test problem. Analogous to Section 3.1 in Ankenman et al. (2010), we then fit a *deterministic* Kriging model to the sampled variances of the simulated function values and use it to predict $\tau(\mathbf{x})$ at any $\mathbf{x} \in \Theta$.

Fig. 5 illustrates the structure of heavy worst case and heavy best case noise for the Branin function, at $x_2 = 0$. Here $R_f = 6$, so the maximum and minimum standard deviation of noise with 55 replications are $0.8 \times 6 = 4.8$ and $0.2 \times 6 = 1.2$, respectively. As shown in the figure, with best case noise, this maximum occurs at $(x_1 = 0, x_2 = 0)$, and the noise drops linearly as $f(\mathbf{x})$ decreases. With worst case noise, we have the opposite behavior.

The noise structure may substantially impact the performance of the algorithms, as it impacts the search pattern in step 2. As an illustration, Fig. 6 shows the structure of MEI for the Camel-back function, with best and worst case noise, at the first iteration of an arbitrary run (so only the n_0 design points have been simulated). With best case noise, the maximum MEI is near the optimal solution; with worst case noise, the search may be driven to non-interesting areas of the search space. The noise structure can thus completely change the search behavior of the algorithms.

Table 6 provides the parameter settings of the different algorithms. Huang et al. (2006) recommend $\beta = 0.84$ for SKO ($\Phi^{-1}(0.84) = 1$); for the β values in MQ and EQI we follow the suggestions of Picheny et al. (2013b). We follow the recommendation of Picheny et al. (2013a) for γ in EQI; we set $n_{inc} = 10$. In TSSO, n_{min} (i.e., the minimum number of replications for simulating an infill point) only affects the way B is shared between the search and replication steps (see Quan et al., 2013). In eTSSO, we chose $n_s = 10$ to match n_{inc} of EQI.

Fig. 5. Heavy worst case and best case noise for rescaled Branin function at $x_2 = 0$ with 55 replications.

(a) Camelback function, the location of the optimum is indicated in red.

(b) MEI for best case heavy noise

(c) MEI for worst case heavy noise

Fig. 6. MEI of Camelback function for the first iteration of a given run, with best and worst case noise. The dark red regions of the figures have the highest values for MEI, and are most likely to be sampled in the next step.

3.5. Implementation details

For best case as well as for worst case noise, we have the following 4 settings: 1) Low budget – Light noise 2) Low budget – Heavy noise 3) High budget – Light noise 4) High budget – Heavy noise. These are applied to each of the 3 analytical test functions,

yielding a total of $2 \times 4 \times 3 = 24$ scenarios. For the inventory problem, we have 2 scenarios: low budget and high budget. Recall that the budget is for the search and replication steps (steps 2 and 3 in Fig. 1). To evaluate the performance of an algorithm for a given scenario, we use 100 macroreplications: we run the algorithm 100 times, each time using a different initial Latin hypercube design. In

Table 7

Average running time of different algorithms for one macroreplication (high budget) in seconds.

Algorithm	(s, S) problem	Camelback	Branin	Hartmann-6
MQ	17	20	19	61
SKO	31	24	31	91
CKG	385	384	379	88734
EQI	133	109	73	256
TSSO	26	24	27	121
eTSSO	4	6	4	17

a given macroreplication, all algorithms start with the same set of initial points (same $\mathbf{x}^i$, $\hat{f}(\mathbf{x}^i)$, and $\text{Var}[\hat{f}(\mathbf{x}^i)]$) and, hence, the same Kriging model (analogous to Picheny et al., 2013b).

All the algorithms were coded in MATLAB. To fit the stochastic Kriging models, we used the code provided on <http://stochastickriging.net/>, for CKG we used the code on <http://people.orie.cornell.edu/pfrazier/src.html>, the TSSO code was based on <http://www.ise.nus.edu.sg/staff/ngsh/download/matlabdocs/SOK/> and the eTSSO code was obtained from the authors of Liu et al. (2014). The remaining algorithms were coded by the authors themselves and compared with the DiceOptim package in R (Picheny & Ginsbourger, 2014) to ensure correctness.

As most of the algorithms allow for revisits, the covariance matrix of the metamodel may become ill-conditioned; we follow the guidelines in Section 4.3 of Picheny and Ginsbourger (2014) for handling this issue. Table 7 shows the average running time of different algorithms for one macroreplication in all the test problems (on a Dell desktop with 3.4 gigahertz Core i7-2600 processor and 8 gigabytes of RAM). For the analytical test functions, we focused on heavy worst case noise; the structure or magnitude of noise had only minor effects on the running times. As evident from the table, CKG and then EQI have the longest running time, while eTSSO requires the smallest computation time followed by MQ. In eTSSO, the number of replications per iteration quickly increases (see Eq. (6)) and the budget is depleted after a relatively small number of iterations: this explains its short running time. CKG requires a lot of time to calculate the expected improvement of all 10000 feasible alternatives in Hartmann-6 and thus, as shown in the table, takes around 24 hours per macroreplication. For the experiments in Section 4, we implemented this algorithm on a server with 5 cores and 100 threads, running each macroreplication on a different thread. However, in many problems, the feasible region contains millions of points (Xu, 2012); the computational requirements may limit the applicability of CKG in these cases.

3.6. Performance measure

Let $\mathbf{x}_m^r$ be the point returned by the algorithm in the identification step at macroreplication $m = 1, 2, \dots, 100$. $GAP_m = f(\mathbf{x}_m^r) - f(\mathbf{x}^*)$ then represents the difference in objective value between the alternative suggested by the algorithm and the optimal solution. The distribution of $GAP = (GAP_1, GAP_2, \dots, GAP_{100})$ is our main performance measure, visualized by means of a boxplot for each of the algorithms.

4. Results

In this section, we provide the results of our analysis. As outlined above, we study the following research questions:

1. Does the structure and magnitude of the noise and the replication budget affect the performance of the algorithms?
2. Do the algorithms that require an estimate of the noise variance function in step 2 (CKG, SKO, EQI) outperform the others that don't need this information?

3. How do the algorithms perform in comparison to our benchmark algorithm (MQ)?
4. How do the algorithms compare in terms of their search strategies? Do they manage to locate good solutions among Θ ?
5. Are the replication strategies of the algorithms effective? Do they help the algorithms to distinguish the superior solutions in the identification step, or do the algorithms experience the identification problem?

In Section 4.1, we discuss the effect of the external factors proposed in Table 5. Section 4.2 details the performance of the algorithms with respect to our performance measure, *GAP*. In Section 4.3 we discuss the effect of the identification step on the performance of the algorithms. As the results were similar for the Camelback, Branin, and the (s, S) problem, we moved the results of the two latter ones to Appendix C and Appendix D, respectively; only the results for Camelback (Figs. 7, and 9) and Hartmann-6 (Figs. 8, and 10) are discussed in the text.

4.1. Effect of external factors

The effect of the external factors on the performance of the algorithms is largely the same for the different test problems. From Figs. 7 and 8, we can derive the following effects:

- *Replication budget*: comparing low and high budget figures reveals that a higher replication budget improves the *GAP* distribution for all algorithms; for the analytical test functions, though, the difference is smaller with heavy noise: simulating more infill points does not significantly improve *GAP* with heavy noise.
- *Noise structure (for the analytical test functions)*: Figs. 7 and 8 show that the performance of all algorithms usually degrades with worst case noise comparing to best case noise, for all scenarios.
- *Noise magnitude (for the analytical test functions)*: comparing light and heavy noise figures, we see that stronger noise has a substantial negative impact on the performance of all algorithms. Moreover, noise magnitude has a stronger influence on the performance of the algorithms than noise structure.

4.2. Comparison with respect to *GAP*

Figs. 7 and 8 show that there is no algorithm that clearly provides a smaller *GAP* than other algorithms in all the tested scenarios. The benchmark algorithm MQ is very competitive in Camelback, Branin, and the (s, S) problem: only CKG and SKO tend to outperform MQ. TSSO and eTSSO give the largest *GAP* on average (especially with best case noise). The results for Hartmann-6 are somewhat different: as evident from Fig. 8, MQ is less competitive (especially with light noise); together with EQI, it usually gives the largest *GAP*. TSSO usually provides the best outcome (except with heavy best case noise), and is similar in performance to CKG.

The above observations show that, among the algorithms that require $\tau^2(\mathbf{x})$ in step 2, CKG and SKO usually outperform the benchmark MQ algorithm, while EQI performs worse. This holds even in the (s, S) inventory problem where $\tau^2(\mathbf{x})$ is estimated via a metamodel, see Appendix D. Analogous to MQ, the TSSO and eTSSO algorithms don't require an estimate for $\tau^2(\mathbf{x})$ in step 2: Figs. 7 and 8 show that TSSO usually provides a better *GAP* than eTSSO (the same holds for the Branin function and the inventory problem, see Figs. C.1 and D.1).

4.3. Effect of the identification step

In the previous section, we mentioned that TSSO usually provides a bad *GAP*: in comparison to other algorithms, the final

Fig. 7. GAP for the Camelback function ($R_f = 7.3$). At each scenario, the boxplots show the distribution of GAP for each distinct algorithm. The median is indicated by a solid circle; the edges of the boxes are the 25th and 75th percentiles.

solutions returned by this algorithm usually have higher objective values. This bad GAP performance can have two reasons: (1) the algorithm fails to discover promising alternatives in the search step, so Θ^T (the set of sampled solutions) does not contain interesting points (i.e., *search problem*); or (2) promising alternatives are present in Θ^T , but the algorithm fails to identify them in the

identification step; it thus returns an inferior solution to the user (i.e., *identification problem*).

For the Camelback function, Fig. 9 shows the GAP performance of algorithms if we had achieved *perfect* identification in the last step, i.e., if there was no identification problem. Algorithms then return the true best point among Θ^T as the final solution (the

Fig. 8. GAP for the Hartmann-6 function ($R_f = 3.3$). At each scenario, the boxplots show the distribution of GAP for each algorithm. The median is indicated by a solid circle; the edges of the boxes are the 25th and 75th percentiles.

Fig. 9. GAP results for the Camelback function with perfect identification ($R_f = 7.3$). At each scenario, the boxplots show the distribution of GAP for each algorithm when the true best point among the sampled points is returned as the final solution. The median is indicated by a solid star. The edges of the boxes are the 25th and 75th percentiles.

search strategies and thus the sampled solutions are unchanged). Comparing to Fig. 7, TSSO is performing much better and as evident from Fig. 9, in the absence of the identification problem, TSSO provides very good performance in all scenarios. Therefore, the bad GAP performance of TSSO in Fig. 7 is caused by the identification problem rather than the search problem; Fig. 9 shows that TSSO

actually has a relatively good search strategy as it usually discovers better alternatives than other algorithms. It thus seems that the replication strategy of TSSO is not very successful: the algorithm usually can't distinguish between the inferior and good solutions in the identification step (the same observations hold for the Branin function, Figs. C.1 and C.2; and the inventory problem, Fig. D.1). For

Fig. 10. GAP results for Hartmann-6 function with perfect identification ($R_f = 3.3$). At each scenario, the boxplots show the distribution of GAP for each algorithm when the true best point among the sampled points is returned as the final solution. The median is indicated by a solid star. The edges of the boxes are the 25th and 75th percentiles.

Hartmann-6, most of the algorithms already achieve perfect identification using the identification criteria in Table 2 (i.e., the GAP boxplots in Figs. 8 and 10 are identical in most cases).

The replication strategy of the algorithms also affects their performance in the search step. Appendix F shows the number of *distinct* infill points (i.e., without counting the revisits) sampled by the algorithms, for the high budget scenarios. As evident from these figures, TSSO always samples 50 distinct infill points as it does not allow for revisits. All other algorithms sample fewer distinct points. This can also contribute to the relatively good search performance of TSSO. EQI and eTSSO usually yield the lowest number of distinct infill points. Especially with heavy noise, EQI becomes very “conservative”, and spends a lot of replications to get more precise estimates for the objective value: the replication budget is depleted after sampling a few infill points. In eTSSO, the number of replications per iteration quickly increases (see Eq. (6)), implying that the budget is also depleted after a relatively small number of iterations.

The above observations lead us to conjecture that if the identification problem of TSSO could be solved, this algorithm would probably provide much better GAP results, and might even be preferred over CKG or SKO (as it does not require $\tau^2(\mathbf{x})$ in step 2). This could be done by either changing the replication strategy used by the algorithm (for instance, using other ranking and selection procedures instead of OCBA), or by modifying the identification criterion. For the Branin and Camelback functions, Figs. C.3 and E.1 in Appendix show, for instance, the change in performance when the identification criterion of TSSO is changed to choose the point with minimum Kriging prediction (as in CKG) instead of the point with minimum sample mean. This modified TSSO (MTSSO) almost always gives a better GAP performance than TSSO, and has a performance similar to MQ; the identification criterion thus has a significant effect on the outcome of Kriging-based algorithms.

Although it is more acute in TSSO, *all* algorithms have identification issues in the inventory problem (Fig. D.1), and the Camelback and Branin functions (compare Fig. 7 vs. 9 and Figs. C.1 vs. C.2); in the latter two functions, the impact on the resulting GAP distribution is more severe with heavy noise, as higher noise also has a negative impact on the search step. The replication strategies of the algorithms are thus, in general, not helpful. While increasing B (replication budget per iteration) might, at first sight, seem the obvious way to mitigate this problem, it is in general not easy to find the right B in practical problems where the optimal solution is unknown. Moreover, a higher B will cause us to sample less infill points. While some algorithms (i.e., eTSSO and EQI) determine B dynamically during the run of the algorithm, our results reveal that these are not very successful.

We thus believe that it is essential to focus on smarter replication strategies. One could include, for instance, a R&S procedure that acts as a “clean-up” in the identification step (see, e.g., the procedure proposed in Boesel, Nelson, & Kim, 2003); the aim of this R&S procedure is to discard the bad alternatives in Θ^T and find a sufficiently good one using as few replications as possible. Alternatively, R&S could be used in the optional replication step (see Fig. 1); Hong and Nelson (2007), for instance, propose a R&S procedure to identify the best alternative among the sampled ones in each iteration. In addition to mitigating the identification problem, this approach might also enable the algorithm to choose better infill points in the search step. While the OCBA technique (Chen et al., 2000) used in TSSO tries to achieve a similar goal, our results indicate that it is not very effective. Although perfect identification (as in Figs. 9 and 10) is unlikely to result from any given R&S method, these figures illustrate the best possible improvement in GAP that we could hope to obtain.

Evidently, R&S consumes replication budget. Note, however, that in our case, all algorithms provide better GAP performance in low

budget scenarios with perfect identification than in their corresponding high budget scenarios without perfect identification (e.g., compare Figs. 7 and 9 for the Camelback function; the same observations hold for the Branin function and the inventory problem). Therefore, reducing the budget of the search step to allow for a suitable R&S could be beneficial. In Hartmann-6, R&S is less useful since the criteria employed by the algorithms usually achieve good identification. Only with heavy noise, R&S might still be beneficial.

Further research is needed on how to combine Kriging-based algorithms with R&S to achieve an appropriate replication strategy. In addition, deciding on how much replication budget to allocate to the R&S step is not straightforward. For instance, when using R&S in a clean-up phase, we should decide on when to stop simulating infill points and move to the R&S stage.

5. Conclusions and future research

In this paper, we have compared six recent Kriging-based algorithms for continuous optimization via simulation with heterogeneous noise. Using Kriging-based algorithms for optimizing systems modeled through stochastic simulation (especially with heterogeneous noise) is relatively new, and provides an active research area: we think that the insights obtained in this article can be useful in improving these algorithms, and give useful suggestions for researchers who wish to employ them for solving engineering/business problems.

Our results showed that not only the magnitude, but also the structure of the noise can affect the performance of the algorithms. More specifically, we observed that with worst case noise (highest noise in the global minimum region) all algorithms performed worse than with best case noise (lowest noise in the global minimum region). We also showed that increasing the replication budget could improve the performance of the algorithms, especially when the noise magnitude was low.

CKG, SKO, and EQI require an estimate for the noise variance function but in practice, finding a reasonable estimate can be challenging. Among these algorithms, CKG and SKO tend to outperform our benchmark algorithm (MQ), while EQI does not. CKG is computationally very intensive, especially when the number of candidate points is high (e.g., 10000 feasible alternatives in the Hartmann-6 function). As, in some optimization problems, the feasible space includes millions of solutions, applying CKG to these problems might be impractical. In these cases, one can use the approximate knowledge gradient policy (Scott, Frazier, & Powell, 2011), which employs an approximation of the expected improvement function to speed up the algorithm. As shown by Scott et al. (2011), this policy performs similarly to SKO, at least in a setting with homogeneous noise.

Analogous to MQ, TSSO and eTSSO do not require information about the noise variance function. We did not observe a big difference between TSSO and eTSSO; TSSO usually even gave better results. The quality of the solutions returned by TSSO was strongly affected by its inability to identify good solutions among the simulated ones in the identification step: if it were successful in identifying the best alternatives, it would give significantly better results than MQ, and would perform as well as CKG or SKO. However, it would likely be preferred over CKG or SKO, since it does not require information with regard to the noise variance function.

Our results indicate that, *all* algorithms suffer from the identification problem; the resulting GAP distributions are worse for settings with heavy noise, as a higher noise also negatively impacts the search step. Consequently, it is essential to combine the algorithms with smarter replication strategies (e.g., using appropriate ranking and selection methods). In our opinion, this is a key area for future research.

Acknowledgments

This research was supported by the [Research Foundation-Flanders \(FWO\)](#) (Grant no. [G.0822.12](#)). We would like to thank Professor Szu Hui Ng, Department of Industrial & Systems Engineering, National University of Singapore, for sharing eTSSO code with us. Some of the computational resources and services used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Hercules Foundation and the Flemish Government – department EWI. We thank the two anonymous reviewers for their constructive comments, that substantially improved the paper.

Appendix A. Kriging expressions

In this appendix we provide formulas for the stochastic Kriging prediction $\hat{f}(\mathbf{x})$ and its estimated variance $s^2(\mathbf{x})$. Stochastic Kriging with a constant trend assumes that the unknown true function can be represented by $f(\mathbf{x}) = \beta + M(\mathbf{x})$, where $\beta \in \mathbb{R}$ is an unknown constant term and M is a realization of a mean zero Gaussian random field ([Ankenman et al., 2010](#)). We assume *spatial correlation*: $M(\mathbf{x}^h)$ and $M(\mathbf{x}^l)$ tend to be close if the distance $\|\mathbf{x}^h - \mathbf{x}^l\|$ is small. As this distance goes to infinity, $\text{Corr}[M(\mathbf{x}^h), M(\mathbf{x}^l)] \rightarrow 0$ and if the distance is zero, $\text{Corr}[M(\mathbf{x}^h), M(\mathbf{x}^l)] = 1$ (the implied covariance is denoted by $\text{Cov}[M(\mathbf{x}^h), M(\mathbf{x}^l)]$). The way the spatial correlation changes with distance $\|\mathbf{x}^h - \mathbf{x}^l\|$ depends on the employed correlation function (As noted in [Section 3.3](#), we used Matern 2.5 in our experiments). The observed goal value in the j th simulation replication at design point $\mathbf{x}^k$, can thus be written as:

$$\tilde{f}_j(\mathbf{x}^k) = \beta + M(\mathbf{x}^k) + \varepsilon_j(\mathbf{x}^k), \quad (7)$$

where, as mentioned in [Section 1](#), $\varepsilon_j(\mathbf{x}^k)$ has mean zero and variance $\tau^2(\mathbf{x}^k)$ and represents the noise inherent in a stochastic simulation model. This noise is assumed to be independent of M , and independent and identically distributed across replications. Since we don't consider common random numbers (CRN), $\text{Corr}[\varepsilon_j(\mathbf{x}^h), \varepsilon_j(\mathbf{x}^l)] = 0$.

Assume we fitted a Kriging metamodel using k design points $(\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^k)$ with estimated function values $\hat{\mathbf{f}} = (\hat{f}(\mathbf{x}^1), \hat{f}(\mathbf{x}^2), \dots, \hat{f}(\mathbf{x}^k))^T$. These estimates are sample means of the simulation output across $(n^1, n^2, \dots, n^k)$ replications. If the Kriging metamodel only has an intercept term (β), we obtain the Kriging prediction at an arbitrary point $\mathbf{x}^i$ as follows ([Ankenman et al., 2010](#)):

$$\hat{f}(\mathbf{x}^i) = \beta + \Sigma_M(\mathbf{x}^i, \cdot)^T [\Sigma_M + \Sigma_\varepsilon]^{-1} (\hat{\mathbf{f}} - \beta \mathbf{1}_k), \quad (8)$$

where $\mathbf{1}_k$ is a $k \times 1$ vector of ones, $\Sigma_M(\mathbf{x}^i, \cdot) = (\text{Cov}[M(\mathbf{x}^i), M(\mathbf{x}^1)], \dots, \text{Cov}[M(\mathbf{x}^i), M(\mathbf{x}^k)])^T$, and Σ_M is the spatial variance covariance matrix of the k design points (with size $k \times k$). Σ_ε is a $k \times k$ variance covariance matrix implied by the sample average simulation noise. Since we don't consider CRN, this matrix is diagonal: $\Sigma_\varepsilon = \text{Diag}\{\tau^2(\mathbf{x}^1)/n^1, \dots, \tau^2(\mathbf{x}^k)/n^k\}$.

To be able to use [Eq. \(8\)](#), we need to estimate β , Σ_M , and Σ_ε (the estimates are denoted by $\hat{\beta}$, $\hat{\Sigma}_M$, and $\hat{\Sigma}_\varepsilon$). For instance, since the noise variance function $\tau^2(\mathbf{x})$ is usually unknown in practice ([Kim & Nelson, 2006](#)), we use $\hat{\Sigma}_\varepsilon = \text{Diag}\{\widehat{\text{Var}}[\tilde{f}(\mathbf{x}^1)], \dots, \widehat{\text{Var}}[\tilde{f}(\mathbf{x}^k)]\}$. See [Ankenman et al. \(2010\)](#) for estimation of the other two parameters.

The mean squared error (MSE) of the Kriging predictor in [Eq. \(8\)](#) is also known as the *Kriging variance* and is estimated as

follows ([Chen & Kim, 2014](#)):

$$\widehat{\text{MSE}}(\hat{f}(\mathbf{x}^i)) = s^2(\mathbf{x}^i) = \hat{\Sigma}_M(\mathbf{x}^i, \mathbf{x}^i) - \hat{\Sigma}_M(\mathbf{x}^i, \cdot)^T [\hat{\Sigma}_M + \hat{\Sigma}_\varepsilon]^{-1} \hat{\Sigma}_M(\mathbf{x}^i, \cdot) + \frac{\hat{\delta}^T \hat{\delta}}{\mathbf{1}_k^T [\hat{\Sigma}_M + \hat{\Sigma}_\varepsilon]^{-1} \mathbf{1}_k}, \quad (9)$$

where $\hat{\delta} = \mathbf{1} - \mathbf{1}_k^T [\hat{\Sigma}_M + \hat{\Sigma}_\varepsilon]^{-1} \hat{\Sigma}_M(\mathbf{x}^i, \cdot)$.

Appendix B. Relation between $\tau(\mathbf{x})$ and $f(\mathbf{x})$ in different scenarios

For the 3 analytical test functions, we assume a linear relationship between the function value and the standard deviation of noise (i.e., noise magnitude). For best case noise, we have:

$$\begin{aligned} (\max_{\mathbf{x} \in \Theta} f(\mathbf{x}) + b) \times a &= \max_{\mathbf{x} \in \Theta} \tau(\mathbf{x}) = 0.6R_f \text{ or } 6R_f \\ &\quad \text{(light noise or heavy noise)} \\ (\min_{\mathbf{x} \in \Theta} f(\mathbf{x}) + b) \times a &= \min_{\mathbf{x} \in \Theta} \tau(\mathbf{x}) = 0.15R_f \text{ or } 1.5R_f \\ &\quad \text{(light noise or heavy noise),} \end{aligned} \quad (10)$$

where $\tau(\mathbf{x})$ is the standard deviation of the noise on $\tilde{f}(\mathbf{x})$ with one replication. From the above equations we can easily obtain a and b (e.g., for Camelback best case low noise: $(6.27 + b) \times a = 0.6 * 7.3$ and $(-1.03 + b) \times a = 0.15 * 7.3 \rightarrow a = 0.45, b = 3.46$; see [Table 8](#) for other scenarios). For worst case noise, we have a similar relation:

$$\begin{aligned} (\min_{\mathbf{x} \in \Theta} f(\mathbf{x}) + b) \times a &= \max_{\mathbf{x} \in \Theta} \tau(\mathbf{x}) = 0.6R_f \text{ or } 6R_f \\ &\quad \text{(light noise or heavy noise)} \\ (\max_{\mathbf{x} \in \Theta} f(\mathbf{x}) + b) \times a &= \min_{\mathbf{x} \in \Theta} \tau(\mathbf{x}) = 0.15R_f \text{ or } 1.5R_f \\ &\quad \text{(light noise or heavy noise)} \end{aligned} \quad (11)$$

Recall that SKO, CKG, and EQI need an estimate of $\tau(\mathbf{x})$: we use $\hat{\tau}(\mathbf{x}) = (\hat{f}(\mathbf{x}) + b) \times a$, as the goal function $f(\mathbf{x})$ is assumed to be unknown. For the inventory test problem, we use a metamodel to estimate the noise, see [Section 3.4](#).

Table 8

Parameters for the relation between $f(\mathbf{x})$ and $\tau(\mathbf{x})$.

		Camelback	Rescaled Branin	Hartmann-6
Best case	Low noise	$a = 0.45,$ $b = 3.46$	$a = 0.45,$ $b = 3.05$	$a = 0.45,$ $b = 4.12$
	High noise	$a = 4.5,$ $b = 3.46$	$a = 4.5,$ $b = 3.05$	$a = 4.5,$ $b = 4.12$
Worst case	Low noise	$a = -0.45,$ $b = -8.704$	$a = -0.45,$ $b = -6.95$	$a = -0.45,$ $b = -1.38$
	High noise	$a = -4.5, b =$ -8.704	$a = -4.5,$ $b = -6.95$	$a = -4.5,$ $b = -1.38$

Appendix C. Results for Branin function

Fig. C1. GAP for the Branin function ($R_f = 6$). At each scenario, the boxplots show the distribution of GAP for each algorithm. The median is indicated by a solid circle; the edges of the boxes are the 25th and 75th percentiles. The results are analogous to the Camelback function: CKG and SKO perform somewhat better than the benchmark algorithm (MQ) while other algorithms don't. TSSO and eTSSO usually have the largest GAP.

Fig. C2. GAP results for the Branin function with perfect identification ($R_f = 6$). At each scenario, the boxplots show the distribution of GAP for each algorithm when the true best point among the sampled points is returned as the final solution. The median is indicated by a solid star. The edges of the boxes are the 25th and 75th percentiles. Analogous to the Camelback function, we see that all algorithms suffer from the identification problem (especially TSSO).

Fig. C3. GAP for the Branin function with MTSSO included ($R_f = 6$). At each scenario, the boxplots show the distribution of GAP for each algorithm. The median is indicated by a solid circle; the edges of the boxes are the 25th and 75th percentiles. MTSSO usually provides smaller GAP than TSSO.

Appendix D. Results for the (s, S) inventory problem

Fig. D1. *GAP* and *GAP* with perfect identification for the (s, S) inventory problem ($R_f = 8584$). The boxplots show the distribution of *GAP* for each algorithm. The median is indicated by a solid circle for *GAP* and by a solid star for *GAP* with perfect identification; the edges of the boxes are the 25th and 75th percentiles. The results are analogous to the Camelback function: CKG and SKO perform somewhat better than the benchmark algorithm (MQ) in terms of *GAP* while other algorithms usually don't (EQI gives better results only with high budget). TSSO and eTSSO usually have the largest *GAP*. Both with high and low budget, the performance of all algorithms suffers from the identification problem (especially TSSO).

Appendix E. MTSSO results for Camelback function

Fig. E1. GAP for the Camelback function with MTSSO included ($R_f = 7.3$). At each scenario, the boxplots show the distribution of GAP for each algorithm. The median is indicated by a solid circle; the edges of the boxes are the 25th and 75th percentiles.

Appendix F. Number of distinct points sampled

Fig. F1. Number of distinct points sampled for the Branin function (High budget). At each scenario, the boxplots show the distribution of the number of distinct points sampled by each algorithm. The median is indicated by a solid circle; the edges of the boxes are the 25th and 75th percentiles.

Fig. F2. Number of distinct points sampled for the Camelback function (High budget). At each scenario, the boxplots show the distribution of the number of distinct points sampled by each algorithm. The median is indicated by a solid circle; the edges of the boxes are the 25th and 75th percentiles.

Fig. F3. Number of distinct points sampled for the Hartmann-6 function (High budget). At each scenario, the boxplots show the distribution of the number of distinct points sampled by each algorithm. The median is indicated by a solid circle; the edges of the boxes are the 25th and 75th percentiles.

Fig. F4. Number of distinct points sampled for the (s, S) inventory problem (High budget). At each scenario, the boxplots show the distribution of the number of distinct points sampled by each algorithm. The median is indicated by a solid circle; the edges of the boxes are the 25th and 75th percentiles.

References

- Ankenman, B., Nelson, B. L., & Staum, J. (2010). Stochastic kriging for simulation metamodeling. *Operations Research*, 58, 371–382.
- Auer, P., Cesa-Bianchi, N., & Fischer, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2–3), 235–256.
- Boesel, J., Nelson, B. L., & Kim, S.-H. (2003). Using ranking and selection to “clean up” after simulation optimization. *Operations Research*, 51(5), 814–825.
- Brochu, E., Cora, V. M., & De Freitas, N. (2010). A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. arXiv preprint arXiv:1012.2599.
- Chen, C.-H., Lin, J., Yücesan, E., & Chick, S. E. (2000). Simulation budget allocation for further enhancing the efficiency of ordinal optimization. *Discrete Event Dynamic Systems*, 10(3), 251–270.
- Chen, X., Ankenman, B. E., & Nelson, B. L. (2012). The effects of common random numbers on stochastic kriging metamodels. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 22(2), 7.
- Chen, X., & Kim, K.-K. (2014). Stochastic kriging with biased sample estimates. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 24(2), 8.
- Cressie, N. (1993). *Statistics for spatial data*. New York: John Wiley & Sons.
- Dixon, L. C. W., & Szegö, G. P. (1978). *Towards global optimisation 2*. The Netherlands: North-Holland Amsterdam.
- Forrester, A., Sobester, A., & Keane, A. (2008). *Engineering design via surrogate modelling: A practical guide*. Chichester: John Wiley & Sons.
- Frazier, P., Powell, W., & Dayanik, S. (2009). The knowledge-gradient policy for correlated normal beliefs. *INFORMS Journal on Computing*, 21(4), 599–613.
- Fu, M. C., & Healy, K. J. (1997). Techniques for optimization via simulation: an experimental study on an (s, s) inventory system. *IIE transactions*, 29(3), 191–199.
- Hong, L. J., & Nelson, B. L. (2007). Selecting the best system when systems are revealed sequentially. *IIE Transactions*, 39(7), 723–734.
- Hong, L. J., Nelson, B. L., & Xu, J. (2015). Discrete optimization via simulation. In *Handbook of simulation optimization* (pp. 9–44). Springer.
- Huang, D., Allen, T. T., Notz, W. I., & Zeng, N. (2006). Global optimization of stochastic black-box systems via sequential kriging meta-models. *Journal of Global Optimization*, 34, 441–466.
- Jalali, H., & Van Nieuwenhuyse, I. (2015). Simulation optimization in inventory replenishment: a classification. *IIE Transactions*, 47, 1–19.
- Jones, D. R. (2001). A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization*, 21(4), 345–383.
- Jones, D. R., Schonlau, M., & Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13, 455–492.
- Kim, S.-H., & Nelson, B. L. (2006). Selecting the best system. *Handbooks in operations research and management science*, 13, 501–534.
- Kleijnen, J. P. (2009). Kriging metamodeling in simulation: A review. *European Journal of Operational Research*, 192, 707–716.
- Kleijnen, J. P. (2014). Simulation-optimization via kriging and bootstrapping: a survey. *Journal of Simulation*, 8(4), 241–250.
- Kleijnen, J. P., van Beers, W., & Van Nieuwenhuyse, I. (2012). Expected improvement in efficient global optimization through bootstrapped kriging. *Journal of global optimization*, 54(1), 59–73.
- Kleijnen, J. P., & Van Beers, W. C. (2005). Robustness of kriging when interpolating in random simulation with heterogeneous variances: some experiments. *European Journal of Operational Research*, 165(3), 826–834.
- Kleijnen, J. P. C. (2015). *Design and analysis of simulation experiments* (2nd). Springer, New York.
- Law, A. M. (2015). *Simulation modeling and analysis*. McGraw-Hill, New York.
- Lemieux, C. (2009). *Monte carlo and quasi-monte carlo sampling*. Springer Science & Business Media.
- Liu, C., Pedrielli, G., & Ng, S. H. (2014). eTSSO: Adaptive search method for stochastic global optimization under finite budget. *Working paper*.
- McCormack, R., & Coates, G. (2015). A simulation model to enable the optimization of ambulance fleet allocation and base station location for increased patient survival. *European Journal of Operational Research*, 247(1), 249–309.
- Naoum-Sawaya, J., Ghaddar, B., Arandia, E., & Eck, B. (2015). Simulation-optimization approaches for water pump scheduling and pipe replacement problems. *European Journal of Operational Research*, 246(1), 293–306.
- Picheny, V., & Ginsbourger, D. (2014). Noisy kriging-based optimization methods: A unified implementation within the DiceOptim package. *Computational Statistics & Data Analysis*, 71, 1035–1053.
- Picheny, V., Ginsbourger, D., Richet, Y., & Caplin, G. (2013a). Quantile-based optimization of noisy computer experiments with tunable precision. *Technometrics*, 55, 2–13.
- Picheny, V., Wagner, T., & Ginsbourger, D. (2013b). A benchmark of kriging-based infill criteria for noisy optimization. *Structural and Multidisciplinary Optimization*, 48, 607–626.
- Preuss, M., Wagner, T., & Ginsbourger, D. (2012). High-dimensional model-based optimization based on noisy evaluations of computer games. In *Lion* (pp. 145–159). Springer.
- Quan, N., Yin, J., Ng, S. H., & Lee, L. H. (2013). Simulation optimization via kriging: a sequential search using expected improvement with computing budget constraints. *IIE Transactions*, 45, 763–780.
- Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. Cambridge, MA: MIT press.
- Saif, A., & Elhedhli, S. (2016). Cold supply chain design with environmental considerations: A simulation-optimization approach. *European Journal of Operational Research*, 251(1), 274–287.
- Scott, W., Frazier, P., & Powell, W. (2011). The correlated knowledge gradient for simulation optimization of continuous parameters using gaussian process regression. *SIAM Journal on Optimization*, 21(3), 996–1026.
- Sun, L., Hong, L. J., & Hu, Z. (2014). Balancing exploitation and exploration in discrete optimization via simulation through a gaussian process-based search. *Operations Research*, 62(6), 1416–1438.
- Xu, J. (2012). Efficient discrete optimization via simulation using stochastic kriging. In *Proceedings of the 2012 winter simulation conference (wsc)* (pp. 1–12). IEEE.
- Yin, J., Ng, S., & Ng, K. (2011). Kriging metamodel with modified nugget-effect: The heteroscedastic variance case. *Computers & Industrial Engineering*, 61, 760–777.