*Proceedings of the 2014 Winter Simulation Conference*
*A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.*

# EVALUATION OF KRIGING-BASED METHODS FOR SIMULATION OPTIMIZATION WITH HOMOGENEOUS NOISE

Hamed Jalali, and Inneke Van Nieuwenhuyse

Research Center for Operations Management
KU Leuven, Naamsestraat 69
Leuven, Belgium

## ABSTRACT

In this poster, we evaluate the effectiveness of four kriging-based approaches for simulation optimization with homogeneous noise: *Augmented Expected Improvement* (AEI), *Approximate Knowledge Gradient* (AKG), the *Two-stage sequential optimization* method (TSSO), and the well-known *Efficient Global Optimization* (EGO) algorithm. We test the performance of these algorithms on test functions with homogeneous noise, assuming given computing budget constraints (i.e., given number of infill points and replication budget). Our results indicate that, as long as we have enough replication budget to implement both stages of the Two-stage algorithm, this method is highly competitive, providing similar or better performance particularly in settings with high noise and many infill points.

## 1 INTRODUCTION

Picheny et al. (2013) compared several kriging-based methods for optimizing functions with homogeneous noise and showed that *Augmented Expected Improvement* (AEI, see Huang et al. (2006)) and *Approximate Knowledge Gradient* (AKG, see Scott et al. (2011)) are the best alternatives. The *Two-stage sequential optimization* method (denoted by TSSO here, see Quan et al. (2013)), which is not included in Picheny et al. (2013), is a recently proposed kriging-based method that can also handle heterogeneous simulation noise without requiring any knowledge about the noise variance function.

This poster evaluates the effectiveness of TSSO for simulation optimization with homogeneous noise by comparing it against AEI, AKG, and the popular *Efficient Global Optimization* (EGO) approach (Jones et al. 1998). To this end, we apply these methods to three well-known test functions: Hartmann-6 and Six-hump Camelback (as in Kleijnen et al. (2012)), and the rescaled Branin function (proposed in Picheny et al. (2013)). While the former functions can be considered as "bumpy", the latter one has a flat valley which makes the convergence to the global minimum difficult.

## 2 TEST SCENARIOS

Each function evaluation (i.e., each replication) is perturbed by a homogeneous Gaussian noise. In EGO, AKG, and AEI, we sample infill points by performing $B$ replications while in TSSO, this budget is used for sampling a new point (the *search* stage) and for performing more replications at already sampled points (the *allocation* stage). A heuristic is employed to divide $B$ between these two stages; for the allocation stage to be active, the budget $B$ needs to be sufficiently high relative to the number of infill points allowed (which we denote by $I$, see Quan et al. (2013) for details).

Throughout the experiment, we vary the maximum number of infill points allowed ($I$), the replication budget ($B$) and, the noise level, yielding 18 scenarios per test function. The initial design is a maximin LHS design of size $10 \times d$ (where $d$ is the dimension of the test function), as suggested by Jones et al. (1998). We perform 100 macroreplications per scenario, each using a different initial design and noise realization.

Analogous to Kleijnen et al. (2012), we use a fixed set of candidate points consisting of 200 points for the Six-hump Camelback and rescaled Branin functions and 500 points for the Hartmann-6 function; our objective is to find the global minimum among these candidate points (referred to as $x^*$, having function value $y^*$) using our 4 kriging-based methods. We use *DACE* to implement EGO, the R package *DiceOptim* for AKG and AEI, and the code provided in http://www.ise.nus.edu.sg/staff/ngsh/download/matlabdocs/SOK/ for TSSO.

## 3 RESULTS

To compare the 4 kriging-based methods for different scenarios, we evaluate the following performance measures for each method: (1) the distribution of the optimality gap (defined as the difference between the function value at the best point returned by the method and $y^*$) across the macroreplications (2) the number of macroreplications that identify $x^*$ as the minimum (i.e., for which the previous criterion equals zero) and (3) the number of infill points required to reach this optimum.

Our results indicate that with regard to the first two performance measures, on average, the Two-stage sequential optimization method is the best alternative, provided that the replication budget ($B$) is high enough to allow the allocation stage to be active. Since the allocation stage is active only when $B$ is sufficiently larger than $I$, opting for high number of infill points requires ample computing budget. Only for the Branin function, TSSO performs slightly worse than AKG or AEI for scenarios with low number of infill points. The superior performance of TSSO is more evident when we have high noise. Indeed, although the performance of all methods suffers from higher noise, we found TSSO to be more robust.

The performance of AKG and AEI are most of the time similar. With low number of infill points, most of the time, EGO surprisingly provides similar results or even outperforms AKG and AEI (we found AKG to be especially sensitive to low number of infill points). As we increase the number of infill points, AKG and AEI give better results, especially when we have high noise.

Increasing the replication budget ($B$) and the maximum number of infill points ($I$) enhance the performance of all algorithms, helping them to return solutions that are closer to $y^*$. In terms of speed (the third performance measure), we didn't find evidence on significant difference between the algorithms; with higher noise all of them became slightly slower to converge to the global minimum.

## REFERENCES

Huang, D., T. T. Allen, W. I. Notz, and N. Zeng. 2006. "Global optimization of stochastic black-box systems via sequential kriging meta-models". *Journal of Global Optimization* 34:441–466.

Jones, D. R., M. Schonlau, and W. J. Welch. 1998. "Efficient global optimization of expensive black-box functions". *Journal of Global Optimization* 13:455–492.

Kleijnen, J. P., W. van Beers, and I. Van Nieuwenhuyse. 2012. "Expected improvement in efficient global optimization through bootstrapped kriging". *Journal of Global Optimization* 54:59–73.

Picheny, V., T. Wagner, and D. Ginsbourger. 2013. "A benchmark of kriging-based infill criteria for noisy optimization". *Structural and Multidisciplinary Optimization* 48:607–626.

Quan, N., J. Yin, S. H. Ng, and L. H. Lee. 2013. "Simulation optimization via kriging: a sequential search using expected improvement with computing budget constraints". *IIE Transactions* 45:763–780.

Scott, W., P. Frazier, and W. Powell. 2011. "The correlated knowledge gradient for simulation optimization of continuous parameters using gaussian process regression". *SIAM Journal on Optimization* 21:996–1026.