Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

MULTI-CRITERIA DECISION MAKING IN SETTINGS WITH UNCERTAINTY

Sebastian Rojas-Gonzalez

Department of Decision Sciences and Information Management
KU Leuven
Naamsestraat 69
Leuven, 3000, BELGIUM

ABSTRACT

We consider multiobjective simulation optimization problems with heteroscedastic noise, where we seek to find the non-dominated set of designs evaluated using noisy simulation evaluations. To perform the search of competitive designs, we propose a metamodel-based scalarization method, which explicitly characterizes both the extrinsic uncertainty of the unknown response surface, and the intrinsic uncertainty inherent in stochastic simulation. To differentiate the designs with the true best expected performance, we propose a multiobjective ranking and selection approach that accounts for correlation between the mean values of alternatives. Empirical results show that the proposed methods only require a small fraction of the available computational budget to find the optimal solutions.

1 INTRODUCTION

In simulation optimization, the goal is to find the values of the decision vector(s) that optimize the objective(s) of interest, evaluated using stochastic simulation. In most settings, the analyst has limited computational budget, and thus, the challenge is to develop a strategy that searches the solution space in a way that is both effective (i.e., it has high probability of finding the optimal decision vectors), and efficient (it does so with limited computational budget). Consequently, different approaches have been developed to provide inexpensive metamodels that seek to provide predictions of the objective(s), using the information obtained from a limited number of evaluations. The use of kriging metamodels in simulation optimization has become increasingly popular. The majority of the algorithms so far use ordinary kriging for constructing the metamodels, assuming that solutions have been sampled with infinite precision. This is a major issue when the simulation problem is stochastic: ignoring the noise in the outcomes may not only lead to an inaccurate metamodel, but also to potential errors in identifying the optimal points among those sampled. Moreover, most algorithms so far have focused on single-objective problems.

In this work, we test the performance of a Bayesian multiobjective simulation optimization algorithm that incorporates stochastic kriging metamodels (Ankenman et al. 2010) to account for the inherent noise in the outputs when deciding which are the best solutions observed so far. We exploit the stochastic kriging information during the *search* and *identification* (i.e., ranking and selection) phases of the algorithm, and evaluate them using standard test problems from the literature.

2 SK-MOCBA

Kriging metamodels can successfully approximate outputs over the entire search space (yielding a global metamodel), while also providing a quantification of the prediction uncertainty through the mean square error (MSE). In general, kriging-based optimization methods perform a sequential procedure to improve the accuracy of the metamodels and thus, facilitate the search of optimal solutions. During the search phase, an *infill criterion* directs the search towards interesting solutions by exploiting the kriging information. The metamodel is sequentially updated by simulating the selected *infill point(s)* in the search space, until the computational budget is depleted or a desired performance level is reached. The well-known *expected improvement* criterion (Jones et al. 1998), developed for deterministic simulation, is commonly used in the literature. In contrast, we consider an infill criterion developed for stochastic simulation.

In SK-MOCBA we first scalarize the objectives into one performance function, and search for the points that will minimize this function at each iteration. By randomly choosing a different set of weights at each iteration, we are able to search for optimal points distributed across the Pareto front (Knowles 2006). For the identification phase, we use MOCBA, a well-known multiobjective ranking and selection (MORS) procedure, in view of maximizing the probability of selecting the true Pareto-optimal points by allocating extra replications on competitive designs. We evaluate the impact of these elements on the search and identification effectiveness, for a set of test functions with different Pareto front geometries, and varying levels of heterogeneous noise. Our results show that the use of stochastic kriging is essential in improving the search efficiency; yet, the allocation procedure appears to lose effectiveness in settings with high noise. This emphasizes the need for further research on multiobjective ranking and selection methods.

3 SK-MORS

Built upon the results obtained with SK-MOCBA in the identification phase, we propose a sequential resampling technique that uses a combination of screening, stochastic kriging metamodels, and hypervolume estimates to decide how to allocate samples, in view of maximizing the probability of selecting the true Pareto-optimal points. In contrast to the common assumption in the literature, we consider MORS problems with heteroscedastic noise and correlation between the mean values of alternatives.

Recent research proposed the *expected hypervolume difference* (Branke et al. 2016) as a good criterion to allocate samples. However, estimating the hypervolume change is computationally expensive. We propose an approximation that simply looks at the effect of replacing the observed mean by the mean predicted by the stochastic kriging metamodels. The proposed measure on its own didn't perform particularly well as criterion to allocate samples to alternatives. Best results were obtained when combining it with a measure of the difference of an alternative's observed means and predicted values, also accounting for the MSE of the predictor. An initial screening step further reduces the computational cost needed to compute the hypervolume differences.

Empirical results show that the proposed method only requires a small fraction of samples compared to the standard *equal* allocation, with the exploitation of the correlation structure being the dominant contributor to the improvement. Future work should include a comparison with other state-of-the-art MORS algorithms and a performance comparison based on probability of correct selection of hypervolume difference.

REFERENCES

Ankenman, B., B. L. Nelson, and J. Staum. 2010. “Stochastic Kriging for Simulation Metamodeling”. *Operations Research* 58(2):371–382.

Branke, J., W. Zhang, and Y. Tao. 2016. “Multiobjective Ranking and Selection Based on Hypervolume”. In *Proceedings of the 2016 Winter Simulation Conference*, edited by T. Roeder, P. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. Chick, 859–870. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

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.

Knowles, J. 2006. “ParEGO: A Hybrid Algorithm with On-line Landscape Approximation for Expensive Multiobjective Optimization Problems”. *IEEE Transactions on Evolutionary Computation* 10(1):50–66.