

Ensemble-Based Modeling

Thomas Bartz-Beielstein, Martina Frieze, Oliver Flasch, Wolfgang Konen, Patrick Koch, and Boris Naujoks

Fakultät für Informatik und Ingenieurwissenschaften, Fachhochschule Köln

E-Mail: { thomas.bartz.beielstein | martina.frieze }@fh-koeln.de

Working Paper

Sequential parameter optimization (SPO) can be described as a tuning algorithm with the following properties [Bartz-Beielstein et al., 2004]: (i) Use the available budget (e.g., simulator runs, number of function evaluations) sequentially, i.e., use information from search-space exploration to guide the search by building one or several meta models, e.g., random forest, linear regression, or Kriging. Choose new design points based on predictions from the meta model(s). Refine the meta model(s) stepwise to improve knowledge about the search space. (ii) Try to cope with noise by improving confidence. Guarantee comparable confidence for search points. (iii) Collect and report tuning process information for exploratory data analysis. (iv) Provide mechanisms both for interactive and automated tuning.

The *SPO toolbox* (SPOT) provides standardized interfaces, which enable the integration of several meta models in a convenient manner [Bartz-Beielstein et al., 2010].¹ Naturally, the question arises, which meta model should be used during the tuning process. Instead of recommending one meta model only, we will analyze an alternative approach: Set up several models in parallel, and provide an effective and efficient *policy* for dynamical model selection.

Goal of this study: How to dynamically select the right meta model amongst an ensemble of meta models.

This is a classical exploration—exploitation problem, which has been discussed in the literature for several decades and in different settings (scheduling, design of clinical trials, search). State-of-the-art policies from dynamical programming [Frazier, 2010] will be compared to basic approaches. Consider k meta models. The following policies are subject of our analysis.

1. The *round-robin* policy, which, at time t , selects model number $\{(t - 1) \bmod k\} + 1$. Probably, this is the simplest strategy.
2. The *greedy-search* policy selects the model with the smallest error.
3. The *epsilon-greedy* policy selects the model with the smallest error. It selects sometimes, with a fixed probability, among other models.
4. The *soft-min decision* policy selects models by probability matching. This policy is related to tournament selection in genetic algorithm.
5. The *epsilon soft-min decision* policy gives an uncertainty bonus to models that have not been selected, which augments their probability of being chosen.
6. The *Gittins index* policy tries to minimize the prediction error (expected outcome) from k meta models. This approach considers k alternative meta models as k -armed bandits [Rudolph, 1997]. It determines dynamic allocation indices (Gittins indices), which depend on the number of times a model has been sampled and its outcomes (rewards). This policy tries to handle

¹SPOT can be downloaded from CRAN, see <http://cran.r-project.org/web/packages/SPOT>

the trade-off between exploration (improve reward estimation) versus exploitation (perform actions on the best model) [Gittins et al., 2011].

Results from pre-experimental runs indicate that ensemble-based policies favor Kriging meta models. However, these results should not be generalized, because our studies were performed on a set of classical test functions (e.g., Rosenbrock, Rastrigin, Griewangk, Ackley; fitness-proportionate noise added). Further results will be presented in the final version of this paper.

References

- [Bartz-Beielstein et al., 2010] Bartz-Beielstein, T., Chiarandini, M., Paquete, L., & Preuss, M., Eds. (2010). *Experimental Methods for the Analysis of Optimization Algorithms*. Berlin, Heidelberg, New York: Springer.
- [Bartz-Beielstein et al., 2004] Bartz-Beielstein, T., Parsopoulos, K. E., & Vrahatis, M. N. (2004). Design and analysis of optimization algorithms using computational statistics. *Applied Numerical Analysis and Computational Mathematics (ANACM)*, 1(2), 413–433.
- [Frazier, 2010] Frazier, P. (2010). Learning with dynamic programming. In *Wiley Encyclopedia of Operations Research and Management Science* (pp. 1–13). Wiley.
- [Gittins et al., 2011] Gittins, J., Glazebrook, K., & Weber, R. (2011). *Multi-armed bandit allocation indices*. Wiley.
- [Rudolph, 1997] Rudolph, G. (1997). Reflections on bandit problems and selection methods in uncertain environments. In T. Bäck (Ed.), *Genetic Algorithms: Proceedings Seventh International Conference (ICGA'97)* (pp. 166–173). San Francisco CA: Morgan Kaufmann.