Targeting Well-Balanced Solutions in Multi-Objective Bayesian Optimization under a Restricted Budget

David Gaudrie[1] Rodolphe le Riche[2] Victor Picheny[3] Benoît Enaux[4] Vincent Herbert[5]
[1]Groupe PSA, david.gaudrie@mpsa.com
[2]CNRS, Mines Saint-Étienne, leriche@emse.fr
[3]MIAT, Université de Toulouse, INRA, victor.picheny@inra.fr
[4]Groupe PSA, benoit.enaux@mpsa.com
[5]Groupe PSA, vincent.herbert@mpsa.com

Key words: Gaussian Processes, Parsimonious Optimization, Computer Experiments, Preference-Based Optimization

Multi-objective optimization aims at finding the Pareto set composed of all the best trade-off solutions between several objectives. When dealing with expensive-to-evaluate black box functions, surrogate-based approaches, in the vein of EGO [1] have proven their effectiveness.

However, for extremely narrow budgets, and/or when the number of objectives is large, uncovering the entire Pareto set becomes out of reach even for these approaches. In this work, we restrict our search to well-chosen targeted parts of the Pareto set. This accelerates the search as only a subset of the objective space is considered. As an end-user would typically choose solutions with equilibrated trade-offs between objectives rather than ones favoring a single objective over the others, we will focus on the central part of the Pareto Front, which corresponds to the most well-balanced solutions.

First, we discuss how to define and estimate the center of the Pareto Front. That estimated point has to fairly represent the topology of the front, in spite of the parsimonious knowledge of the objective space. Then, three infill criteria which will guide the optimization, including the *Expected Hypervolume Improvement* [2] are studied and tailored through some of their hyperparameters to enable them to target.

To assess performance, a benchmark built from real-world airfoil aerodynamic data, with variable dimension and number of objectives is used. Compared with standard techniques, the proposed methodology leads to a faster and a more precise convergence towards the Pareto Front in the targeted region.

References

[1] Jones, D. R., Schonlau, M., & Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. Journal of Global optimization, 13(4), 455-492

[2] Emmerich, M. T., Giannakoglou, K. C., & Naujoks, B. (2006). Single-and multiobjective evolutionary optimization assisted by gaussian random field metamodels. IEEE Transactions on Evolutionary Computation, 10(4), 421-439.