

TOWARDS FRUGAL BAYESIAN OPTIMIZATION FOR ENGINEERING DESIGN

Gaston Plat⁽¹⁾, Paul Saves⁽²⁾, Nathalie Bartoli⁽³⁾, Thierry Lefebvre⁽³⁾, Joseph Morlier⁽⁴⁾

⁽¹⁾Université de Toulouse, Institut Supérieur de l'Aéronautique et de l'Espace ISAE-SUPAERO, Toulouse, 31055, France; DTIS, ONERA, Université de Toulouse, 31000 Toulouse, France; Fédération ENAC ISAE-SUPAERO ONERA, Université de Toulouse, 31000 Toulouse, France;

⁽²⁾Université Toulouse Capitole, IRIT UMR 5505 CNRS, Toulouse, France

⁽³⁾DTIS, ONERA, Université de Toulouse, 31000 Toulouse, France; Fédération ENAC ISAE-SUPAERO ONERA, Université de Toulouse, 31000 Toulouse, France

⁽⁴⁾Université de Toulouse, Institut Supérieur de l'Aéronautique et de l'Espace ISAE-SUPAERO, Toulouse, 31055, France

Keywords: Computer experiments, grid computing, energy-aware Bayesian optimization, Gaussian processes, eco-design

Abstract

Blackbox simulations are fundamental to modern engineering and industrial applications. In particular, multidisciplinary design optimization of complex systems, such as aircraft, often necessitates high-fidelity solvers. However, they are computationally intensive.

To mitigate these prohibitive costs, fast-to-evaluate surrogate models have been adopted across diverse domains, ranging from aerospace engineering and socio-technical transportation systems to meta-modeling in deep learning. Consequently, there is growing interest in Bayesian optimization that leverages Gaussian process surrogates for gaining efficiency.

Recent advances in mixed-variable and hierarchical design space techniques have been integrated into the open-source software SMT 2.0 and validated against an extensive suite of academic and industrial benchmarks. While mixed-variable and hierarchical design spaces enable the modeling of increasingly complex problems, expensive simulations remain a bottleneck within optimization loops. To address these intractable execution times, parallel computing has become an essential solution. However, large-scale tasks often demand energy-intensive exascale infrastructure, whereas smaller tasks may be more efficiently handled by regional clusters or standard workstations.

This paper evaluates a resource-aware optimization approach that distributes computer experiment workloads across a heterogeneous network of computing nodes to minimize the energy consumption of computations. Given that hierarchical hardware architectures and time-dependent power grids' availability and consumption greatly vary, the cost and carbon footprint of resource-intensive

experiments fluctuate based on hardware configuration, time of day, and location. We model hierarchical spaces to extend the standard optimization workflow into a system architecture optimization problem with multiple objectives, specifically targeting computational energy consumption.

We first formulate energy-minimization toy problems of varying complexities and compare energy-aware Bayesian optimizers against the NSGA-II evolutionary algorithm. The results demonstrate the existence of significant Pareto optimal trade-offs between objective performance and energy consumption. For example, in a deep learning case study (MNIST digit classification), we show that a 5% improvement in accuracy (91% to 96%) can result in a 2800% increase in training energy consumption. Furthermore, we find that more complex hierarchical relations between variables yield a richer set of trade-offs. Finally, the approach is successfully applied to a revisited classical Haftka's 10-bar truss problem.