Preference-Aware Constrained Multi-Objective Bayesian Optimization (Student Abstract)

Alaleh Ahmadianshalchi[1], Syrine Belakaria[2], Janardhan Rao Doppa[1]

[1]School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA, USA
[2]Stanford University, Stanford, CA, USA
a.ahmadianshalchi@wsu.edu, syrineb@stanford.edu, jana.doppa@wsu.edu

Abstract

We consider the problem of constrained multi-objective optimization over black-box objectives, with user-defined preferences, with a largely infeasible input space. Our goal is to approximate the optimal Pareto set from the small fraction of feasible inputs. The main challenges include huge design space, multiple objectives, numerous constraints, and rare feasible inputs identified only through expensive experiments. We present PAC-MOO, a novel preference-aware multi-objective Bayesian optimization algorithm to solve this problem. It leverages surrogate models for objectives and constraints to intelligently select the sequence of inputs for evaluation to achieve the target goal.

Introduction

A large number of engineering design problems such as electric power systems design (Wang et al. 2018) and design of analog circuits (Zhou et al. 2020) involve making design choices to optimize multiple objectives. The common challenges in such constrained multi-objective optimization (MOO) problems include the following. **1)** The objective functions are unknown and expensive. **2)** The objectives are conflicting in nature and all of them cannot be optimized simultaneously. **3)** All constraints need to be evaluated using expensive experiments and satisfied. **4)** Only a small fraction of the input design space is feasible (i.e., satisfies all constraints). Additionally, in several real-world applications, the practitioners have specific preferences over the objectives. For example, the designer prefers efficiency over settling time when optimizing analog circuits.

Bayesian optimization (BO) is an efficient framework to solve black-box optimization problems with expensive functions (Shahriari et al. 2016). The key idea behind BO is to learn surrogate models of the expensive objective function and intelligently select the sequence of inputs for evaluation using an acquisition function. We propose a novel and efficient information-theoretic approach referred to as **P***reference*-**A***ware* **C***onstrained* **M***ulti*-**O***bjective Bayesian* **O***ptimization (PAC-MOO)*. PAC-MOO builds surrogate models for both output objectives and constraints based on the training data from past function evaluations. The selected input design at each iteration maximizes the information gain about the constrained optimal Pareto front while factoring in the designer preferences over objectives. We develop a tractable acquisition function to select inputs for evaluation, increasing the chances of finding feasible inputs, and incorporating preferences over objectives.

Problem Setup and PAC-MOO Algorithm

Constrained MOO w/ Preferences. Constrained MOO is the problem of optimizing $\mathbf{K} \geq 2$ real-valued objective functions $\{f_1(x), \cdots, f_{\mathbf{K}}(x)\}$, while satisfying $\mathbf{L}$ black-box constraints of the form $c_1(x) \geq 0, \cdots, c_{\mathbf{L}}(x) \geq 0$ over the given design space $\mathfrak{X}$. A function evaluation with the candidate parameters $\mathbf{x} \in \mathfrak{X}$ generates two vectors, one consisting of objective values and one consisting of constraint values $\mathbf{y} = (y_{f_1}, \cdots, y_{f_{\mathbf{K}}}, y_{c_1}, \cdots, y_{c_{\mathbf{L}}})$ where $y_{f_j} = f_j(x) \forall j \in \{1, \cdots, K\}$ and $y_{c_i} = C_i(x) \forall i \in \{1, \cdots, L\}$. We define an input vector $\mathbf{x}$ as feasible if and only if it satisfies all constraints. The input vector $\mathbf{x}$ *Pareto-dominates* another input vector $\mathbf{x}'$ if $f_j(\mathbf{x}) \leq f_j(\mathbf{x}')$ $\forall j$ and there exists some $j \in \{1, \cdots, K\}$ such that $f_j(\mathbf{x}) < f_j(\mathbf{x}')$. The designer/practitioner can define input preferences over multiple black-box functions through the notion of preference specification, which is defined as a vector of scalars $\mathbf{p} = \{p_{f_1}, \cdots, p_{f_K}, p_{c_1}, \cdots, p_{c_L}\}$ with $0 \leq p_i \leq 1$ and $\sum_{i \in \mathcal{I}} p_i = 1$ such that $\mathcal{I} = \{f_1, \cdots, f_K, c_1, \cdots, c_L\}$. Higher values of p_i mean that the corresponding objective function f_i is highly preferred.

Overview of PAC-MOO. PAC-MOO algorithm is an instance of the BO framework, which takes as input the input space $\mathfrak{X}$, preferences over objectives p, expensive objective functions and constraints evaluator, and produces a Pareto set of candidate inputs as per the preferences. In each iteration, PAC-MOO selects a candidate input design $\mathbf{x}_t \in \mathfrak{X}$ to perform a function evaluation. Consequently, the surrogate models for both objective functions and constraints are updated based on training data from the function evaluations.

The MESMOC acquisition function (Belakaria, Deshwal, and Doppa 2020a) maximizes the information gain between the next candidate input for evaluation $\mathbf{x}$ and the optimal constrained Pareto front:

$$AF(i, x) = \sum_{s=1}^{S} \frac{\gamma_s^i(\mathbf{x})\phi(\gamma_s^i(\mathbf{x}))}{2\Phi(\gamma_s^i(\mathbf{x}))} - \ln \Phi(\gamma_s^i(\mathbf{x})) \quad (1)$$

$$\alpha(\mathbf{x}) \simeq \sum_{i \in \mathcal{I}} AF(i,x), \mathcal{I} = \{c_1, \cdots, c_L, f_1, \cdots, f_K\} \quad (2)$$

Where ϕ and Φ are the p.d.f and c.d.f of a standard normal distribution respectively. We provide The definition of γ and the complete derivation of this acquisition function in (Ahmadianshalchi, Belakaria, and Doppa 2023). The acquisition function proposed in Equation 1 resulted in a function in the form of a summation of an entropy term defined for each of the objective functions and constraints as $AF(i,x)$. Given this expression, the algorithm will select an input while giving the same importance to each of the functions and constraints. We propose to inject preferences from the designer into our algorithm by associating different weights to each of the objectives. A principled approach is to assign appropriate preference weights resulting in a convex combination of the individual components of the summation $AF(i,x)$. Let p_i be the preference weight associated with each individual component. The preference-based acquisition function is defined below.

$\alpha_{pref}(\mathbf{x}) \simeq \sum_{i \in \mathcal{I}} p_i \times AF(i,x) \; s.t \sum_{i \in \mathcal{I}} p_i = 1$

Intuitively, the algorithm should first aim at identifying feasible input configurations by maximizing the probability of satisfying all the constraints. We define a special case of our acquisition function for such challenging scenarios as $\alpha_{prob}(x) = \prod_{i=1}^{L} Pr(c_i(x) \geq 0)$ This acquisition function enables an efficient feasibility search due to its exploitation characteristics (Gardner et al. 2014). Given that the probability of constraint satisfaction is binary, the algorithm will be able to quickly prune unfeasible regions of the input space and move to other promising regions until it identifies feasible input configurations.

Results and Discussion

A detailed explanation of the experimental setup, real-world problems, and the pseudo-code for PAC-MOO are included in (Ahmadianshalchi, Belakaria, and Doppa 2023).

We utilize a modified version of the OSY problem (Osyczka and Kundu 1995) where each dimension of the input space is expanded to 1.5 times its original size; We include an additional constraint flagging any input outside the original input space as infeasible. The modified OSY problem exhibits a significantly reduced rate of feasible points. This problem is initialized with 12 random initial points and comprises 6 input dimensions, 2 objective functions, and 7 constraints.

In Figure 1, PAC-MOO-0 refers to assigning equal preferences to all black-box functions. PAC-MOO-1, 2, and 3 assign higher preference values to the preferred objective function. Our implementation of PAC-MOO is included in https://github.com/Alaleh/PAC-MOO.

As shown in Figures 1a, 1c, and 1e, PAC-MOO with no preferences typically outperforms all the baseline methods. This is attributed to the efficient information-theoretic acquisition function and the exploitation approach to finding feasible regions in the input space. At least one version of USEMOC (Belakaria, Deshwal, and Doppa 2020b) outperforms all evolutionary baselines demonstrating that BO methods have an advantage over evolutionary algorithms. The performance of PAC-MOO with preference may slightly decline

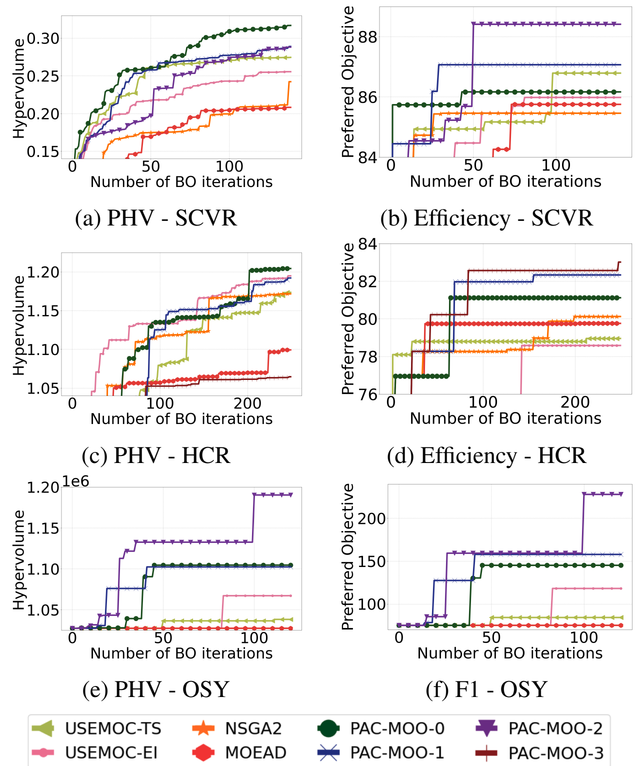

(a) PHV - SCVR (b) Efficiency - SCVR (c) PHV - HCR (d) Efficiency - HCR (e) PHV - OSY (f) F1 - OSY

Figure 1: PAC-MOO results on multiple benchmarks

in some cases in terms of the PHV because the PHV metric evaluates the quality of the general Pareto front, while PAC-MOO focuses on specific regions of the Pareto front through preference specification. Although this behavior is expected, we have observed that the PHV of PAC-MOO with preferences remains competitive with PAC-MOO-0.

Figures 1b, 1d, and 1f show the results for maximum discovered feasible value of the preferred objective as a function of the number of BO iterations. PAC-MOO with preferences outperforms all baseline methods, by finding feasible input designs with higher values of the preferred objective function. In some cases, the improvement in maximum efficiency of uncovered circuit configurations for PAC-MOO with preferences comes at the expense of loss in hypervolume metric as shown in Figure 1a and Figure 1c. In problems involving a small number of objective functions (e.g., the OSY problem), the increase in the value of the preferred objective can outweigh the negative effects of emphasizing specific regions of the Pareto front through objective preferences.

Conclusion

PAC-MOO is a specialized BO algorithm that addresses large input spaces, costly experiments, rare feasible inputs, and preference-driven objectives. Its key innovations include an efficient information gain based acquisition function, effective exploration of feasible input space, and integration of preferences via convex combination of acquisition functions. It can be used in a variety of complex engineering design optimization problems.

Acknowledgements

The authors gratefully acknowledge the support from National Science Foundation (NSF) grants IIS-1845922, SII-2030159, and CNS-2308530.

References

Ahmadianshalchi, A.; Belakaria, S.; and Doppa, J. R. 2023. Preference-Aware Constrained Multi-Objective Bayesian Optimization. arXiv:2303.13034.

Belakaria, S.; Deshwal, A.; and Doppa, J. R. 2020a. Max-value Entropy Search for Multi-Objective Bayesian Optimization with Constraints. *arXiv preprint arXiv:2009.01721*.

Belakaria, S.; Deshwal, A.; and Doppa, J. R. 2020b. Uncertainty aware search framework for multi-objective Bayesian optimization with constraints. *arXiv preprint arXiv:2008.07029*.

Gardner, J. R.; Kusner, M. J.; Xu, Z. E.; Weinberger, K. Q.; and Cunningham, J. P. 2014. Bayesian Optimization with Inequality Constraints. In *ICML*, volume 2014, 937–945.

Osyczka, A.; and Kundu, S. 1995. A new method to solve generalized multicriteria optimization problems using the simple genetic algorithm. *Structural optimization*, 10: 94–99.

Shahriari, B.; Swersky, K.; Wang, Z.; Adams, R. P.; and de Freitas, N. 2016. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE*, 104(1): 148–175.

Wang, J.; Yang, W.; Du, P.; and Li, Y. 2018. Research and application of a hybrid forecasting framework based on multi-objective optimization for electrical power system. *Energy*, 148: 59–78.

Zhou, Z.; Belakaria, S.; Deshwal, A.; Hong, W.; Doppa, J. R.; Pande, P. P.; and Heo, D. 2020. Design of Multi-Output Switched-Capacitor Voltage Regulator via Machine Learning. In *DATE*.