One Step Preference Elicitation in Multi-Objective Bayesian Optimization

Juan Ungredda
J.Ungredda@warwick.ac.uk
University of Warwick
Coventry, UK

Juergen Branke
juergen.branke@wbs.ac.uk
University of Warwick
Coventry, UK

Mariapia Marchi
marchi@esteco.com
ESTECO SpA
Trieste, Italy

Teresa Montrone
montrone@esteco.com
ESTECO SpA
Trieste, Italy

ABSTRACT

We consider a multi-objective optimization problem with objective functions that are expensive to evaluate. The decision maker (DM) has unknown preferences, and so the standard approach is to generate an approximation of the Pareto front and let the DM choose from the generated non-dominated designs. However, especially for expensive to evaluate problems where the number of designs that can be evaluated is very limited, the true best solution according to the DM's unknown preferences is unlikely to be among the small set of non-dominated solutions found, even if these solutions are truly Pareto optimal. We address this issue by using a multi-objective Bayesian optimization (BO) algorithm (see, e.g., [5]) and allowing the DM to select a preferred solution from a predicted continuous Pareto front just once before the end of the algorithm rather than selecting a solution after the end. This allows the algorithm to understand the DM's preferences and make a final attempt to identify a more preferred solution. We demonstrate the idea using ParEGO [2], and show empirically that the found solutions are significantly better in terms of true DM preferences than if the DM would simply pick a solution at the end.

CCS CONCEPTS

• **Computing methodologies** → **Gaussian processes**; **Search methodologies**; *Genetic algorithms*; • **Applied computing** → **Multi-criterion optimization and decision-making**.

KEYWORDS

Preference Elicitation, Gaussian Processes, Bayesian Optimization

ACM Reference Format:
Juan Ungredda, Juergen Branke, Mariapia Marchi, and Teresa Montrone. 2021. One Step Preference Elicitation in Multi-Objective Bayesian Optimization. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion), July 10–14, 2021, Lille, France.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3449726.3459515

GECCO '21 Companion, July 10–14, 2021, Lille, France

ACM ISBN 978-1-4503-8351-6/21/07.
https://doi.org/10.1145/3449726.3459515

1 PROBLEM DEFINITION

We assume a space of possible *solutions* $\mathbf{x} \in X \subset \mathbb{R}^D$ where the objective function is an arbitrary black box $\mathbf{f} : X \rightarrow \mathbb{R}^K$ that returns a deterministic vector *output* $\mathbf{y} \in \mathbb{R}^K$. The (unknown) DM preference over the outputs can be characterized by a utility function $U : \mathbb{R}^K \rightarrow \mathbb{R}$. Thus, of all solutions in X, the DM's most preferred solution is $\mathbf{x}^* = \arg\max_{\mathbf{x}\in X} U(\mathbf{f}(x))$.

There is a budget of B objective function evaluations. It is possible to engage with the DM once after $B - p$ evaluations to learn about the DM's preferences. The algorithm then continues running for the final p iterations. At the end, the DM chooses a preferred solution $\mathbf{x}_B$ from the sampled solution set Γ according to $\mathbf{x}_B = \arg\max_{x\in\Gamma} U(\mathbf{f}(x))$. The quality of the final solution set Γ is determined by the utility of the chosen sample, and we choose to minimize the Opportunity Cost (OC) or regret:

$$OC = U(\mathbf{f}(\mathbf{x}^*)) - U(\mathbf{f}(\mathbf{x}_B))$$

2 ALGORITHM

2.1 Statistical Model over Simulator and Utility

We follow ParEGO where the DM's utility can be described with a Tchebychev utility $U_{\boldsymbol{\theta}}(\mathbf{x})$ with parameters $\boldsymbol{\theta} \in \Theta$. The utility may be modelled by a Gaussian Process (GP) with mean function $\mu_U(\mathbf{x}) : X \rightarrow \mathbb{R}$ and a covariance function $k_U(\mathbf{x}, \mathbf{x}') : X \times X \rightarrow \mathbb{R}$. However, to show a Pareto front approximation to the DM, we propose to use an independent GP to model each objective function $y_j = f_j(\mathbf{x}), \forall j = 1, \ldots, K$, defined by a mean function $\mu_j(\mathbf{x}) : X \rightarrow \mathbb{R}$ and a covariance function $k_j(\mathbf{x}, \mathbf{x}') : X \times X \rightarrow \mathbb{R}$.

We use the popular squared exponential kernel and set the prior mean to zero and we estimate hyper-parameters by maximising the marginal likelihood. Further details can be found in [3].

2.2 ParEGO with p-Step Preference Elicitation

During the first $B - p - 1$ iterations ParEGO translates a multi-objective problem into a single-objective problem using a Tchebychev function with randomly sampled $\boldsymbol{\theta} \in \Theta$. To focus more directly on the region interesting to the DM, at step $B - p$ we fit a GP for each objective function $f_j(x)$ to produce a response surface over each output using the mean posterior μ^m. Then, we use NSGA-II to produce a Pareto front approximation based on the response surfaces and show this approximation to the DM. At this point, the

DM picks their most preferred solution x_p. Assuming that the DM's utility function is based on the Tchebychev utility, we can estimate the parameters $\boldsymbol{\theta}$ as

$$\frac{\hat{\theta}_i}{\hat{\theta}_j} = \frac{f_j(x_p)}{f_i(x_p)} \qquad \forall i, j = 1 \dots K$$

Finally, we compute the last p optimization steps by using expected improvement with the estimated parameters $EI_{\hat{\theta}}(x)$.

3 RESULTS AND DISCUSSION

We compare the proposed approach against using standard ParEGO without exploiting information from a DM in the final p optimization steps. To plot the convergence of the opportunity cost (OC) over iterations, we want to show, at each iteration i, the performance if the algorithm would have stopped there. Thus, to determine the OC at iteration i, we take the set of non-dominated solutions generated up to iteration $i - p$, create an approximation of the Pareto front, select the most preferred solution from that approximated front by using the true DM's preferences, and finally execute iterations $i-p+1, \dots, i$ of the optimization taking the preference information into account. The final OC at iteration i is then the OC of this final solution set and is based on the true DM utility function.

3.1 Experimental setup

The true underlying parameters are generated randomly for every replication of a run using a different random seed. In all experiments, NSGA-II is run for 300 generations with a population size of 100 to produce a Pareto front approximation. We use the POL function [1] defined over $X = [-\pi, \pi]^2$ and also has $K = 2$ objectives, and the HOLE function [4] defined over $X = [-1, 1]^2$ with $K = 2$ objectives with parameter $b > 0$.

3.2 Synthetic Experiments

Figure 1 (**first row**) displays the benefit that can be gained from showing the DM the approximated Pareto front at the final stage of optimization, with $p = 1$ or $p = 2$. Results look similar, although there seems to be a more noticeable difference between the settings of $p = 1$ and $p = 2$ on the POL problem. Figure 1 (**middle row**) shows the results when the DM's true utility model is linear, meaning there is a mismatch between the DM's true utility model (unknown to the algorithm) and the learned Tchebychev model used in ParEGO. Results show a reduced benefit from asking the DM one step before the end of the run. Lastly, Figure 1 (**bottom row**) considers the final OC after 100 iterations for larger values of p. Best results for both synthetic functions were found when the DM is shown the approximated front approximately 20 iterations before the end of the run. Asking the DM very early yields rather poor results, presumably because the approximated Pareto front shown to the DM at that time is far away from the true Pareto front, and the information learned from the interaction is therefore not very helpful. However, after 80 evaluations, the approximated Pareto front seems of sufficient quality to get a meaningful estimate of the DM's preferences and guide the remaining optimization steps in a reasonable direction.

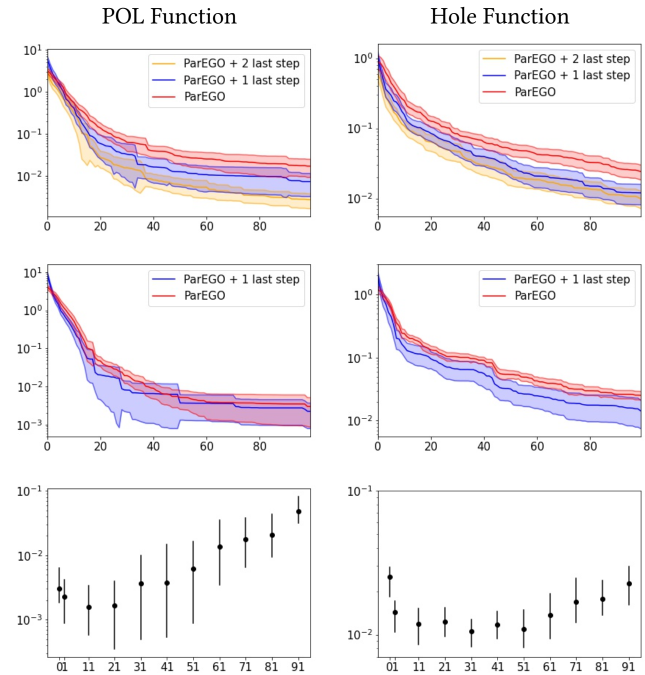

Figure 1: Mean and 95% CI for the OC over iterations (first and middle row) and final OC over p (bottom row).

4 CONCLUSION

For the case of expensive multi-objective optimization, we show how the surrogate models generated by BO can be used not only to speed up optimization, but also to show an approximated continuous Pareto front to the DM once before the end of optimization. The information on the most preferred solution can be used to focus the final iterations of the algorithm to try to find this predicted most preferred solution, or even a dominating solution. We argue that the additional cognitive effort for the DM should be small, but the benefit in terms of true utility to the DM may be significant. We demonstrate the proposed approach on a few test problems.

Future directions may include to explore this idea with other BO algorithms, and also to turn this into a fully interactive approach with multiple interactions with the DM during the optimization.

REFERENCES

[1] C.Poloni. 1997. Hybrid GA for multiobjective aerodynamic shape optimization. In *Genetic Algorithms in Engineering and Computer Science*. Wiley, 397–414.
[2] J. Knowles. 2006. ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation* 10, 1 (2006), 50–66.
[3] C. E. Rasmussen and C. K. I. Williams. 2006. *Gaussian Processes for Machine Learning*. MIT Press.
[4] E. Rigoni. 2004. *Hole functions problem*. Technical Report 2004-02. ESTECO.
[5] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. 2016. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proc. IEEE* 104, 1 (2016), 148–175.