An evaluation of BoTorch for Structured Bayesian Optimisation and Human-in-the-loop Approaches

Umer Hasan

December 3, 2025

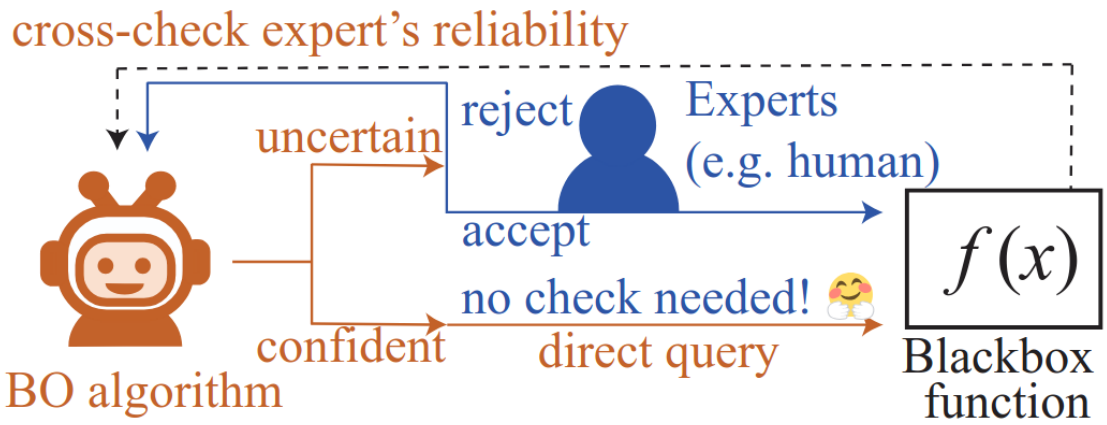

Figure 1: The binary labeller scenario. [2]

Primal-dual problem where human feedback is the regulariser with a primal-dual weight λ_t at the t-th iteration:

$$min_{x \in \chi} f(x) \text{ st } g_t(x) \leq 0$$

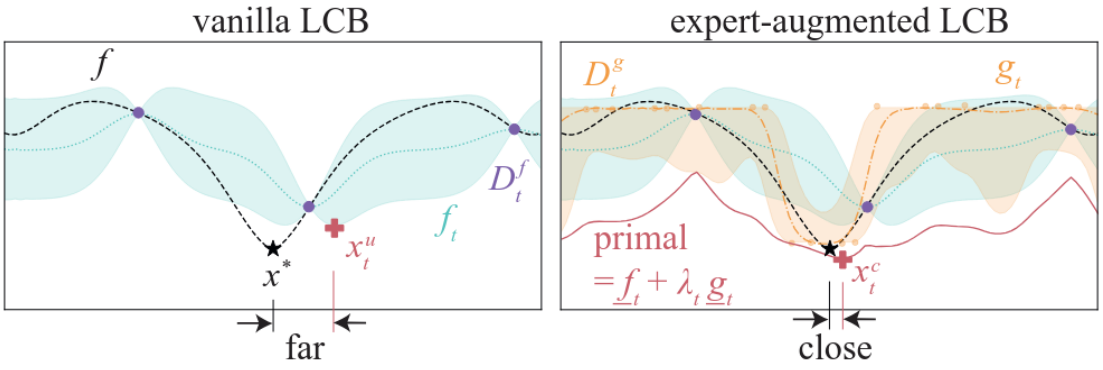

Figure 2: Expert-augmented LCB navigates to a closer point x_t^c by mixing f_t and g_t with $f_t + \lambda_t g_t$. D_t^f is the set of sample points of the objective function and D_t^g is the human feedback. [2]

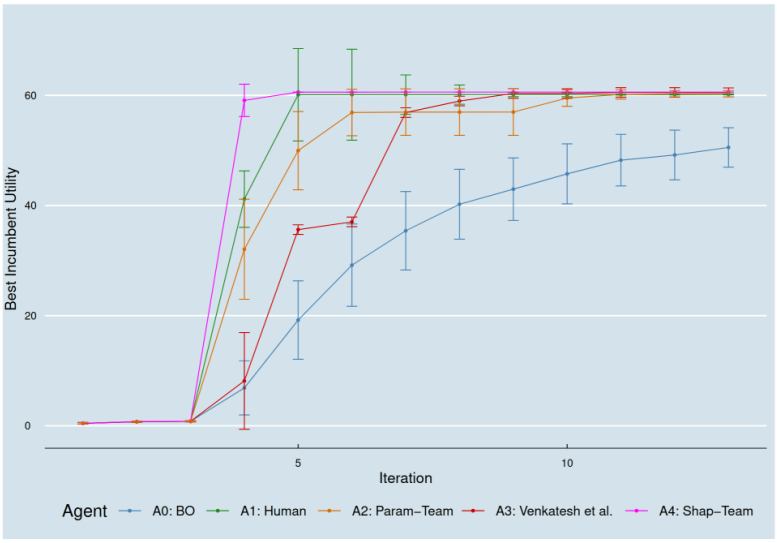

Figure 11: Results of Agents A0-A4 (see Table 4) in human-AI collaborative BO for simulated exosuit personalization (individual 2) with 10 iterations and 3 initial samples each. Error bars indicate 95% confidence intervals; $k = 2$ for A3, $\beta = 2$ for A2 and A4.

Figure 3: Shapley values from game theory are used to explain why BO chose the parameters. Appended to the BO output is a corresponding vector ϕ of parameter contributions. [1]

Research in Bayesian optimisation (BO) usually focusses on black box functions where no inductive bias is supplied to the system. As users, we often have detailed domain knowledge so are in a position to guide optimisers to do efficient work. This is where the advantage of structured Bayesian optimisation (SBO) lies.[1]

There are some competitors to BoTorch: *Trieste* backed by Tensorflow, *Dragonfly* for scalable BO. Often, Dragonfly implementations use BoTorch. What is the gap in BoTorch? Is direct implementation better than using any framework?

Progress so far is mostly reading abstracts: Principled Bayesian Optimisation in Collaboration with Human Experts (NeurIPS 2024 spotlight).

12 Dec - Finish Survey Report on SBO.

15-17 Dec - Implementation of a new SBO strategy in BoTorch

18-19 Dec - Critical analysis after using BoTorch

22-24 Dec - Finish literature review

27-30 Dec - Any remaining results / final reviews

1 Jan - New Year New Me.

References

[1] Julian Rodemann et al. *Explaining Bayesian Optimization by Shapley Values Facilitates Human-AI Collaboration*. 2024. arXiv: 2403.04629 [cs.LG]. URL: https://arxiv.org/abs/2403.04629.

[1] In this work we will use the terms grey box BO, human-in-the-loop BO, SBO interchangeably. For the most part, we will refer to it as SBO.

[2] Wenjie Xu et al. *Principled Bayesian Optimisation in Collaboration with Human Experts*. 2024. arXiv: 2410.10452 [cs.LG]. URL: https://arxiv.org/abs/2410.10452.