Parallel cost-aware optimization of multidimensional black-box functions

Author
Oliver Sihlovec, O.Sihlovec@student.tudelft.nl

Responsible Professor
Matthijs Spaan, m.t.j.spaan@tudelft.nl

Supervisor
Joery de Vries, J.A.deVries@tudelft.nl

01 Background

- Blackbox functions are functions that have unknown analytical expression, are noisy and expensive to evaluate in terms of time and/or resources
- Bayesian optimizers estimate global optima by inferring posterior distributions given previous observations and are commonly utilized in Hyperparameter Optimization or Algorithm selection
- Most optimizers are short-sighted and do not consider dynamic evaluation costs
- Parallel cost-aware optimizers could lead to faster convergence and become more robust against environment with delayed evaluation outputs

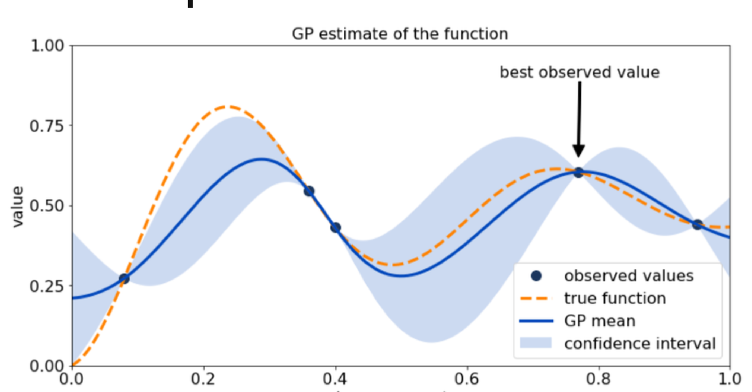

Figure 1. Gaussian Process of Bayesian Optimization

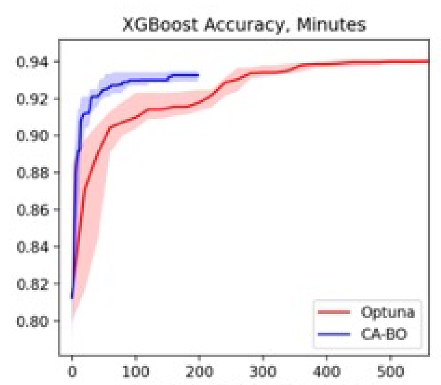

Figure 2.Comparison of cost (un)aware optimization

02 Research Questions

- Do parallel implementations enhance the optimality of multi-objective Bayesian optimizers?
- Can parallel expected improvement (qEI) be adapted in order to achieve cost-awareness for multi-objective Bayesian optimization?
- Is the multi-timestep variant of cost-aware Bayesian optimization resilient to environment misspecification?

03 Methodology

- The research is based on work of Abdolshah et al. [1], who propose a cost-aware heuristic combined with UCB acquisition function
- UCB policy is greedy and is replaced using the Parallel Expected Improvement (qEI) proposed by Wang et al. [2]
- This acquisition function outputs a set of data points that jointly minimize the cost and is thus multi-timestep by design
- The two implementations are then benchmarked in various environments and input/objective spaces

04 Results

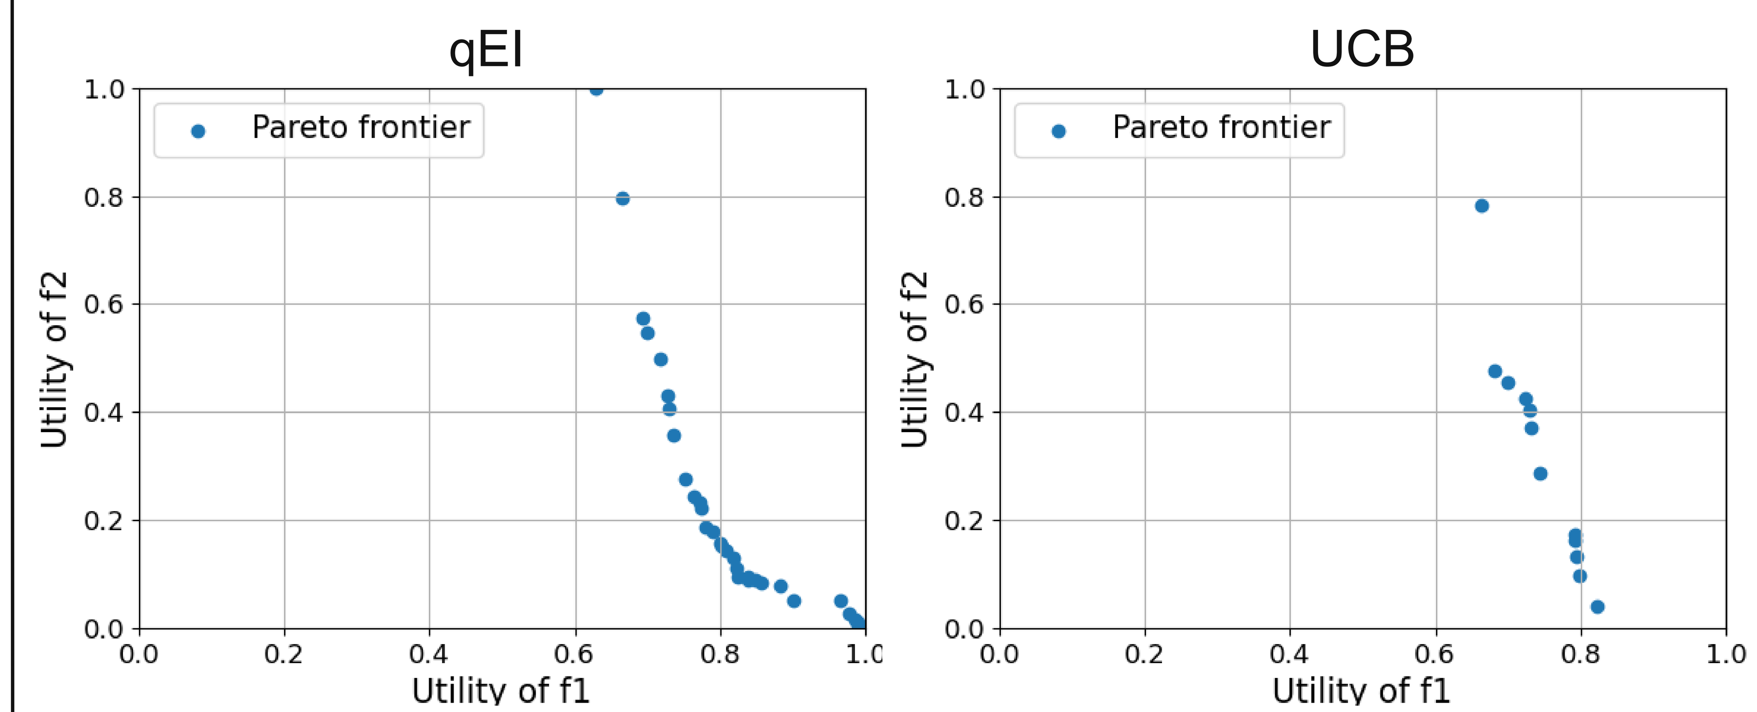

Figure 3. Comparison of optimization performance of qEI (left) and UCB (right) acquisition functions in terms of objective utilities. The more area the Pareto frontier envelops, the better.

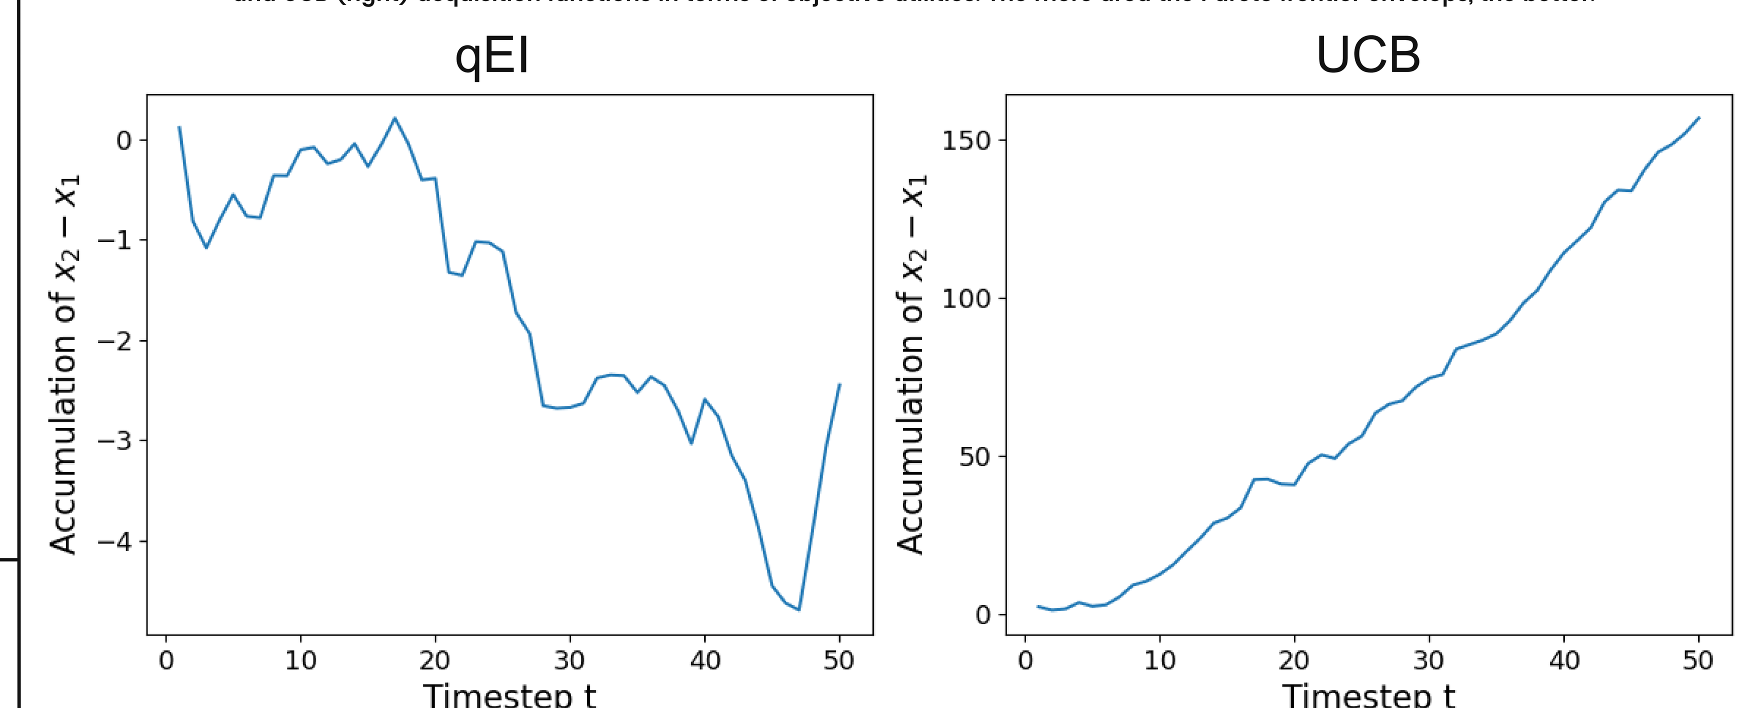

Figure 4. Comparison of cost-awareness of qEI (left) and UCB (right) acquisition functions in terms of cumulative differences per dimensions of selected input. Positive difference corresponds to prioritization of cheaper inputs and vice-versa. Thus, larger positive difference is better.

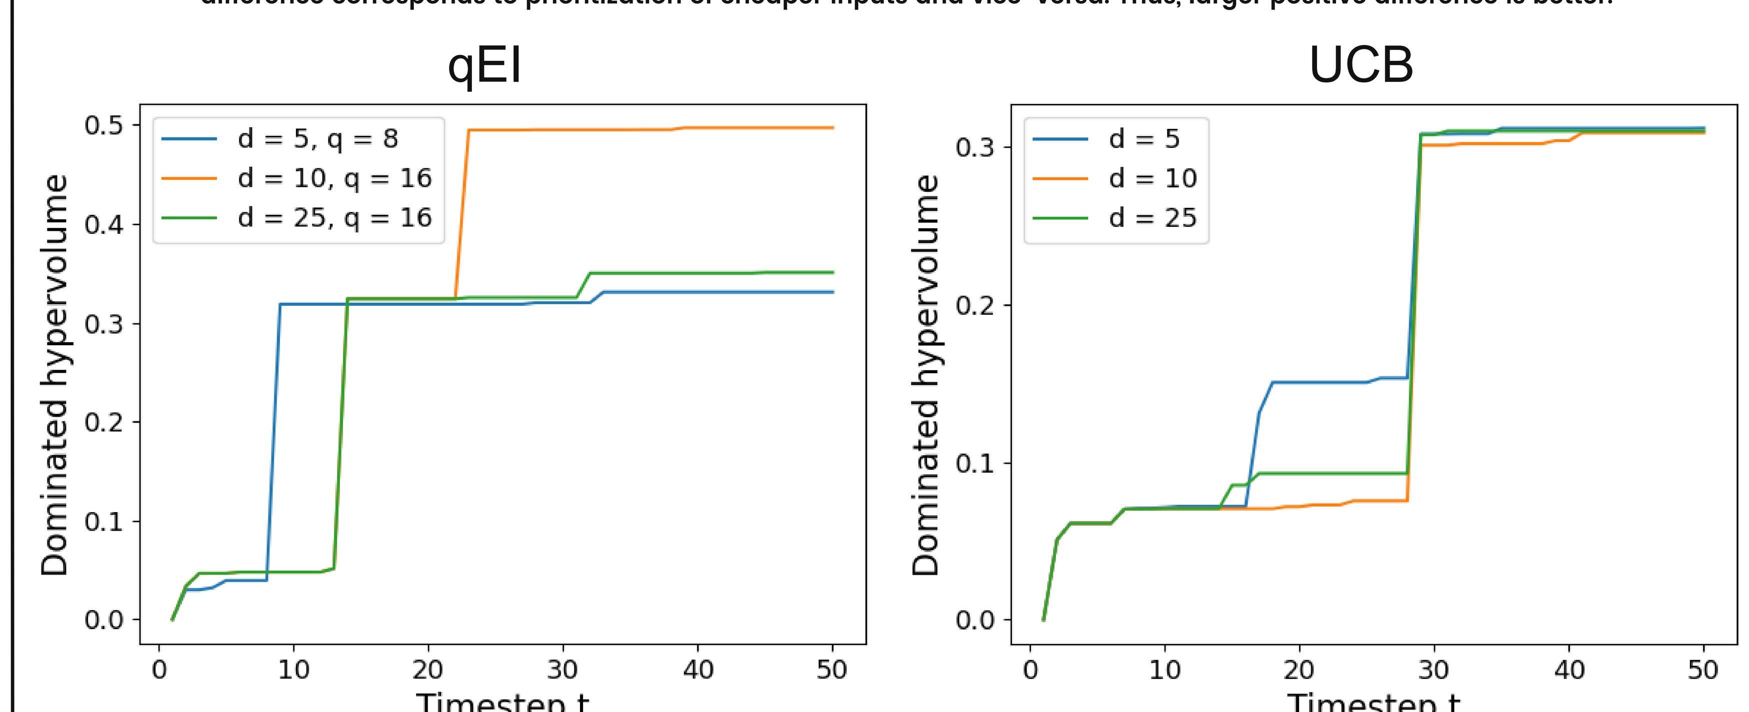

Figure 5. Comparison of optimization performance in a delayed environment of qEI (left) and UCB (right) acquisition functions with respect to dominated hypervolume of the resulting Pareto frontier. Higher dominated hypervolume is better. d represents the response delay in terms of number of timesteps, q denotes the number of look-ahead steps. Since the UCB strategy is myopic, the q parameter was omitted.

05 Conclusion

- qEI achieves more optimal results than UCB due to querying significantly more inputs within the same number of timesteps
- qEI is incompatible with cost-aware heuristic of Aboldshah et al. [2], possibly due to its input sensitivity or cost-per-batch aggregation
- qEI is able to diminish the effect of a delayed environment, but fine-tuning the batch size for a given delay is crucial, as suboptimal choices of control parameters cause the method to regress to a judicious random strategy

06 Limitations and Future Work

- The short time-frame of the project was a limiting factor. Running the experiment for more timesteps, repetition of the experiments and construction of confidence intervals or scalability analysis would augment the robustness of the research.
- Experimentation with less discriminatory parallel methods could be fruitful, as such alternatives could facilitate the cost-aware heuristic more effectively
- A different approach to cost-awareness is also worth entertaining. For example, it is possible to employ Bayesian optimization to estimate the costs of unseen samples based on previous observations and attempt to minimize them as a secondary objective [3]

References

[1] Majid Abdolshah et al. Cost-aware Multi-objective Bayesian optimisation. 2019. arXiv: 1909. 03600 [cs.LG]

[2] Jialei Wang et al. Parallel Bayesian Global Optimization of Expensive Functions. 2019. arXiv: 1602.05149.

[3] Eric Hans Lee et al. Cost-aware Bayesian Optimization. 2020. arXiv: 2003.10870 [cs.LG]