qPOTS: Efficient Batch Multiobjective Bayesian Optimization via Pareto Optimal Thompson Sampling

28th International Conference on Artificial Intelligence and Statistics (AISTATS), 2025.

Ashwin Renganathan [1] Kade Carlson[1]

Computational complex engineered Systems Design Lab (CSDL)
[1]The Pennsylvania State University

Introduction

Mathematically, we consider the simultaneous constrained optimization of K objectives

$$\max_{\mathbf{x}\in\mathcal{X}} \{f_1(\mathbf{x}), \ldots, f_K(\mathbf{x})\} \quad \text{s.t. } c_i(\mathbf{x}) = 0,\ i \in \mathcal{E} \quad c_i(\mathbf{x}) \geq 0,\ i \in \mathcal{I} \tag{1}$$

- $f_i, c_i,\ \forall i$ are expensive stochastic zeroth-order oracles.
- We are interested in a **derivative-free** and **sample efficient** solution to (1).

State of the art and gaps

- **Evolutionary algorithms**, e.g., NSGA-II Deb et al. [2000], are the de facto method in the absence of derivatives. They are **accurate but not sample efficient**.
- **Bayesian optimization** use statistical surrogates (e.g., Gaussian process models) with Bayesian decision theory. They are **sample efficient but not very accurate.**
- Existing work in multiobjective BO
  - **(1) Scalarization.** Zhang et al. [2009], Zhou et al. [2012], Daulton et al. [2021], Paria et al. [2020] **(2) Hypervolume based.** Couckuyt et al. [2014], Emmerich [2008], Yang et al. [2019], Daulton et al. [2020], Picheny [2015], Daulton et al. [2021]. **(3) Entropy based.** Hvarfner et al. [2022], Tu et al. [2022], Suzuki et al. [2020], Belakaria et al. [2019]

*We propose a **best of both worlds** solution: we combine Bayesian optimization with evolutionary algorithms for improved sample efficiency and accuracy.*

Thompson sampling for Bayesian optimization

- Thompson sampling posits that we make decisions according to the probability that they are optimal.
- That is, $\mathbf{x}_{n+1} \sim p_{\mathbf{x}_*}(\mathbf{x})$, where $p_{\mathbf{x}_*}(\mathbf{x})$ is the probability density of the maximizers of f.
- In the context of Bayesian optimization, we could use the surrogate Y to define this density as

$$p_{\mathbf{x}^*}(\mathbf{x}) = \int \delta\left(\mathbf{x} - \arg\max_{\mathbf{x}\in\mathcal{X}} Y(\mathbf{x})\right) p(Y|\mathcal{D}_n)dY, \tag{2}$$

qPOTS : Batch Pareto optimal Thompson sampling

Our approach extends Thompson sampling to the multiobjective setting. *Whereas TS makes decisions according to the probability that they are optimal, qPOTS makes decisions according to the probability that they are Pareto optimal.*

$$p_{\mathcal{X}^*}(\mathbf{x}) = \int \delta\left(\mathbf{x} - \arg\max_{\mathbf{x}\in\mathcal{X}} \{Y_1(\mathbf{x}), \ldots, Y_K(\mathbf{x})\}\right) p(Y_1|\mathcal{D}_n)\ldots p(Y_K|\mathcal{D}_n)dY_1\ldots dY_K.$$

- In practice, we fit K GP models and draw K GP posterior sample paths $\{Y_1(\cdot,\omega), \ldots, Y_K(\cdot,\omega)\}$.
- Then, solve the following **cheap** multiobjective optimization problem using, e.g., **evolutionary approaches**, to sample new acquisitions.

$$\mathbf{x}_{n+1} \in X^* = \arg\max_{\mathbf{x}\in\mathcal{X}} \{Y_1(\mathbf{x},\omega), \ldots, Y_K(\mathbf{x},\omega)\}. \tag{3}$$

- Thereby we blend EAs with BO.

Extension to constrained problems

Handling constraints is straightforward in qPOTS.

- Let there be a total of C constraints $\{c_1, \ldots, c_C\}$, and $\mathcal{I} \cup \mathcal{E} = [C]$.
- We fit a total of $K + C$ posterior GPs $Y_1, \ldots, Y_K, Y_{K+1}, \ldots, Y_{K+C}$.
- For all $i \in [C]$, we define the constrained qPOTS to choose points according to

$$p_{\mathcal{X}^*}(\mathbf{x}) = \int \delta\Big(\mathbf{x} - \arg\max_{\mathbf{x}\in\mathcal{X}} [\{Y_1(\mathbf{x}), \ldots, Y_K(\mathbf{x})\} \times \mathbb{1}_{\{\bigcap_i Y_{k+i\in\mathcal{E}}(\mathbf{x})=0\}} \times \mathbb{1}_{\{\bigcap_i Y_{k+i\in\mathcal{I}}(\mathbf{x})\geq 0\}}]\Big) p(Y_1|\mathcal{D}_n)\ldots p(Y_K|\mathcal{D}_n)dY_1\ldots dY_K,$$

- where $\mathbb{1}$ denotes the indicator function for feasibility.

A quick demo

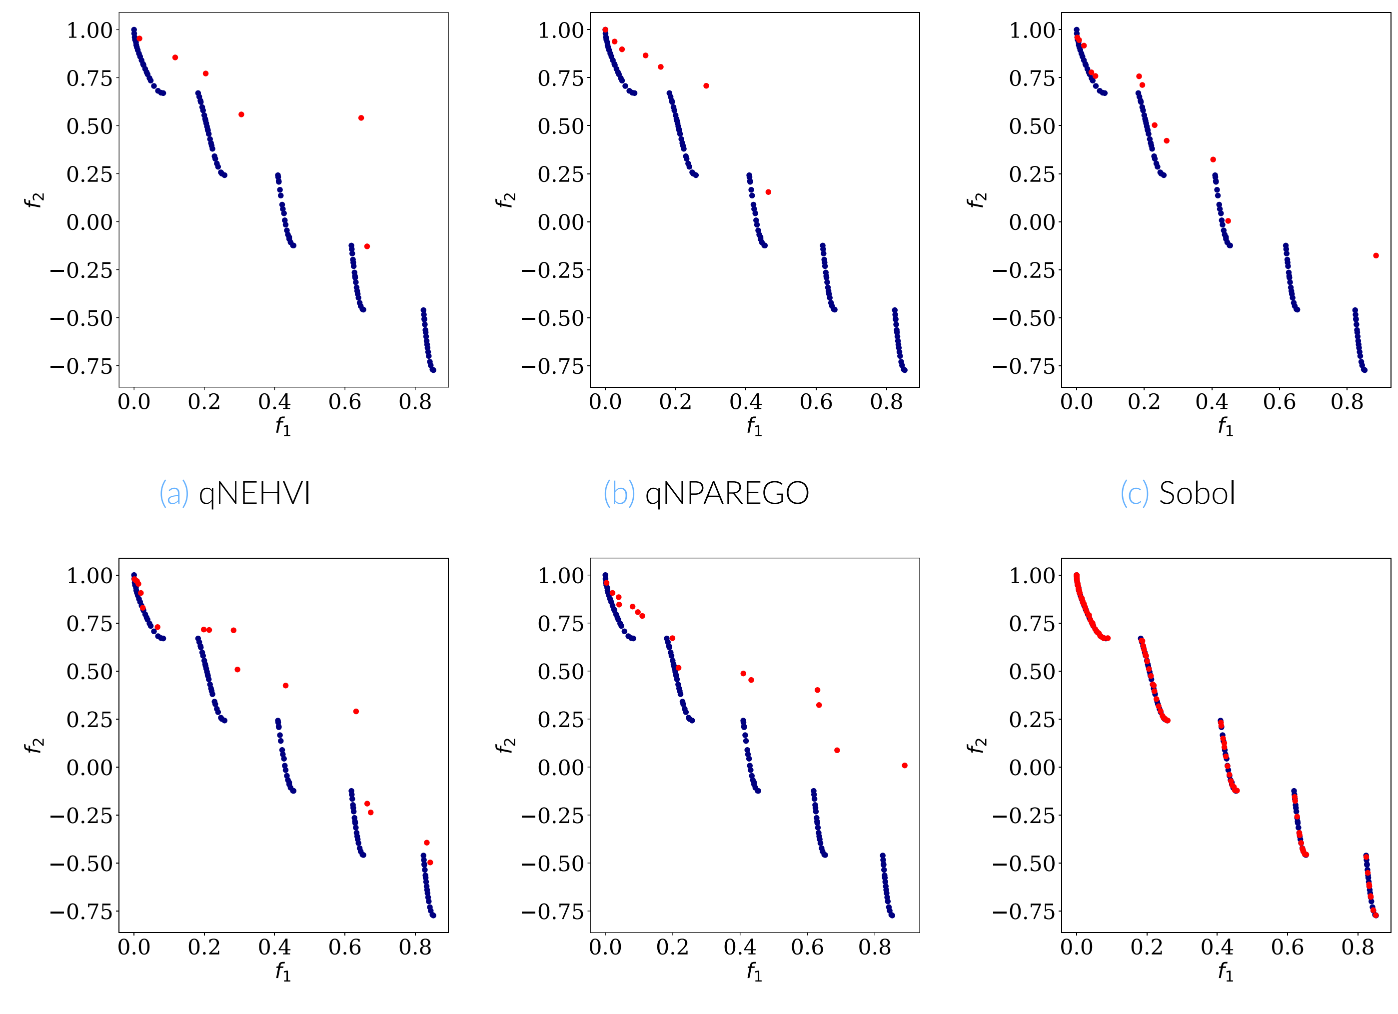

(a) qNEHVI (b) qNPAREGO (c) Sobol (d) PESMO (e) MESMO (f) qPOTS (Ours)

Figure 1. Demonstration on the ZDT3 ($K = 2$) test function that has disjoint Pareto frontiers. Navy are the true Pareto frontier and red are the predicted Pareto frontier. With 20 seed points and only an additional 204 acquisitions, in batches of $q = 4$, notice that qPOTS is quickly able to resolve the true Pareto frontier while other methods in the state of the art struggle to come any close.

Experiments

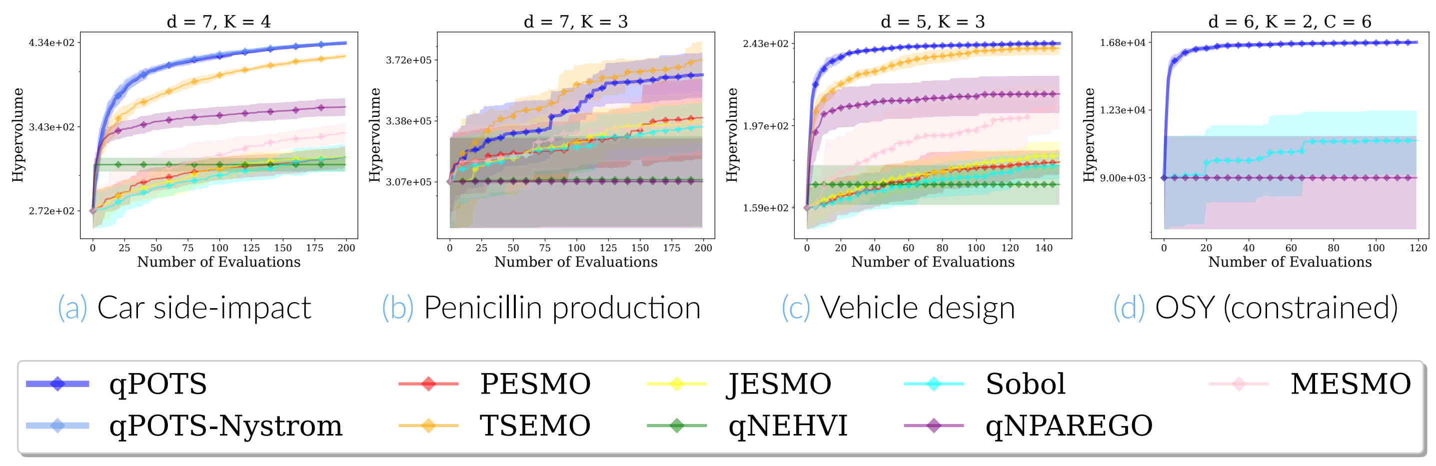

(a) Car side-impact (b) Penicillin production (c) Vehicle design (d) OSY (constrained)

Figure 2. **Sequential acquisition.** Hypervolume Vs. iterations for sequential ($q = 1$) acquisition; plots show mean and ± 1 standard deviation out of 10 repetitions. qPOTS outperforms all competitors or is amongst the best. Bottom right shows constrained case on the OSY problem Osyczka and Kundu [1995]. Additional experiments are provided in the supplementary material.

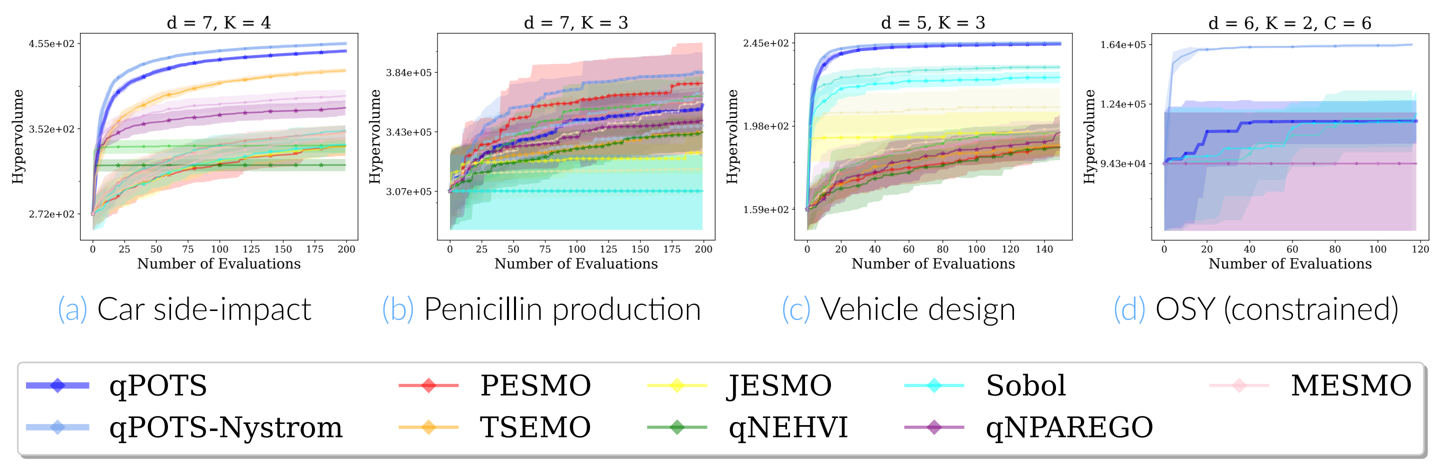

(a) Car side-impact (b) Penicillin production (c) Vehicle design (d) OSY (constrained)

Figure 3. **Batch acquisition.** Hypervolume Vs. iterations for batch ($q > 1$) acquisition; plots show mean and ± 1 standard deviation out of 10 repetitions. qPOTS outperforms all competitors, but the benefit is more pronounced in the batch case. Additional experiments, including constrained problems, are shown in the supplementary material.

Conclusions

- qPOTS presents a new method for multiobjective Bayesian optimization with several improvements over the state of the art.

(a) Paper (b) Software