

Bayesian Optimization with Preference Exploration using a Monotonic Neural Network Ensemble

Hanyang Wang ¹ Juergen Branke ¹ Matthias Poloczek ²


¹University of Warwick

²Amazon

Multi-objective Bayesian Optimiziation with Posterior Selection

WARWICK

In Multi-objective Bayesian Optimization with Posterior Selection, a Decision Maker chooses the point that maximizes their posterior expected utility.

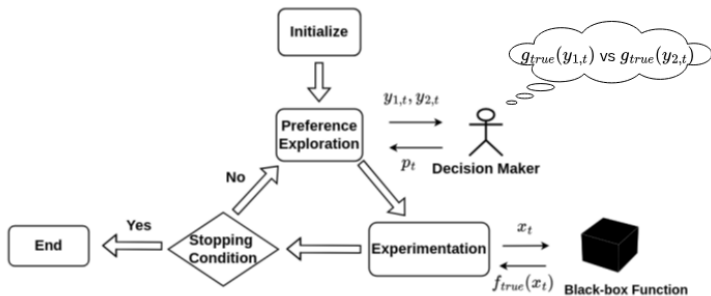


Bayesian Optimization with Preference Exploration

The goal is to find

$$x^* = \arg \max_{x \in \mathcal{X}} g_{\text{true}}(f_{\text{true}}(x)),$$

where $f_{\text{true}} : \mathbb{R}^d \rightarrow \mathbb{R}^k$ maps inputs to objectives, and $g_{\text{true}} : \mathbb{R}^k \rightarrow \mathbb{R}$ denotes the decision maker's utility.



- ▶ The utility function is typically **monotonically increasing** : For example, both a higher treatment effect and a higher tolerability indicate a better drug.
- ▶ Standard GP models **ignore this monotonicity**.
- ▶ We propose a model called **Monotonic Neural Network Ensemble** (MoNNE) that exploits this property to improve BO performance.




Fig. 1: Models are trained with pairwise comparisons.

Experimental Results

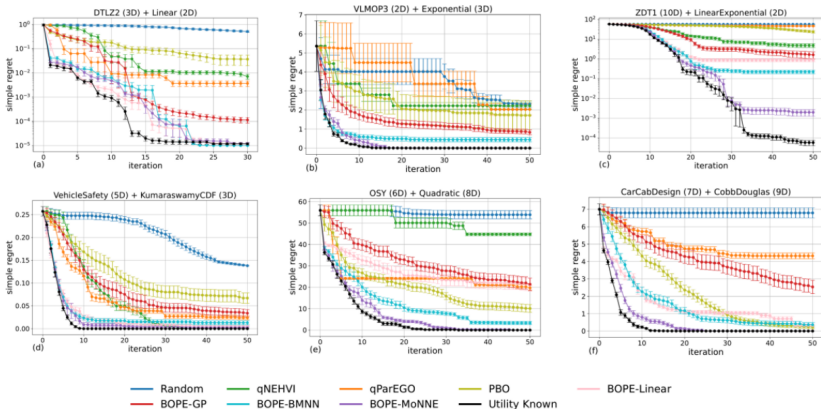


Fig. 2: BOPE-MoNNE Achieves the Best Performance Among All Benchmark Algorithms.

By incorporating monotonicity information into the surrogate model, our algorithm achieves superior performance.

For more details, visit our poster:

Location: Exhibit Hall C, D, E

Time: Fri, Dec 5 12:30–3:30 a.m. GMT

Thanks!