An Entropy Search Portfolio for Bayesian Optimization

Bobak Shahriari*, Ziyu Wang†, Matthew W. Hoffman‡, Alexandre Bouchard-Côté*, Nando de Freitas*†§◦

*University of British Columbia, †University of Oxford, ‡University of Cambridge, §CIFAR, ◦Google DeepMind

Bayesian optimization

Setting: Find $\mathbf{x}^\star \in \arg\max_{\mathbf{x}\in\mathcal{X}} f(\mathbf{x})$ by sequentially

- selecting inputs $\mathbf{x}_t$ and
- observing a random output Y_t such that $f(\mathbf{x}_t) = \mathbb{E}[Y_t|\mathbf{x}_t]$.

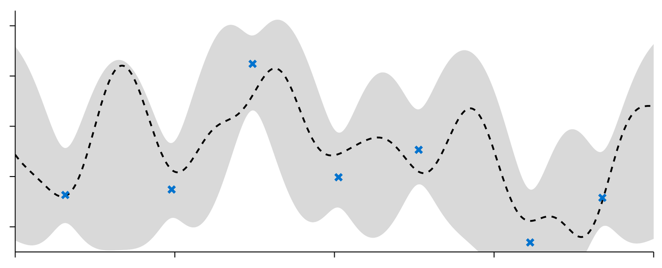

Given observations $\mathcal{D}_t = (\mathbf{x}_{1:t}, y_{1:t})$ we can construct a Bayesian posterior over the unknown function f. This posterior is then used to select $\mathbf{x}_{t+1}$ by maximizing a so-called acquisition function. Finally we make a recommendation $\tilde{\mathbf{x}}_T$.

The acquisition function α is function that is derived from the posterior at each time step. It is designed to trade-off exploration and exploitation.

Problem: No one acquisition function outperforms all others in all instances and it is difficult to know *a priori* which one to use.

Visualization of the problem

Expected improvement [3]:

$$\alpha_t(\mathbf{x}) = \mathbb{E}\big[f(\mathbf{x}) - f(\mathbf{x}_t^+)\big]$$

where $\mathbf{x}_t^+$ is the current best observed input.

Probability of improvement [3]:

$$\alpha_t(\mathbf{x}) = \mathbb{P}\big[f(\mathbf{x}) - f(\mathbf{x}_t^+)\big]$$

Thompson sampling [4] optimizes a sample from the GP posterior. We use a random set of features to fix a sample $\tilde{f}$ to optimize [2]:

$$\alpha_t(\mathbf{x}) = \tilde{f}(\mathbf{x}) = \boldsymbol{\Phi}(\mathbf{x})^\top \mathbf{w}$$

where $\mathbf{w}$ are random weights and $\boldsymbol{\Phi}(\mathbf{x})$ consists of random Fourier features evaluated at $\mathbf{x}$.

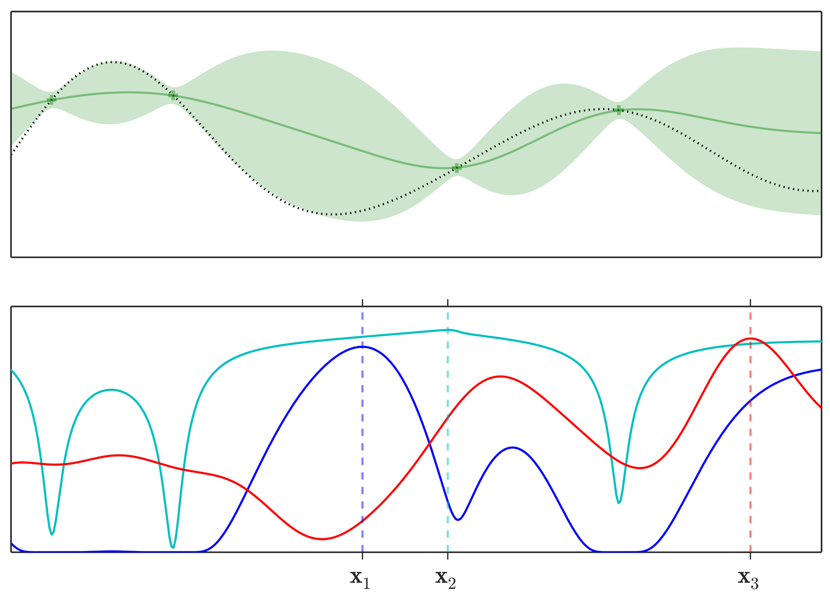

The three acquisition functions behave differently; it is unclear which will produce the best final recommendation.

Portfolio Bayesian optimization

We address this issue by considering a portfolio of acquisition functions and select between them using a information-based meta-criterion.

for $t = 1, \ldots, T$ **do**
  collect candidates $\{\mathbf{x}_t^k = \arg\max_{\mathbf{x}} \alpha_{t-1}^k(\mathbf{x})\}_{k=1}^K$ — consider a candidate from each acquisition function
  select point $\mathbf{x}_t = \arg\max_{\mathbf{x}_t^{1:K}} u_{t-1}(\mathbf{x})$ — pick candidate with highest meta-criterion
  observe $y_t \sim p(\cdot|\mathbf{x}_t)$ — observe data then update the model
  update model posterior $p(f|\mathcal{D}_{t-1})$
end for
return $\tilde{\mathbf{x}}_T$

Note: The $t-1$ subscript on α and u indicate an implicit dependence on the current posterior. We write $\mathbf{x}_k := \mathbf{x}_t^k$ when the time index is superfluous.

Entropy Search Portfolio

Let $\mathbf{x}^\star$ be a random variable whose distribution (induced by the GP posterior) represents our represents our belief on the location of the maximizer. We propose to use the expected information gained about $\mathbf{x}^\star$ as a meta-policy [5, 1].

$$u(\mathbf{x}_k|\mathcal{D}) = \mathbb{E}_{p(y|\mathcal{D},\mathbf{x}_k)}\Big[H\big[\mathbf{x}^\star|\mathcal{D}\cup\{\mathbf{x}_k, y\}\big]\Big]$$

Entropy estimation

Strategy: First discretize the distribution and compute the discrete Shannon entropy.

We sample the discretization $\mathbf{z}^{(i)} \sim p(\mathbf{x}^\star|\mathcal{D})$ approximately, using Thompson sampling [2].

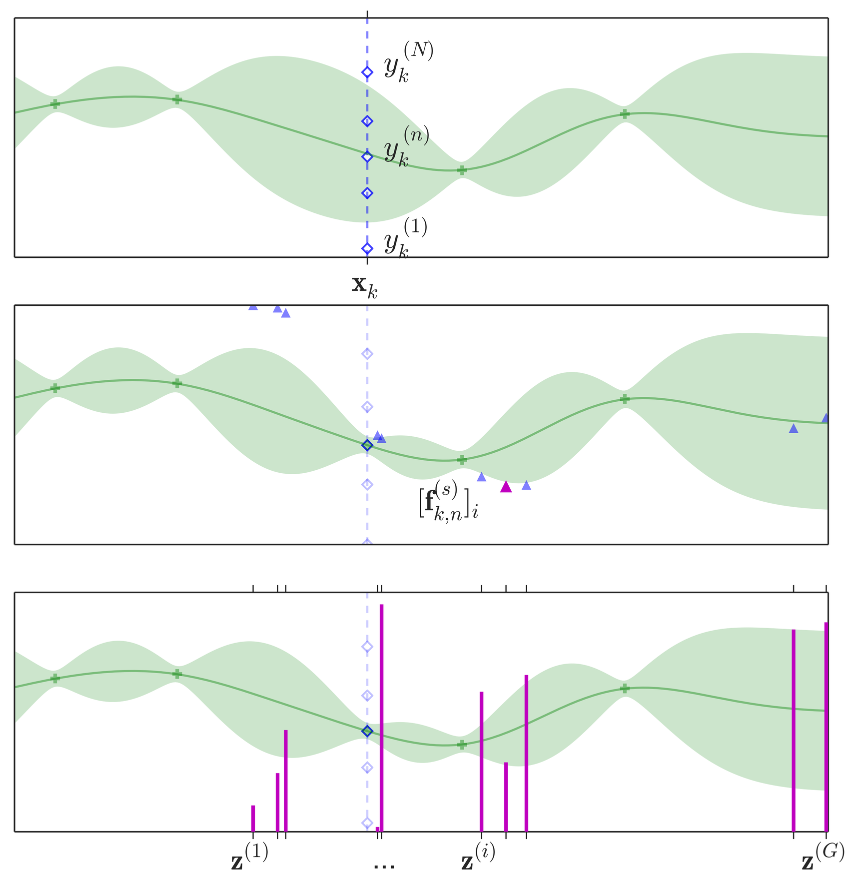

- For each candidate $\mathbf{x}_k$, exactly simulate
$$y_k^{(n)} \sim p(y|\mathcal{D}, \mathbf{x}_k).$$
- For each simulation $y_k^{(n)}$ (◇), condition and sample the resulting GP at the points $\mathbf{z}^{(1:G)}$
$$\mathbf{f}_{k,n}^{(s)} \sim p(\mathbf{f}|\mathcal{D}\cup(\mathbf{x}_k, y_k^{(n)}), \mathbf{z}^{(1:G)}).$$
- For all the samples $\mathbf{f}_{k,n}^{(s)}$ (▲) we collect the maxima (▲) to produce a discrete distribution $\hat{\mathbf{p}}_{kn}$ (bars).
- Compute the meta-criterion for candidate $\mathbf{x}_k$
$$u(\mathbf{x}_k|\mathcal{D}) \approx -\frac{1}{N}\sum_{n=1}^{N}\sum_{i=1}^{G} \hat{p}_{ikn} \log \hat{p}_{ikn}.$$

Experiments

We evaluated ESP on 6 experiments. We compare ESP to the baseline random portfolio (RP), which selects among the candidates uniformly at random, and GPHedge, which uses past performance to establish the quality of acquisition functions.

The dashed lines represent portfolios with 9 additional random selectors—acquisition functions that select points in the space uniformly at random. By polluting the portfolio with random candidates, we compare the robustness of the portfolios to poor acquisition functions.

We point out that:

- no acquisition function was best in all experiments;
- even RP performs well in practice;
- ESP often outperforms both other portfolios and its own constituent acquisition functions;
- ESP is more robust to the addition of poor acquisition functions.

Performance of ESP: median over 20 runs.

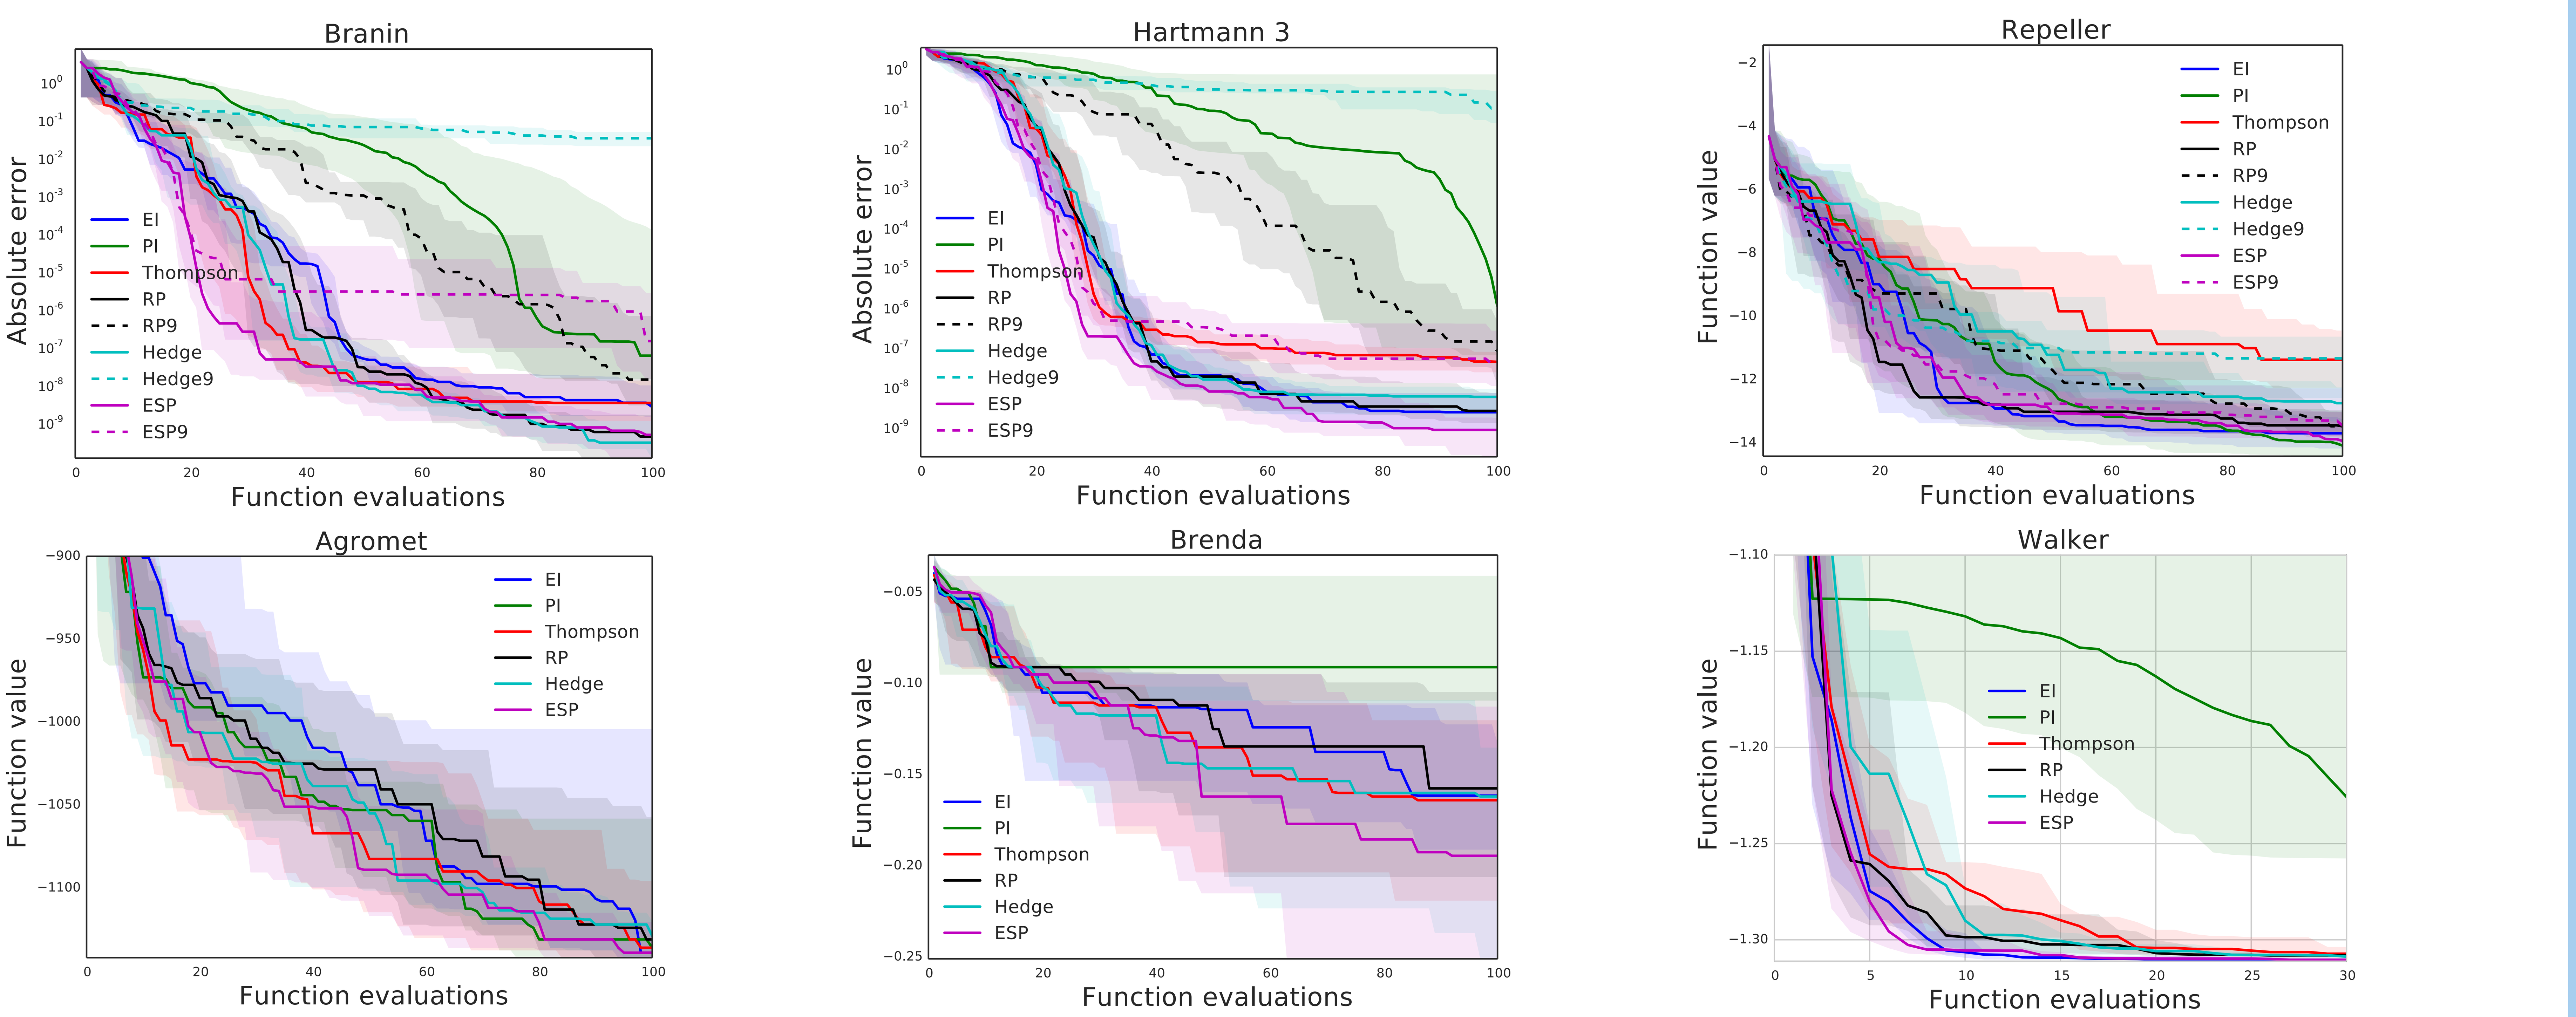

References

[1] P. Hennig and C. Schuler. *Entropy search for information-efficient global optimization. The Journal of Machine Learning Research*, 98888:1809–1837, 2012.

[2] J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani. *Predictive entropy search for efficient global optimization of black-box functions.* In *Advances in Neural Information Processing Systems*. 2014.

[3] D. Jones, M. Schonlau, and W. Welch. *Efficient global optimization of expensive black-box functions. J. of Global optimization*, 13(4):455–492, 1998.

[4] W. R. Thompson. *On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika*, 25(3/4):285–294, 1933.

[5] J. Villemonteix, E. Vazquez, and E. Walter. *An informational approach to the global optimization of expensive-to-evaluate functions. J. of Global Optimization*, 44(4):509–534, 2009.