

ProBO: Versatile Bayesian Optimization Using Any Probabilistic Programming Language

W. Neiswanger et al. 2019

Paper review by Sean Parker

Structure

- Background / Motivation
- Overview of ProBO
- Key contributions
- Experiments & Evaluation
- Review

Bayesian Optimization (BO)

- Aim: Optimize the function $f(x)$
- Restricted to sampling the function at points x
- *Surrogate model* used to approximate objective function
- Uses *acquisition function* to sample areas of interest
 - MPI, EI, UCB, TS

Probabilistic Programming Languages (PPLs)

- Often built upon existing languages
 - PyMC3/PyMC4 (Python)
 - Edward (Tensorflow)
 - Pyro (PyTorch)
- Each PPL uses a different inference strategy + posterior representations
 - MCMC, SMC
 - VI, EI

Probabilistic Programming Languages (PPLs)

- Domain-specific languages
- Inference on probabilistic models
- Assumptions encoded over variables of the model
- Output: Probability Distribution

PPL: Example

Coin toss

- Calculate the bias of a coin:
 - Bernoulli distribution with latent variable θ
 - $P(x_i = 1 | \theta) = \theta$ and $P(x_i = 0 | \theta) = 1 - \theta$
 - Infer θ based on previous results of coin toss - $P(\theta | x_1, x_2, \dots, x_N)$

Motivation

- Models built in PPL is optimised using BO techniques in that PPL
- BOPP - BO in specific PPL to estimate latent variables
- BOAT - Custom framework, uses exact inference & expected improvement

Key contributions

- General abstraction for PPL programs
- ProBO system implementation*
- Evaluation of ProBO using BO models, implemented in various PPLs

Probabilistic Programs Abstraction

- Three core PPL operations:
 - `inf(D)` - returns `post` (PPL dependent)
 - `post(s)` - returns a sample from the posterior distribution
 - `gen(x, z, s)` - returns sample from generative distribution

ProBO Algorithm

Algorithm 1 ProBO($\mathcal{D}_0, \text{inf}, \text{gen}$)

```
1: for  $n = 1, \dots, N$  do
2:    $\text{post} \leftarrow \text{inf}(\mathcal{D}_{n-1})$ 
3:    $x_n \leftarrow \text{argmin}_{x \in \mathcal{X}} a(x, \text{post}, \text{gen})$ 
4:    $y_n \sim s(x_n)$ 
5:    $\mathcal{D}_n \leftarrow \mathcal{D}_{n-1} \cup (x_n, y_n)$ 
6: Return  $\mathcal{D}_N$ .
```

- ▷ Run inference algorithm to compute post
- ▷ Optimize acquisition using post and gen
 - ▷ Observe system at x_n
 - ▷ Add new observations to dataset

- Goal: Return $x^* = \text{argmin}_{x \in \mathcal{X}} \mathbb{E}_{y \sim s(x)} [f(y)]$
- Algorithm:
 - Invoke the PPLs inference procedure via $\text{inf}()$
 - Get new x by optimising acquisition function
 - Observe system at x
 - Add new observation to dataset

ProBO - Computation Cost

- `inf()` cost dependent on PPLs inference algorithm
 - e.g. MCMC algorithms - $O(n)$ per iteration
- `inf()` only executed **once per query**
- Acquisition optimisation executed 100s times per query
 - `post()` & `gen()` cheaply implemented - $O(1)$

Acquisition function optimisation

- $\text{post}()$ & $\text{gen}()$ not analytically differentiable
- Authors explored zeroth-order optimisation of a_{MF}
 - $\text{post}()$ & $\text{gen}()$ called M_f times
- Any zeroth-order optimisation algorithm can be used

Algorithm 6 $a_{MF}(x, \text{post}, \text{gen})$

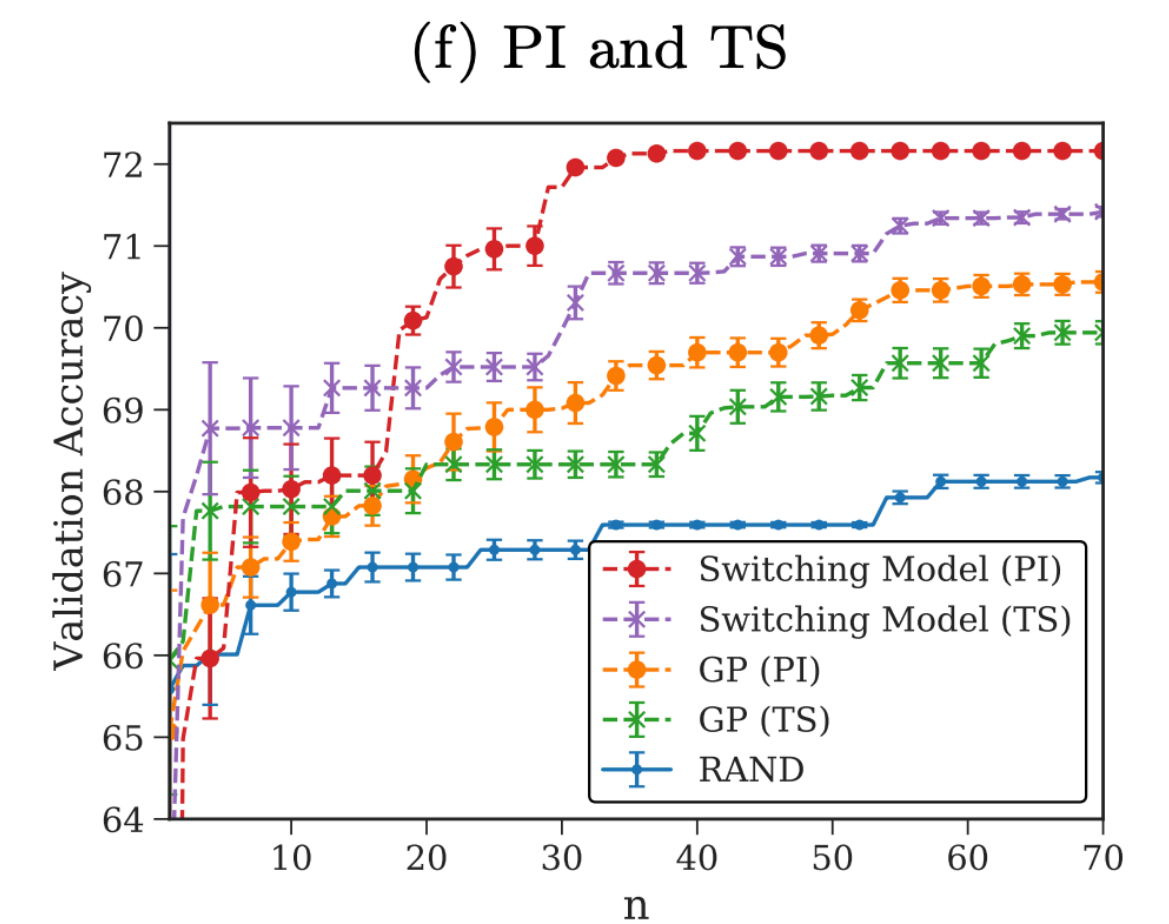
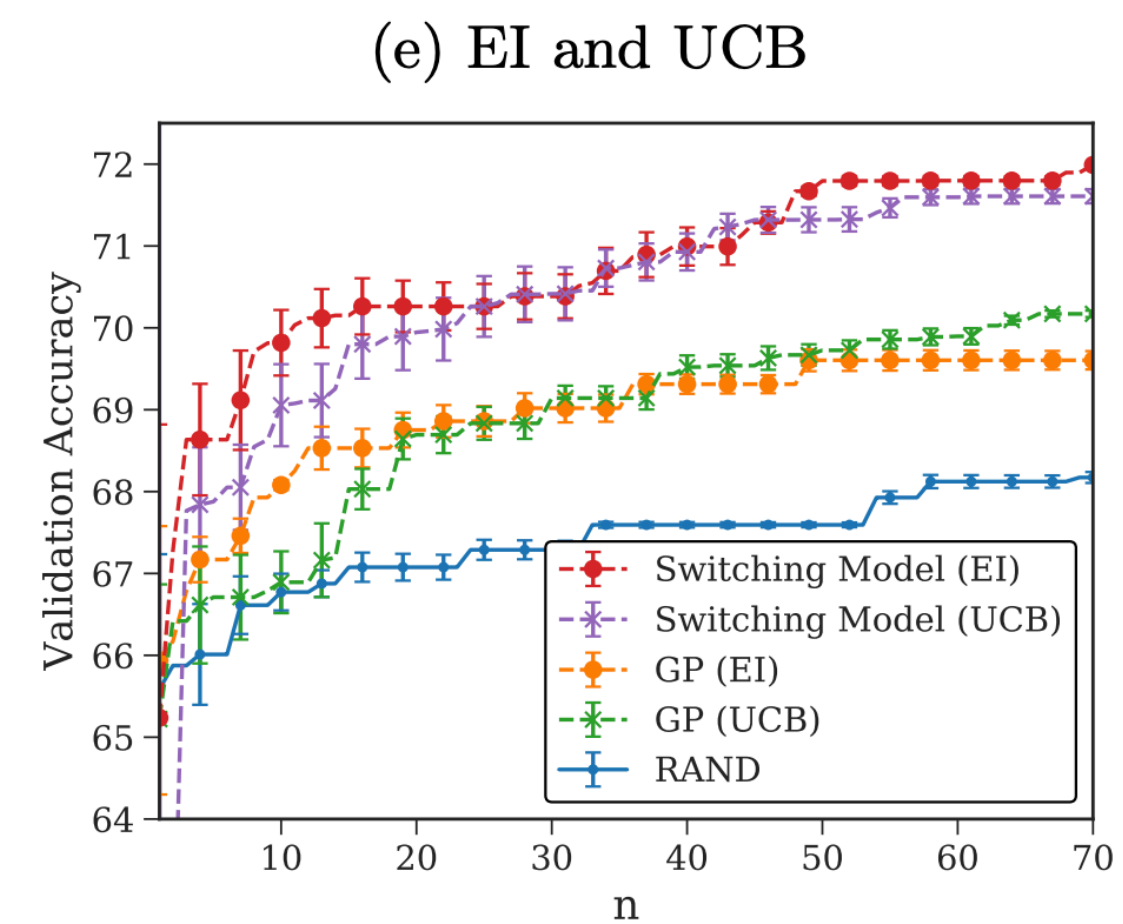
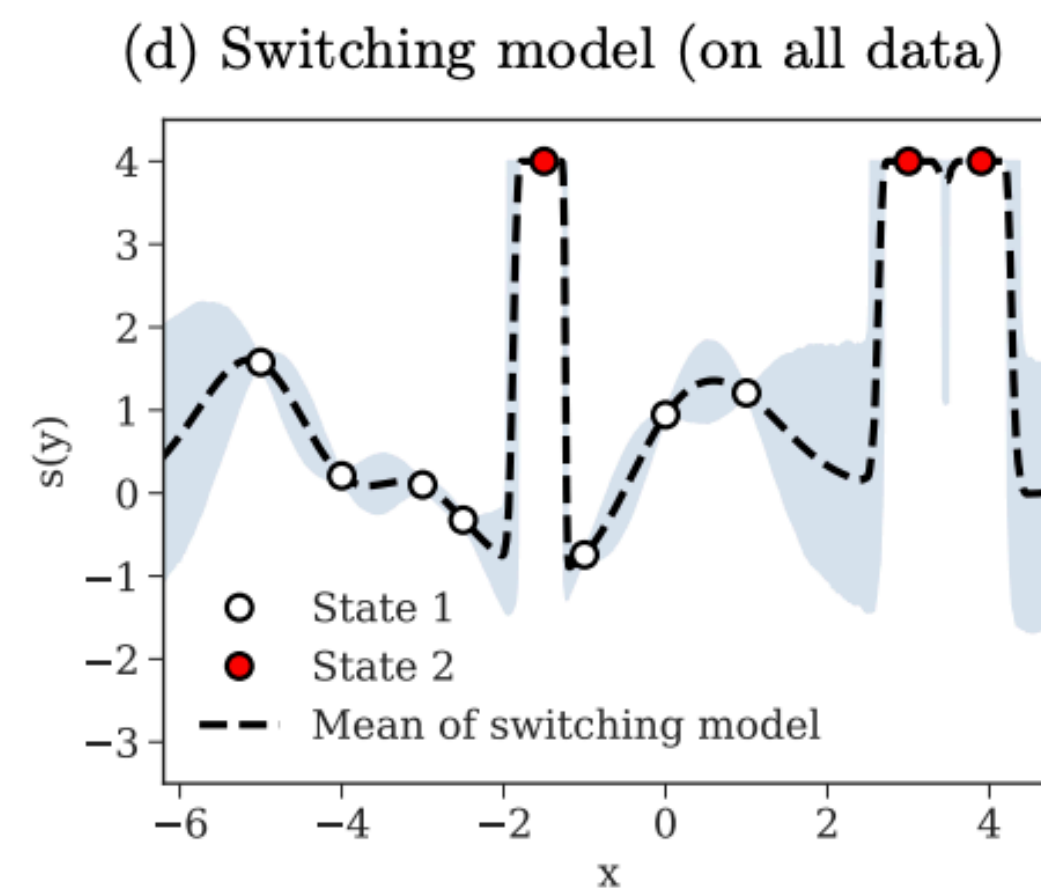
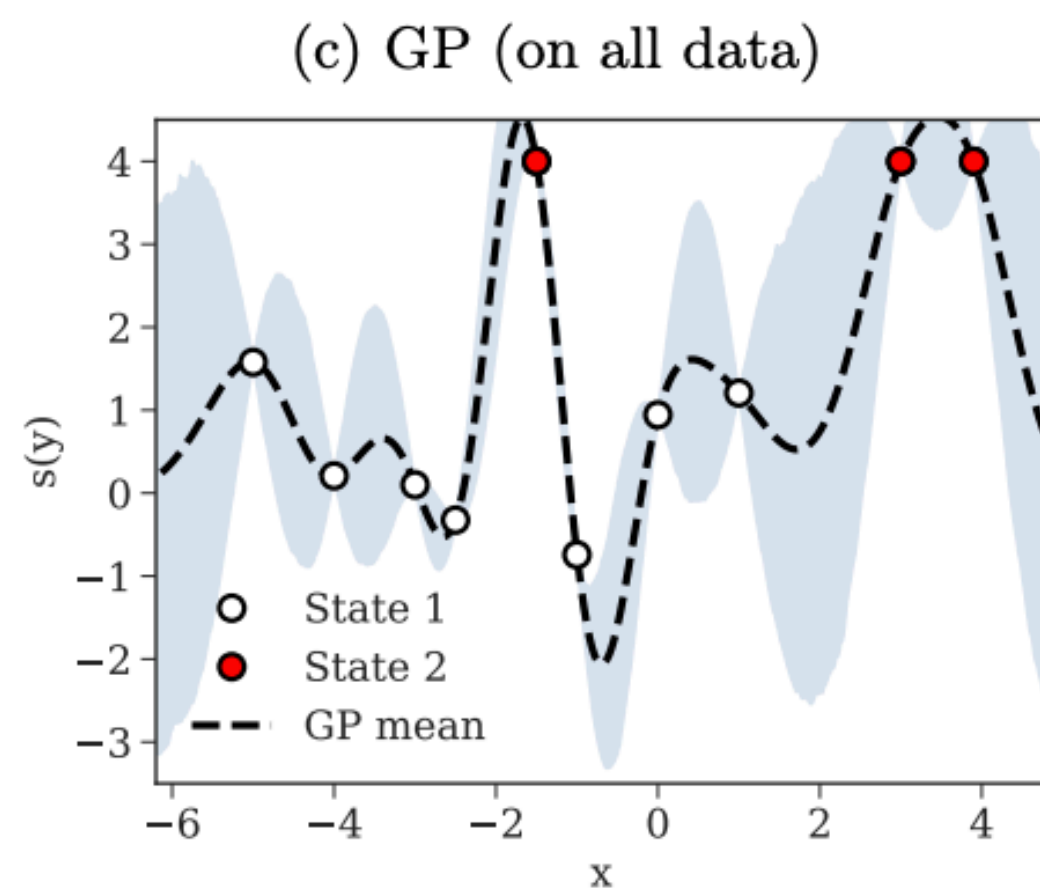
```
1:  $a_{\min} \leftarrow$  Min value of  $a$  seen so far
2:  $\ell = -\infty, f = 1$ 
3: while  $\ell \leq a_{\min}$  do
4:    $\ell \leftarrow$  LCB-bootstrap( $\text{post}, \text{gen}, M_f$ )
5:    $f \leftarrow f + 1$ 
6: Return  $a(x, \text{post}, \text{gen})$  using  $M = M_f$ 
```

Algorithm 7 LCB-bootstrap($\text{post}, \text{gen}, M_f$)

```
1:  $y_{1:M_f} \leftarrow$  Call post and gen  $M_f$  times
2: for  $j = 1, \dots, B$  do
3:    $\tilde{y}_{1:M_f} \leftarrow$  Resample( $y_{1:M_f}$ )
4:    $a_j \leftarrow \lambda(\tilde{y}_{1:M_f})$   $\triangleright$  See text for details
5: Return LCB( $a_{1:B}$ )
```

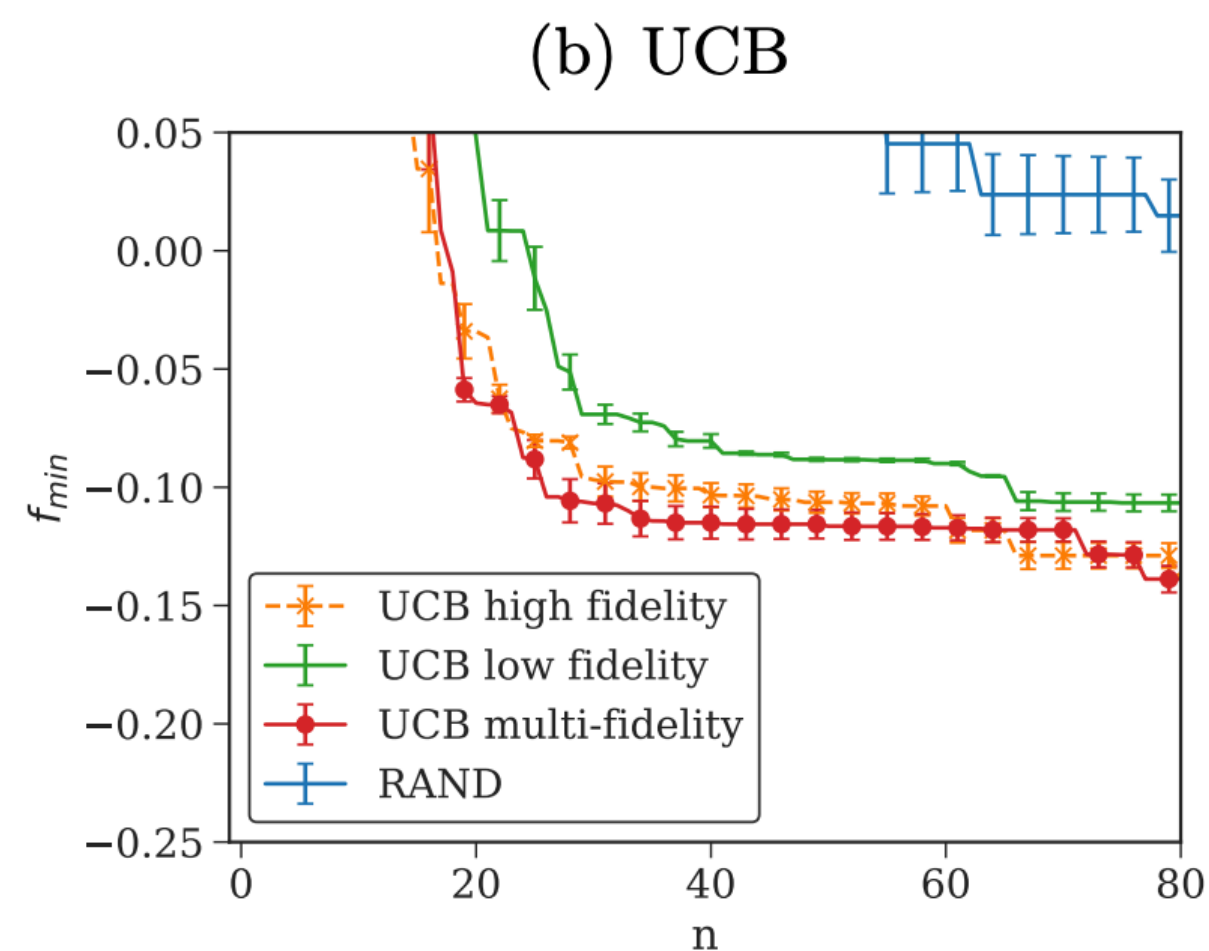
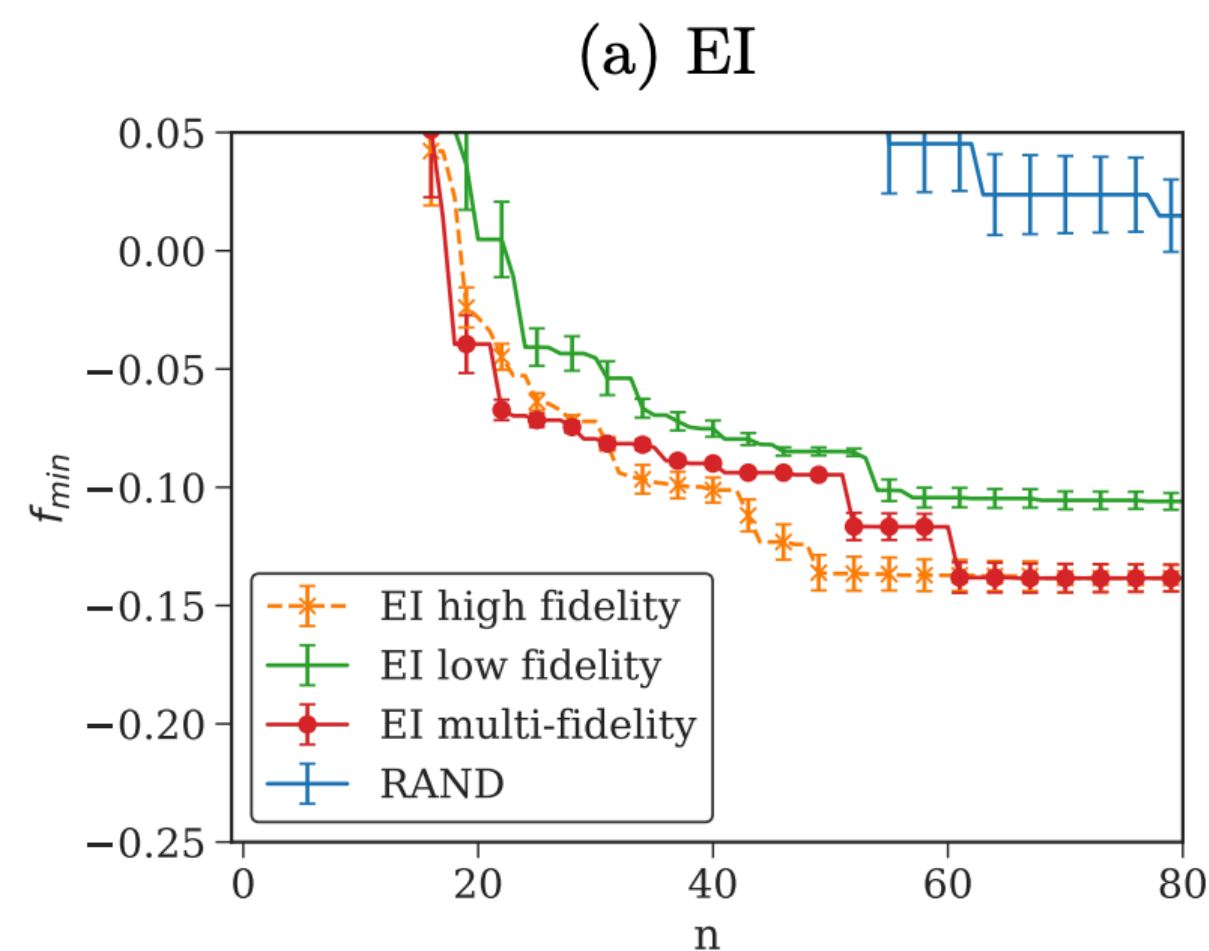
Evaluation

- Optimisation of MLP hyperparameters
 - “Switching model” is ProBO using a dynamic value of M_f



Evaluation

- 3x better performance than high-fidelity in terms of calls to $\text{gen}()$
- High and multi fidelity have comparable performance
- Converges to very similar value



(c) Calls to gen

PPL acquisition method $a(x)$	Avg. number $\text{gen}/a(x)$
EI high-fidelity	1000
EI multi-fidelity	347.89
EI low-fidelity	10
UCB high-fidelity	1000
UCB multi-fidelity	324.65
UCB low-fidelity	10

Review

- Difficult paper to understand
- Implementation of ProBO not provided
 - Questions remain of how ProBO is integrated into existing PPLs
- Good idea for providing uniform way of performing BO across PPLs