ICASSP-2018

CLUSTERING-GUIDED GP-UCB FOR BAYESIAN OPTIMIZATION

Jungtaek Kim and Seungjin Choi

Department of Computer Science and Engineering, POSTECH

POSTECH POHANG UNIVERSITY OF SCIENCE AND TECHNOLOGY

Introduction and Motivation

- Bayesian optimization is a powerful method for finding extrema of an objective function [1, 3].
- One of acquisition functions, GP-upper confidence bound (GP-UCB) determines where next to sample from the true function, balancing exploration and exploitation.
- We first present a geometric interpretation of GP-UCB.
- We develop GP-UCB to clustering-guided method, called as clustering-guided GP-UCB (CG-GPUCB).

Background

- Bayesian optimization [1]
  - ✓ Global optimization for black-box function.

Algorithm 1 Bayesian Optimization with GP regression

Require: Initial data $\mathcal{D}_{1:I} = \{(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_I, y_I)\}$, and $T \in \mathbb{N} > 0$

1: **for** $t = 1, 2, \ldots, T$ **do**

2: Find $\mathbf{x}_{I+t}$ that maximizes the acquisition function over the current GP: $\mathbf{x}_{I+t} = \arg\max_{\mathbf{x}} a(\mathbf{x}|\mathcal{D}_{1:I+t-1})$.

3: Sample the objective function: $y_{I+t} = f(\mathbf{x}_{I+t}) + \epsilon_{I+t}$.

4: Augment the data: $\mathcal{D}_{1:I+t} = \{\mathcal{D}_{1:I+t-1}, (\mathbf{x}_{I+t}, y_{I+t})\}$.

5: Update the GP, computing $\mu_{I+t}(\mathbf{x}), \sigma^2_{I+t}(\mathbf{x})$:

6: **end for**

7: **return** $\mathbf{x}^* = \arg\max_{\mathbf{x} \in \{\mathbf{x}_1, \ldots, \mathbf{x}_{I+T}\}} \mu_{I+T}(\mathbf{x})$

- GP-UCB [2]
  - ✓ Linear combination of posterior mean and standard deviation function.
  - ✓ Trade-off hyperparameter controls tightness of confidence bound.

$$a(\mathbf{x}|\mathcal{D}_{1:t}) = -\mu(\mathbf{x}) + \kappa\sigma(\mathbf{x})$$

Selected References

[1] E. Brochu, V. M. Cora, and N. de Freitas, "A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning," 2010, arXiv preprint arXiv:1012.2599.

[2] N. Srinivas, A. Krause, S. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," in Proceedings of the International Conference on Machine Learning (ICML), Haifa, Israel, 2010, pp. 1015–1022.

[3] J. Snoek, H. Larochelle, and R. P. Adams, "Practical Bayesian optimization of machine learning algorithms," in Advances in Neural Information Processing Systems (NIPS), Lake Tahoe, NV, USA, 2012, vol. 25, pp. 2951–2959.

Clustering-Guided GP-UCB

- Geometric view of GP-UCB

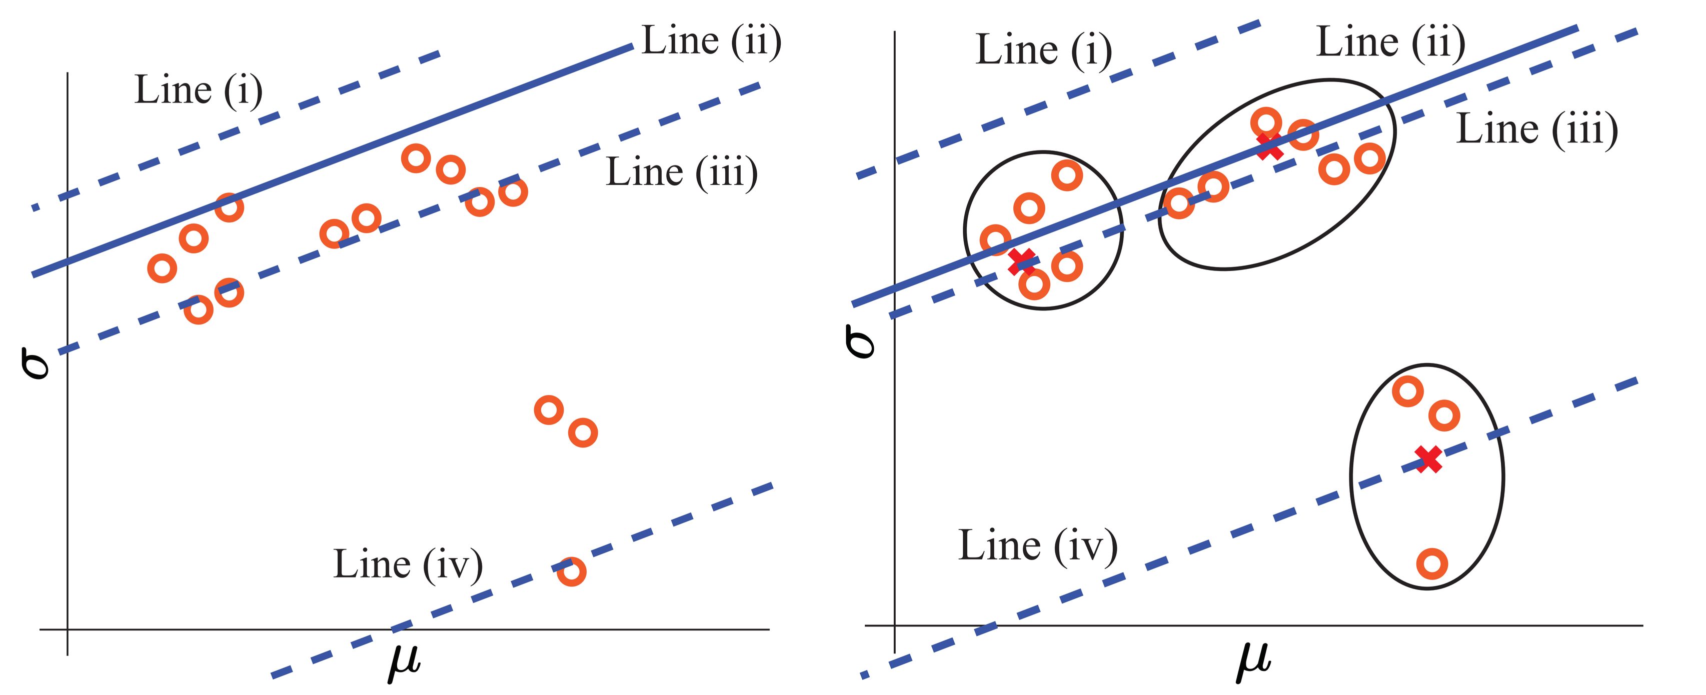

$$\sigma(\mathbf{x}) = -\frac{1}{\kappa}\left(\mu(\mathbf{x}) + a(\mathbf{x}|\mathcal{D}_{1:t})\right)$$

- Proposed method: Clustering-guided GP-UCB

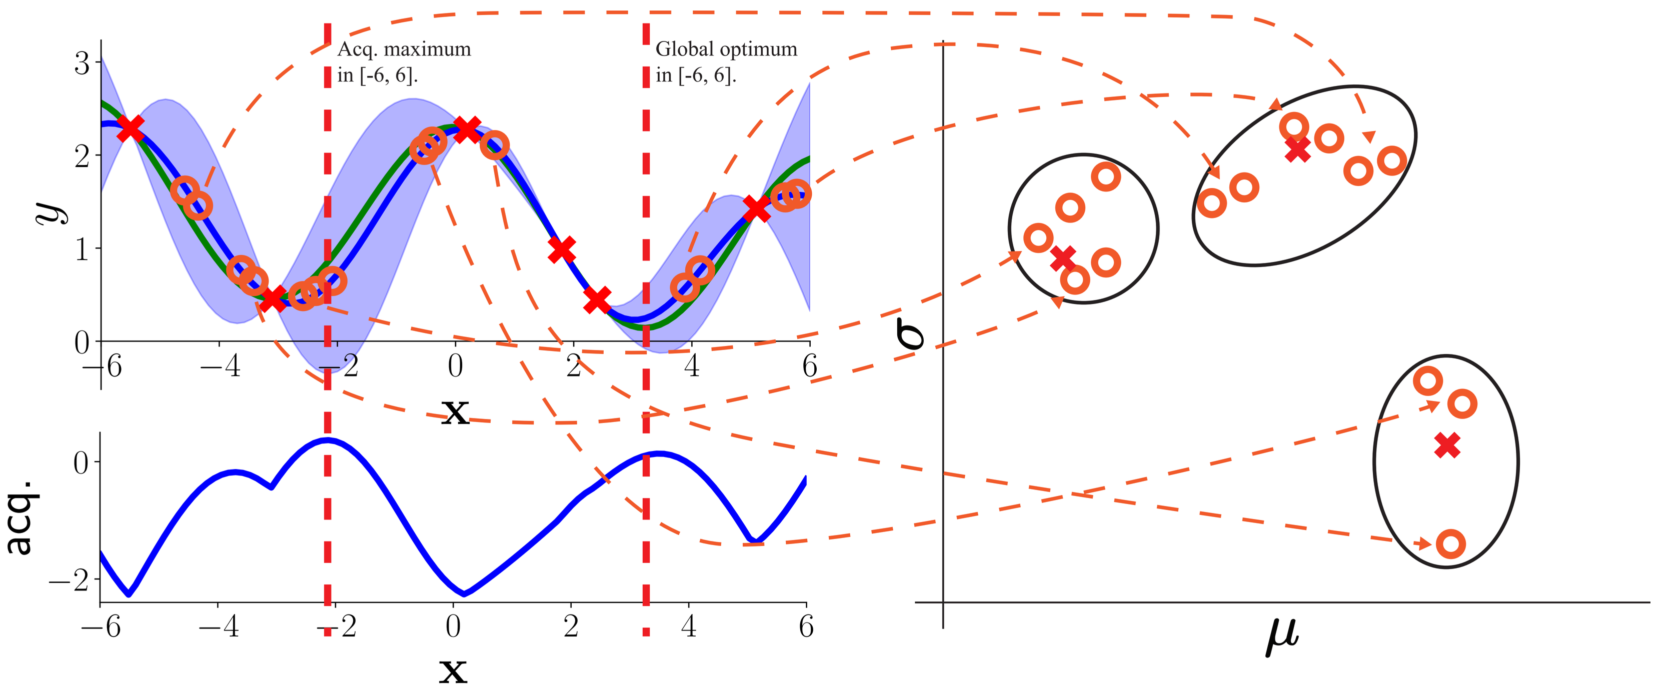

- ✓ The points on posterior mean-standard deviation space are grouped, and a query point is selected on the chosen cluster.
- ✓ We select a query point using one of two criteria, CG-GPUCB-NN and CG-GPUCB².

$$\mathbf{x}_{t+1} = \arg\min_{\mathbf{x} \in \mathcal{C}_{i^*}} \|[\mu(\mathbf{x}), \sigma(\mathbf{x})]^\top - \mathbf{c}_{i^*}\|_2^2$$

$$\mathbf{x}_{t+1} = \arg\max_{\mathbf{x} \in \mathcal{C}_{i^*}} a(\mathbf{x}|\mathcal{D}_{1:t})$$

Algorithm 2 Bayesian Optimization with CG-GPUCB

Require: Initial data $\mathcal{D}_{1:I} = \{(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_I, y_I)\}$, $T \in \mathbb{N} > 0$, and $K \in \mathbb{N} > 0$ (number of clusters)

1: **for** $t = 1, 2, \ldots, T$ **do**

2: Calculate centers $\mathbf{c}_i$ of K clusters determined in the μ-σ space, given the current GP.

3: Find the best cluster $\mathcal{C}_{i^*}$ via (9).

4: Find $\mathbf{x}_{I+t}$ by (10) or (11).

5: Sample the objective function: $y_{I+t} = f(\mathbf{x}_{I+t}) + \epsilon_{I+t}$.

6: Augment the data: $\mathcal{D}_{1:I+t} = \{\mathcal{D}_{1:I+t-1}, (\mathbf{x}_{I+t}, y_{I+t})\}$.

7: Update the GP via (4) & (5).

8: **end for**

Experimental Results

- Our method was conducted on synthetic function, benchmarks, and hyperparameter optimization (HPO).
- We showed the effect of hyperparameters for CG-GPUCB-NN and CG-GPUCB².
- HPO for logistic regression (LR) and deep convolutional networks (DCN) was tested.

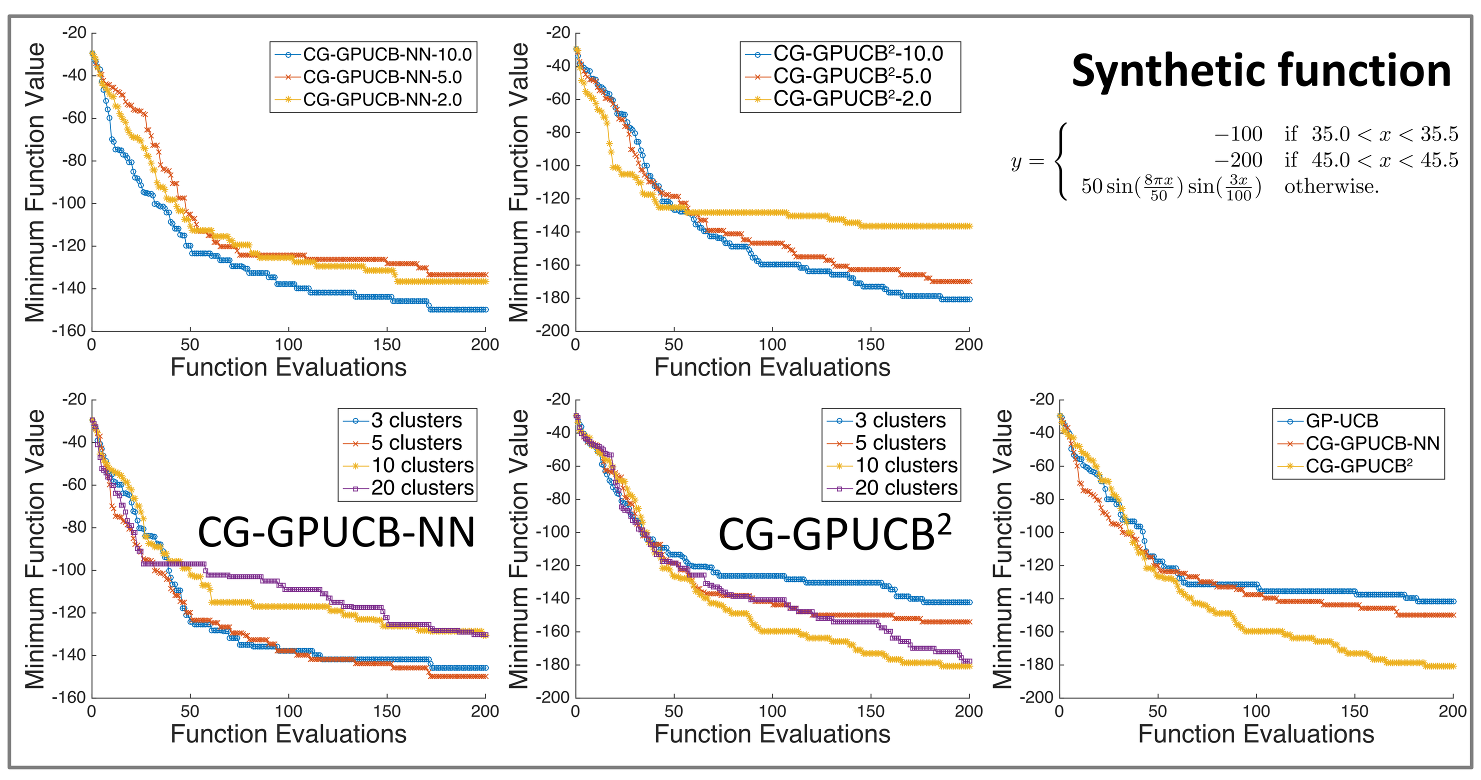

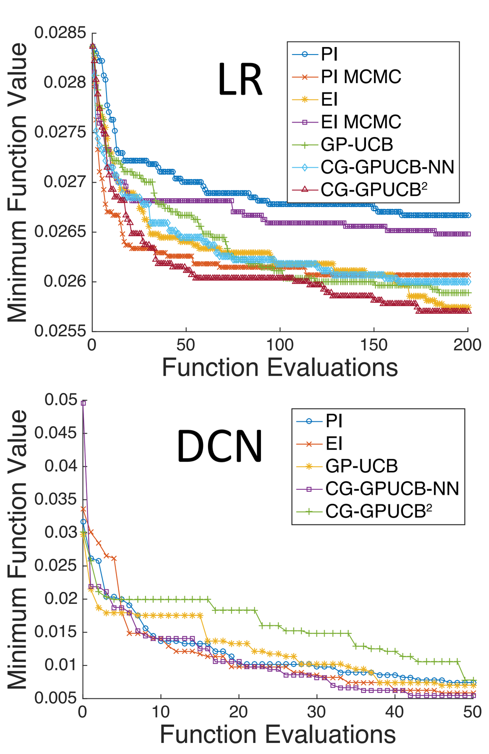

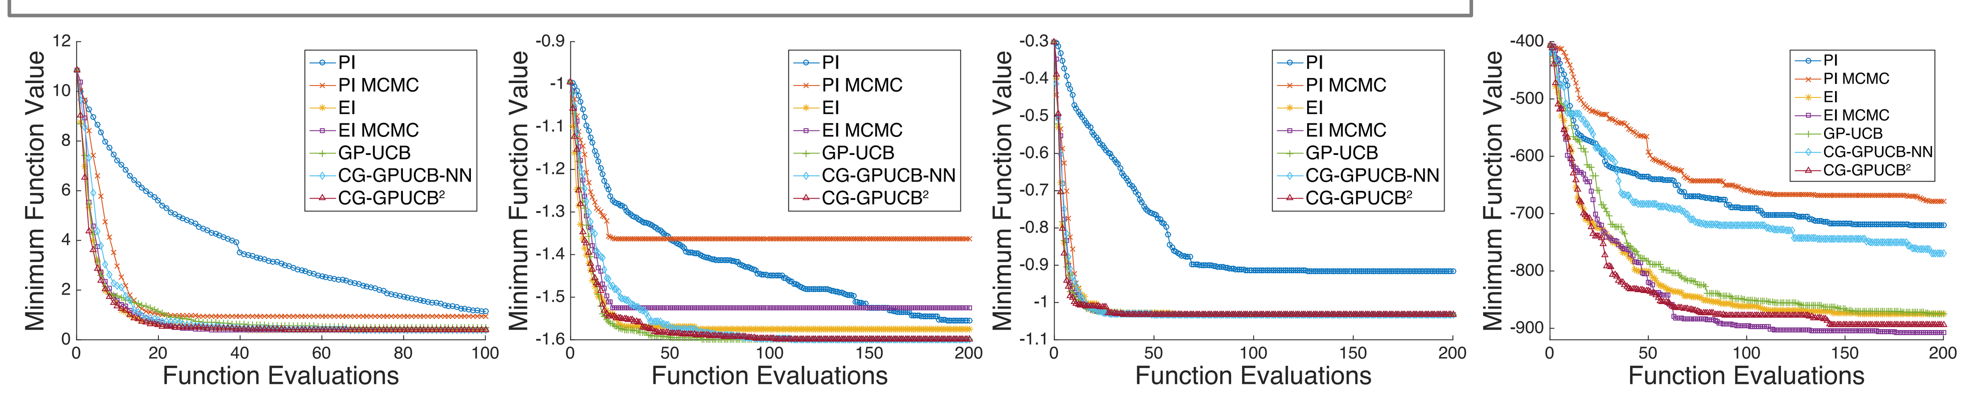

Branin-Hoo Cosines Sixhumpcamel Eggholder

Conclusion

- We presented our own geometric interpretation of GP-UCB and new acquisition function, CG-GPUCB.
- Our method outperformed in non-smooth function with sharp peaks and valleys.