Ada-BKB

Scalable Gaussian Process Optimization on Continuous Domain by Adaptive Discretizations

Marco Rando[1] Luigi Carratino[1] Silvia Villa[2] Lorenzo Rosasco[134]

[1] *MaLGa - DIBRIS, University of Genova, Italy.* [2] *MaLGa - DIMA, University of Genova, Italy.* [3] *CBMM - Massachusetth Institute of Technology, USA.* [4] *Istituto Italiano di Tecnologia, Genova, Italy.*

Problem Setup

Given a compact $X \subset \mathbb{R}^d$ and $f : X \to \mathbb{R}$

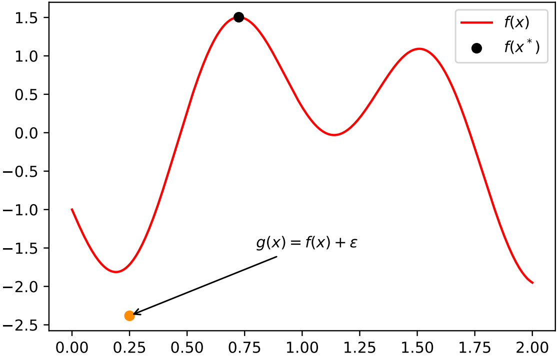

Goal: $x^* \in \arg\max_{x \in X} f(x)$

BKB: Budgeted Kernel Bandit

(Calandriello et al., 2019)

Assume $f \in \mathcal{H}$ and let $\tilde{k}(x, x') = k(x_i, S_t) K_{S_t}^{\dagger} k(S_t, x')$

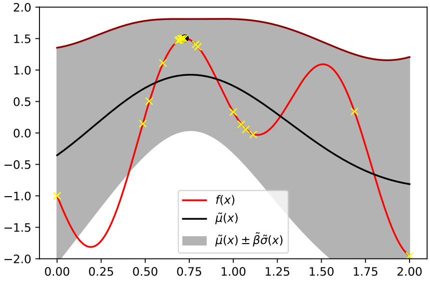

For $t = 1, \cdots, T$

- $x_t = \arg\max_{x \in X} \tilde{\mu}_t(x) + \tilde{\beta}_t \tilde{\sigma}_t(x)$
- $y_t = f(x_t) + \epsilon_t$
- update model

$$\tilde{\mu}_t(x) = \tilde{k}(x, X_t)(\tilde{K} + \lambda I)^{-1} y_t$$

$$\tilde{\sigma}_t(x) = \frac{1}{\lambda}(k(x, x) - \tilde{k}(x, X_t)(\tilde{K} + \lambda I)^{-1} \tilde{k}(X_t, x))$$

BKB: Budgeted Kernel Bandit

Regret and Computational Cost

Smaller computational cost than GP-UCB(Srinivas et al., 2010) and similar cumulative regret.

Continuous search space

It would require either to discretize the space or to solve the inner optimization problem (at each time step) i.e.

- either huge discretizations(Srinivas et al., 2012).
- or expensive inner optimization tasks without guarantees.

Adaptive Discretizations

(Bubeck et al., 2011; Munos, 2014; Shekhar and Javidi, 2018)

Update the discretization of X during the optimization process

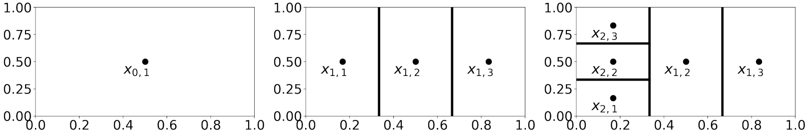

- function is evaluated only at partition centroids
- a centroid is selected using an UCB I
- the discretization is composed by partition centroids

Ada-BKB: Adaptive discretization + BKB

$$I_t(x_{h,i}) = \min\{\tilde{u}_t(x_{h,i}), \tilde{u}_t(\text{parent}(x_{h,i})) + V_{h-1}\} + V_h$$

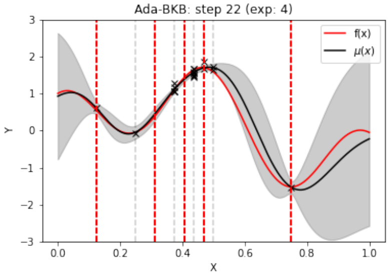

where $V_h \geq \sup_{x,x' \in X_{h,i}} |f(x) - f(x')|$ for all i.

For $t = 1, \cdots, T$:

- $x_{h,i} = \arg\max_{x \in L} I_t(x)$
- if $\tilde{\beta}_t \tilde{\sigma}_t(x_{h,i}) \leq V_h$ and $h < h_{\max}$, expand
- else $y_t = f(x_{h,i}) + \epsilon$
- update model

Pruning rule and early stopping

Let l_t^* be the highest lower bound observed at timestep t, partitions that may contain a global maximizer are

$$\{x_{h,i} | \tilde{u}_t(x_{h,i}) + V_h \geq l_t^*\}$$

Thus, we can erase every $x_{h,i}$ s.t. $\tilde{u}_t(x_{h,i}) + V_h < l_t^*$.

Early stopping

if $|L| = 0$ or $L = \{x_{h_{\max},i}\}$ and $|L| = 1$, we can interrupt the execution of the algorithm

Theoretical Results

Cumulative regret and Computational Cost

Similar regret guarantees and smaller time cost,

$$\text{(Regret)} \qquad \mathcal{O}(\sqrt{T} d_{\text{eff}} \log T) \quad \text{or} \quad \mathcal{O}\left(\sqrt{T d_{\text{eff}} \log T \frac{N^{h_{\max}} - 1}{N - 1}}\right)$$

$$\text{(Computational Cost)} \qquad \mathcal{O}(T^2 d_{\text{eff}}^2 h_{\max})$$

Empirical Results

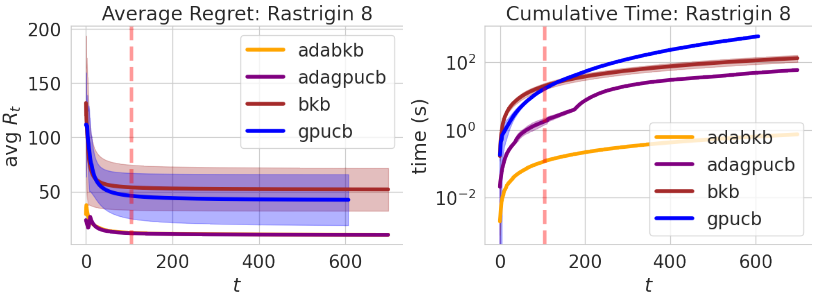

Thank you!

Thank you for your Attention!

Contact: marco.rando@edu.unige.it

References I

Bubeck, S., Munos, R., Stoltz, G., and Szepesvári, C. (2011). X-armed bandits. *Journal of Machine Learning Research*, 12:1655–1695.

Calandriello, D., Carratino, L., Lazaric, A., Valko, M., and Rosasco, L. (2019). Gaussian process optimization with adaptive sketching: Scalable and no regret. In *Conference on Learning Theory*, pages 533–557. PMLR.

Munos, R. (2014). From bandits to monte-carlo tree search: The optimistic principle applied to optimization and planning. *Foundations and Trends in Machine Learning*, 7(1):1–129.

Shekhar, S. and Javidi, T. (2018). Gaussian process bandits with adaptive discretization. *Electronic Journal of Statistics*, 12(2):3829 – 3874.

Srinivas, N., Krause, A., Kakade, S., and Seeger, M. (2012). Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory - TIT*, 58:3250–3265.

References II

Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. (2010). Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, pages 1015–1022.