Ada-BKB: Scalable Gaussian Process Optimization on Continuous Domains by Adaptive Discretizations

Marco Rando
University of Genova

Luigi Carratino
University of Genova

Silvia Villa
University of Genova

Lorenzo Rosasco
University of Genova / IIT/ MIT

Motivation

Gaussian process optimization (GP-Opt) is a successful class of algorithms to optimize a black-box function through sequential evaluations. However, for functions with continuous domains, GP-Opt has to rely on either a fixed discretization of the space, or the solution of a non-convex optimization sub-problem at each evaluation. We introduce Ada-BKB (Adaptive Budgeted Kernelized Bandit), a no-regret GP-Opt algorithm with adaptive discretizations for functions on continuous domains, that provably runs in $O(T^2 d_{\text{eff}}^2)$.

Problem Setup

Given a compact $X \subset \mathbb{R}^d$ and a function $f : X \to \mathbb{R}$, we want to find

$$x^* \in \operatorname*{argmax}_{x \in X} f(x).$$

Goal: to get small cumulative regret $R_T = \sum_{t=0}^{T} f(x^*) - f(x_t)$.

BUT: only perturbed function values $y = f(x) + \epsilon \qquad \epsilon \sim \mathcal{N}(0, \sigma^2)$

Main assumption: $f \in \mathcal{H}$ reproducing kernel Hilbert space (RKHS)

Background

BKB

Given an discrete set X, a kernel k and an RKHS $\mathcal{H}$.

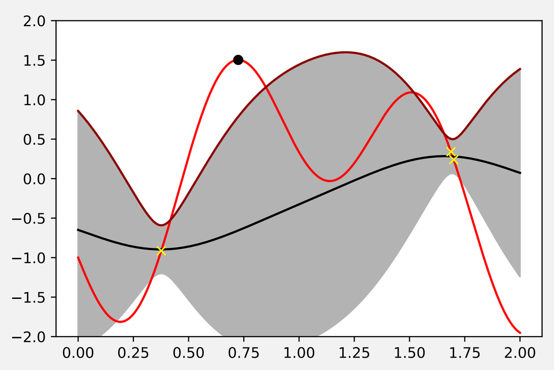

At each time step

- choose $x_t = \operatorname*{argmax}_{x \in X} \tilde{u}_t(x)$
- observe $y_t = f(x_t) + \epsilon_t$
- update the upper bound

With $\tilde{u}_t$ defined as in eq. (1).

Time complexity: $\mathcal{O}(T d_{\text{eff}}^2 |X|)$.

Problem: guarantees only with $|X| = \mathcal{O}(T^d)$[3].

Adaptive Discretizations

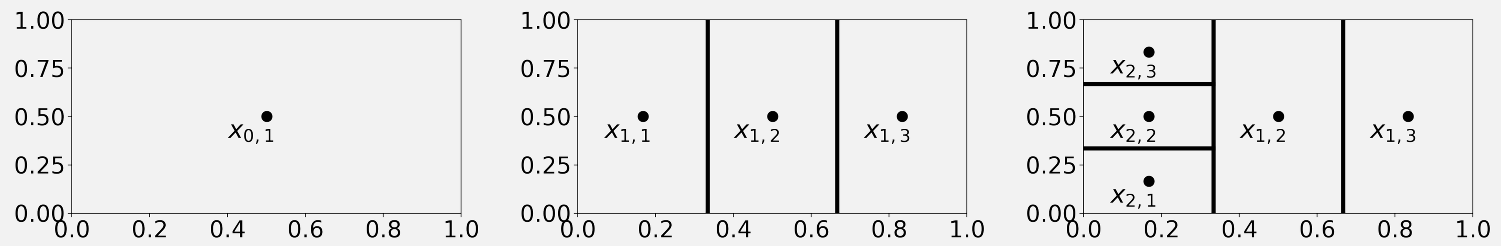

The point $x_{h,i}$ indicates the centroid of the partition $X_{h,i}$. Given a centroid $x_{h,i}$, the splitting procedure produces N partitions with centroids $x_{h+1,j}$. The centroid $x_{h,i}$ is called *parent* of $x_{h+1,j}$.

$$x_{h,i} = parent(x_{h+1,j})$$

Our Contribution

- a new GP optimization algorithm for continuous X: Ada + BKB
- new early-stopping condition and pruning rule
- theoretical guarantees on cumulative regret and computational time

Main Results

Cumulative Regret and Computational Time

Ada-BKB gets a cumulative regret,

$$R_T \leq \mathcal{O}\left(\sqrt{T d_{\text{eff}}(T) \log(T) \frac{N^{h_{max}} - 1}{N - 1}}\right)$$

Ada-BKB has time complexity,

$$\mathcal{O}(T^2 d_{\text{eff}}^2(T))$$

Other algorithms

Algorithm	Computational Cost	Cumulative Regret
GP-UCB	$\mathcal{O}(T^3 A)$	$\mathcal{O}(\sqrt{T}\gamma_T)$
BKB	$\mathcal{O}(TAd_{\text{eff}}^2)$	$\mathcal{O}(\sqrt{T}\gamma_T \log(T))$
AdaGP-UCB	$\mathcal{O}(T^4(N-1)h_{\text{max}})$	$\mathcal{O}(\sqrt{T}\gamma_T)$
GP-ThreDS	$\mathcal{O}(T^4)$	$\mathcal{O}(\sqrt{T}\gamma_T \log^2 T)$

Experimental results

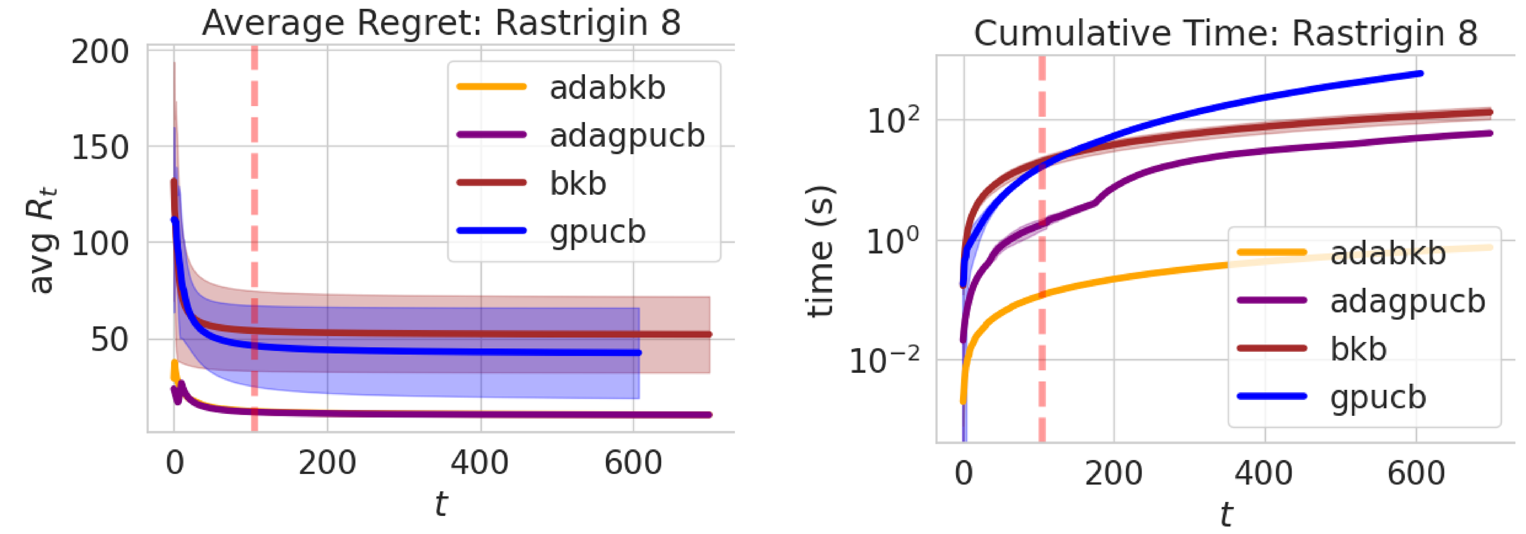

Left: average regret. Our algorithm obtain similar performance in average regret to AdaGP-UCB[2] but in less time. GP-UCB and BKB obtain worse regret because the offline discretizations don't contain a good suboptimal optimizer.

Right: computational savings. The red vertical line indicates that the early stopping condition is reached and we could interrupt the execution at that time step.

Ada-BKB

Let $x_{0,1}$ be the centroid of the first partition, let $L_0 = \{x_{0,1}\}$ and let $\tau = 0$

Algorithm

Algorithm 1 Ada-BKB

while $t \leq T$ **do**
  $x_{h,i} = \operatorname*{argmax}_{x_i \in L_\tau} I_t(x_i)$
  if $\tilde{\beta}_t \tilde{\sigma}_{t-1}(x_{h,i}) \leq V_h$ and $h_t < h_{\text{max}}$ **then**
    $L_{\tau+1} = (L_\tau \setminus \{x_{h,i}\}) \cup \text{expand}(x_{h,i})$
  else
    $y_t = f(x_{h,i}) + \epsilon_t$ (with ϵ_t noise)
    compute $\tilde{\mu}_{t+1}, \tilde{\sigma}_{t+1}, I^*_{t+1}$
    $L_{\tau+1} = L_\tau$
    $t = t + 1$
  $L_{\tau+1} = \text{prune}(L_\tau)$
  if $|L_{\tau+1}| == 0$ or $L_{\tau+1} == \{x_{h_{\text{max}},i}\}$ **then**
    break
  $\tau = \tau + 1$

where, for every i

$$I_t(x_{h,i}) = \min\{\tilde{u}_t(x_{h,i}), \tilde{u}_t(\text{parent}(x_{h,i}))\} + V_h \quad V_h \geq \sup_{x,x' \in X_{h,i}} |f(x) - f(x')|$$

with $X_{h,i}$ partition and

$$\begin{aligned} \tilde{u}_t(x) &= \tilde{\mu}_t(x) + \tilde{\beta}_t \tilde{\sigma}_t(x) \\ \tilde{\mu}_t(x) &= \tilde{k}(x, X_t)(\tilde{K}_t + \lambda I)^{-1} y_t \\ \tilde{\sigma}_t(x) &= k(x,x) - \tilde{k}(x, X_t)(\tilde{K}_t + \lambda I)^{-1} \tilde{k}(X_t, x) \end{aligned} \tag{1}$$

where $K_{i,j} = \tilde{k}(x_i, x_j)$ and, let S_t be a subset of X_t

$$\tilde{k}_t(x, x') = k(x, S_t) K_{S_t}^{\dagger} k(S_t, x')$$

A threshold on the number of expansion h_{max} is introduced to avoid infinite expansions.

Pruning rule & Early stopping

Let the highest lower confidence bound (LCB) be defined as

$$I_t^* = \max_{x \in X_t} \tilde{\mu}_t(x) - \tilde{\beta}_t \tilde{\sigma}_t(x).$$

After each iteration, the pruning rule deletes every $x \in L_\tau$ s.t.

$$I_t(x) < I_t^*$$

If after pruning $|L_\tau| == 0$ or $L_\tau = \{x_{h_{max},j}\}$ then we can interrupt the execution of the algorithm (**early stopping**).

References

[1] Daniele Calandriello, Luigi Carratino, Alessandro Lazaric, Michal Valko, and Lorenzo Rosasco. In *Conference on Learning Theory*, pages 533–557. PMLR, 2019.

[2] Shubhanshu Shekhar and Tara Javidi. *Electronic Journal of Statistics*, 12(2):3829 – 3874, 2018.

[3] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory - TIT*, 58:3250–3265, 05 2012.