Improved Regret Bounds for Agnostic Gaussian Process Bandits using Local Polynomial Estimators

Shubhanshu Shekhar[1] ECE Department, UCSD
Tara Javidi[2] ECE Department, UCSD

REALML Workshop at ICML 2020

[1] shshekha@eng.ucsd.edu
[2] tjavidi@eng.ucsd.edu

Problem Setup

We consider the problem of maximizing a black-box function f which can be accessed via noisy and expensive evaluation queries.

- We assume that $f : [0,1]^D \mapsto \mathbb{R}$ has bounded norm in the RKHS of a kernel K, i.e., $\boldsymbol{\|f\|_{\mathcal{H}_K} \leq B}$.
- Observation model: $\boldsymbol{y = f(x) + \eta}$ with $\eta \sim \sigma^2$-sub-Gaussian.
- Budget $= n$ evaluations
- Goal: Select queries $\{X_1, X_2, \ldots, X_n\}$ sequentially to minimize Simple ($\mathcal{S}_n$) and Cumulative ($\mathcal{R}_n$) regret:

$$\mathcal{S}_n = f(x^*) - f(X_n), \qquad \mathcal{R}_n = \sum_{t=1}^{n} f(x^*) - f(X_t)$$

$\texttt{LP-GP-UCB}$ Algorithm

Key Idea: Augment the global GP surrogate (as used in GP-UCB algorithm) with Local Polynomial (LP) estimators over non-uniform partitions of the domain to construct tighter UCBs.

Repeat the following steps for all times $t \geq 1$:

- Maintain a partition $\mathcal{P}_t = \{E_1, E_2, \ldots, E_{m_t}\}$.
- Use Local Polynomial (LP) Estimators along with global GP surrogate to construct UCB.
- Select query point $x_t \sim \texttt{Unif}(E_t)$ where $E_t = \arg\max_{E \in \mathcal{P}_t} UCB(E)$.
- Update the partition $\mathcal{P}_t$ and the GP posterior.

Theoretical Results

Key Observation: For Matérn kernels with $\nu > 0$, we can show an embedding of the RKHS into the space $\mathcal{C}^{k,\alpha}$ for $k = \lceil \nu - 1 \rceil$ and $\alpha = \nu - k$.

- For Matérn kernels, the `LP-GP-UCB` algorithms achieves uniformly tighter bounds on both $\mathcal{S}_n$ and $\mathcal{R}_n$, for all $\nu > 0$.
- In particular, the bounds on $\mathcal{S}_n$ (resp. $\mathcal{R}_n$) match the algorithm independent lower bounds for $\nu \leq D(D+1)$ (resp. $\nu \leq 1$).
- Besides Matérn kernels, we also obtain the first explicit in n regret bounds for some other important kernels such as Rational-Quadratic and Gamma-Exponential.

Empirical Results-Branin Function

We use the 2-dim Branin function (g_B) to construct an objective function (f) on $\mathbb{R}^8$ as follows:

$$f(x) = \sum_{i=1}^{4} c_i g_B \left(x[2i-1:2i] \right),$$

where $c_1 = 1.0$ and $c_i = 0.1$ for $i = 2, 3, 4$.

Informally, f has 2 *active* dimensions and an *ambient* dimension of 8.