Bayesian Optimization under Heavy-tailed Payoffs

Sayak Ray Chowdhury

Joint work with

Aditya Gopalan

Department of ECE,
Indian Institute of Science

NeurIPS, Dec. 2019

Black-box optimization

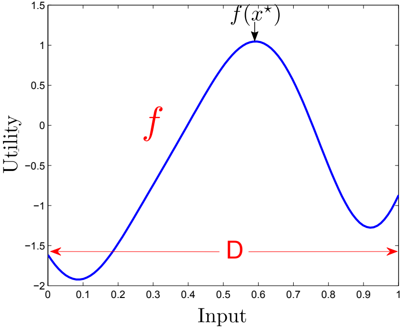

Problem: Maximize an unknown utility function $f : D \to \mathbb{R}$ by

- Sequentially querying f at inputs $x_1, x_2, \ldots, x_T$ and
- Observing noisy function evaluations: $y_t = f(x_t) + \epsilon_t$

Want: Low cumulative regret: $\sum_{t=1}^{T} \Big(f(x^\star) - f(x_t) \Big)$

Heavy-tailed noise

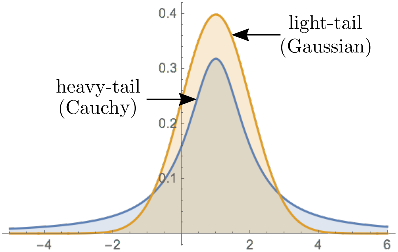

eg. Student's-t, Pareto, Cauchy etc.

Motivation:

- Significant chance of very high/low values
- Corrupted measurements
- Bursty traffic flow distributions
- Price fluctuations in financial and insurance data

- Existing works assume light-tailed noise (e.g. Srinivas et. al '11, Hernandez-Lobato et al.'14, ...)
- Question: Bayesian optimization algorithms with guarantees under heavy-tailed noise?

Algorithm 1: Truncated GP-UCB (TGP-UCB)

Unknown function f modeled by a Gaussian Process $f \sim \mathsf{GP}(0, k)$
At round t:

Algorithm 1: Truncated GP-UCB (TGP-UCB)

Unknown function f modeled by a Gaussian Process $f \sim \mathrm{GP}(0, k)$
At round t:

1. Choose the query point x_t using current GP posterior and a suitable parameter β_t:

$$x_t = \operatorname*{argmax}_{x \in D} \mu_{t-1}(x) + \beta_t \sigma_{t-1}(x)$$

Algorithm 1: Truncated GP-UCB (TGP-UCB)

Unknown function f modeled by a Gaussian Process $f \sim \text{GP}(0, k)$

At round t:

1. Choose the query point x_t using current GP posterior and a suitable parameter β_t:

$$x_t = \operatorname*{argmax}_{x \in D} \mu_{t-1}(x) + \beta_t \sigma_{t-1}(x)$$

2. Truncate the observed payoff y_t using a suitable threshold b_t:

$$\hat{y}_t = y_t \mathbb{1}_{|y_t| \leq b_t}$$

Algorithm 1: Truncated GP-UCB (TGP-UCB)

Unknown function f modeled by a Gaussian Process $f \sim \text{GP}(0, k)$

At round t:

1. Choose the query point x_t using current GP posterior and a suitable parameter β_t:

$$x_t = \underset{x \in D}{\operatorname{argmax}}\, \mu_{t-1}(x) + \beta_t \sigma_{t-1}(x)$$

2. Truncate the observed payoff y_t using a suitable threshold b_t:

$$\hat{y}_t = y_t \mathbb{1}_{|y_t| \le b_t}$$

3. Update GP posterior (μ_t, σ_t) with new observation $(x_t, \hat{y}_t)$:

$$\begin{aligned} \mu_t(x) &= k_t(x)^T (K_t + \lambda I)^{-1} [\hat{y}_1, \ldots, \hat{y}_t]^T \\ \sigma_t^2(x) &= k(x,x) - k_t(x)^T (K_t + \lambda I)^{-1} k_t(x) \end{aligned}$$

Regret bounds

Assumption on heavy-tailed payoffs:

$$\mathbb{E}\left[|y_t|^{1+\alpha}\right] < +\infty \quad \text{for} \quad \alpha \in (0,1]$$

Algorithm	Payoff	Regret
GP-UCB (Srinivas et. al)	sub-Gaussian	$O\left(\gamma_T T^{\frac{1}{2}}\right)$
TGP-UCB (this paper)	Heavy-tailed	$O\left(\gamma_T T^{\frac{2+\alpha}{2(1+\alpha)}}\right)$

- $\alpha = 1 \quad \Rightarrow \quad$ Regret $\tilde{O}\left(T^{\frac{3}{4}}\right)$

Regret bounds

Assumption on heavy-tailed payoffs:

$$\mathbb{E}\left[|y_t|^{1+\alpha}\right] < +\infty \quad \text{for} \quad \alpha \in (0,1]$$

Algorithm	Payoff	Regret
GP-UCB (Srinivas et. al)	sub-Gaussian	$O\left(\gamma_T T^{\frac{1}{2}}\right)$
TGP-UCB (this paper)	Heavy-tailed	$O\left(\gamma_T T^{\frac{2+\alpha}{2(1+\alpha)}}\right)$

- $\alpha = 1 \quad \Rightarrow \quad$ Regret $\tilde{O}\left(T^{\frac{3}{4}}\right)$
- We also give a $\Omega\left(T^{\frac{1}{1+\alpha}}\right)$ regret lower bound for any algorithm

Regret bounds

Assumption on heavy-tailed payoffs:

$$\mathbb{E}\left[|y_t|^{1+\alpha}\right] < +\infty \quad \text{for} \quad \alpha \in (0,1]$$

Algorithm	Payoff	Regret
GP-UCB (Srinivas et. al)	sub-Gaussian	$O\left(\gamma_T T^{\frac{1}{2}}\right)$
TGP-UCB (this paper)	Heavy-tailed	$O\left(\gamma_T T^{\frac{2+\alpha}{2(1+\alpha)}}\right)$

- $\alpha = 1 \quad \Rightarrow \quad$ Regret $\tilde{O}\left(T^{\frac{3}{4}}\right)$
- We also give a $\Omega\left(T^{\frac{1}{1+\alpha}}\right)$ regret lower bound for any algorithm
- Question: Can we achieve $\tilde{O}\left(T^{\frac{1}{1+\alpha}}\right)$ regret scaling?

Regret bounds

Assumption on heavy-tailed payoffs:

$$\mathbb{E}\left[|y_t|^{1+\alpha}\right] < +\infty \quad \text{for} \quad \alpha \in (0,1]$$

Algorithm	Payoff	Regret
GP-UCB (Srinivas et. al)	sub-Gaussian	$O\left(\gamma_T T^{\frac{1}{2}}\right)$
TGP-UCB (this paper)	Heavy-tailed	$O\left(\gamma_T T^{\frac{2+\alpha}{2(1+\alpha)}}\right)$

- $\alpha = 1 \quad \Rightarrow \quad$ Regret $\tilde{O}\left(T^{\frac{3}{4}}\right)$
- We also give a $\Omega\left(T^{\frac{1}{1+\alpha}}\right)$ regret lower bound for any algorithm
- Question: Can we achieve $\tilde{O}\left(T^{\frac{1}{1+\alpha}}\right)$ regret scaling?
- Ans: **YES**

Algorithm 2: Adaptively Truncated Approximate GP-UCB

Idea: UCB with Kernel approximation + Feature adaptive truncation:

$$x_t = \text{argmax}_{x \in D} \, \tilde{\mu}_{t-1}(x) + \beta_t \tilde{\sigma}_{t-1}(x)$$

Algorithm 2: Adaptively Truncated Approximate GP-UCB

Idea: UCB with Kernel approximation + Feature adaptive truncation:

$$x_t = \text{argmax}_{x \in D}\, \tilde{\mu}_{t-1}(x) + \beta_t \tilde{\sigma}_{t-1}(x)$$

Kernel approximation:

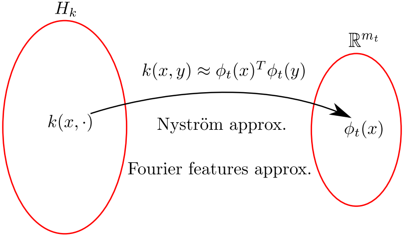

Compute:

- $V_t = \sum_{s=1}^{t} \phi_t(x_s)\phi_t(x_s)^T + \lambda I$ (m_t rows and m_t columns)
- $U_t = V_t^{-\frac{1}{2}}[\phi_t(x_1), \ldots, \phi_t(x_t)]$ (m_t rows and t columns)

Algorithm 2: Adaptively Truncated Approximate GP-UCB

Idea: UCB with Kernel approximation + Feature adaptive truncation:

$$x_t = \text{argmax}_{x \in D}\, \tilde{\mu}_{t-1}(x) + \beta_t \tilde{\sigma}_{t-1}(x)$$

Kernel approximation:

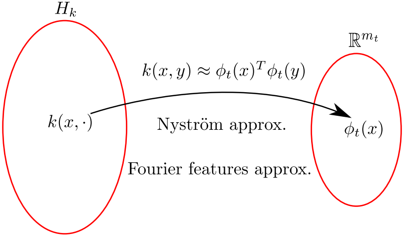

Compute:

- $V_t = \sum_{s=1}^{t} \phi_t(x_s)\phi_t(x_s)^T + \lambda I$ (m_t rows and m_t columns)
- $U_t = V_t^{-\frac{1}{2}}[\phi_t(x_1), \ldots, \phi_t(x_t)]$ (m_t rows and t columns)

Feature adaptive truncation:

$$\begin{bmatrix} u_{11} & u_{12} & \cdots & u_{1t} \\ u_{21} & u_{22} & \cdots & u_{2t} \\ \vdots & \vdots & \ddots & \vdots \\ u_{m_t1} & u_{m_t2} & \cdots & u_{m_tt} \end{bmatrix} \odot \begin{bmatrix} y_1 & y_2 & \cdots & y_t \\ y_1 & y_2 & \cdots & y_t \\ \vdots & \vdots & \ddots & \vdots \\ y_1 & y_2 & \cdots & y_t \end{bmatrix}$$

Hadamard product

Algorithm 2: Adaptively Truncated Approximate GP-UCB

Idea: UCB with Kernel approximation + Feature adaptive truncation:

$$x_t = \text{argmax}_{x \in D}\, \tilde{\mu}_{t-1}(x) + \beta_t \tilde{\sigma}_{t-1}(x)$$

Kernel approximation:

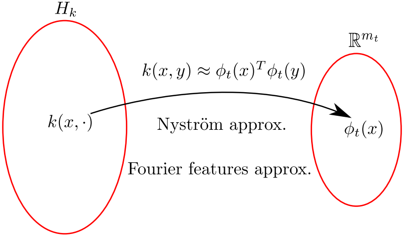

Compute:

- $V_t = \sum_{s=1}^{t} \phi_t(x_s)\phi_t(x_s)^T + \lambda I$ (m_t rows and m_t columns)
- $U_t = V_t^{-\frac{1}{2}} [\phi_t(x_1), \ldots, \phi_t(x_t)]$ (m_t rows and t columns)

Feature adaptive truncation:

$$r_{12} = u_{12} y_2 \mathbf{1}_{|u_{12}y_2| \le b_t}$$

$$\begin{bmatrix} u_{11}y_1 & u_{12}y_2 & \cdots & u_{1t}y_t \\ u_{21}y_1 & u_{22}y_2 & \cdots & u_{2t}y_t \\ \vdots & \vdots & \ddots & \vdots \\ u_{m_t 1}y_1 & u_{m_t 2}y_2 & \cdots & u_{m_t t}y_t \end{bmatrix} \quad \begin{bmatrix} r_{11} & r_{12} & \cdots & r_{1t} \\ r_{21} & r_{22} & \cdots & r_{2t} \\ \vdots & \vdots & \ddots & \vdots \\ r_{m_t 1} & r_{m_t 2} & \cdots & r_{m_t t} \end{bmatrix}$$

Find row sums

$r_1, r_2, \ldots, r_{m_t}$

Algorithm 2: Adaptively Truncated Approximate GP-UCB

Idea: UCB with Kernel approximation + Feature adaptive truncation:

$$x_t = \text{argmax}_{x \in D} \, \tilde{\mu}_{t-1}(x) + \beta_t \tilde{\sigma}_{t-1}(x)$$

Kernel approximation:

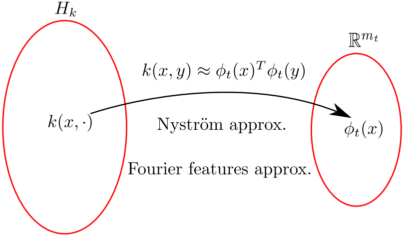

Compute:

- $V_t = \sum_{s=1}^{t} \phi_t(x_s)\phi_t(x_s)^T + \lambda I$ (m_t rows and m_t columns)
- $U_t = V_t^{-\frac{1}{2}}[\phi_t(x_1), \ldots, \phi_t(x_t)]$ (m_t rows and t columns)

Approximate posterior GP:

$$\begin{aligned} \tilde{\mu}_t(x) &= \phi_t(x)^T V_t^{-1/2} [r_1, \ldots, r_{m_t}]^T \\ \tilde{\sigma}_t^2(x) &= k(x, x) - \phi_t(x)^T \phi_t(x) + \lambda \phi_t(x)^T V_t^{-1} \phi_t(x) \end{aligned}$$

where $r_i = \sum_{s=1}^{t} u_{is} y_s \mathbb{1}_{|u_{is} y_s| \leq b_t}$ (u_i is the i^{th} row of U_t)

See you at the poster session

Bayesian Optimization under Heavy-tailed Payoffs

Poster #11

Tue Dec 10th 05:30 – 07:30 PM @ East Exhibition Hall B + C

Acknowledgements:

1. Tata Trusts travel grant
2. Google India Phd fellowship grant
3. DST Inspire research grant