Gaussian Process Optimization with Mutual Information

Emile Contal[1] Vianney Perchet[2] Nicolas Vayatis[1]

[1]CMLA, UMR CNRS 8536, ENS Cachan, France
[2]LPMA, Université Paris Diderot, France

Abstract

We introduce the Gaussian Process Mutual Information algorithm (GP-MI) for sequential global optimization using Gaussian processes. The upper bounds we derive on the cumulative regret for this algorithm improve by an exponential factor the previously known bounds for algorithms like GP-UCB. We confirm the empirical efficiency of this algorithm on synthetic and real tasks against the natural competitor GP-UCB, and also the Expected Improvement heuristic (EI).

Keywords Optimization – Gaussian Process – Cumulative Regret – Exploration/Exploitation tradeoff

Background

1. Motivations

Examples of Optimization Problems

- Minimization of energy consumption in engineering
- Maximization of benefits in marketing
- Minimization of validation error in machine learning

Related Work

- Experimental design — Fedorov [1972]
- Multiarmed bandit — Bubeck et al. [2011]
- Active learning — Carpentier et al. [2011]
- Bayesian optimization — Srinivas et al. [2012]

2. Problem Statement

Sequential Optimization

Let $f : \mathcal{X} \to \mathbb{R}$ where $\mathcal{X} \subset \mathbb{R}^d$ is compact and convex. We consider the problem of finding the maximum of f denoted by: $f(x^\star) = \max_{x\in\mathcal{X}} f(x)$, via sequential queries $f(x_1), f(x_2), \dots$

At iteration T we choose x_{T+1} using the previous noisy observations $Y_T = \{y_1, \dots, y_T\}$, where $\forall t \leq T : y_t = f(x_t) + \epsilon_t$ and $\epsilon_t \overset{\text{iid}}{\sim} \mathcal{N}(0, \eta^2)$.

Objective

The efficiency of a policy is measured via the cumulative regret:

$$R_T = \sum_{t<T} f(x^\star) - f(x_t)\,.$$

3. Gaussian Processes

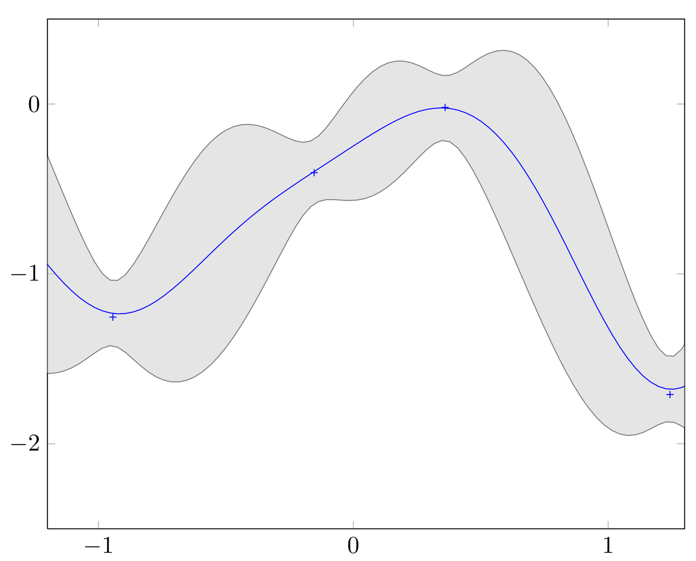

Definition

$f \sim \mathcal{GP}(m, k)$ with mean $m : \mathcal{X} \to \mathbb{R}$ and kernel $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}^+$, when for all $x_1, \dots, x_n$ we have that $\big(f(x_1), \dots, f(x_n)\big)$ is a multivariate Gaussian $\mathcal{N}([m(x_i)]_i, [k(x_i, x_j)]_{i,j})$.

Bayesian Inference

Given Y_T, the posterior distribution $\Pr[f \mid Y_T]$ is a GP with mean μ_{T+1} (prediction) and covariance σ^2_{T+1} (uncertainty) computed by Bayesian inference.

Algorithm

4. Mutual Information

Information Gain

The information gain on f at X_T is given by: $I_T(X_T) = \frac{1}{2}\operatorname{logdet}(\mathbf{I} + \eta^{-2}\mathbf{K}_T)$.

We define $\gamma_T = \max_{X_T \subseteq \mathcal{X} : |X_T| = T} I_T(X_T)$ the maximum information gain by T queries.

Empirical Lower Bound

For GPs with bounded variance: Srinivas et al. [2012]

$$\widehat{\gamma}_T = \sum_{t=1}^{T} \sigma_t^2(x_t) \leq C_1 \gamma_T \text{ where } C_1 = \frac{2}{\log(1+\eta^{-2})}\,.$$

5. GP-MI

$\widehat{\gamma}_0 \leftarrow 0$
for $t = 1, 2, \dots$ **do**
  Compute μ_t and σ_t^2 using Bayesian inference
  $\phi_t(x) \leftarrow \sqrt{\alpha}\left(\sqrt{\sigma_t^2(x) + \widehat{\gamma}_{t-1}} - \sqrt{\widehat{\gamma}_{t-1}}\right)$
  $x_t \leftarrow \operatorname{argmax}_{x\in\mathcal{X}} \mu_t(x) + \phi_t(x)$; $\widehat{\gamma}_t \leftarrow \widehat{\gamma}_{t-1} + \sigma_t^2(x_t)$
  Query at x_t and observe y_t
end

6. Theorem

$\forall \delta > 0$ and $T > 1$, set $\alpha = \log\frac{2}{\delta}$.

$$\Pr\left[R_T \leq 5\sqrt{\alpha C_1 \gamma_T} + 4\sqrt{\alpha}\right] \geq 1 - \delta\,,$$

where $C_1 = \frac{2}{\log(1+\eta^{-2})}$.

7. Corollary

- For linear kernel: $\mathcal{O}(\sqrt{d \log T})$
- For RBF kernel: $\mathcal{O}\big(\sqrt{(\log T)^{d+1}}\big)$
- For Matérn kernel: $\mathcal{O}\big(\sqrt{T^a \log T}\big)$, where $a < \frac{d^2}{2\nu + d^2} < 1$ and ν is the Matérn parameter.

8. Impact of δ on the Regret

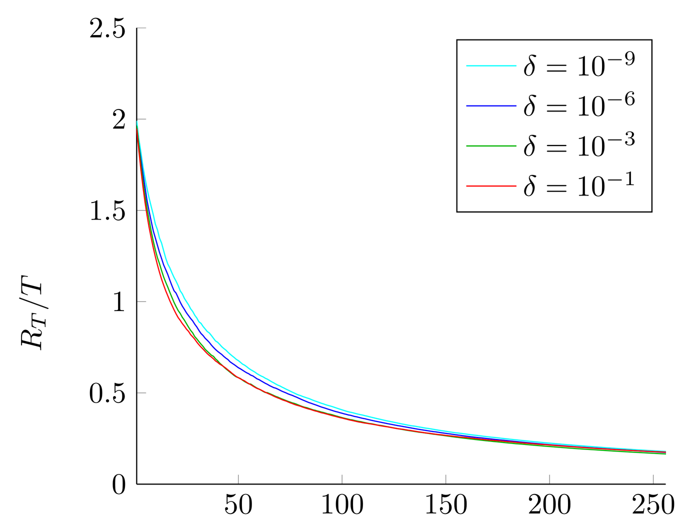

Small impact of the value of δ on the mean average regret with respect to the iteration of the GP-MI algorithm.

Experiments

9. Empirical Average Regret

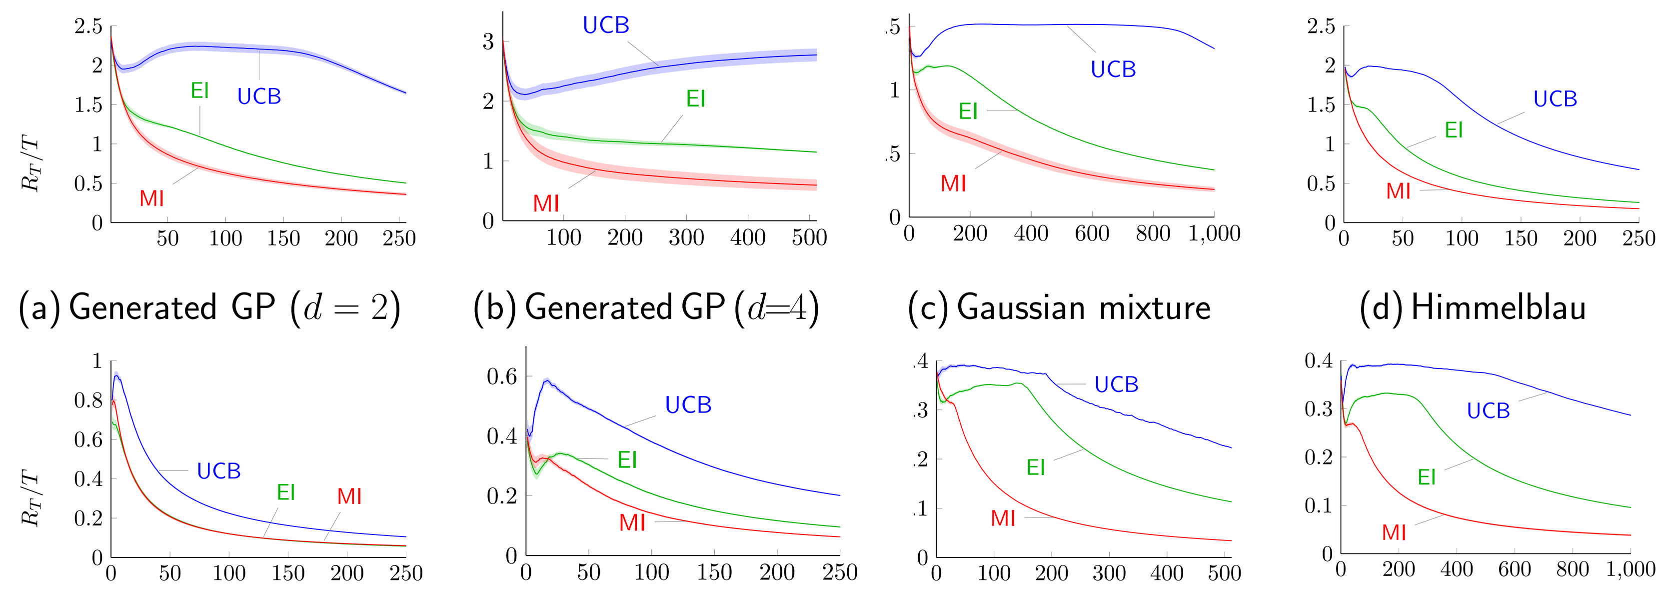

(a) Generated GP ($d = 2$) (b) Generated GP ($d=4$) (c) Gaussian mixture (d) Himmelblau
(e) Branin (f) Goldstein (g) Tsunamis (h) Mackey-Glass

10. Implementation

Exact Inference for Gaussian likelihood

Numerical complexity in $\mathcal{O}(T^2)$ using the Cholesky sequential updates of the covariance matrix. Osborne [2010]

Algorithms for non-Gaussian likelihood

For other likelihood functions (e.g. Laplacian or Student's t), one can use the EP algorithm or Monte Carlo sampling. Kuss et al. [2005]

11. Open Questions and Discussion

- Theoretical performance for simple regret
- Include kernel learning procedure
- Calibration of δ

References

Bubeck, Munos, Stoltz, and Szepesvári. X-armed bandits. *JMLR*, 12:1655–1695, 2011.

Carpentier, Lazaric, Ghavamzadeh, Munos, and P. Auer. Upper-confidence-bound algorithms for active learning in multi-armed bandits. In *ALT*, volume 6925, pages 189–203. Springer-Verlag, 2011.

Fedorov. *Theory of Optimal Experiments*. Academic Press, 1972.

Kuss, Pfingsten, Csató, and Rasmussen. Approximate inference for robust Gaussian process regression. *MPI Bio. Cyb., Tech. Rep.*, 136, 2005.

Osborne. *Bayesian Gaussian processes for sequential prediction, optimisation and quadrature*. PhD thesis, Oxford University New College, 2010.

Srinivas, Krause, Kakade, and Seeger. Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. *IEEE Information Theory*, 58(5):3250–3265, 2012.