Quantum Bayesian Optimization

Zhongxiang Dai* [1] Gregory Kang Ruey Lau* [1,2] Arun Verma[1] Yao Shu[1] Bryan Kian Hsiang Low[1] Patrick Jaillet[3]

* Equal contribution [1] Department of Computer Science, National University of Singapore [2] CNRS@Create [3] Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology

Background: Bayesian Optimization (BO) and Quantum speed-up of classical algorithms

- BO aims to maximize a noisy black-box reward function given sequential queries to the function, with the smallest possible regret

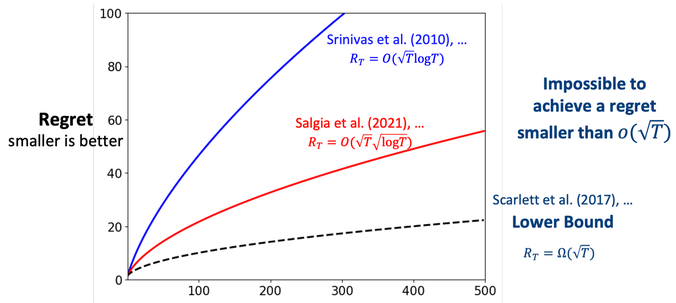

- Quantum computing allows for a superposition of states and operations acting on these states that enables speed-up of classical algorithms
- Quantum Monte Carlo (QMC) enables a quadratic speed-up of the number of samples needed vs. classical Monte Carlo mean estimation

Algorithm: Quantum-Gaussian Process-Upper Confidence Bound (Q-GP-UCB)

- Given a quantum oracle that encodes the noisy reward distribution for x_s
- For stage $s = 1,2,...m$, query oracle N_s times for chosen x_s and use QMC to efficiently estimate y_s up to error ϵ_s
- Update the GP posterior given all past stages of (x_s, y_s) and weights ϵ_s, which will inform the choice of x_{s+1} for the next stage

Algorithm 1 Q-GP-UCB

1: **for** stage $s = 1, 2, \ldots$ **do**
2: $x_s = \arg\max_{x\in\mathcal{X}} \tilde{\mu}_{s-1}(x) + \beta_s \tilde{\sigma}_{s-1}(x)$ (3).
3: $\epsilon_s = \tilde{\sigma}_{s-1}(x_s)/\sqrt{\lambda}$.
4: *(a)* $N_{\epsilon_s} = \frac{C_1}{\epsilon_s} \log(\frac{2m}{\delta})$ (for bounded noise), or
(b) $N_{\epsilon_s} = \frac{C_2\sigma}{\epsilon_s} \log_2^{3/2}(\frac{8\sigma}{\epsilon_s}) \log_2(\log_2 \frac{8\sigma}{\epsilon_s}) \log \frac{2m}{\delta}$ (for noise with variance bounded by σ^2).
5: If $\sum_{k=1}^{s} N_{\epsilon_k} > T$, terminate the algorithm.
6: Run the QMC$(\mathcal{O}_{x_s}, \epsilon_s, \delta/(2m))$ algorithm to query the quantum oracle of x_s for the next N_{ϵ_s} rounds, and obtain y_s as an estimate of $f(x_s)$.
7: Update the weighted GP posterior (3) using (x_s, y_s).

- Unlike standard GP-UCB, our algo. uses a weighted GP posterior, along with the QMC speed-up to estimate noisy y_s, to yield lower regret
- We keep track of ϵ_s, which is related to the QMC estimation error, the weight for the GP update ($\frac{1}{\epsilon_s^2}$), and # of queries to perform N_{ϵ_s}

Theoretical analysis: Based on weighted info. gain analysis and tight confidence ellipsoid

Upper bound on weighted information gain $\tilde{\gamma}_m \leq \gamma_{T^2}$

↓

Upper bound on total number of stages m $m = O((\log T)^{d+1})$

↓

Confidence ellipsoid $|f(x) - \tilde{\mu}_{s-1}(x)| \leq \beta_s \tilde{\sigma}_{s-1}(x)$

↓

Regret bound $R_T = O((\log T)^{3(d+1)/2} \log((\log T)^{d+1}))$

Theoretical results: Regret upper bounds for Q-GP-UCB lower than classical BO lower bounds

Squared exponential kernel

$R_T = O((\log T)^{3(d+1)/2} \log((\log T)^{d+1})) = O(\text{poly} \log T)$

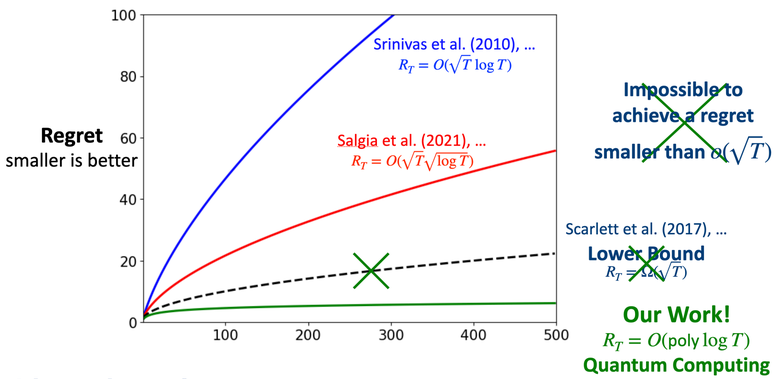

Linear kernel

$R_T = O(d^{3/2}(\log T)^{3/2} \log(d \log T))$

- We improve the regret upper bound for quantum linear bandits [1] by $\sqrt{d}$, given our tighter confidence ellipsoid

Matern kernel

$R_T = O(T^{3d/(2\nu+d)})$

- Regret is smaller than classical lower bound when $\nu > 2d$ (i.e., sufficiently smooth function)

Empirical results: Lower cumulative regret, even on real quantum computer

Synthetic data simulations

Q-GP-UCB has lower cumulative regret compared to classical GP-UCB when tested on synthetic dataset with Bernoulli and Gaussian reward distributions

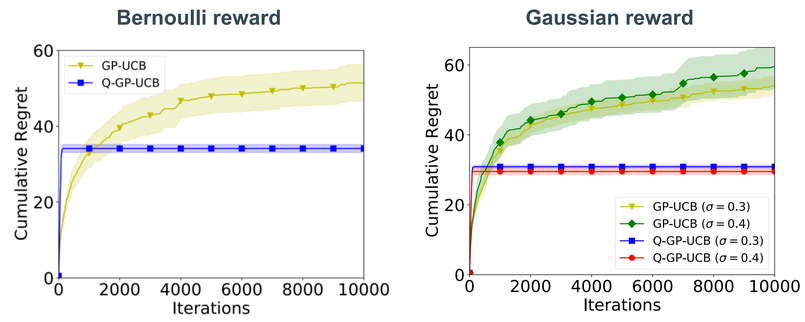

AutoML simulations

Better results (vs. both GP-UCB and quantum linear bandits) also hold for AutoML simulation task to tune SVM hyperparameters for classification on diabetes dataset

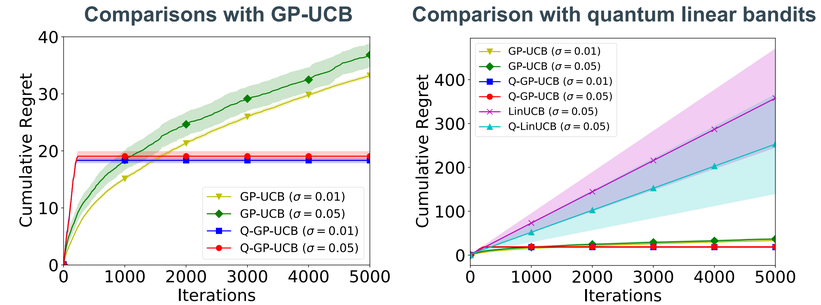

Real quantum computer

Better results hold even when running Q-GP-UCB on a noisy, real quantum computer (7 Qubit IBM quantum computer)

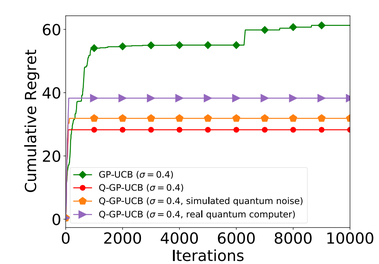

- Running Q-GP-UCB on real quantum computers incurs noise from hardware implementation
- Our Q-GP-UCB run on a real quantum computer still significantly outperforms classical GP-UCB

Acknowledgement

This research/project is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-PhD/2023-01-039J). DesCartes: this research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme.

[1] Zongqi Wan, Zhijie Zhang, Tongyang Li, Jialin Zhang and Xiaoming Sun. Quantum multi-armed bandits and stochastic linear bandits enjoy logarithmic regrets. In AAAI 2023.