

Pseudo-Bayesian Optimization

Thursday, June 20, 2024 4:00 PM (30 minutes)

Bayesian Optimization aims to optimize expensive black-box functions using minimal function evaluations. Its key idea is to strategically model the unknown function structure via a surrogate model and, importantly, quantify the associated uncertainty that allows a sequential search of query points to balance exploitation-exploration. While Gaussian process (GP) has been a flexible and favored surrogate model, its scalability issues have spurred recent alternatives whose convergence properties are nonetheless more opaque. Motivated by these dilemmas, we propose an axiomatic framework, which we call Pseudo-Bayesian Optimization, that elicits the minimal requirements to guarantee black-box optimization convergence beyond GP-based methods. The design freedom in our framework subsequently allows us to construct algorithms that are both scalable and empirically superior. In particular, we show how using simple local regression, together with an uncertainty quantifier that adapts the “randomized prior” idea in reinforcement learning, not only guarantees convergence but also consistently outperforms state-of-the-art benchmarks in examples ranging from high-dimensional synthetic experiments to realistic hyperparameter tuning and robotic applications.

Primary author: CHEN, Haoxian (Columbia University)

Co-author: Prof. LAM, Henry (Columbia University)

Presenter: CHEN, Haoxian (Columbia University)

Session Classification: Parallel session: Learning and optimization