Optimum-statistical Collaboration Towards Efficient Black-box Optimization

Wenjie Li[1], Chi-Hua Wang[1], and Guang Cheng[1]

[1]Department of Statistics, Purdue University

Introduction

Background

- Black-box optimization has gained more and more attention nowadays because of its popularity in application [5, 4]
- Existing works are scattered, and only focus on one specific type of local-smoothness assumption

Contributions

In this paper, we introduce the optimum-statistical collaboration framework with its analysis. In terms of resolution, such an analysis uses a general local smoothness assumption. In terms of uncertainty, the framework inspires us to propose a better quantifier.

Definitions

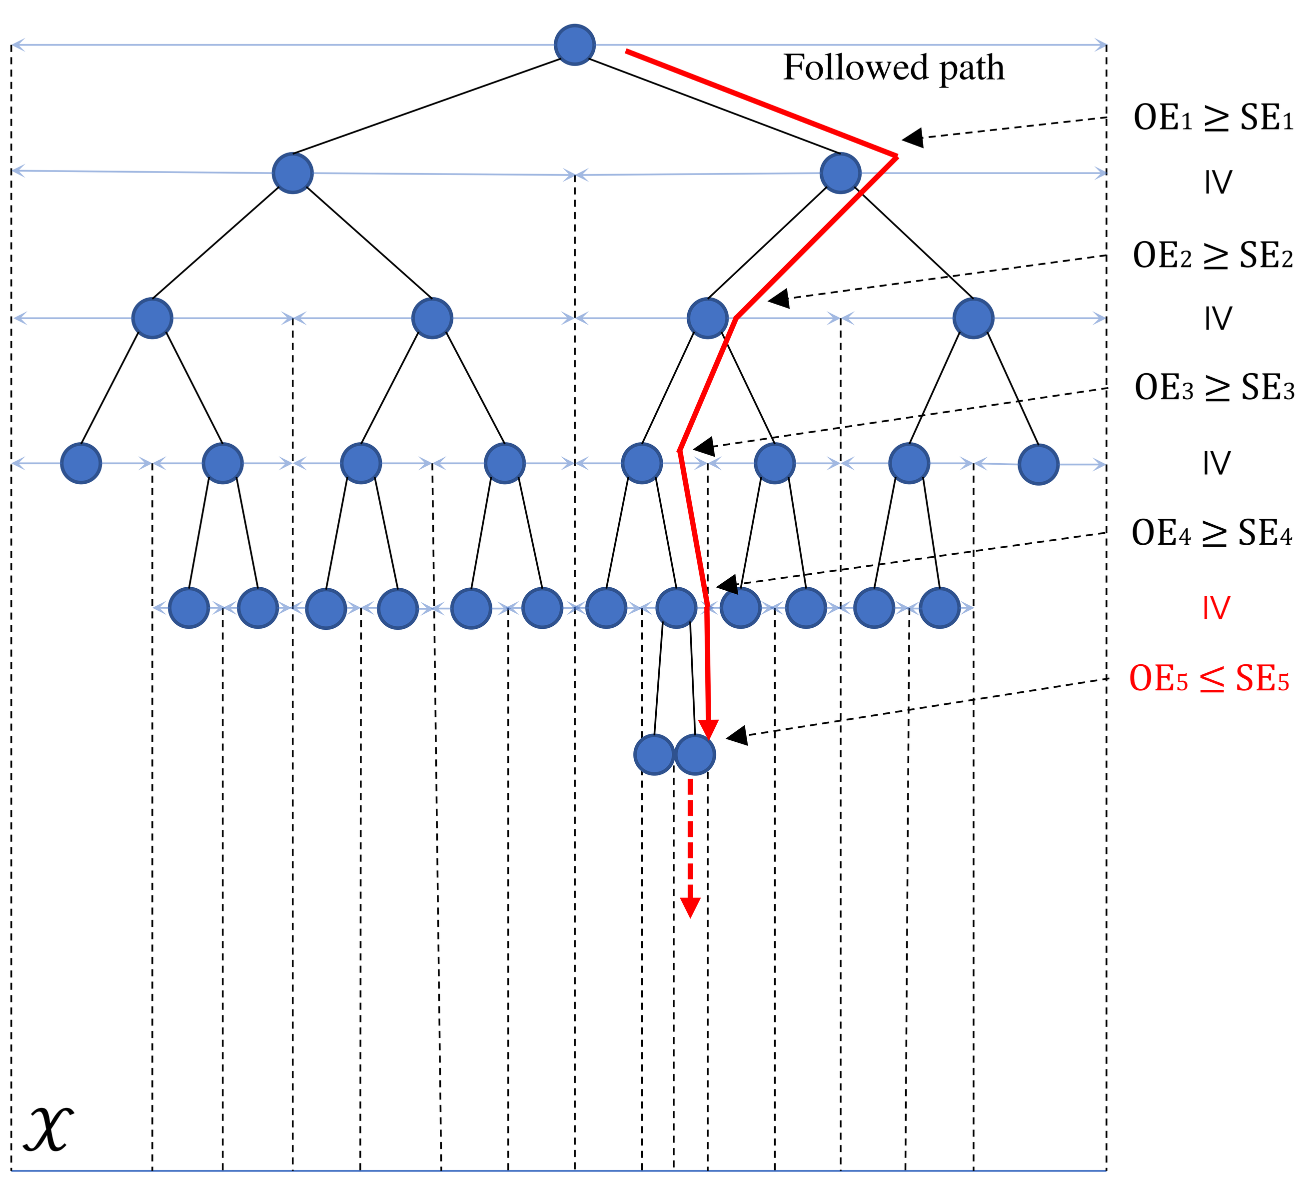

Definition 1. (Resolution Descriptor) Define $\mathtt{OE}_h$ to be the *resolution* for each level h, which is a function that upper-bounds the change of f around the optimum and measures the current optimization error, i.e., for any global optimum x^*,

$$\forall h \geq 0, \forall x \in \mathcal{P}_{h,i_h^*}, f(x) \geq f(x^*) - \mathtt{OE}_h \qquad (\mathtt{OE})$$

where $\mathcal{P}_{h,i_h^*}$ is the node on level h in the partition that contain the global optimum x^*.

Definition 2. (Uncertainty Quantifier) Define $\mathtt{SE}_{h,i}(t)$ to be the *uncertainty estimate* for each node $\mathcal{P}_{h,i}$ at time t, which is a function that upper-bounds the current statistical error of the average $\widehat{\mu}_{h,i}(t)$ of the rewards obtained with high probability, i.e., the event

$$\mathcal{A}_t = \left\{ \forall (h,i), |\widehat{\mu}_{h,i}(t) - f(x_{h,i})| \leq \mathtt{SE}_{h,i}(t) \right\} \qquad (\mathtt{SE})$$

is a high probability event.

Algorithm and Results

- We propose the following general optimum-statistical collaboration algorithm for black-box optimization with unspecified resolution and uncertainty, and unspecified algorithm policy.

Algorithm 1 Optimum-statistical Collaboration

Input: Hierachical partition $\mathcal{P}$, resolution descriptor $\mathtt{OE}_h$, Uncertainty quantifier $\mathtt{SE}_{h,i}(t)$.
Step 1: Refresh the confidence at some specific time steps to update all $\mathtt{SE}_{h,i}(t)$ in the tree.
Step 2: Find the optimal node $\mathcal{P}_{h_t,i_t}$ at time t that satisfies $\mathtt{OE}_{h_t} \leq \mathtt{SE}_{h_t,i_t}(t)$ and pull $\mathcal{P}_{h_t,i_t}$.
Step 3: If $\mathtt{OE}_{h_t} \geq \mathtt{SE}_{h_t,i_t}(t)$ after the pull, expand $\mathcal{P}_{h_t,i_t}$ and explore deeper.

- A **more general** local smoothness (resolution descriptor) for different functions and partitions.

$$\forall h \geq 0, \forall x \in \mathcal{P}_{h,i_h^*}, f(x) \geq f(x^*) - \phi(h) \qquad (\text{GLS})$$

where $\phi(h)$ is a function of the level h, e.g., $\phi(h) = \nu\rho^h, 1/h$

- A **more efficient** variance adaptive uncertainty quantifier.

$$\mathtt{SE}_{h,i}(t) \equiv c\sqrt{\frac{2\mathbb{V}_{h,i}(t)\log(1/\tilde{\delta}(t^+))}{T_{h,i}(t)}} + \frac{3bc^2\log(1/\tilde{\delta}(t^+))}{T_{h,i}(t)} \qquad (\mathtt{VHCT})$$

Regret Bounds

Theorem 3: General Regret Bound

uppose that under a sequence of probability events $\{\mathcal{E}_t\}_{t=1,2,\ldots}$, at each time t, the designed policy to select the optimal node $\mathcal{P}_{h_t,i_t}$ in Algorithm 1 satisfies $f^* - f(x_{h_t,i_t}) \leq a \cdot \max\{\mathtt{SE}_{h_t,i_t}(t),\ \mathtt{OE}_{h_t}\}$, where $a > 0$ is an absolute constant. Then for any $\overline{H} \in [1, H(n))$ we have the following bound on the expected regret

$$\mathbb{E}[R_n] \leq 2aC \sum_{h=1}^{\overline{H}} (\mathtt{OE}_{h-1})^{-\bar{d}} \int_1^{T_h(n)} \max_i \mathtt{SE}_{h,i}(s)ds + \sum_{t=1}^{n} \mathbb{P}(\mathcal{E}_t^c)$$
$$+ a\sum_{\overline{H}+1}^{H(n)} \sum_{i \in \mathcal{I}_h(n)} \int_1^{T_{h,i}(n)} \mathtt{SE}_{h,i}(s)ds + \sqrt{2n\log(4n^3)} + \frac{1}{4n^2}$$

where $\bar{d} := d(a, C, \mathtt{OE}_{h-1})$ is the near-optimality dimension defined with respect to a, C, and $\mathtt{OE}_{h-1}$, and $T_h(n) = \max_i T_{h,i}(n)$

Example Regret Bounds for Different $\phi(h)$

- Regret bound when $\phi(h) = \nu\rho^h$

$$\mathbb{E}[R_n] \leq 2\sqrt{2n\log(4n^3)} + C_1 V_{\max}^{\frac{1}{d_1+2}} n^{\frac{d_1+1}{d_1+2}} (\log n)^{\frac{1}{d_1+2}} + C_2 n^{\frac{2d_1+1}{2d_1+4}} \log n$$

- Regret bound when $\phi(h) = 2/h$

$$\mathbb{E}[R_n] \leq 2\sqrt{2n\log(4n^3)} + \bar{C}_1 V_{\max}^{\frac{1}{2d_2+3}} n^{\frac{2d_2+2}{2d_2+3}} (\log n)^{\frac{1}{2d_2+3}} + \bar{C}_2 n^{\frac{2d_2+1}{2d_2+3}} \log n$$

Experiments

- We empirically compare the proposed `VHCT` algorithm with the existing **anytime** blackbox optimization algorithms, including `T-HOO` (the truncated version of `HOO` [2]), `HCT` [1], `POO` [3], and `PCT` (`POO` + `HCT`, [6]). We use the Garland function and the Double-sine function as the blackbox objectives.

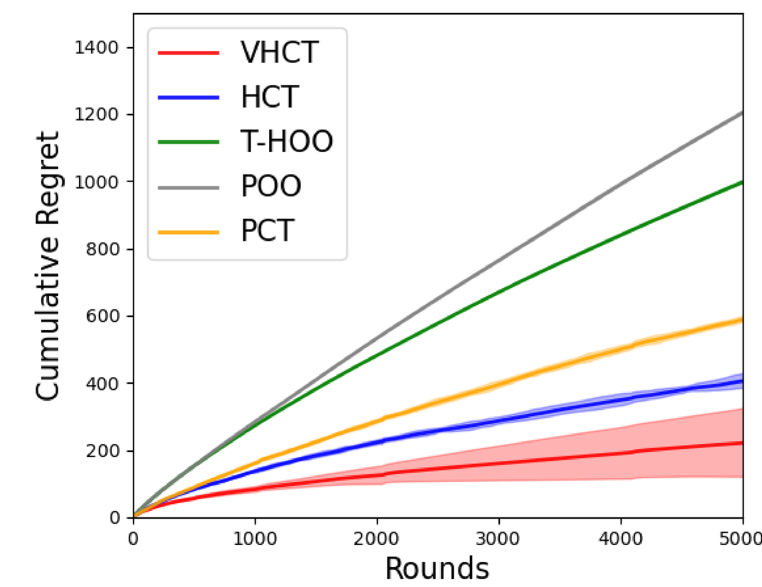

Figure 1: Garland/Low-noise

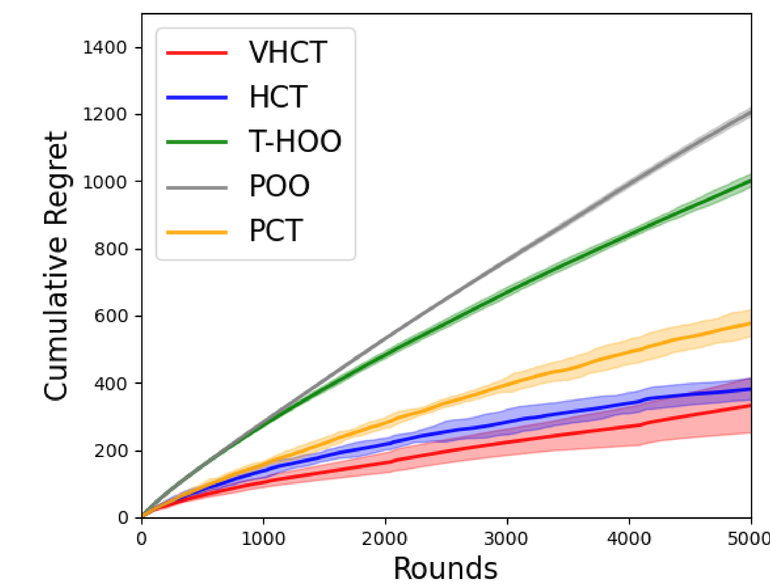

Figure 2: Garland/Moderate-noise

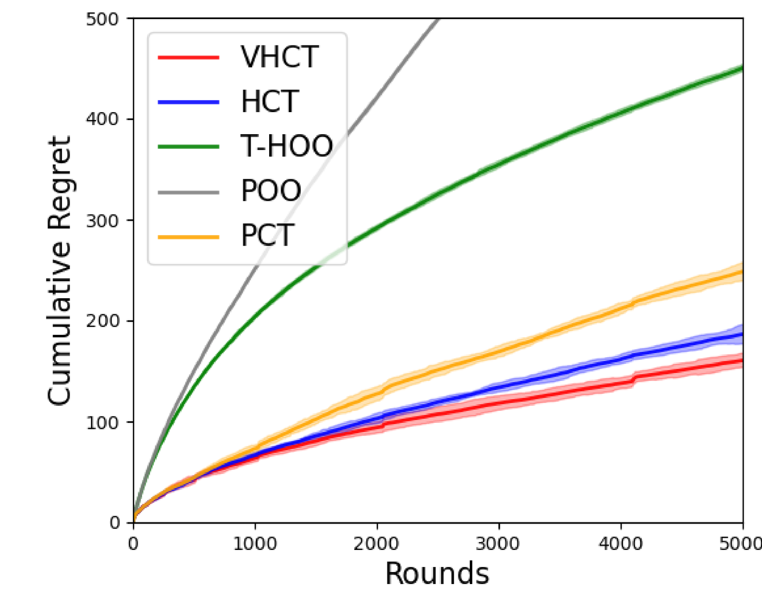

Figure 3: DoubleSine/Low-noise

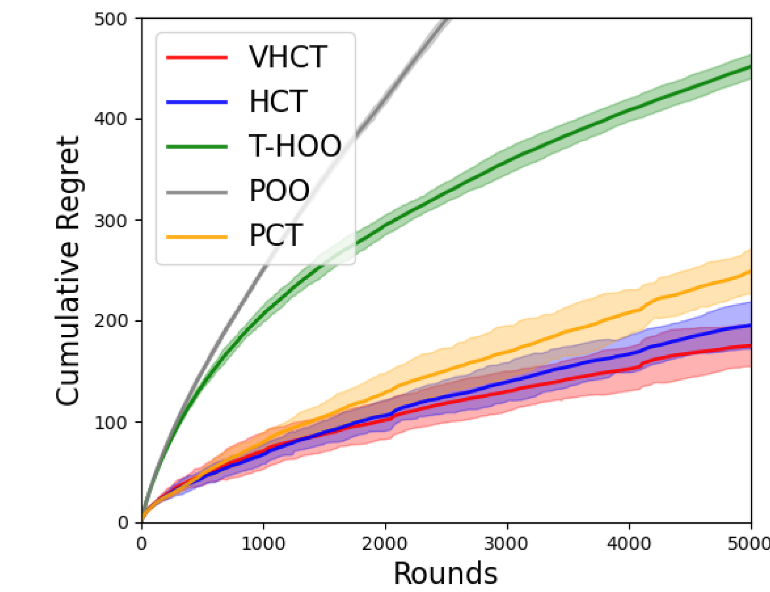

Figure 4: DoubleSine/Moderate-noise

- As shown in the figures, `VHCT` converges much faster than any other algorithms in both the low-noise setting and the moderate-noise setting.

References

[1] Mohammad Gheshlaghi Azar, Alessandro Lazaric, and Emma Brunskill. “Online stochastic optimization under correlated bandit feedback”. In: *International Conference on Machine Learning*. PMLR. 2014, pp. 1557–1565.

[2] Sébastien Bubeck et al. “<i>X</i>-Armed Bandits”. In: *Journal of Machine Learning Research* 12.46 (2011), pp. 1655–1695.

[3] Jean-Bastien Grill et al. “Black-box optimization of noisy functions with unknown smoothness”. In: *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2015.

[4] Dragan Komljenovic et al. “Asset Management in Electrical Utilities in the Context of Business and Operational Complexity”. In: *World Congress on Resilience, Reliability and Asset Management*. July 2019.

[5] Lisha Li et al. “Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization”. In: *Journal of Machine Learning Research* 18.185 (2018), pp. 1–52.

[6] Xuedong Shang, Emilie Kaufmann, and Michal Valko. “General parallel optimization a without metric”. In: *Algorithmic Learning Theory*. 2019, pp. 762–788.