Bayesian Optimization Meets Bayesian Optimal Stopping

Zhongxiang Dai [1] Haibin Yu [1] Bryan Kian Hsiang Low [1] Patrick Jaillet [2]

[1] Department of Computer Science, National University of Singapore

[2] Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology

Machine Learning (ML) models have achieved unprecedented level of performances

Problem:

Choice of **hyperparameters**

- **Reason**: high computational cost

Solution:

Bayesian Optimization (BO)

- Sample efficiency: requires a small number of function evaluations
- Theoretical guarantee: Gaussian Process-Upper Confidence Bound (GP-UCB)

- Many ML models require iterative training
  - e.g. stochastic gradient descent (neural network)

wasted computation

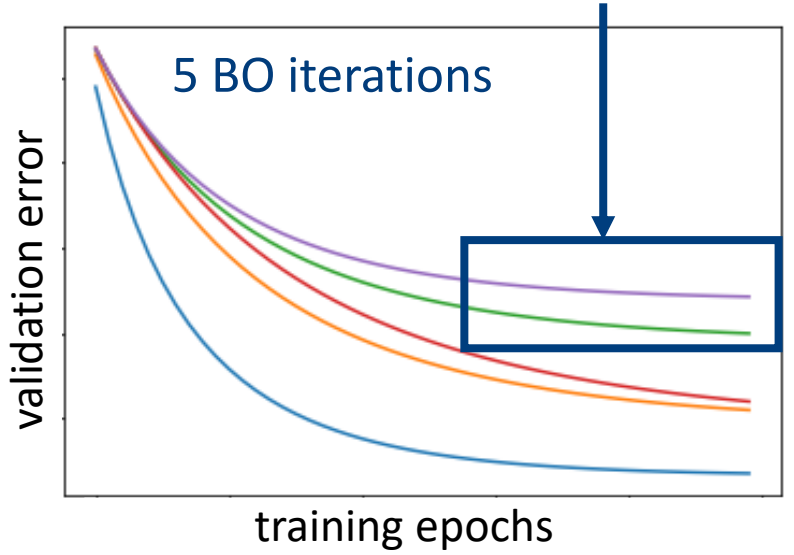

- Many ML models require iterative training
  - e.g. stochastic gradient descent (neural network)

Can we early-stop unpromising hyperparameters to save resource?

wasted computation

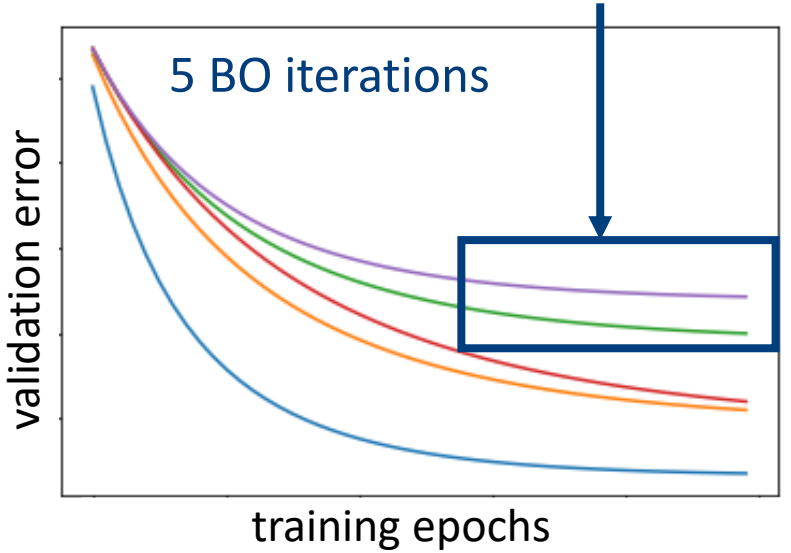

- Many ML models require iterative training
  - e.g. stochastic gradient descent (neural network)

wasted computation

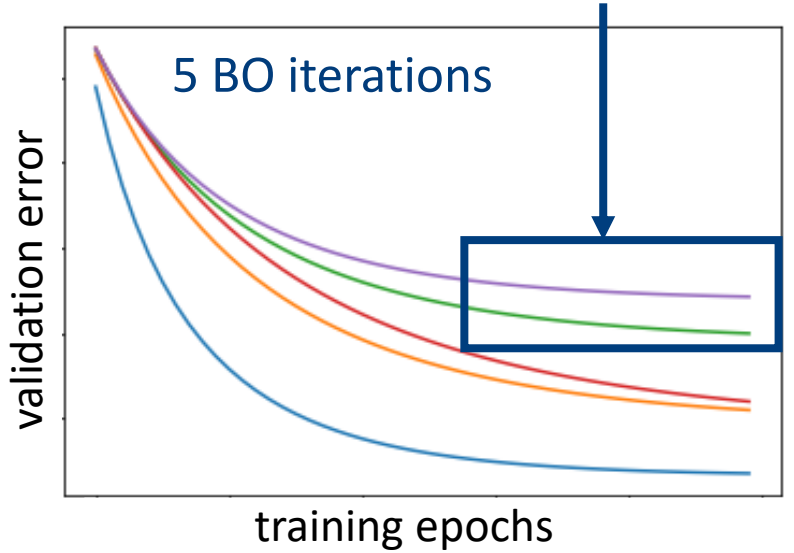

Can we early-stop unpromising hyperparameters to save resource?

a principled optimal stopping rule is needed

Bayesian Optimal Stopping (BOS)

- **Contribution**
  - We propose **BO-BOS**, which unifies BO (GP-UCB) and BOS to achieve epoch-efficient hyper-parameter tuning
  - We show that BO-BOS is asymptotically no-regret and performs effectively in practice

The BO-BOS Algorithm

- for $t = 1,2 \ldots$
  - choose the hyperparameter $\mathbf{x}_t$ to query by maximizing the GP-UCB acquisition function

The BO-BOS Algorithm

- for $t = 1,2 \ldots$
  - choose the hyperparameter $\mathbf{x}_t$ to query by maximizing the GP-UCB acquisition function
  - Formulate & solve the **BOS** problem

The BO-BOS Algorithm

- for $t = 1,2 \ldots$
  - choose the hyperparameter $\mathbf{x}_t$ to query by maximizing the GP-UCB acquisition function
  - Formulate & solve the **BOS** problem
  - Train the ML model using $\mathbf{x}_t$, early-stop the training **if BOS outputs the stopping decision**

BOS outputs the stopping decision if $\mathbf{x}_t$ is believed to end up performing worse than the currently found optimum

The BO-BOS Algorithm

- for $t = 1,2 \ldots$
  - choose the hyperparameter $\mathbf{x}_t$ to query by maximizing the GP-UCB acquisition function
  - Formulate & solve the **BOS** problem
  - Train the ML model using $\mathbf{x}_t$, early-stop the training **if BOS outputs the stopping decision**
  - Update Gaussian process belief of the objective function

BOS outputs the stopping decision if $\mathbf{x}_t$ is believed to end up performing worse than the currently found optimum

Theoretical Analysis

- BO aims to minimize the simple regret: $S_T \triangleq f(\mathbf{z}^*) - \max_{t \in \{1,\ldots,T\}} f(\mathbf{z}_t)$

expectation over the uncertainty introduced by BOS

$$\mathbb{E}[S_T] \leq \frac{\kappa\sqrt{TC_1\beta_T\gamma_T}}{T} + \frac{\sum_{t=1}^{T}\eta_t}{T} + \frac{1}{T}Nb\sqrt{\log\frac{da}{\delta'}\tau_T}$$

- Matches the simple regret of GP-UCB, thus asymptotically vanishes

- Asymptotically vanish if the BOS parameters are chosen with consideration of the **exploration vs. exploitation trade-off**

Experiments

- Hyperparameter Tuning
  - Convolutional Neural Networks

CIFAR-10 dataset

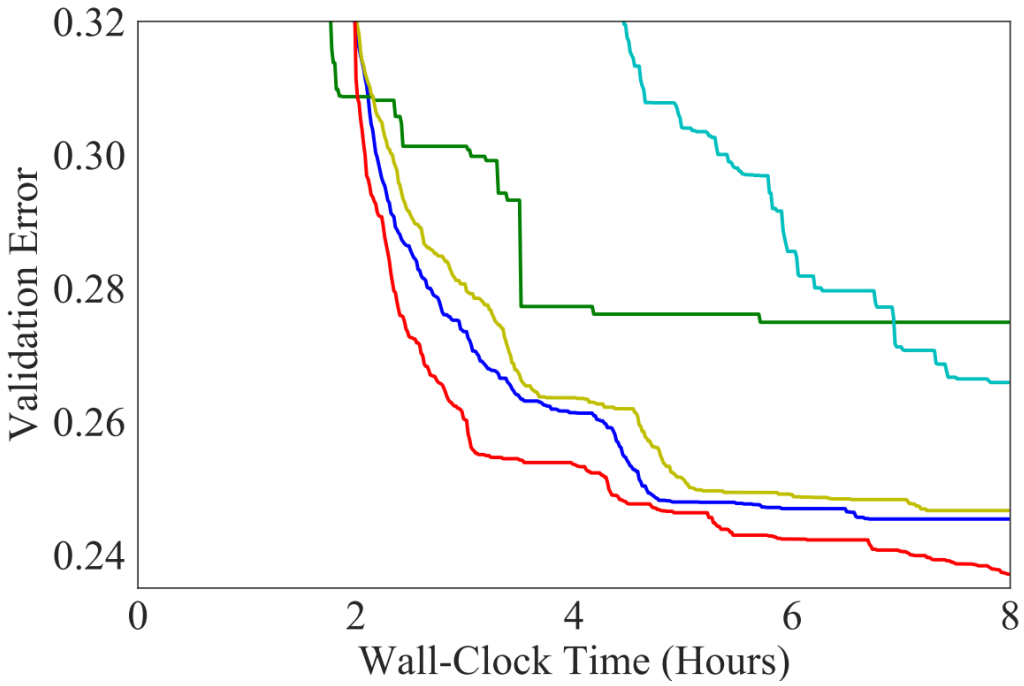

SVHN dataset

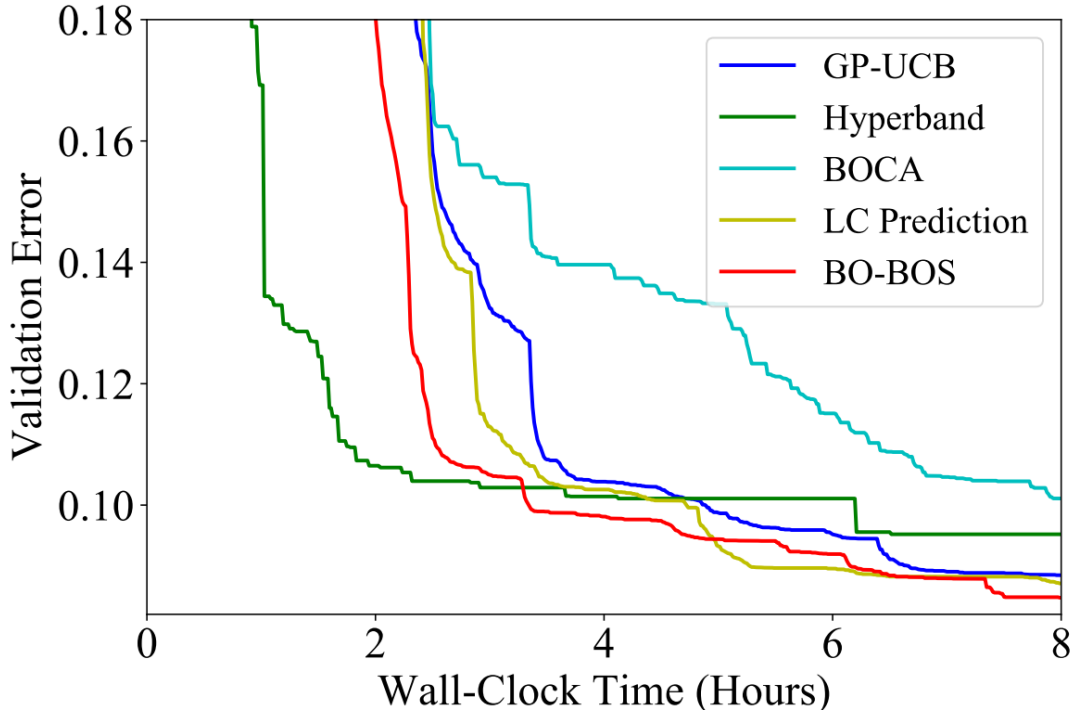

Experiments

- Beyond Hyperparameter Tuning:
  - Policy search for Reinforcement Learning (Swimmer-v2 from OpenAI, Mujoco)
  - Joint hyperparameter tuning and feature selection

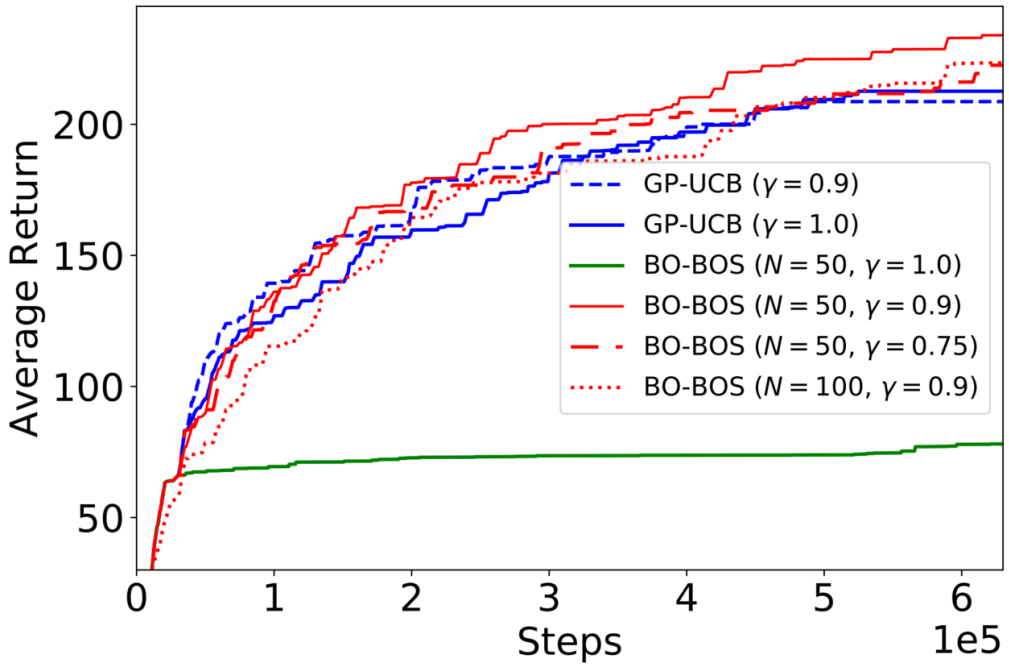

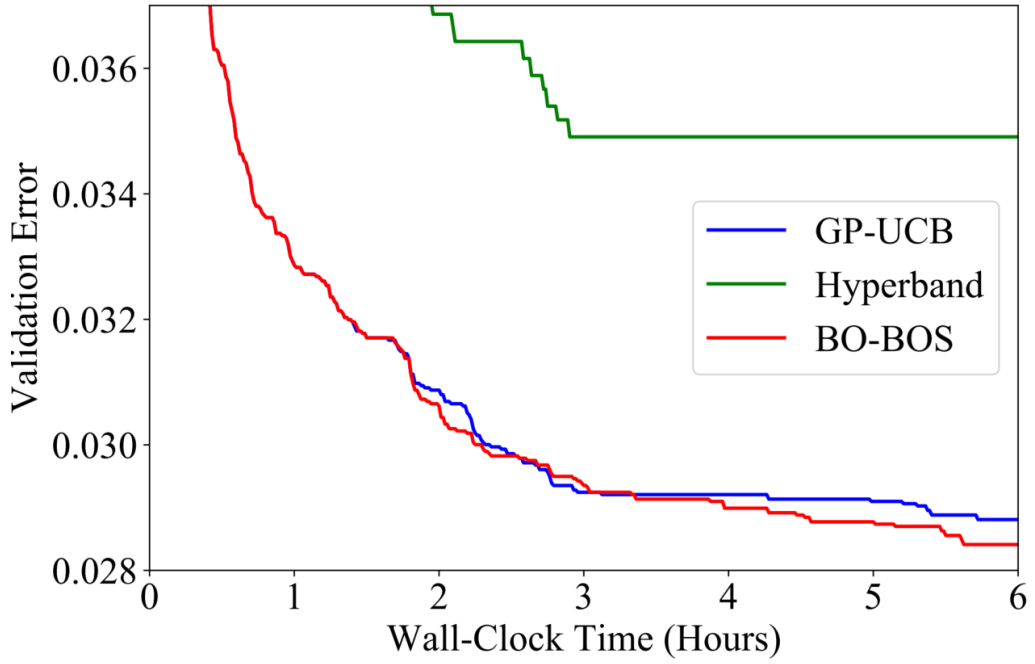

Thank you!

For more information, please visit our poster at **Pacific Ballroom #228**