Not All Explorations Are Equal: Harnessing Heterogeneous Profiling Cost for Efficient MLaaS Training

Jun Yi[1], Chengliang Zhang[2], Wei Wang[2], Cheng Li[3], Feng Yan[1]

[1]University of Nevada, Reno

[2]Hong Kong University of Science and Technology

[3]University of Science and Technology of China

Resource Acknowledgement

aws

OIT | Cyberinfrastructure

MLaaS Training

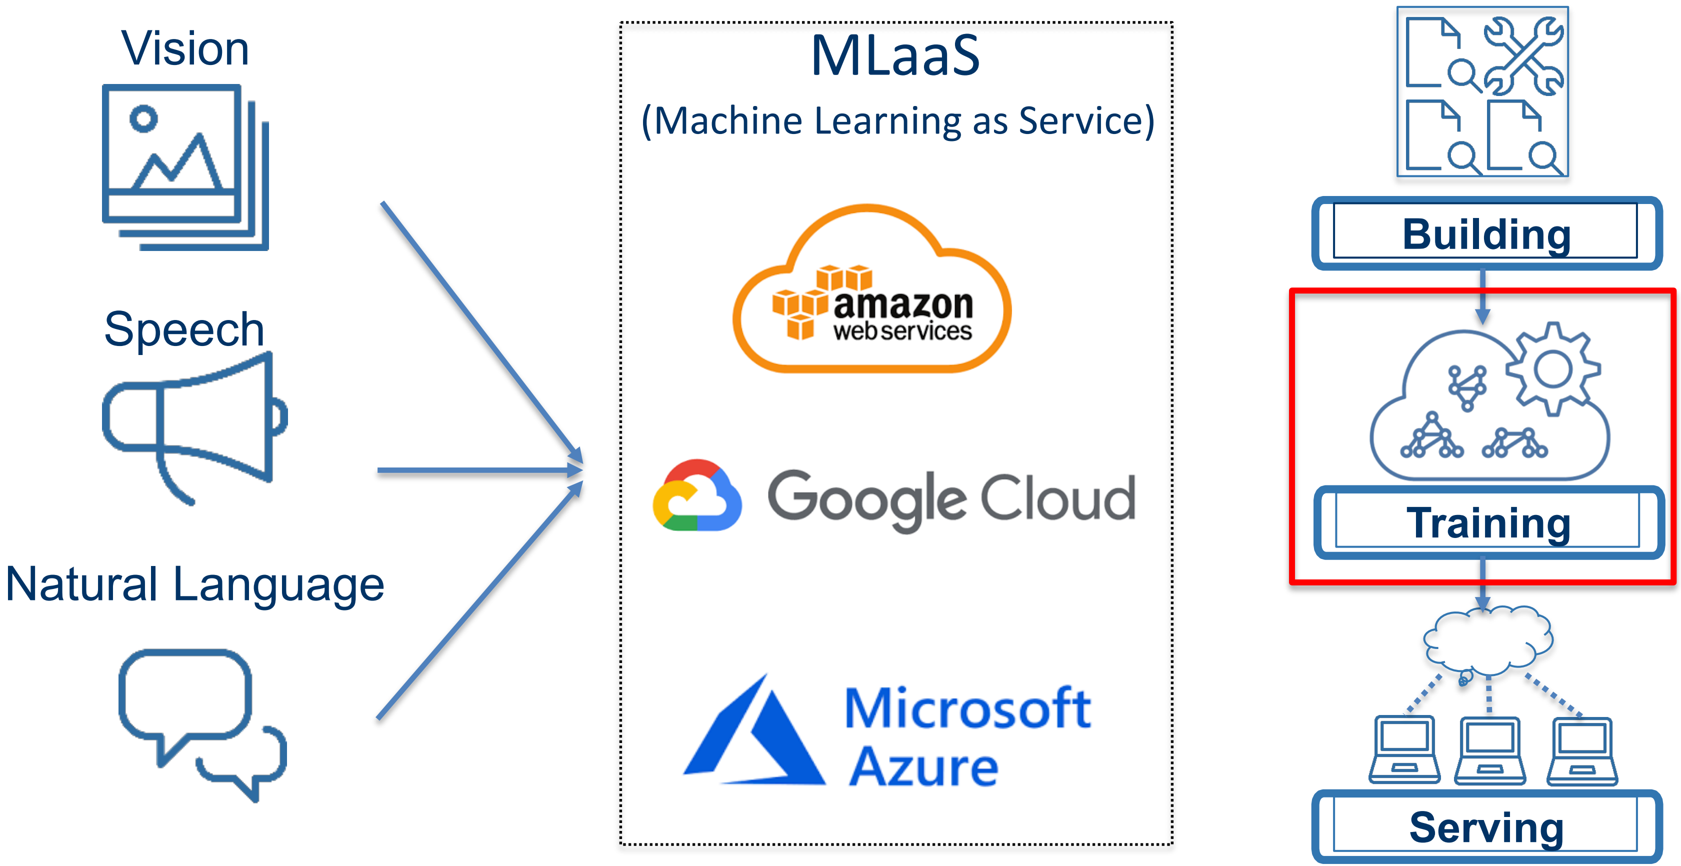

Problem

Practical MLaaS training scenarios:

- *Scenario-1*: Training project without time or cost limit
- *Scenario-2*: Training project with time limitation
- *Scenario-3*: Training project with cost limitation

How to deploy MLaaS training jobs in Cloud?

Scale-up (more capable instance)

VS scale-out (more instances)

E.g., use many cheapest instances (40 c5.4xlarge) or a few costly instances (9 p2.xlarge)?

Neither case is optimal (see the right figure)

Up to 42.5X in hourly cost

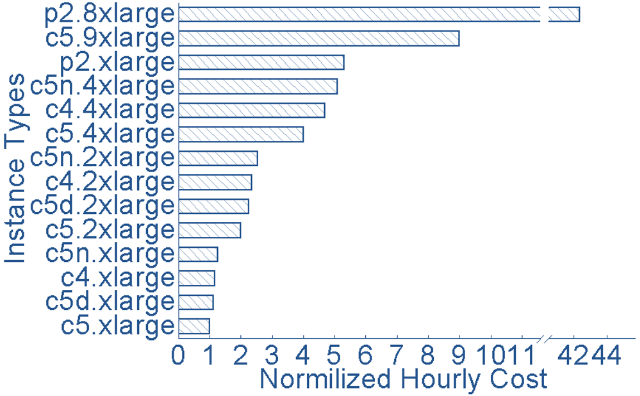

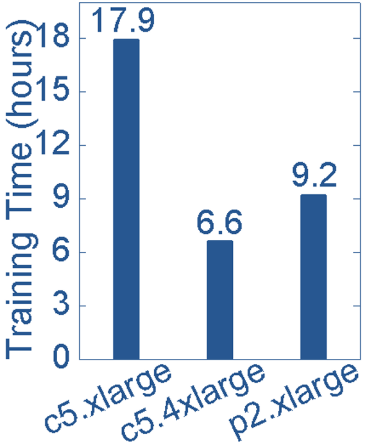

Challenges and Existing Work

Challenges: large deployment scheme search space (62 scale-up & 50 scale-out->3100 schemes)

Existing Work

Analytical Modeling (assumptions on model/hardware)	• Limited applicability (fast-evolving ML models) • Poor fit for cloud (increasing diversified hardware)	❖ [SIGKDD '15] *Performance modeling and scalability optimization of distributed deep learning systems* ❖ [ICLR, '17] *Paleo: A performance model for deep neural networks.*
Reinforcement Learning	• Requires extensive training samples and high computing resources	❖ [Nature '15] Human-level control through deep reinforcement learning.
Pareto-Optimization	• Falls short in performance	❖ [CCGRID '17] Predicting cloud performance for hpc applications: A user-oriented approach.
Conventional Bayesian Optimization (BO) (assume uniform profiling cost of every point)	• Assume uniform exploration cost • Lack of ML-specific insights	❖ [NSDI '17] CherryPick: Adaptively Unearthing the Best Cloud Configurations for Big Data Analytics. ❖ [ICDCS '18] Arrow: Low-level augmented bayesian optimization for finding the best cloud vm.

Key Observations and Main Idea

Conventional BO:

- For problems with unknow objective function
- Start with random initial points
- Select next points based on acquisition function
- Acquisition function optimizes expected improvement, probability of improvement, confidence bound, etc.

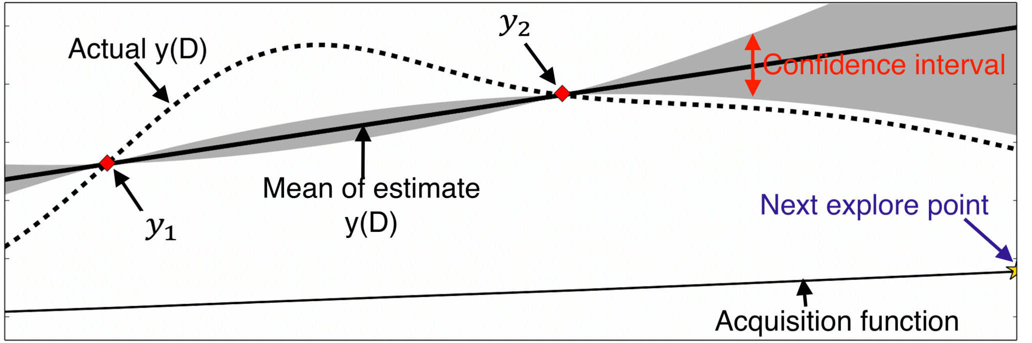

Key Observations

- ➢ Heterogenous exploration cost
  - Some schemes (i.e., large scale-out, high-end GPU instance) are more costly to explore than others
- ➢ No ML-specific prior is adopted in deployment optimization
  - Speedup trend of scale-out follows a concave-shape curve

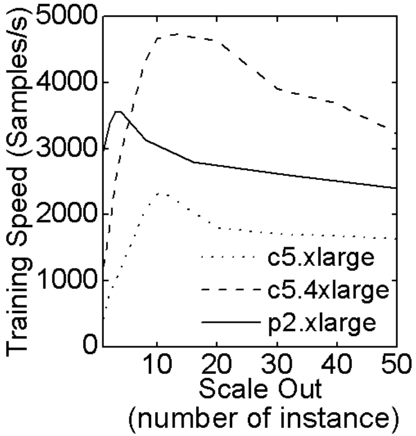

Main Idea: Heterogenous cost-aware and ML prior aware BO

HeterBO Overview

- Problem formulation

$$\text{minimize} \quad T(D)/C(D)$$
$$\text{subject to} \quad D \in D(m,n)$$

$T(D)$ - Total Time; $C(D)$ - Total Cost
$D(m,n)$ - Possible schemes; m - Instance type
n – Number of selected Instance type

- Search process

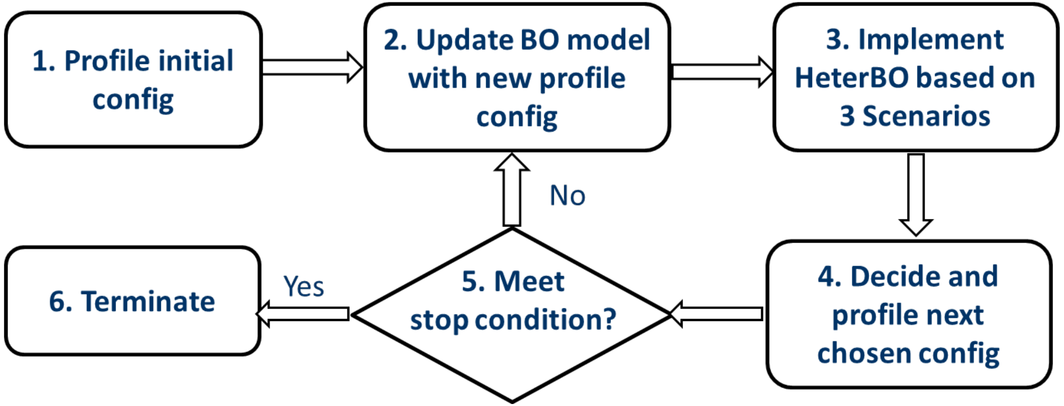

- Key Components
  - ➢ Prior function: Gaussian Process (flexibility and tractability)
  - ➢ Acquisition function: EI (Expected Improvements) with constraints (profiling cost) -> T(rue)EI
  - ➢ Heterogeneous search cost aware: avoid randomly jumping into high profile cost regions
  - ➢ ML-specific aware: detects down slope of the concave-shape -> avoid high overheads

HeterBO Example

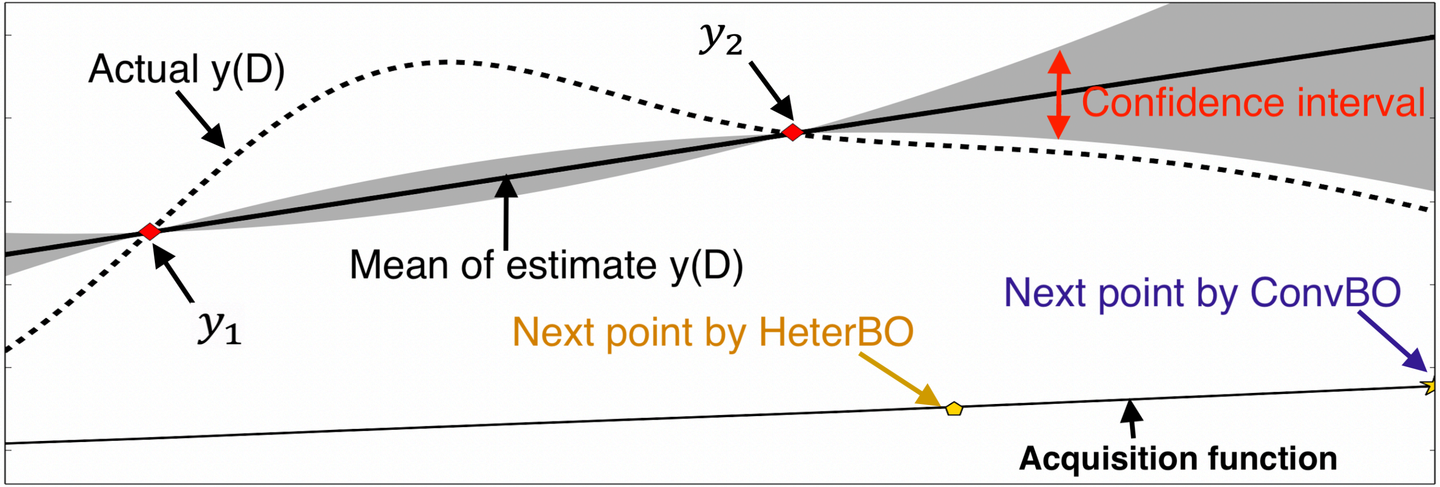

- y_1 and y_2 are profiled points
- Not select the maximum point in acquisition function as next point (i.e., ConvBO)
- HeterBO considers the ***user constraints*** and ***heterogeneous search cost*** when selecting next point (35% less profiling cost)

MLCD System Design

MLaaS training Cloud Deployment system (MLCD):

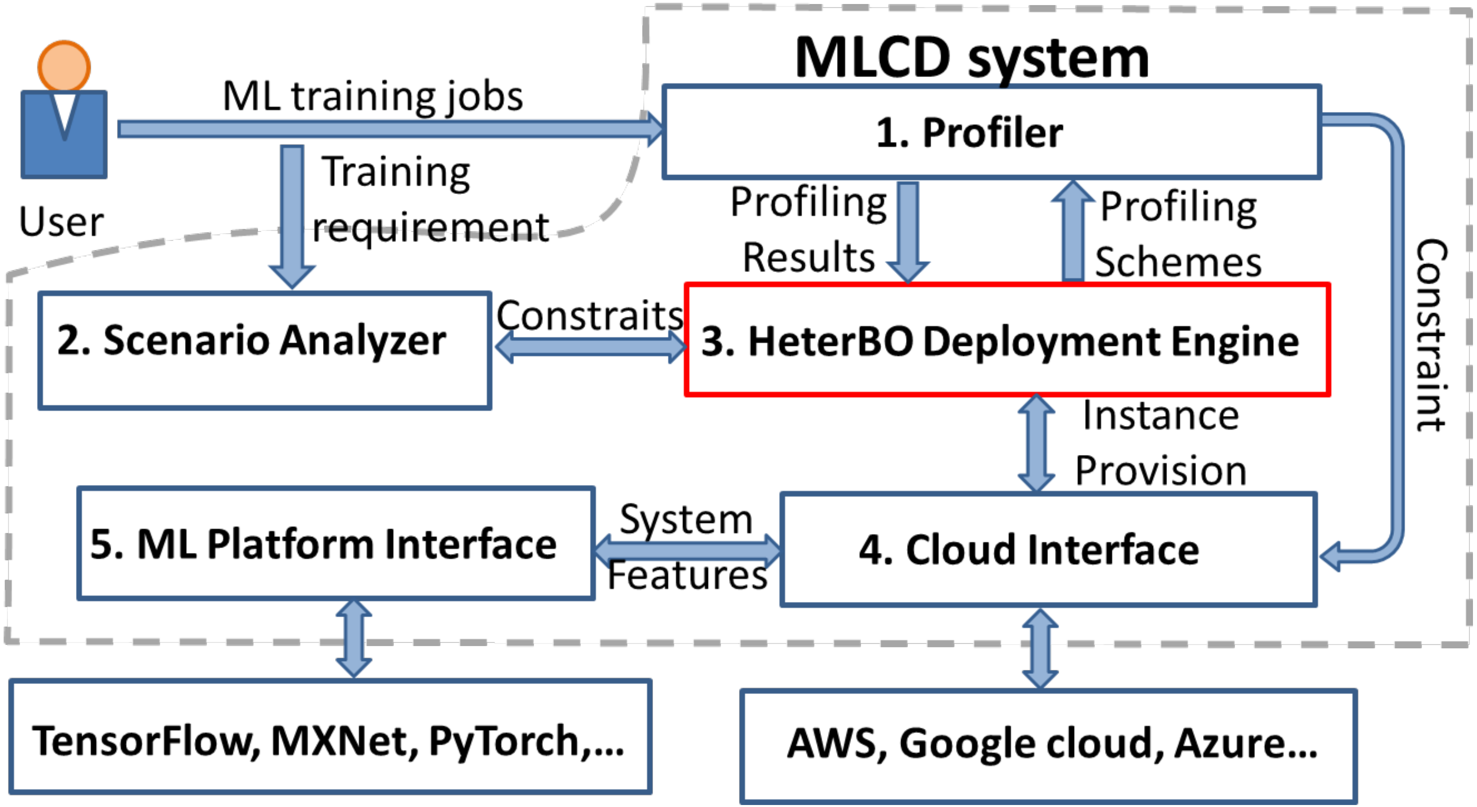

Eval Setup & Performance Comparison

Testbed
- AWS CPU, GPU instances

ML platforms
- TensorFlow and MXNet

ML Models
- AlexNet, ResNet, Inception-v3, CharCNN, BERT

HeterBO vs. Existing Approaches using TensorFlow

Limited monetary budget ($80) scenario

Total Cost

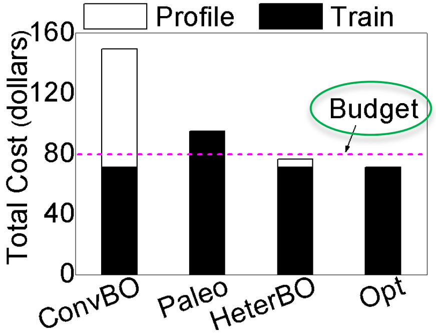

Total Time

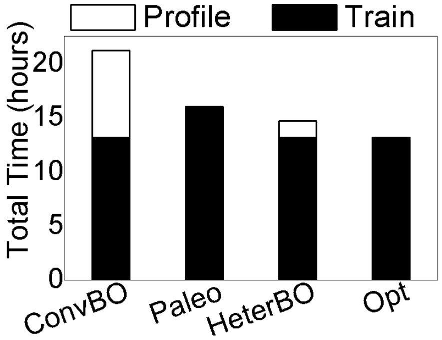

HeteBO costs under budget (ConvBO/Paleo not)

36.4% and **12.5%** better than ConvBO and Paleo in Total Time

Limited total time (20 hours) scenario

Total Cost

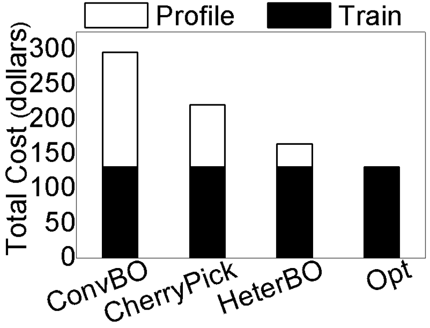

Total Time

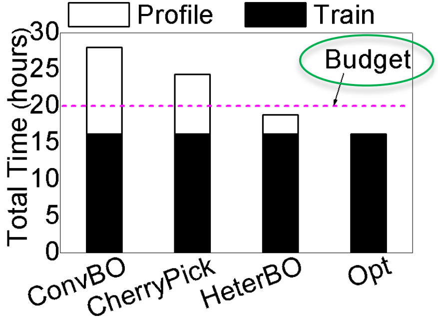

HeterBO finishes on time (ConvBO/CheeryPick not)

44.8% and **28.9%** better than ConvBO and CheeryPick in Total Cost

Robustness and Adaptivity

Total cost vs Budget

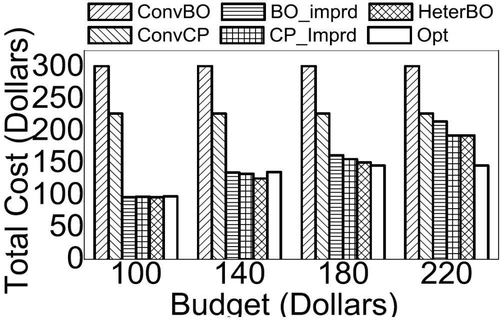

HeterBO outperforms SOTA by up to **3.1×**

Total Time vs Budget

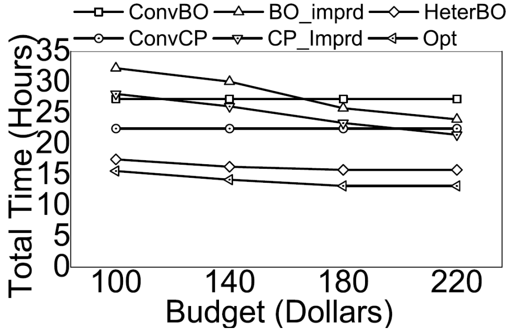

HeterBO outperforms SOTA by up to **2.34×**

Char-RNN using TensorFlow

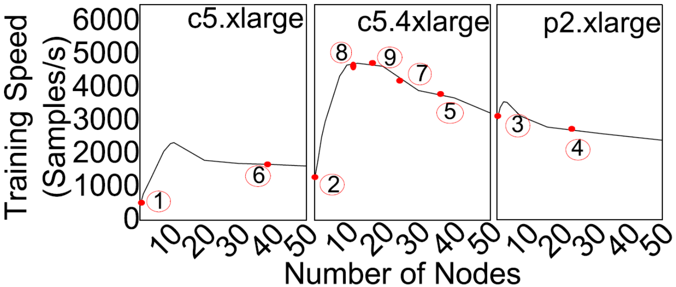

HeterBO found optimal within budget $120

BERT using MXNet

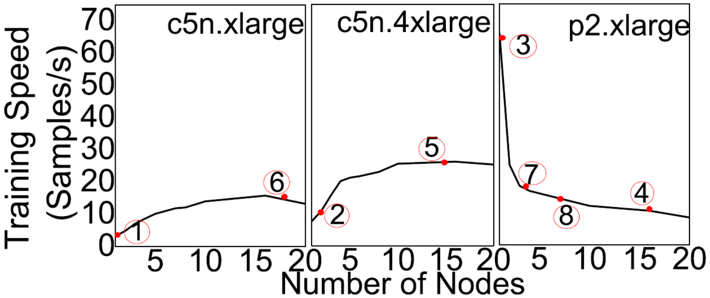

HeterBO found optimal within budget $120

Takeaway:

Not all explorations are equal: heterogeneous exploration cost + machine learning specific prior

→ A fully-automated MLaaS training Cloud Deployment system (**MLCD**) driven by **HeterBO** search method

Jun Yi

junyi@nevada.unr.edu, https://www.cse.unr.edu/~jyi/

https://www.youtube.com/channel/UCMgXRQdpjlmc5GLkGV0Av8g?view_as=subscriber

NSF

aws

OIT | Cyberinfrastructure

Research Grants Council of Hong Kong
香港 研究資助局

N S F C
National Natural Science Foundation of China