BOAT

BespOke Auto-Tuners

with Structured Bayesian Optimization
Dalibard · Schaarschmidt · Yoneki | WWW 2017

Uli Prantz
LSDP · Probabilistic Programming · 11 November 2026

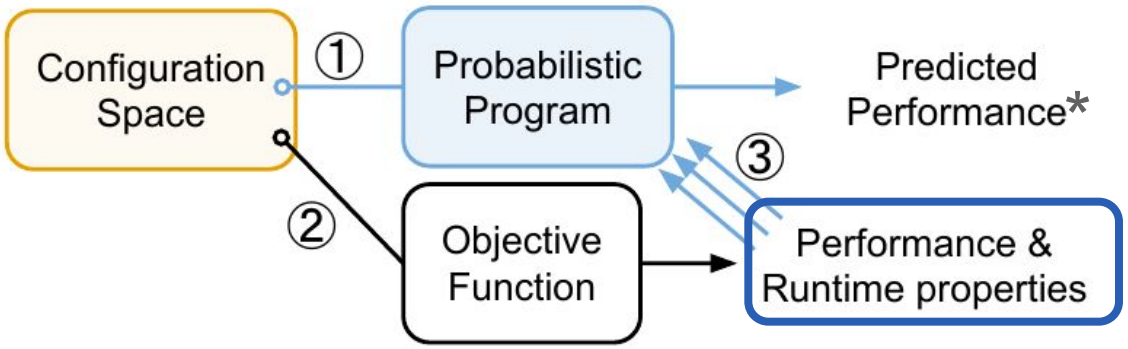
Tuning distributed neural-network training

Synchronous stochastic gradient descent (SGD)

1. Which machines **compute gradients**?
2. Which machines act as **parameter servers**?
3. How much work goes to each **CPU or GPU**?

10 heterogeneous machines · **30-32 configuration parameters**

BOAT learns **more** from each system run



Structured Bayesian optimisation (SBO)

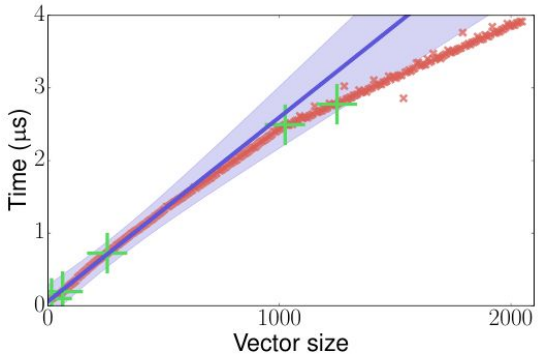
1. Suggest a configuration
2. Run the real system
3. Update the model

One costly run can supply **many useful measurements**

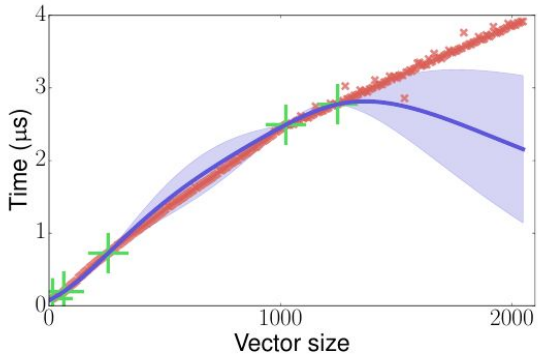
*Predictions include uncertainty

Known trend + learned correction

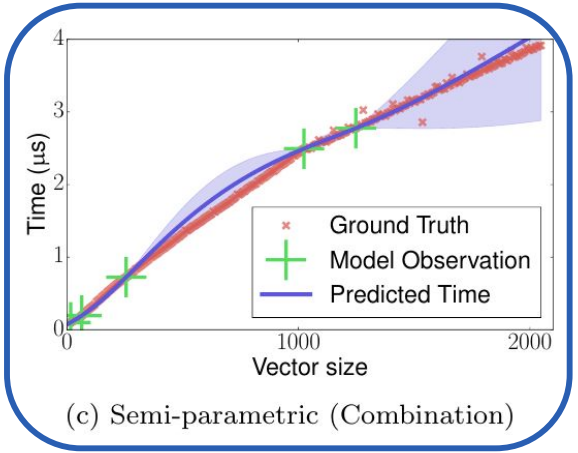
One component = **parametric trend** + **Gaussian process (GP)**



(a) Parametric (Linear regression)



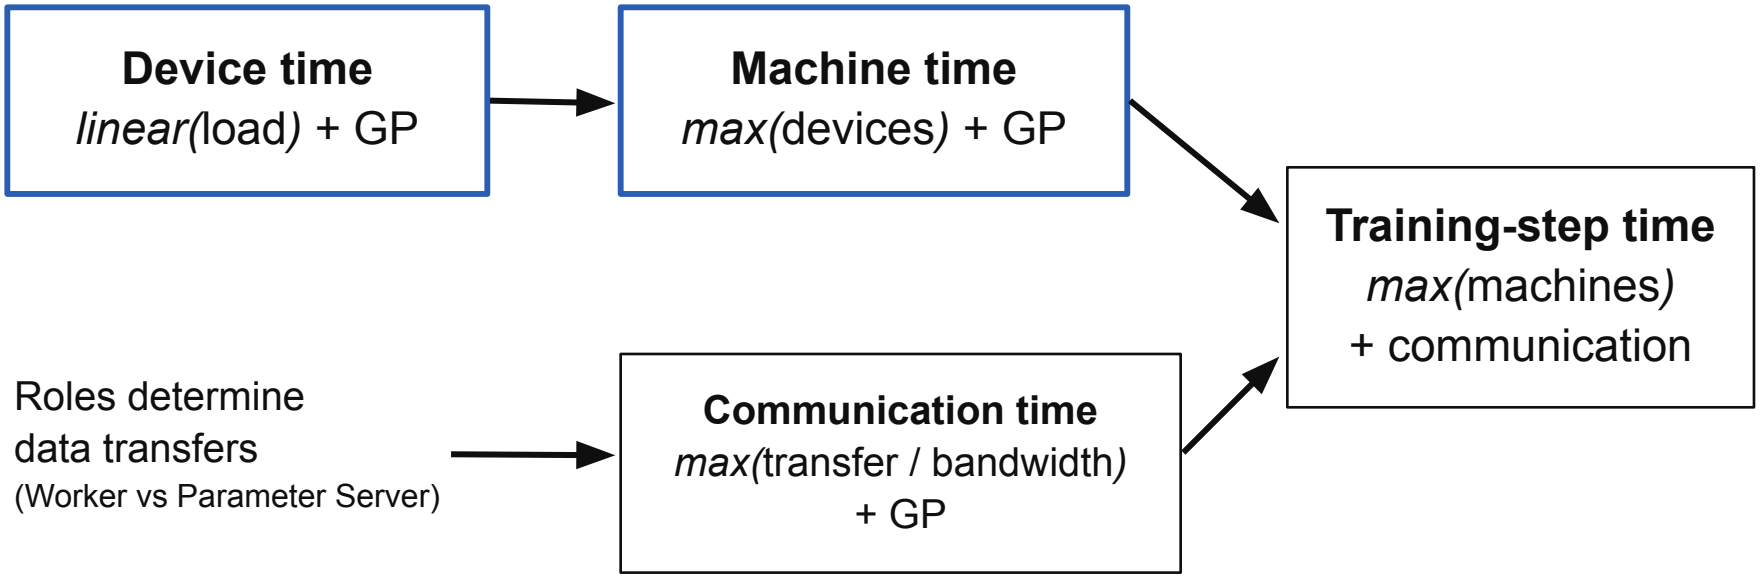
(b) Non-parametric (Gaussian process)



(c) Semi-parametric (Combination)

- The GP learns deviations from the trend
- The **mean and kernel** determine how it extrapolates

Modelling the inputs of a distributed training step



Measured **device and machine times** update the corresponding *components*

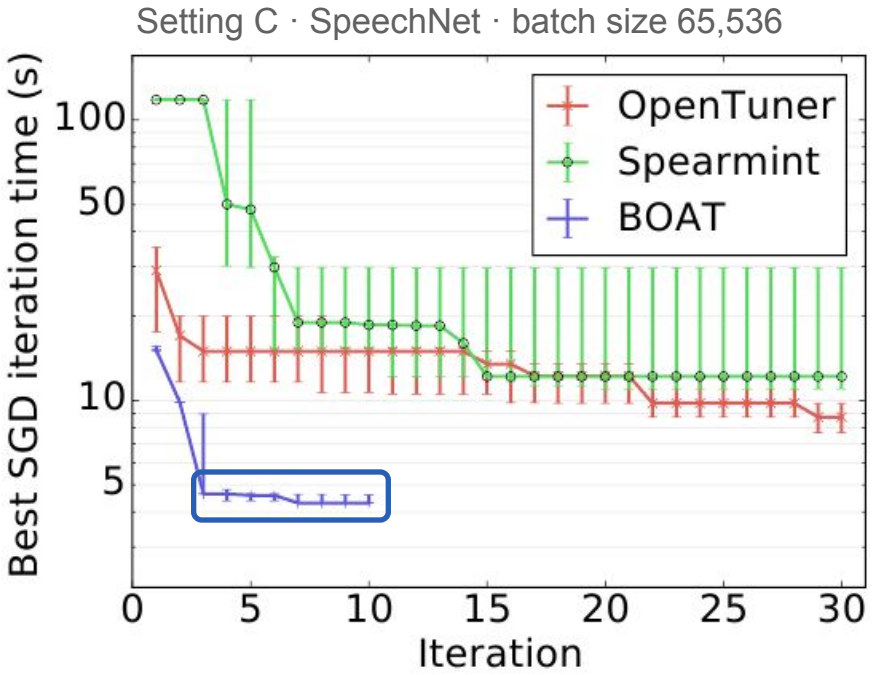
Evaluation setup

Amazon EC2 · TensorFlow 0.8 · fixed NN and batch per tuning task

1. **3 clusters:** 10 machines with 0, 1 or 2 GPUs
2. 3 NNs × 4 batch sizes × 3 clusters = **36 settings**
3. 20 SGD steps per trial; **average the last 10**

10 BOAT trials; each tuning run finished within **2 hours**

Better configurations after fewer trials



BOAT

4.31 s after **10 trials**

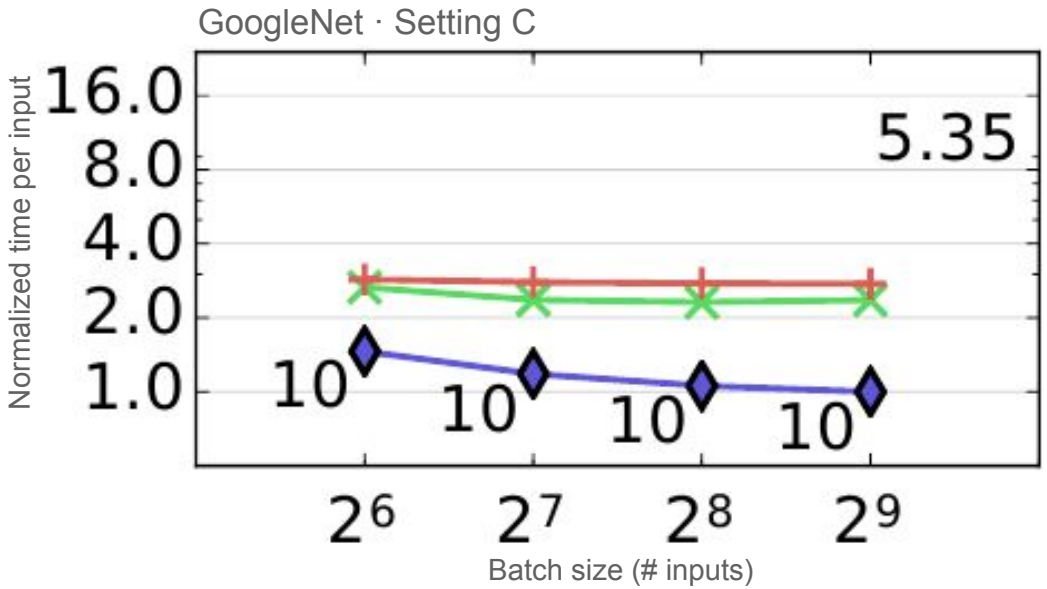
OpenTuner

8.71 s after **30 trials**

Median + min/max
3 runs
log y-axis

Complete tuner comparison → sources of the gain are not isolated

Load balance matters even with 10 workers

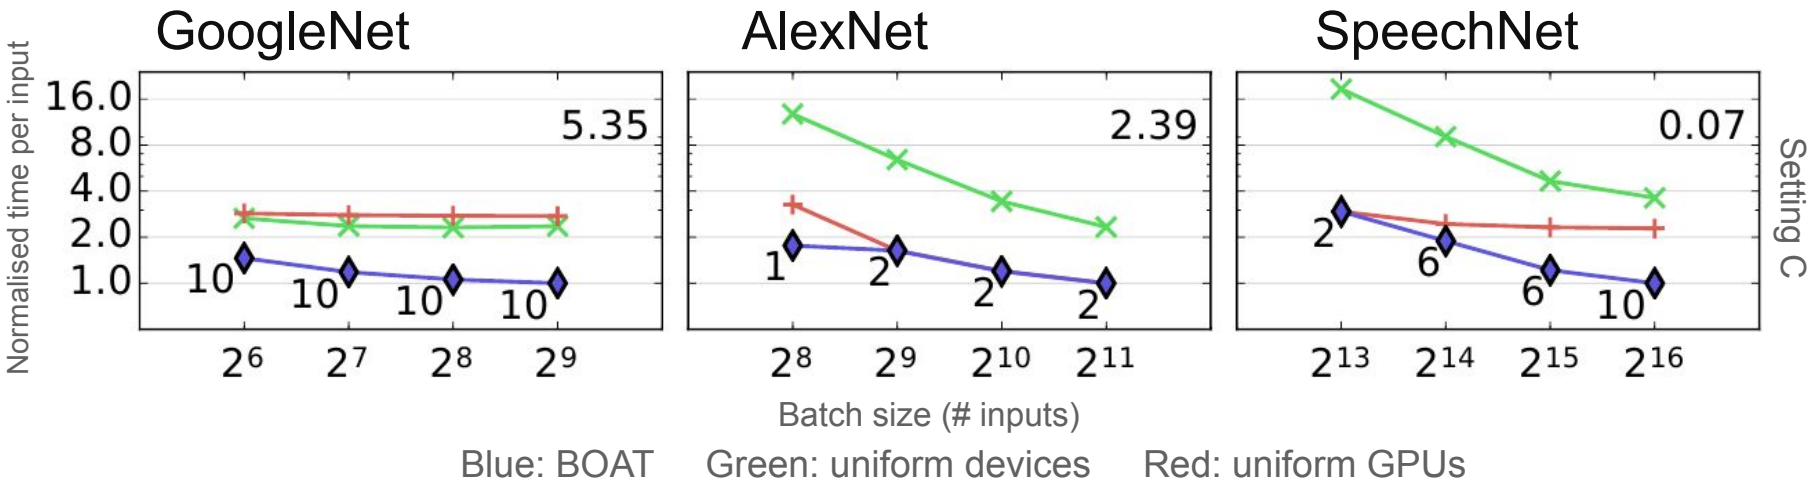


All **10** machines used as workers

Equal work per device is slower than the optimised allocation

Blue: BOAT Green: uniform devices Red: uniform GPUs

Useful worker counts depend on the workload



Blue: BOAT Green: uniform devices Red: uniform GPUs

10 workers **1-2** workers **2 → 10** workers

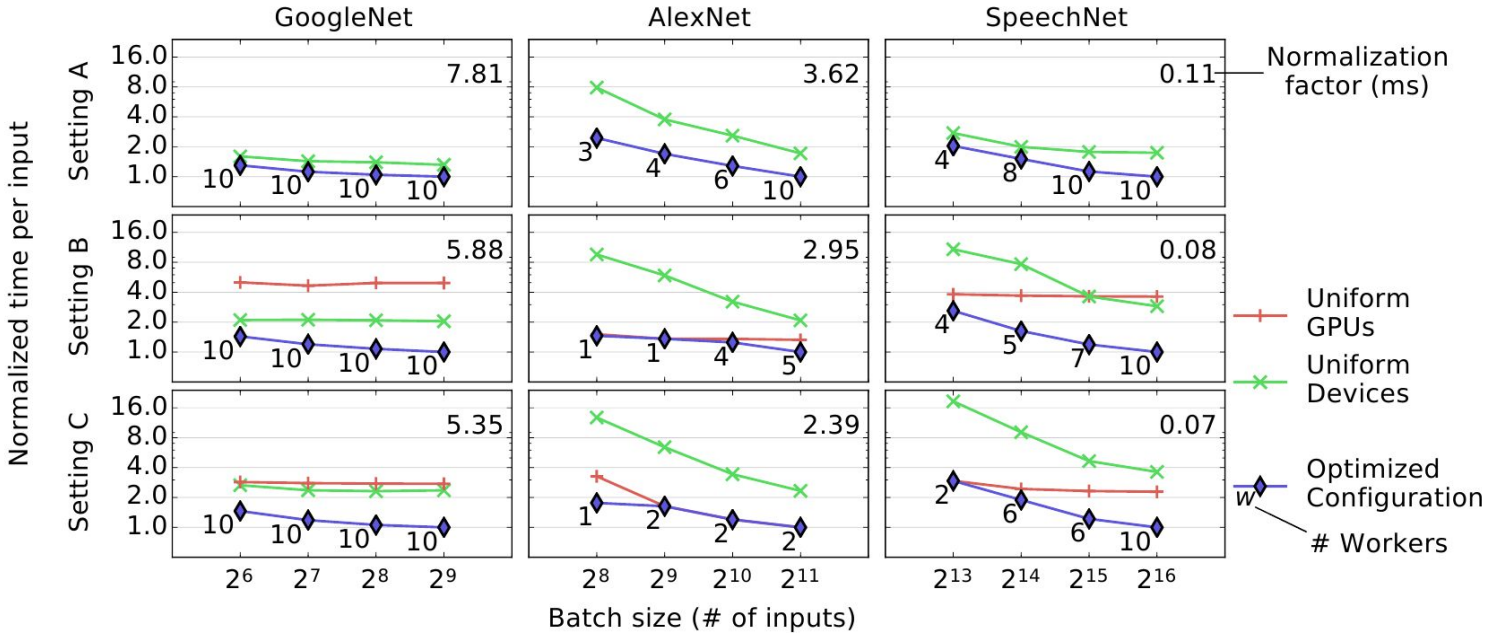
→ More computation can amortise the **cost of communication**

What BOAT contributes

1. **Domain knowledge** simplifies what must be learned
2. **Internal timings** are valuable information of every expensive trial
3. BOAT makes this usable in a **system-specific auto-tuner**

Backup: all 36 training settings

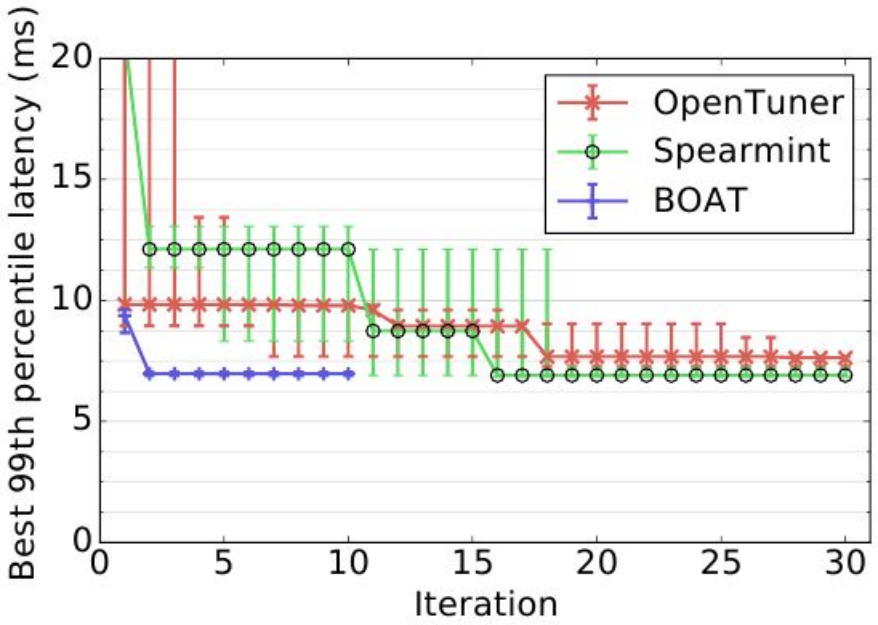
3 hardware settings × 3 networks × 4 batch sizes



Each panel has its **own normalisation**; point labels count workers

Backup: garbage-collection tuning

Cassandra · Java garbage collection (GC) · 99th-percentile latency



3 configuration flags
15 minutes per trial

BOAT: good values within **2 trials**

Spearmint: about **16**

Median, min and max across 3 tuning runs

Backup: novelty and modelling effort

2006: parametric trends + GP residuals

2015: additive structure in Bayesian optimisation

2017: BOAT connects models, measurements and tuning

2021: RocksDB study reports substantial integration work