ARTIFACT EVALUATED usenix ASSOCIATION AVAILABLE

ARTIFACT EVALUATED usenix ASSOCIATION FUNCTIONAL

ARTIFACT EVALUATED usenix ASSOCIATION REPRODUCED

Hydro: Surrogate-based Hyperparameter Tuning Service in Datacenters

Qinghao Hu[1,2], Zhisheng Ye[2,3], Meng Zhang[1,2], Qiaoling Chen[2,4], Peng Sun[2,5], Yonggang Wen[1], Tianwei Zhang[1]

1 NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE

2 上海人工智能实验室 Shanghai Artificial Intelligence Laboratory

3 北京大學 PEKING UNIVERSITY

4 NUS National University of Singapore

5 商汤 sensetime

Background

What is Hyperparameter Tuning?

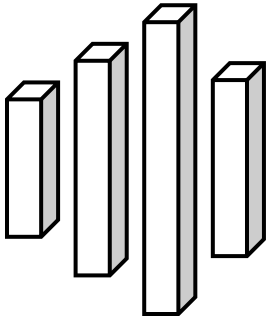

ResNet, GPT...

```
General:
    learning_rate=0.01
    batch_size=256
    weight_decay=0.01

Optimizer=SGD(momentum=0.5)
  /Adam(betas=(0.9, 0.99))

LR_Scheduler=Step(gamma=0.1)
  /CosineAnnealing(T_max=10)
```

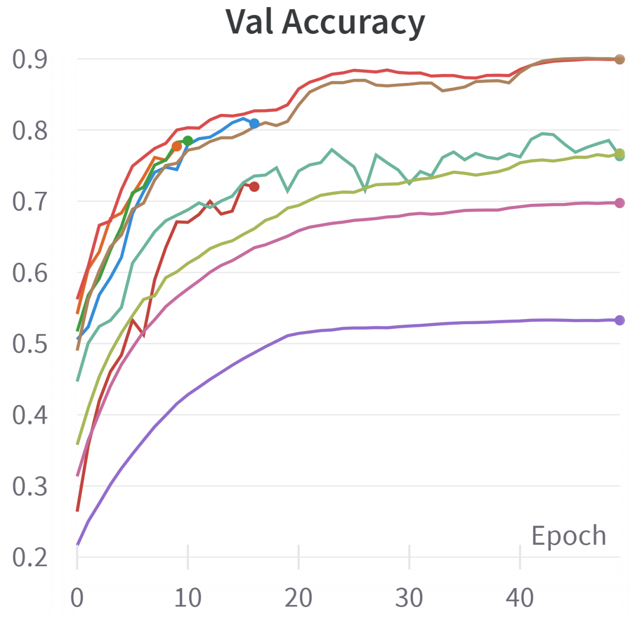

Model ▶ **Hyperparameter Recipes** ▶ **Best Configuration**

Effect of Hyperparameter Tuning

PyTorch v1.10 released an updated version of their official model weights

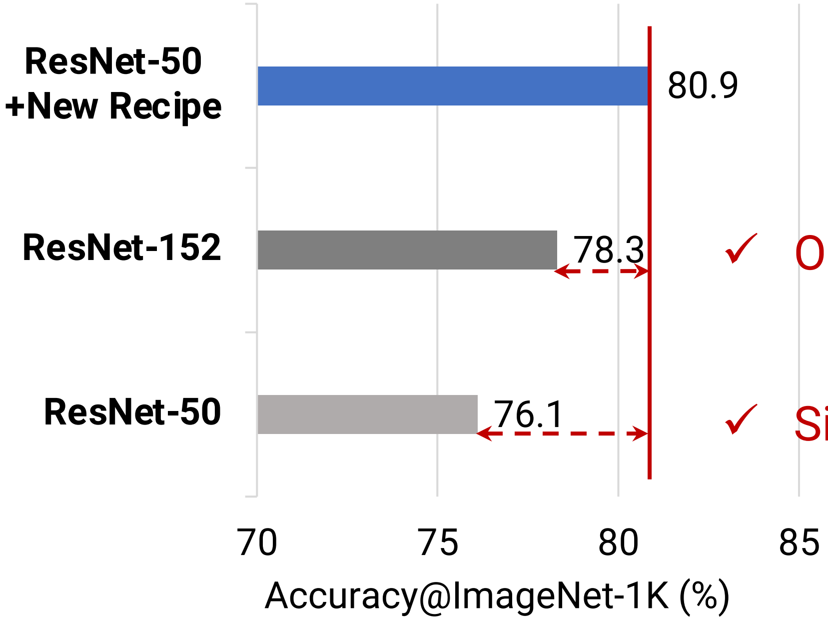

✓ Outperform Larger Model

✓ Significant 4.8% Improvement

Source: PyTorch Blog

Challenge 1: High Tuning Cost of Large Models

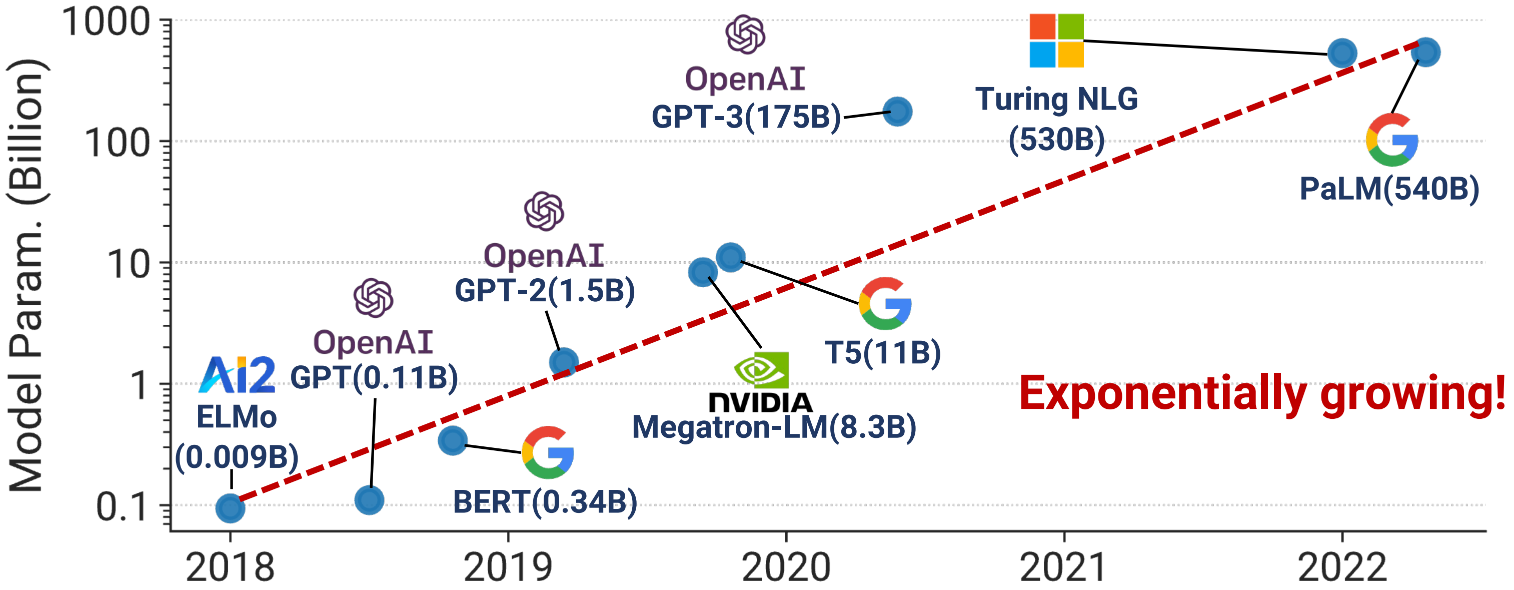

The cost of tuning large models is unacceptable

→ lead to subpar model quality

Challenge 2: Inefficient Resource Usage

Tuning jobs consume substantial resources from enterprise & institute clusters

Microsoft — 90% of models require tuning, 75 trials in median [1]

Alibaba — 65% of jobs repeatedly run ≥5 times [2]

商汤 sensetime — 90% of jobs are repetitive for tuning or debugging [3]

VECTOR INSTITUTE — 46% of GPU hours contribute to singe-GPU tuning jobs [4]

Source: [1] Themis (NSDI '20) [2] MLaaS (NSDI '22) [3] Lucid (ASPLOS '23) [4] HFTA (MLSys '21)

Challenge 2: Inefficient Resource Usage

GPUs are significantly underutilized

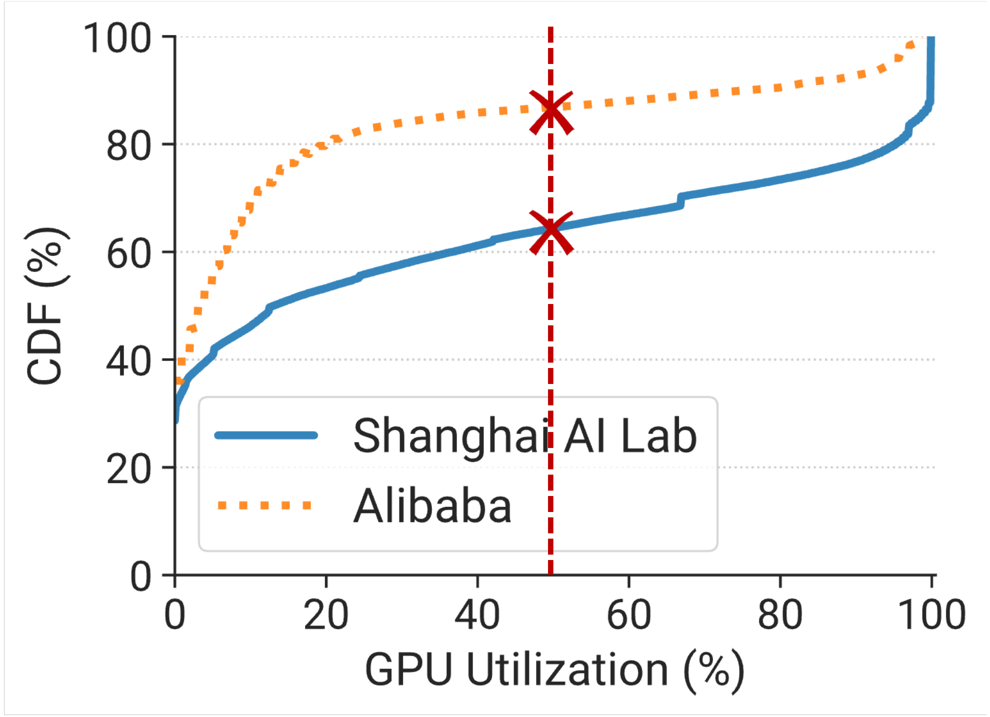

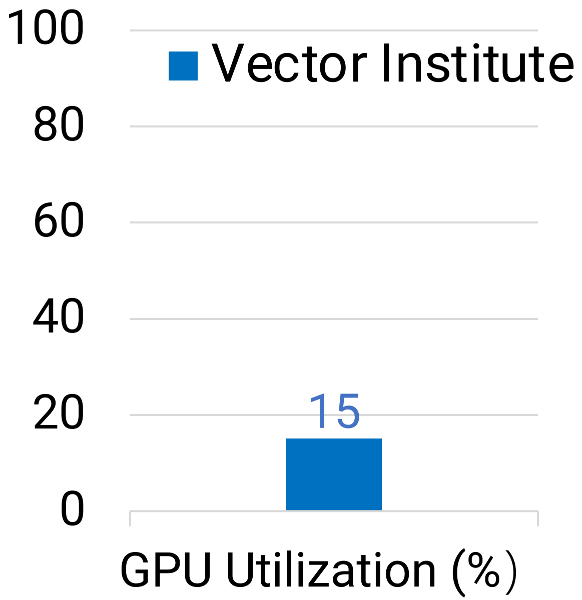

Source: [1] Themis (NSDI '20) [2] MLaaS (NSDI '22) [3] Lucid (ASPLOS '23) [4] HFTA (MLSys '21)

Hydro

Job-level ***Hydro Tuner***

Automatically generate surrogate models for tuning by applying transfer theory and model fusion

Datacenter-level ***Hydro Coordinator***

Leverage idle bubble resources of pretraining jobs via interleaving training

Key Mechanism: Surrogate-based Tuning

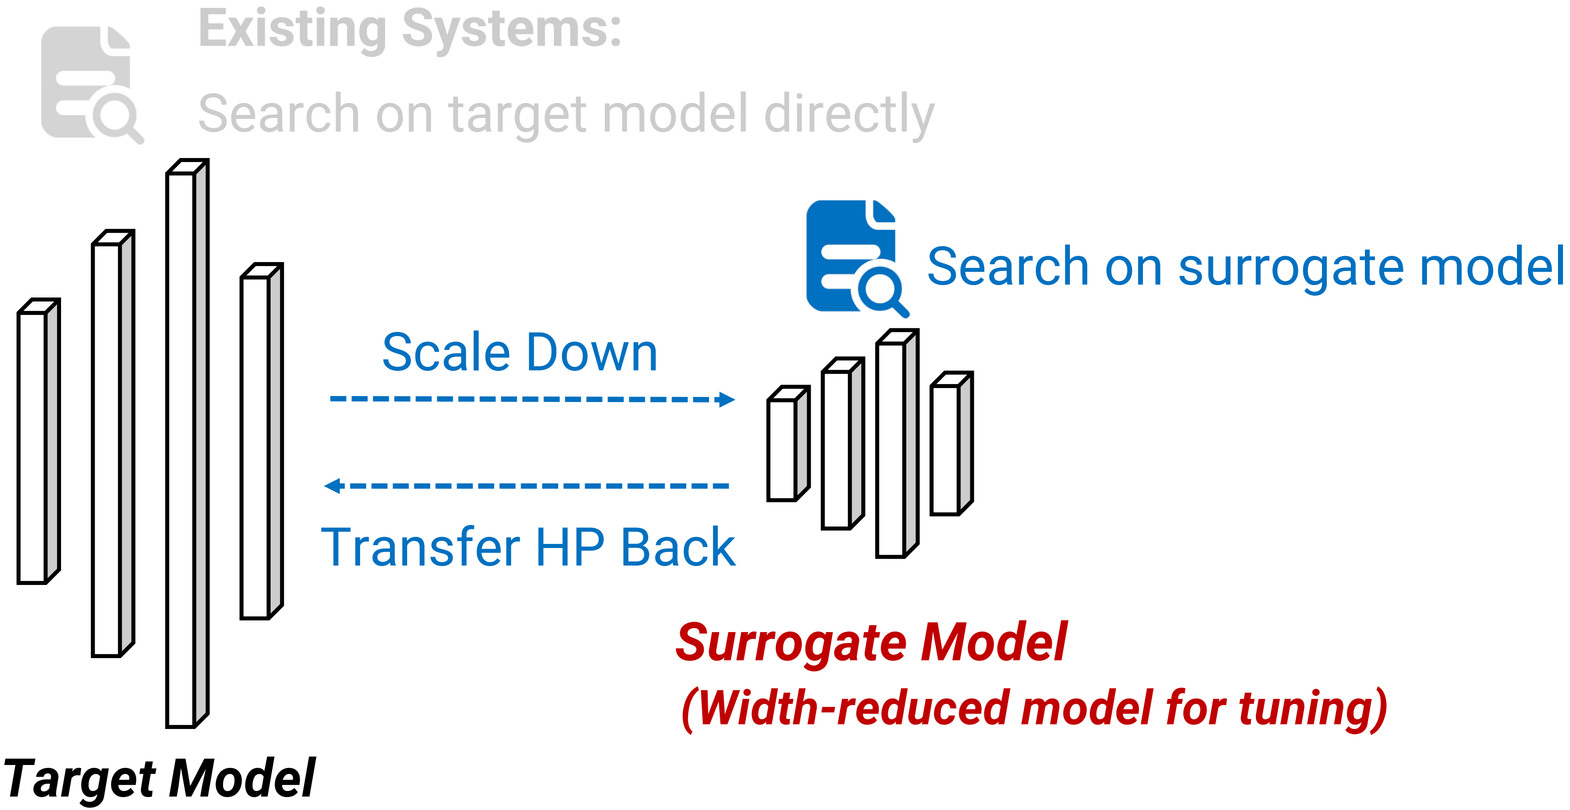

Can Hyperparameters be Transferred?

Toy Example: 2-layer Transformer model over WikiText-2 dataset using Adam

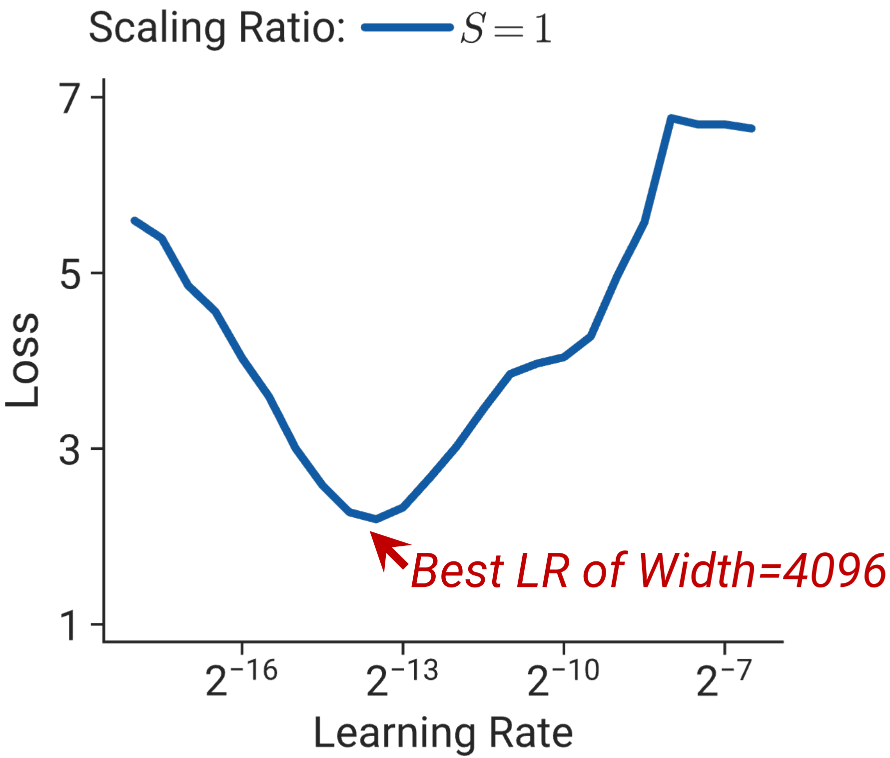

Can Hyperparameters be Transferred?

Toy Example: 2-layer Transformer model over WikiText-2 dataset using Adam

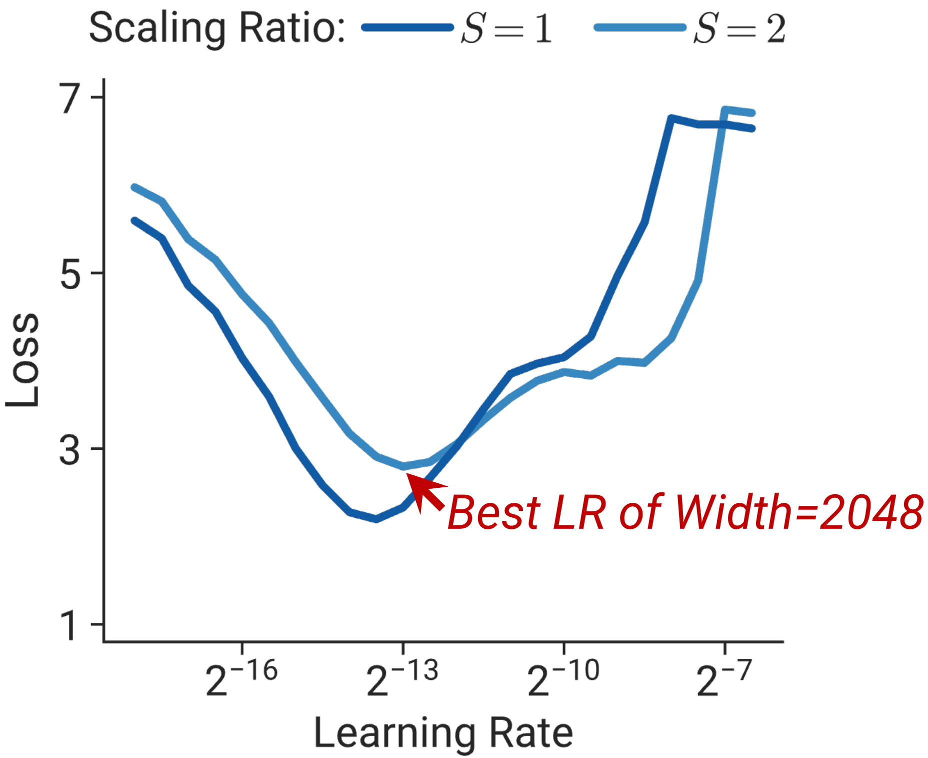

Can Hyperparameters be Transferred?

Toy Example: 2-layer Transformer model over WikiText-2 dataset using Adam

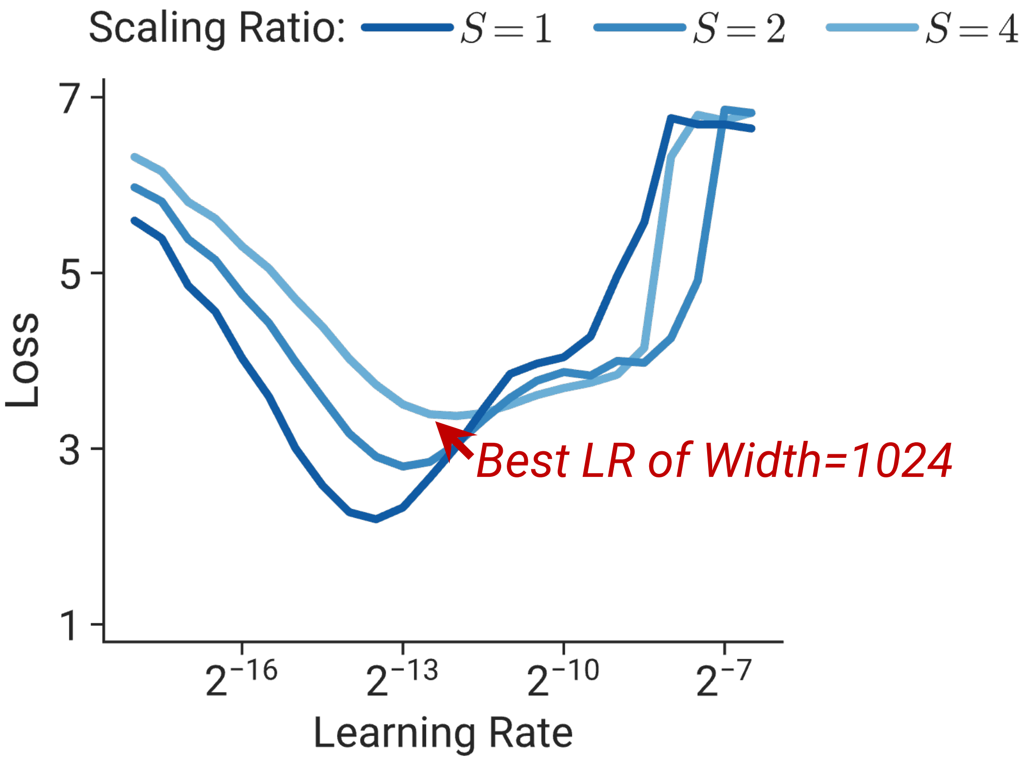

Can Hyperparameters be Transferred?

Toy Example: 2-layer Transformer model over WikiText-2 dataset using Adam

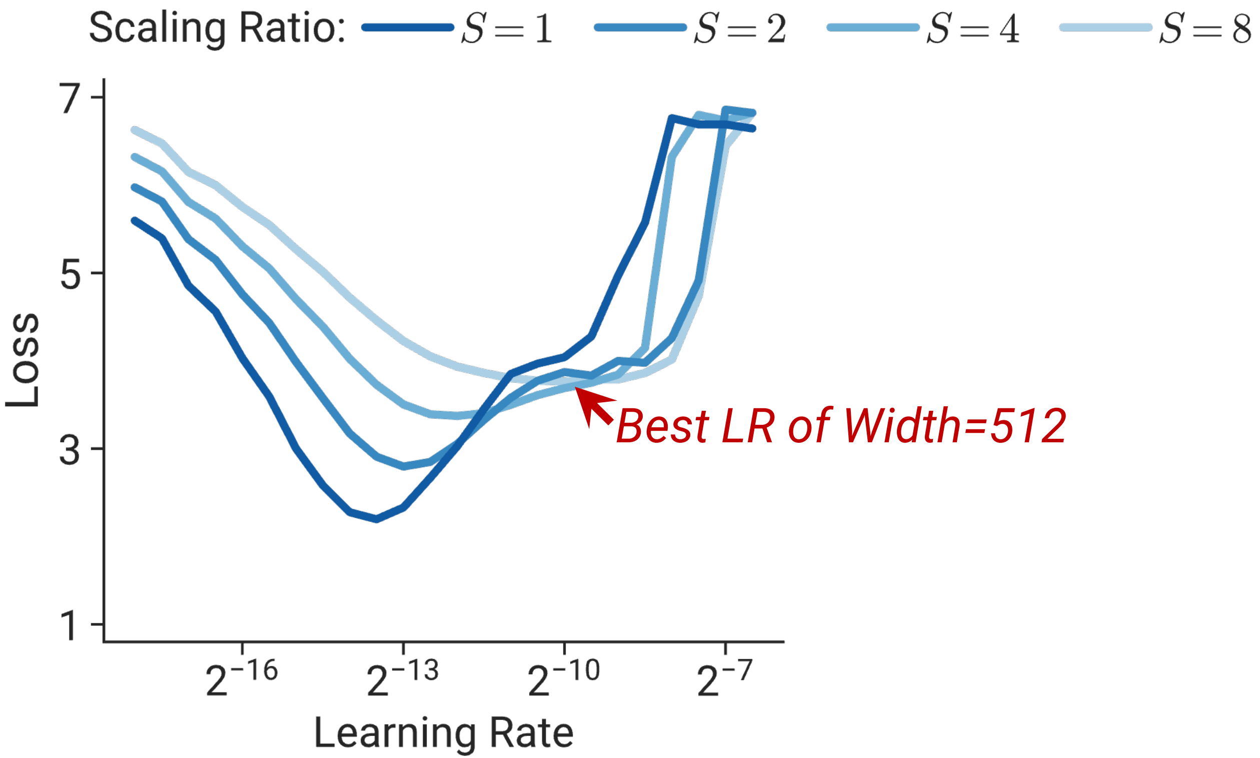

Can Hyperparameters be Transferred?

Toy Example: 2-layer Transformer model over WikiText-2 dataset using Adam

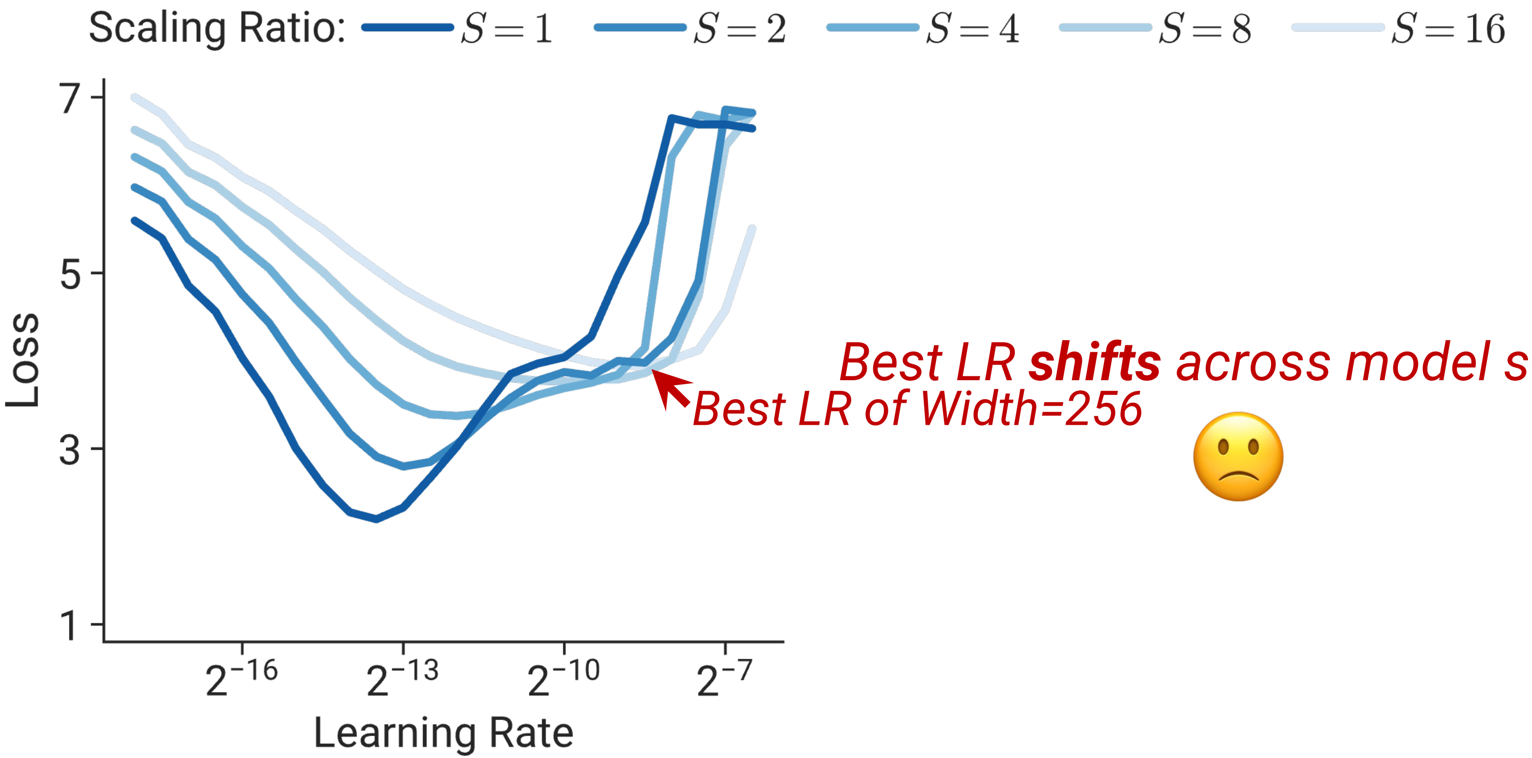

*Best LR **shifts** across model scales* 🙁

Hydro Makes Hyperparameters Transferable

Applying Hydro on the same Transformer model

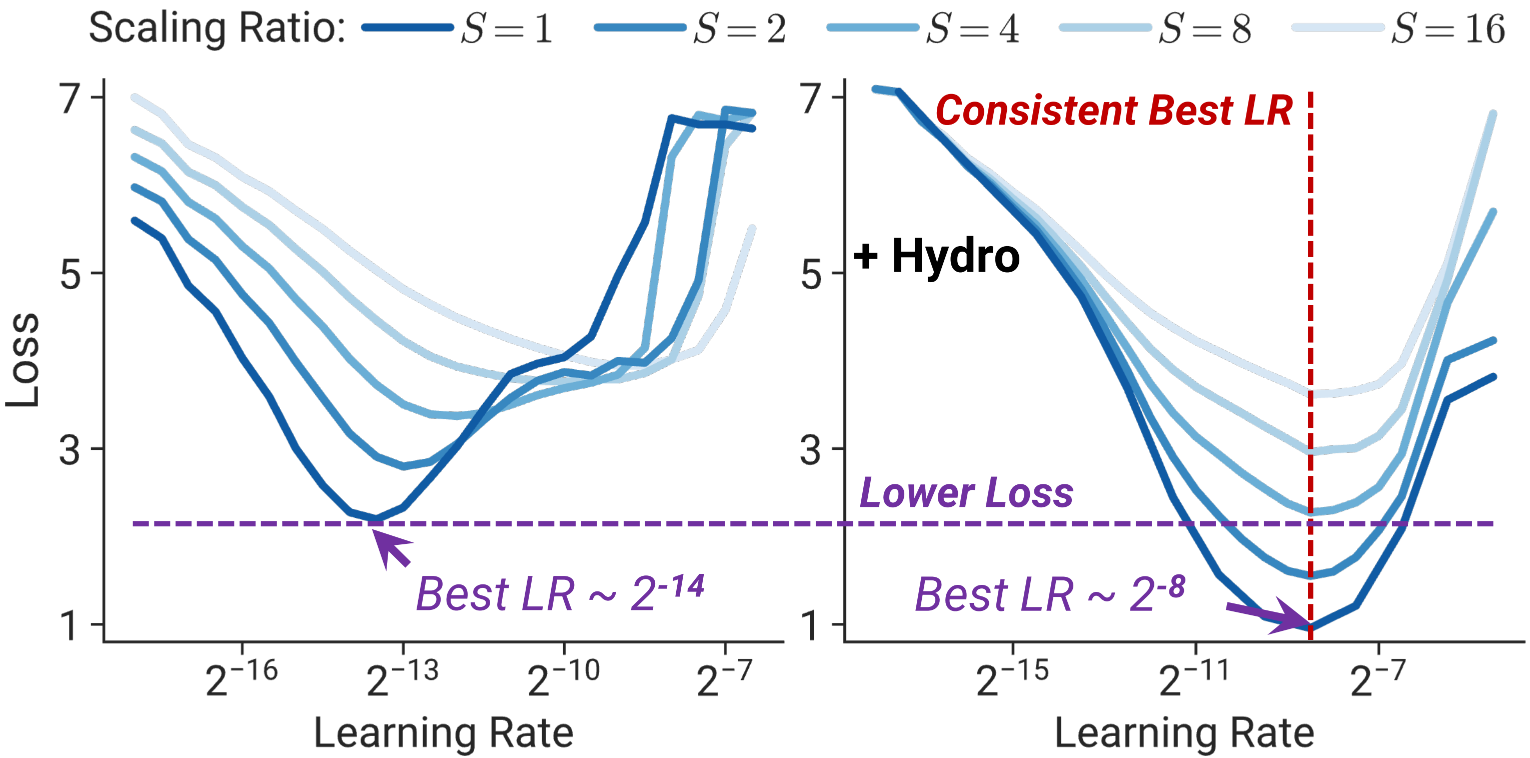

Underlying Theory: Maximum Update (MU) Parametrization[1]

Theoretically enabling maximal feature learning for **infinite-width** neural networks

1-hidden-layer MLP:

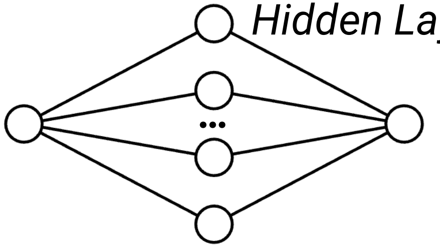

Optimizer: *SGD with* $lr = 1$

Common Practice: Initialization: $U \sim \mathcal{N}(0,1), \quad V \sim \mathcal{N}(0, 1/w)$

Learning rates: $\eta_U = 1, \quad \eta_V = 1$

MU Parametrization: Initialization: $U \sim \mathcal{N}(0,1), \quad V \sim \mathcal{N}(0, 1/w^2)$

Learning rates: $\eta_U = w, \quad \eta_V = 1/w$

Avoid Output Layer Blow-up

MU Parametrization: Intuitive Insights[1]

Theoretical: Maximal feature learning for infinite-width neural networks

Impact in Practice

Empirical: Hyperparameter transfer across model scales (in terms of width)

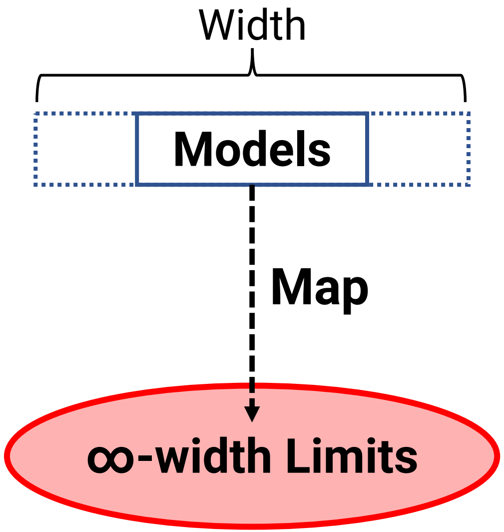

Correspond models with different scales to their ∞ limits

Benefits:

- Solve the unbalanced training issue (e.g., output layer update much faster) via layer-wise *lr* adjustment
- Ensure consistent magnitude updates for each layer during training regardless of its width

Problem: Manually implementing MU parametrization is burdensome and error-prone

Source: [1] Tuning Large Neural Networks via Zero-Shot Hyperparameter Transfer (NeurIPS '21)

Hydro Tuner

MU parametrization theory + system support to jointly accelerate tuning

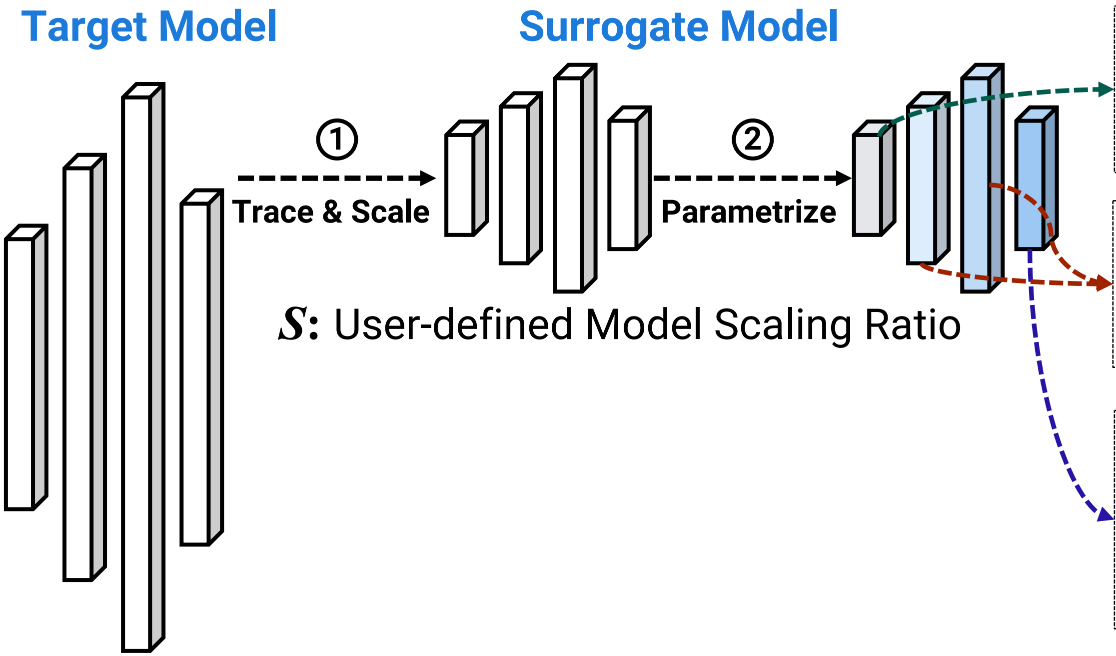

Input Layer:
1. Init Variance Multiply S
2. If SGD Optimizer, Layer LR Divide S

Hidden Layer:
1. Init Variance Multiply S
2. SGD & Adam Optimizer LR Multiply S

Output Layer:
1. Zero-Variance Initialization
2. Layer Input Multiply S
3. If SGD Optimizer, Layer LR Divide S

Hydro Tuner: Scaling Effect

Example: WideResNet-50

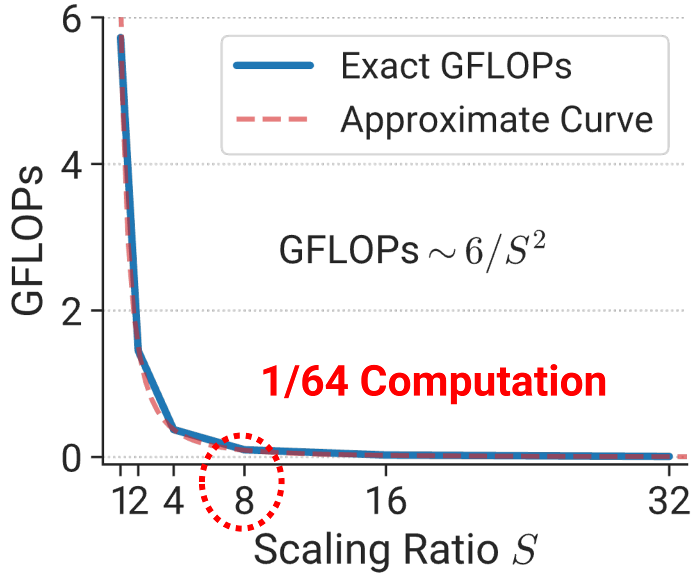

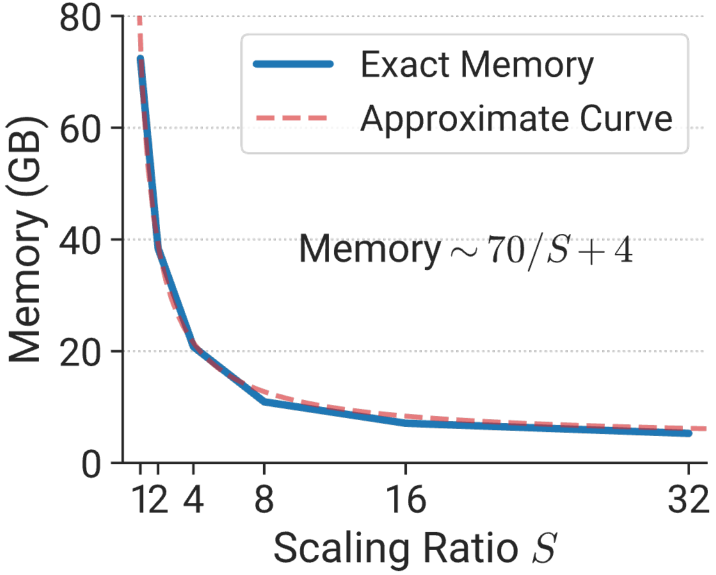

Hydro Tuner: Scaling Effect

However... For Small Model: ResNet-18

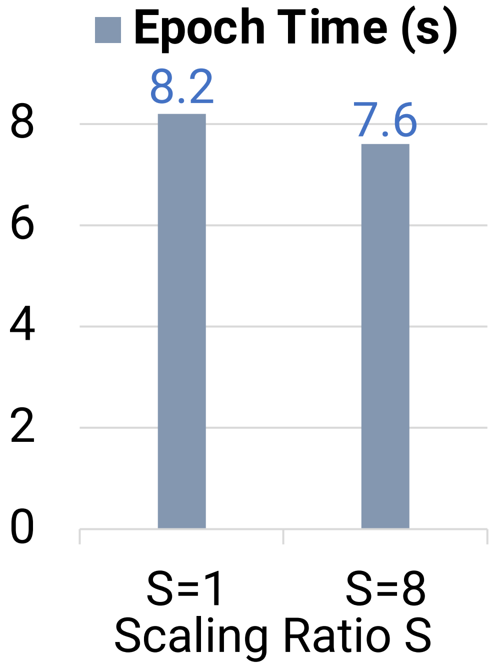

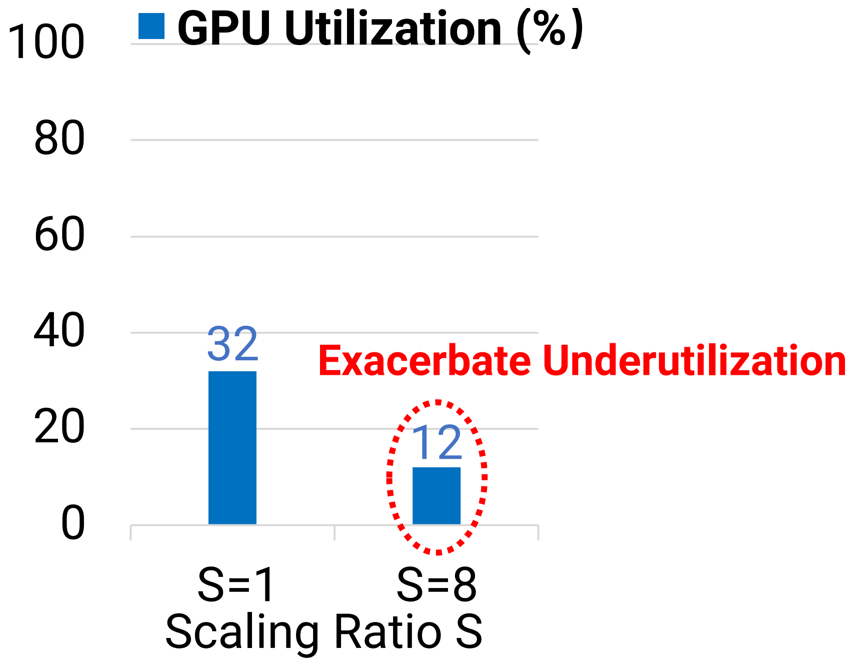

Hydro Tuner: Model Fusion

Hydro further enables inter- and intra-trial fusion to improve hardware efficiency

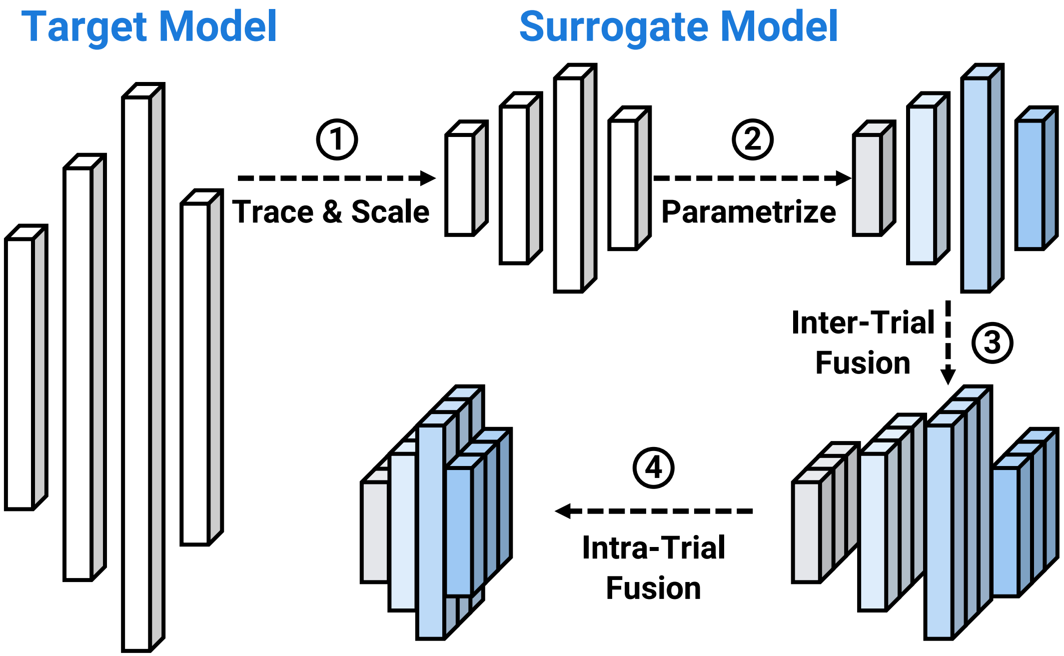

Hydro Tuner: Inter-trial Fusion

Hydro extends the application scope of HFTA[1] & automizes the fusion process

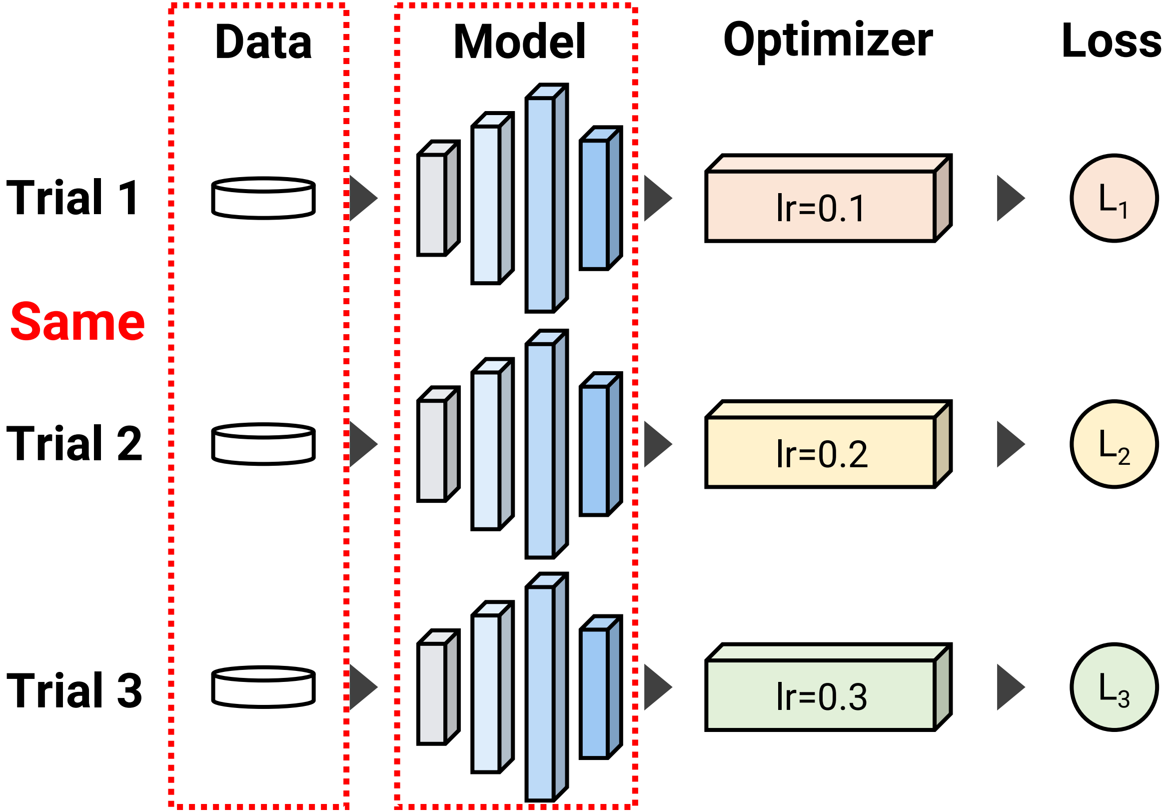

Source: [1] Horizontally Fused Training Array: An Effective Hardware Utilization Squeezer for Training Novel Deep Learning Models (MLSys '21)

Hydro Tuner: Inter-trial Fusion

Hydro extends the application scope of HFTA[1] & automizes the fusion process

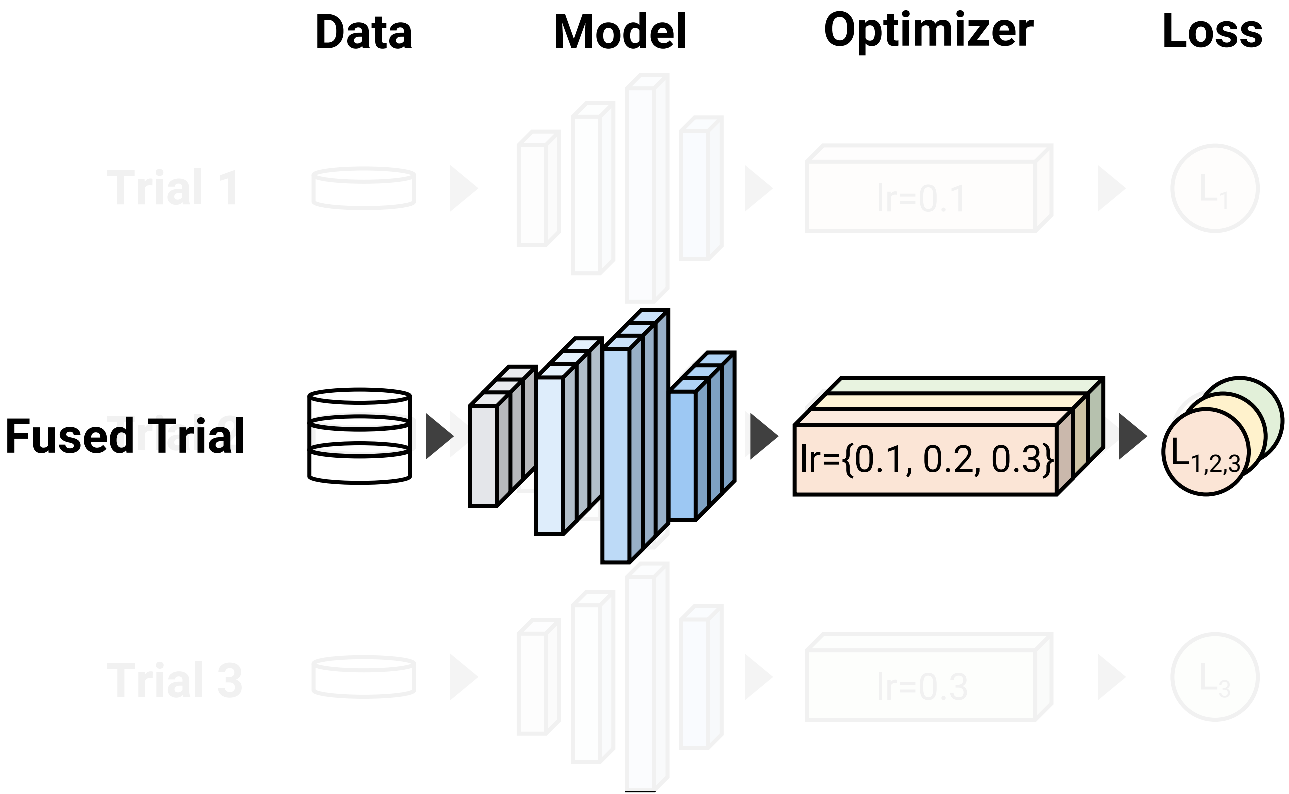

Source: [1] Horizontally Fused Training Array: An Effective Hardware Utilization Squeezer for Training Novel Deep Learning Models (MLSys '21)

Effect of Scaling + Inter-trial Fusion

Example: ResNet-18 (Scaling=8, CIFAR-10 Batch_Size=256) on **A100 80GB**

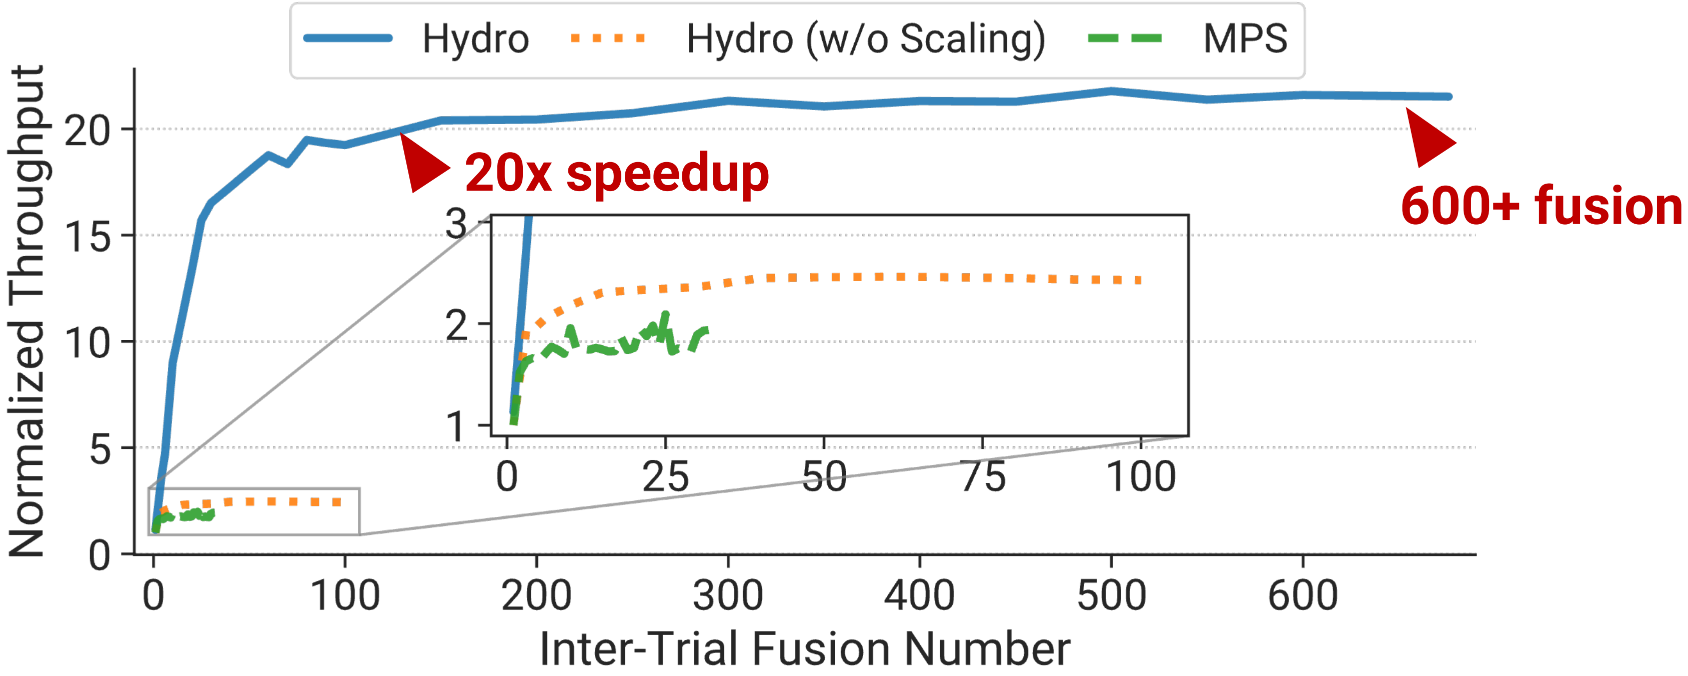

Hydro

Job-level ***Hydro Tuner***

Automatically generate surrogate models for tuning by applying transfer theory and model fusion

Datacenter-level ***Hydro Coordinator***

Leverage idle bubble resources of pretraining jobs via interleaving training

Resource Contention between LLM Pretraining and Tuning Jobs

Large Language Model (LLM) pretraining jobs occupy massive resources

→ **Long queuing delay of tuning jobs**

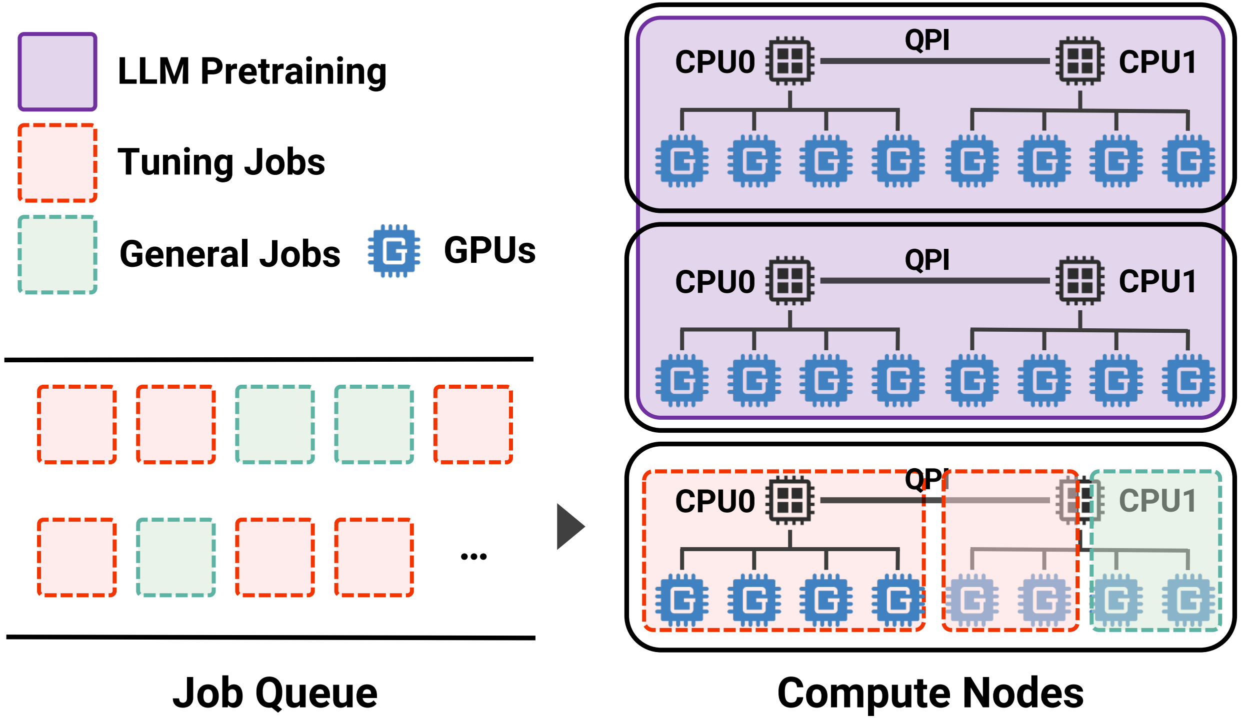

Opportunity: Co-exist LLM Pretraining Jobs

Massive Resources: Long-term occupy hundreds ~ thousands of GPUs

Pipeline Parallelism is commonly applied but introduces **bubbles**

Example: 1F1B[1] Pipeline Schedule

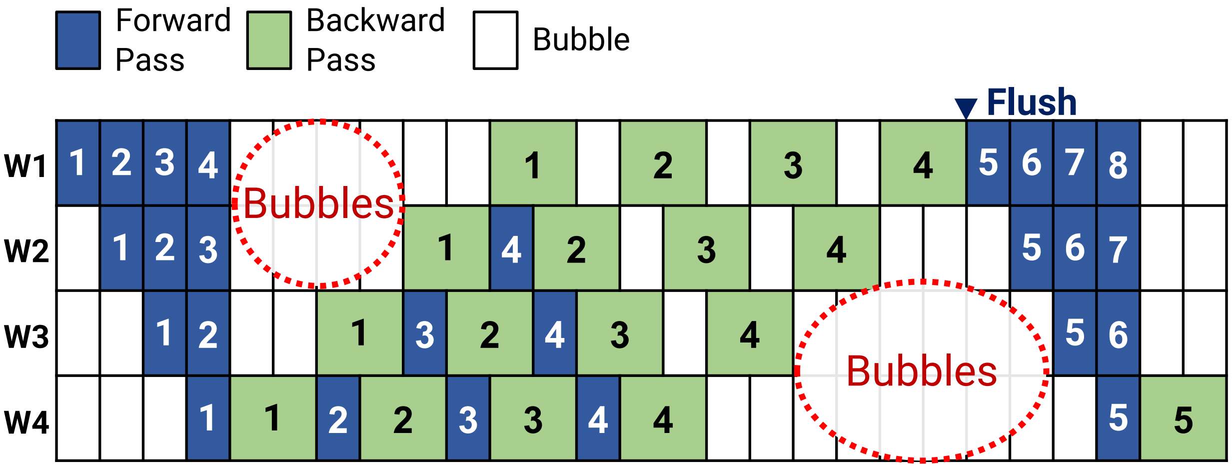

GPU Memory:

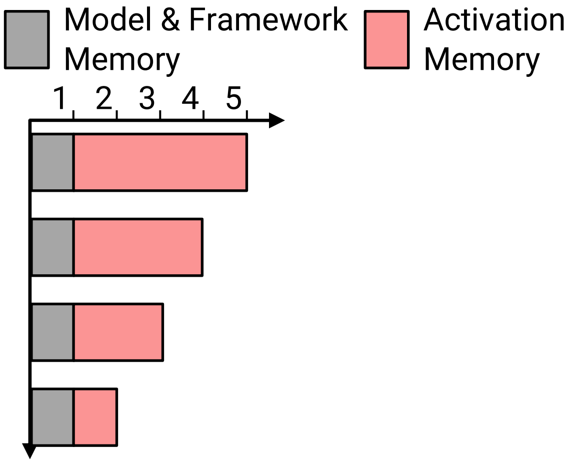

✗ **Wasted spare resources**

✗ **Unbalanced memory footprint**

Source: [1] PipeDream (SOSP '19)

Hydro Coordinator: Leverage Bubble Resources

Solution: Interleaving Hydro trials with a LLM pretraining job

Hydro Trials +1F1B Workload

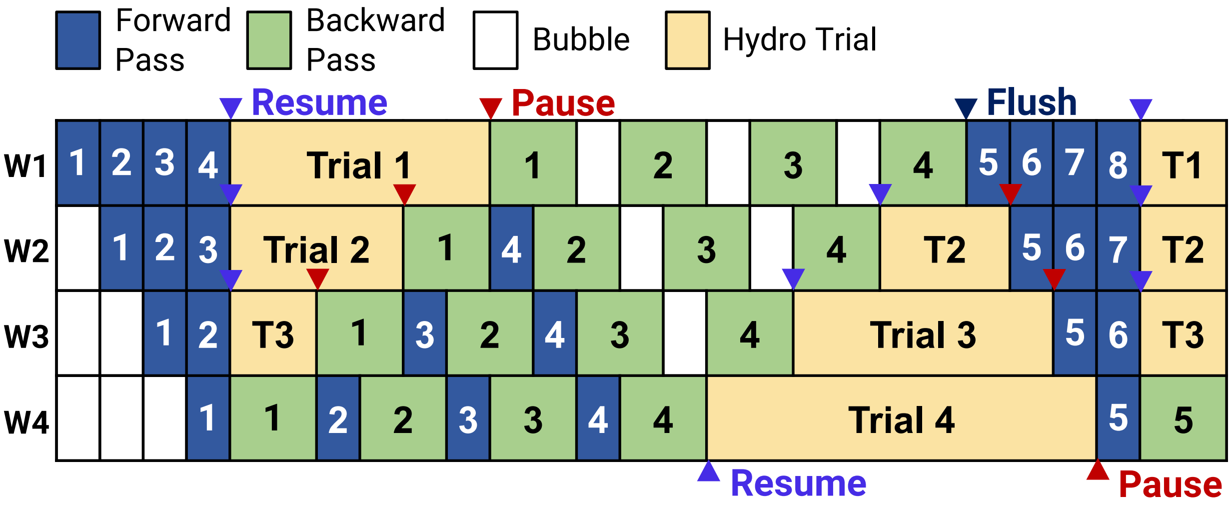

GPU Memory:

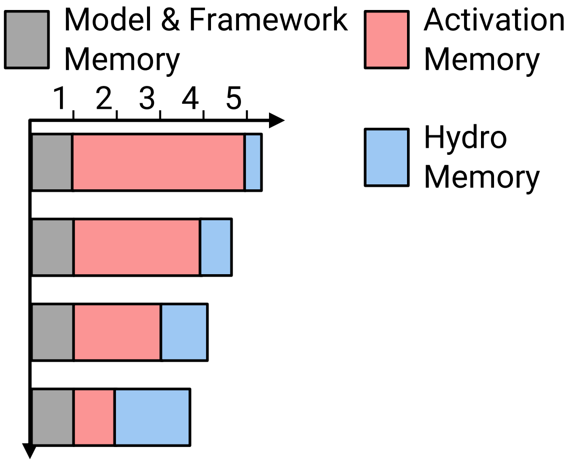

✓ No interference

✓ Resilient trial sizes

Hydro Coordinator: Leverage Bubble Resources

Solution: **Interleaving** Hydro trials with a LLM pretraining job

Why Hydro Trials are suitable for interleaving?

1. Throughput Insensitive

Tuning jobs are more tolerant of partial trials slowdown

2. Deterministic and Scaled Memory Footprint

Memory is profiled and greatly reduced via model scaling

3. Elastic and Opportunistic Trial Placement

Trials can adjust the fusion number to fit the remaining memory

GPU Memory:

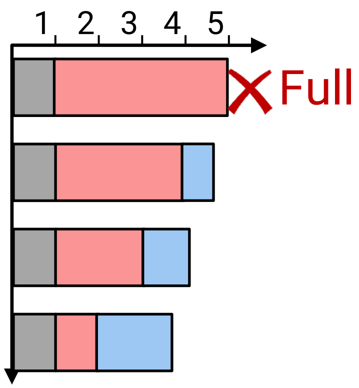

Effect of Hydro Coordinator

SM Activity of a GPT model with 4 pipeline stages (over 4x8 A100 GPUs)

+ Hydro Trials (fuse 16x ResNet-18 models) with interleaving training

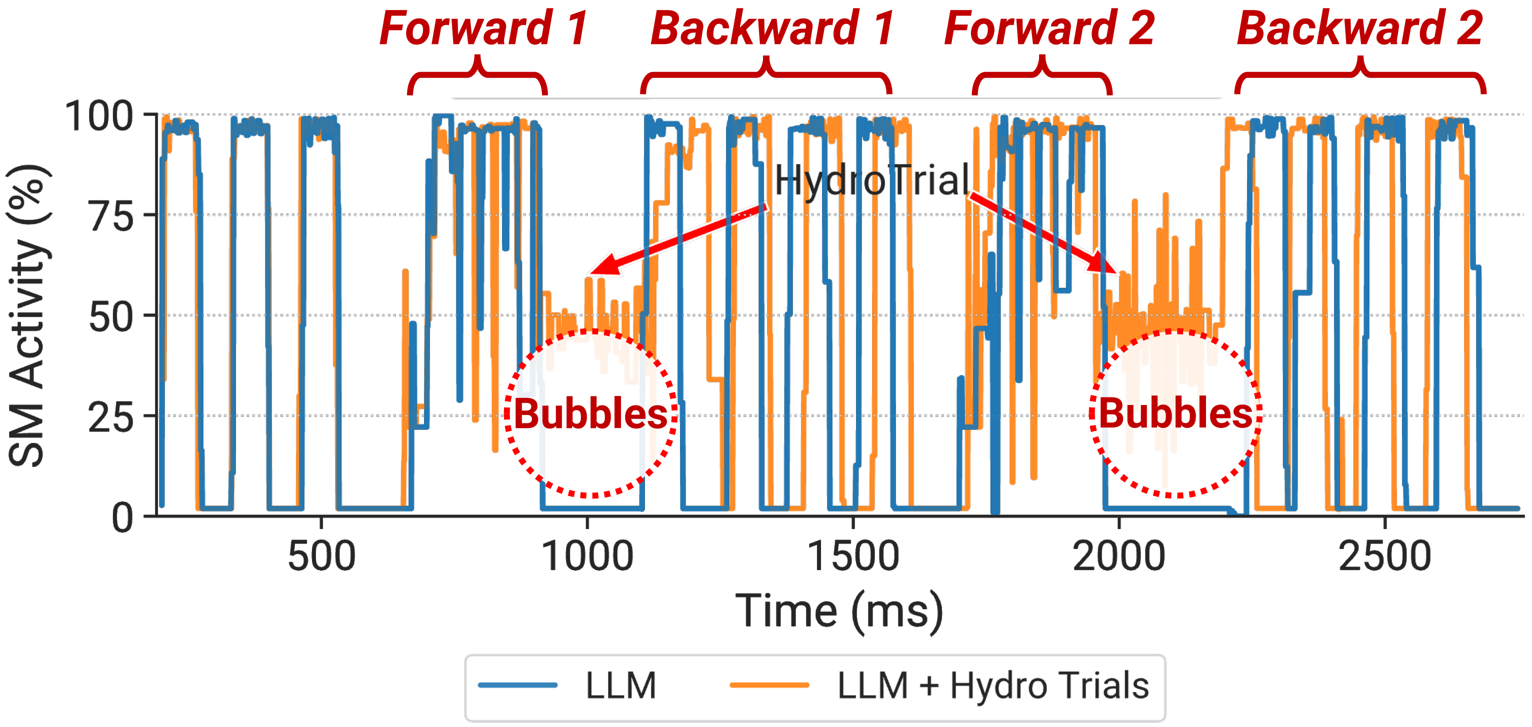

Evaluation

Evaluation: Intuitive Study of Surrogate-Based Tuning

10 trials of ResNet-18 on CIFAR-10: **A-J:** [batch_size, learning_rate, momentum]

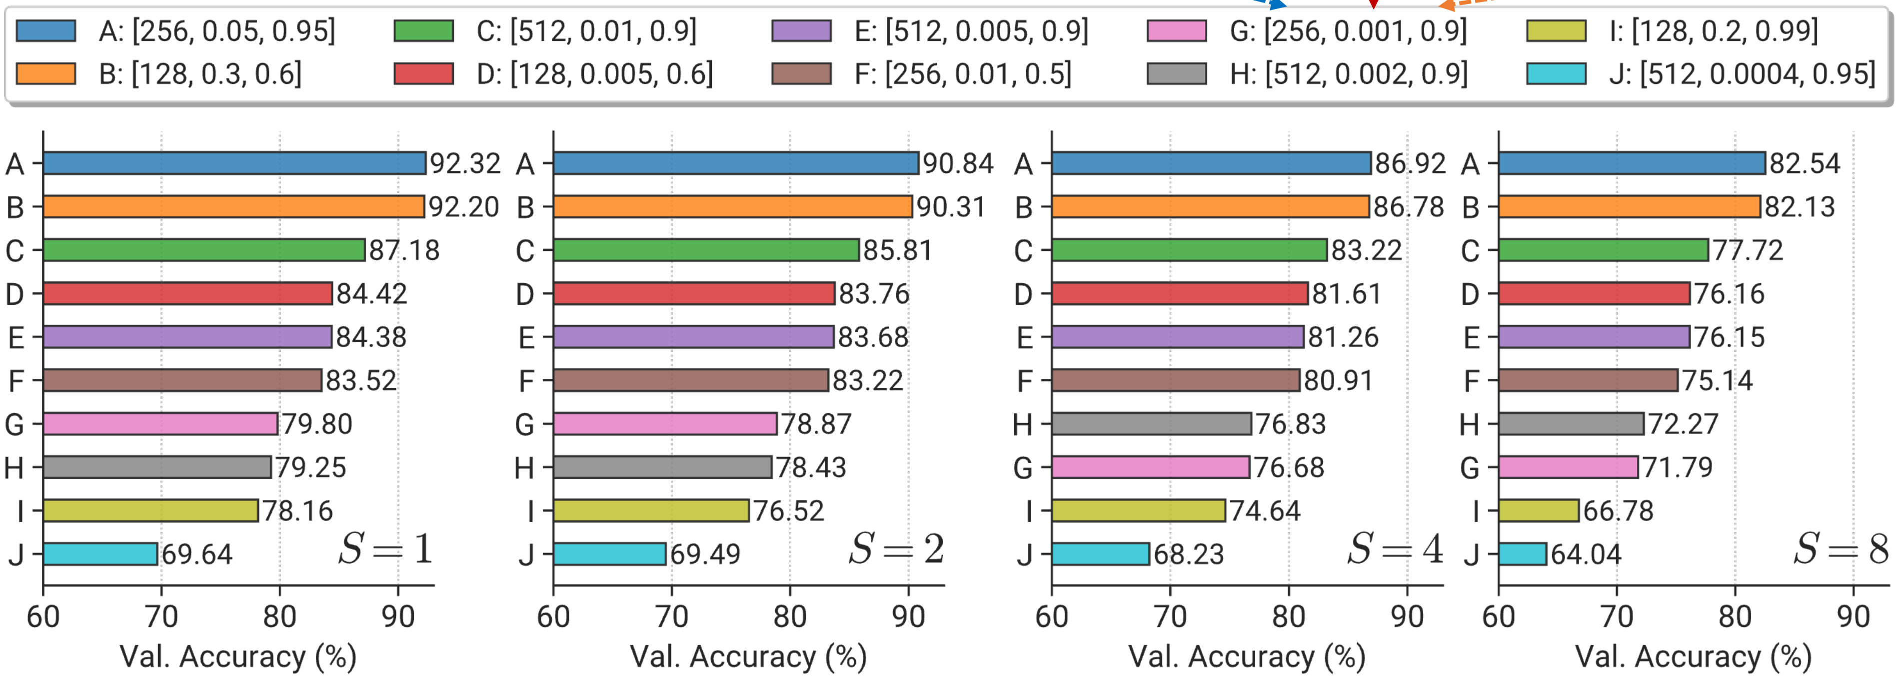

Hyperparameter ranking transfers well across different scaling ratios

Evaluation: End-to-End Experiments

Testbed: A100 GPU cluster of Shanghai AI Laboratory

Baseline: Ray Tune[1] system + FIFO algorithm

Task	Search Space	Model	#GPUs	#Trial	Acceleration	Quality
Language Modeling	**lr**: $(10^{-5}, 10^{-1})$ **gamma**: (0.01, 0.9)	GPT-3 XL	128	100	**78.5 ×**	−0.48 ppl*
		Transformer	8	200	**8.7 ×**	−0.15 ppl
Image Classification	**lr**: $(10^{-4}, 10^{0})$ **momentum**: (0.5, 0.999) **batchsize**: [128, 256,512] **gamma**: (0.01, 0.9)	WideResNet-50	32	200	**20.3 ×**	+1.18% acc*
		MobileNetV3-L	16	500	**12.3 ×**	+0.05% acc
		VGG-11	8	500	**10.8 ×**	+0.09% acc
		ResNet-18	8	1000	**16.2 ×**	+0.02% acc

* Compared with the official hyperparameter setting as the model quality baseline

Source: [1] Tune: A Research Platform for Distributed Model Selection and Training (ICML AutoML '18)

Hydro

Job-level

Joint optimization of theory and system techniques

Datacenter-level

Leverage idle bubble resources of pretraining jobs

https://github.com/S-Lab-System-Group/Hydro

qinghao.hu@ntu.edu.sg