Train your ML models at scale

Gal Oshri

Product Manager
Amazon Web Services

Rise of large-scale models

"a picture of a very clean living room"

2017

StackGAN,
Zhang et al.

2022

Stable Diffusion,
Rombach et al.

Large-scale models lead to better results

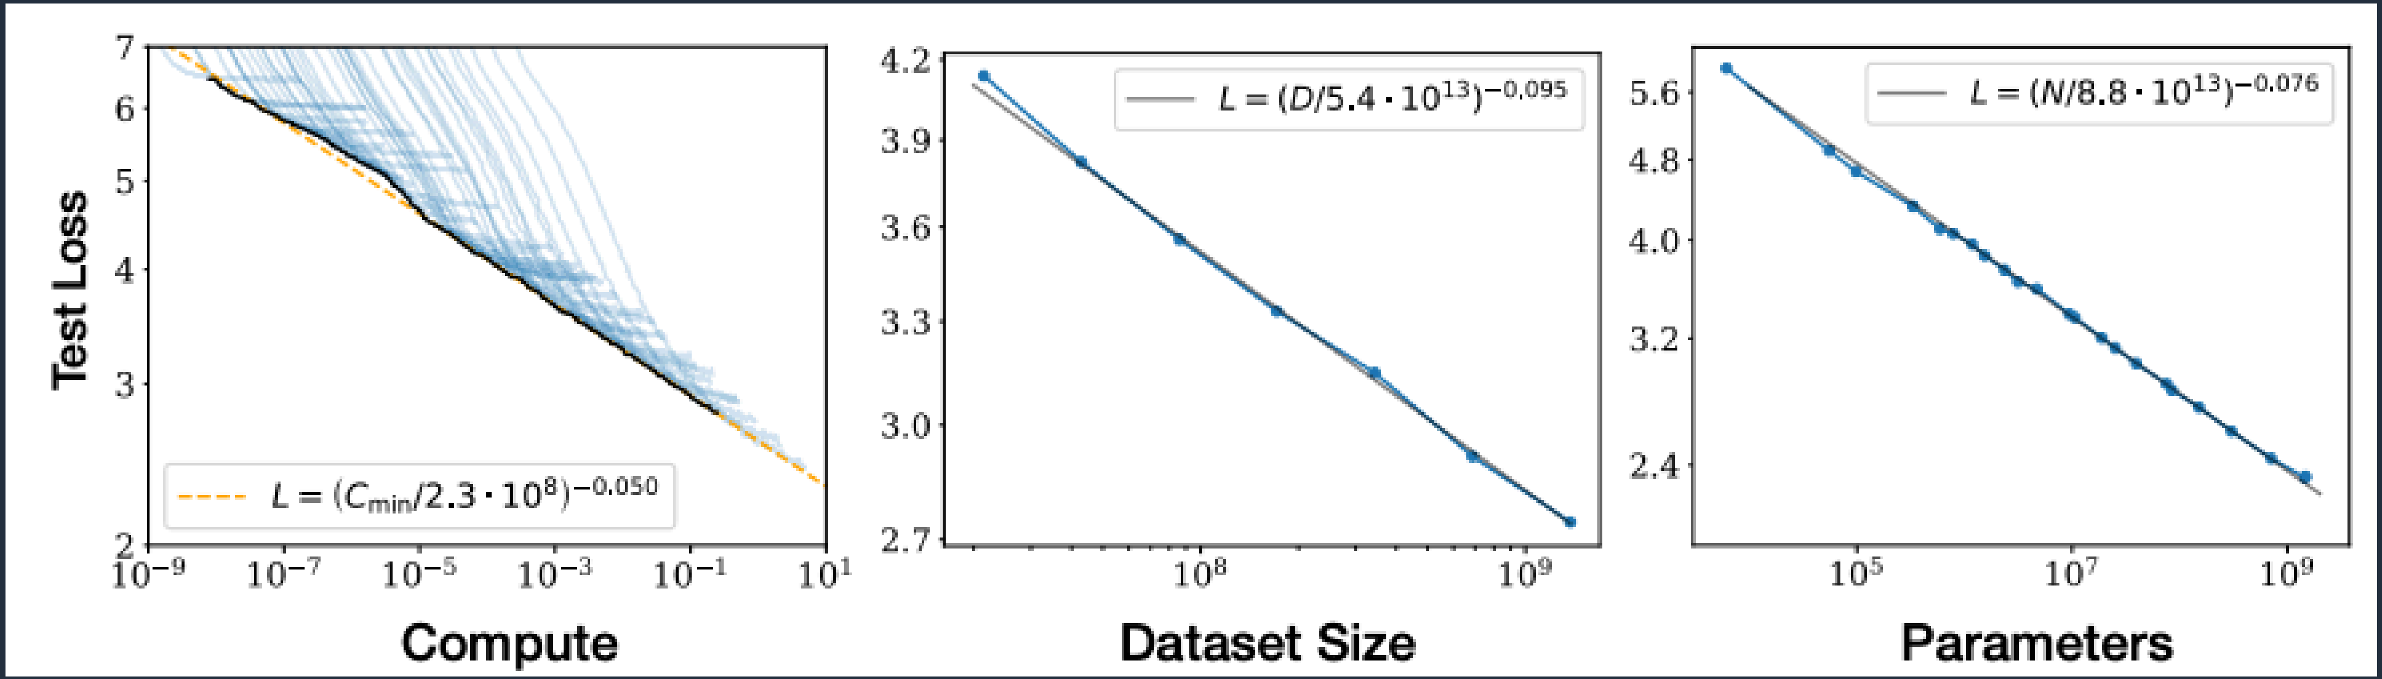

Scaling Laws for Neural Language Models
Kaplan et al., 2020

Challenges with training large-scale models

Hardware

Health checks

Orchestration

Data

Scaling up

Cost

SageMaker accelerates large-scale model training

Large-scale training on SageMaker

OPTIMIZED DISTRIBUTED TRAINING LIBRARIES & FRAMEWORKS

TensorFlow	PyTorch	Hugging Face	SageMaker Distributed Training Libraries	Bring your own library (e.g. DeepSpeed, Megatron)

AMAZON SAGEMAKER TRAINING

Large Scale Cluster Orchestration	NCCL Health Checks	Resilient training	SageMaker Compiler	Warm pools	SSH to container
Data loading	Debugger	Profiling	Experiment tracking	Hyperparameter optimization	Pay for what you use

ML COMPUTE INSTANCES & ACCELERATORS

NVIDIA GPUS A100, V100, K80, T4, A10	AWS Nitro	400/800 Gbps EFA Networking	CPU instances	AWS Trainium

Amazon SageMaker ephemeral training clusters

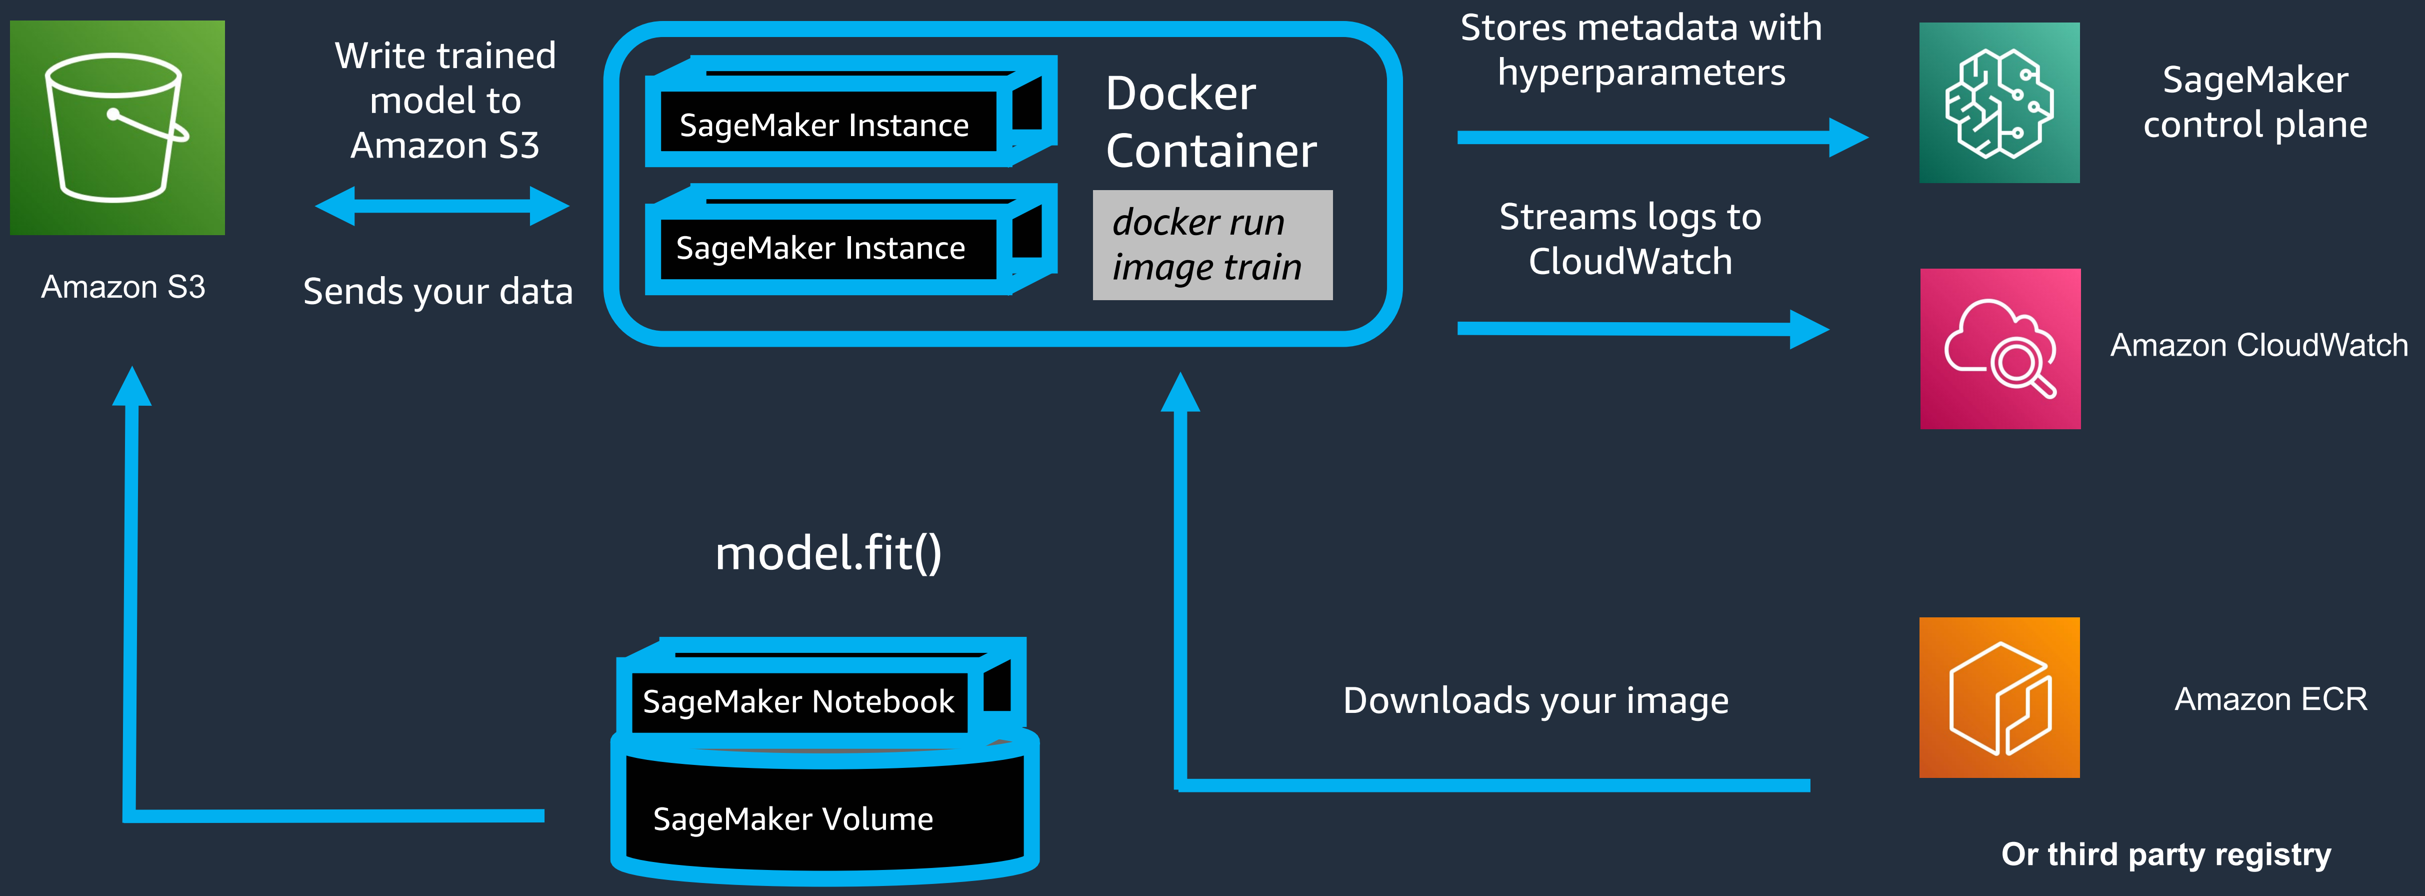

Fully managed by SageMaker

Train with your own deep learning model

```
from sagemaker.pytorch import PyTorch

estimator = PyTorch(entry_point = './cifar10.py',
                    role = role,
                    framework_version = '1.13',
                    py_version = 'py38',
                    instance_count = 1,
                    instance_type = 'ml.g5.xlarge',
                    hyperparameters = {'epochs': 50, 'batch_size': 32},
                    metric_definitions = [{'Name': 'train:loss', 'Regex': 'loss: (.*)'}])

estimator.fit("s3://bucket/path/to/training/data")
```

Accelerate local ML code conversion to training jobs

```
from sagemaker.remote_function import remote

@remote(instance_type="ml.g4dn.xlarge",dependencies = "./environment.yml")
def train_hf_model(
    train_input_path,test_input_path,s3_output_path = None,
    *,epochs = 1, train_batch_size = 32, eval_batch_size = 64,
    warmup_steps = 500,learning_rate = 5e-5
    ):
    model_name = "distilbert-base-uncased"
    model = AutoModelForSequenceClassification.from_pretrained(model_name)
    ...
    return os.path.join(s3_output_path, model_dir), eval_result
```

Replicate experimental results by default

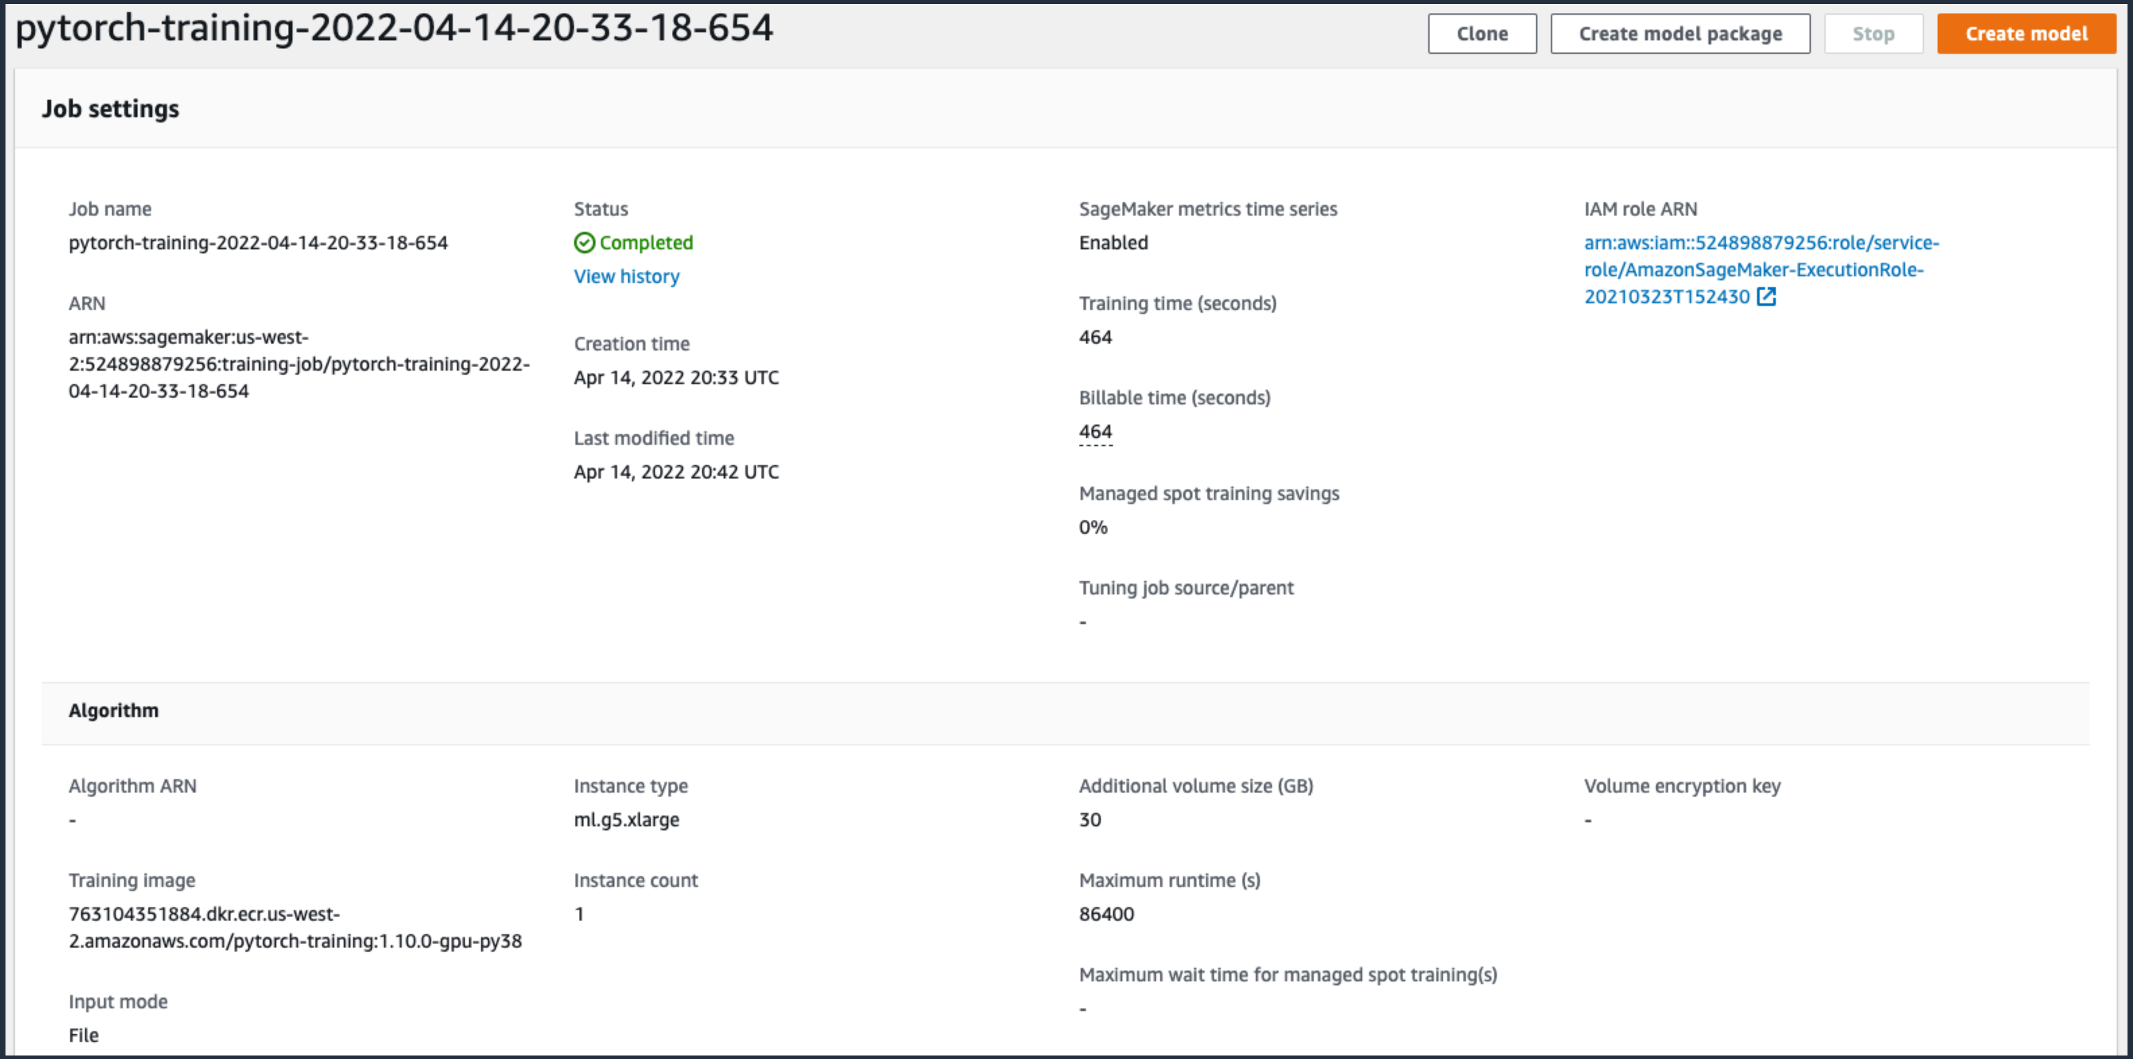

Hosted TensorBoard

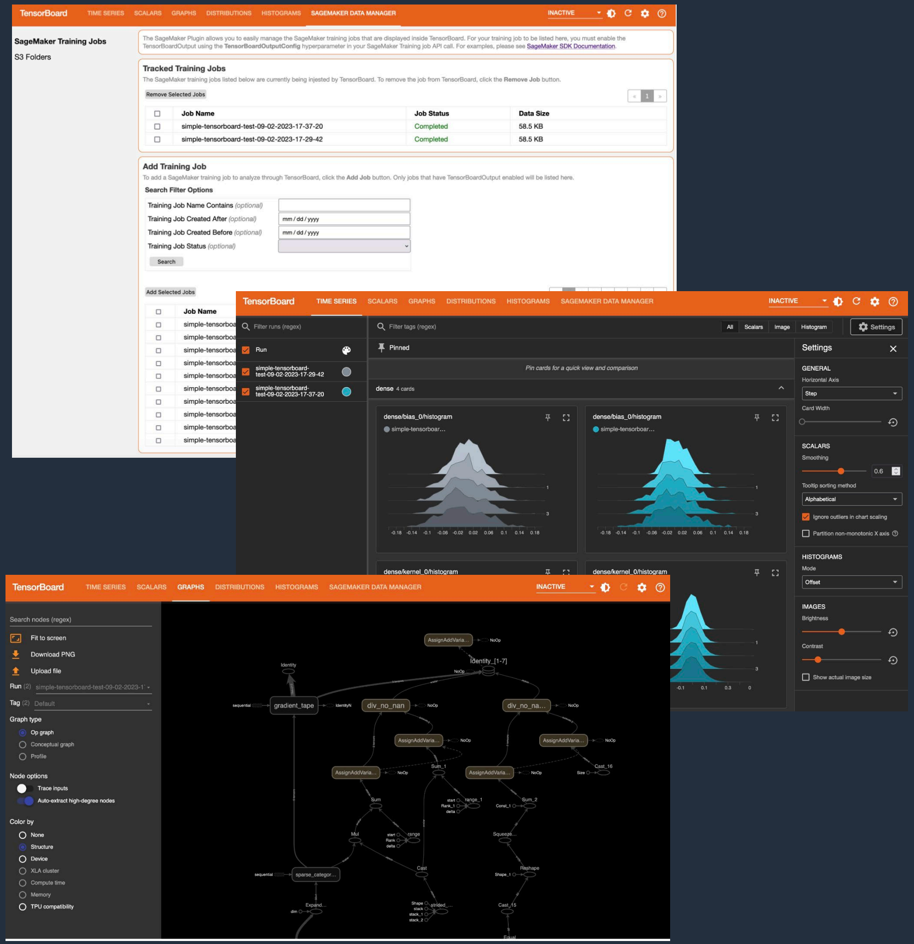

Automated TensorBoard app management
provides a pre-configured and managed TensorBoard interface, which can save data scientists time by eliminating the need to install and configure TensorBoard on their own.

Scalable
designed to meet the needs of large-scale distributed machine learning workloads by having the UI automatically hosted on memory optimized r5 instances

Automated data management
easily configure TB data uploaded to S3 with the training job by passing an API argument (TensorBoardOutputConfig)

Security
provides a secure and reliable environment for storing and processing machine learning data

LG's Tilda, the AI artist powered by EXAONE

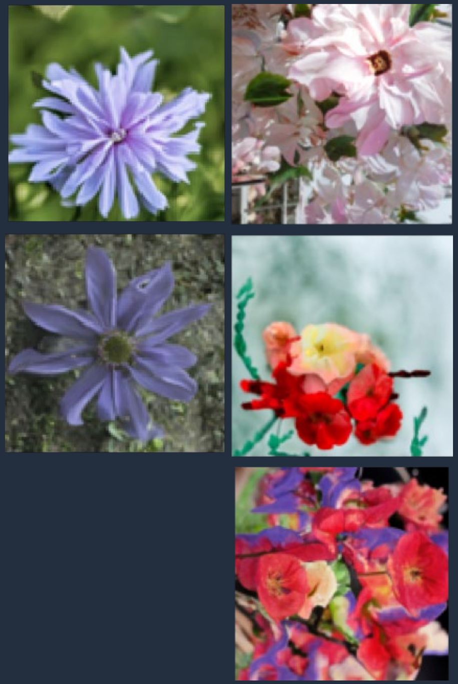

"... we could optimize distributed training and were able to train the model faster by 59% ..."

Seung Hwan Kim , Vice president and vision lab leader at LG AI Research

Sharded data parallelism with tensor parallelism

```
smp_parameters = {
    "ddp": True,
    "fp16": True,
    "prescaled_batch": True,
    "sharded_data_parallel_degree": 4,
    "tensor_parallel_degree": 4
}

pytorch_estimator = PyTorch(
    entry_point="your_training_script.py",
    role=role,
    instance_type="ml.p4d.24xlarge",
    distribution={
        "smdistributed": {
            "modelparallel": {
                "enabled": True,
                "parameters": smp_parameters,
            }
        },
        ...
    },
    ...
)
```

Falcon

Technology Innovation Institute trains the state-of-the-art Falcon LLM 40B foundation model on Amazon SageMaker

by Dr. Ebtesam Almazrouei, Olivier Cruchant, and Will Badr | on 07 JUN 2023 | in Amazon SageMaker, Artificial Intelligence, Intermediate (200) | Permalink | Comments | Share

This blog post is co-written with Dr. Ebtesam Almazrouei, Executive Director–Acting Chief AI Researcher of the AI-Cross Center Unit and Project Lead for LLM Projects at TII.

Getting Started

Product page

aws.amazon.com/sagemaker/train/

Technical Documentation

docs.aws.amazon.com/sagemaker/latest/dg/how-it-works-training.html

SageMaker examples on GitHub

github.com/aws/amazon-sagemaker-examples

Training LLMs on Amazon SageMaker: Best Practices

https://aws.amazon.com/blogs/machine-learning/training-large-language-models-on-amazon-sagemaker-best-practices/

Thank you!

Gal Oshri
@galoshri