AWS re:Invent

NOV. 29 – DEC. 3, 2021 | LAS VEGAS, NV

AIM418

Distributed training on AWS

Aditya Bindal
Sr. Product Manager Lead, AWS Deep Learning
AWS

Can Karakus
Sr. Applied Scientist, AWS Deep Learning
AWS

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

Distributed training libraries on Amazon SageMaker

Benchmarks

Demo

Q&A

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Model parallelism

Automated and efficient model partitioning

Minimal code change

Optimized for training large models

Distributed training on Amazon SageMaker

https://aws.amazon.com/sagemaker/distributed-training/

Data parallelism

Reduced training time

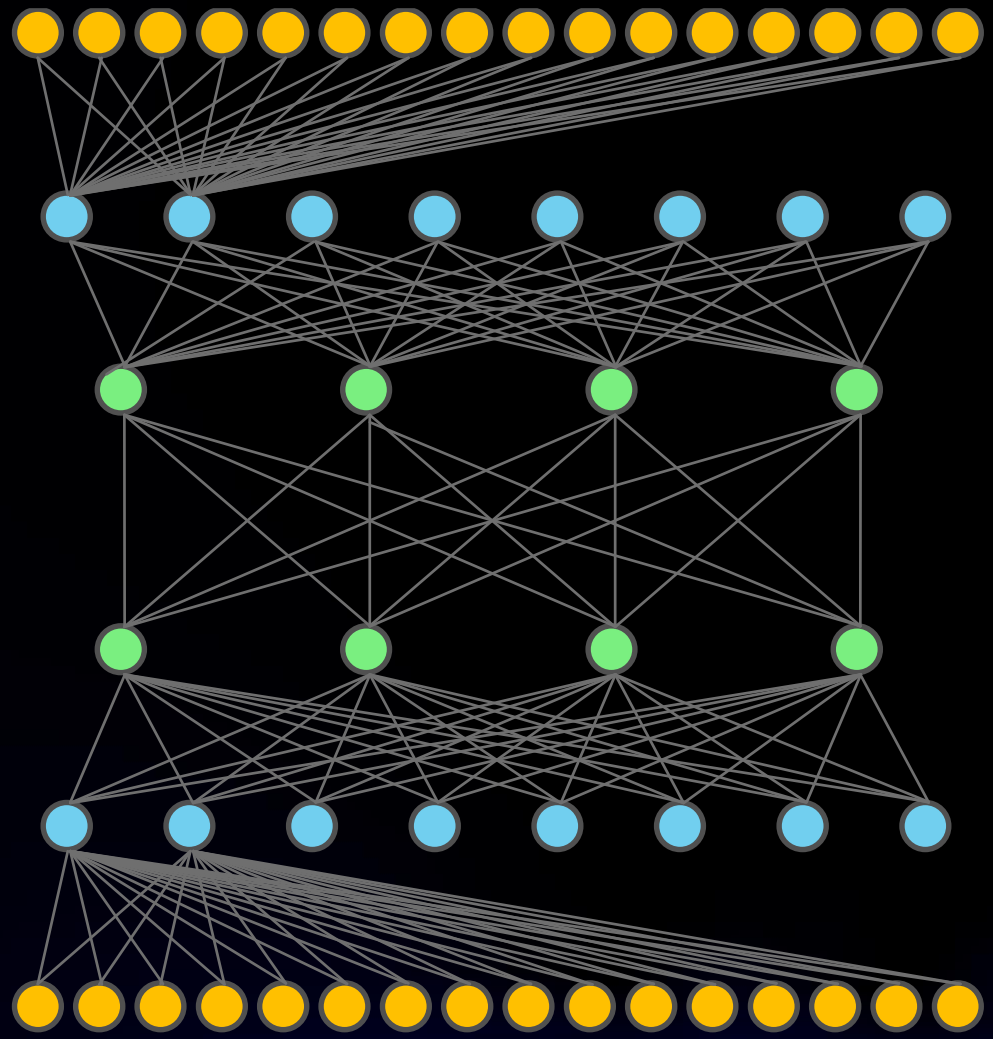
Optimized for AWS network

Support for popular ML framework APIs

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

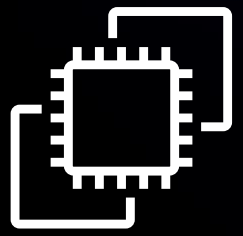
DataParallel in SageMaker

- Library for distributed training of deep learning models in TensorFlow and PyTorch
- Accelerates training for network-bound workloads
- Built and optimized for AWS network topology and hardware
- 20%–40% faster and cheaper; best performance on AWS

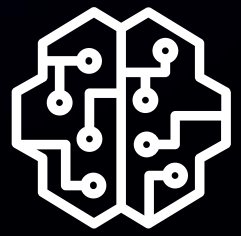


© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

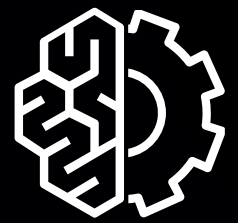
Model parallelism on SageMaker



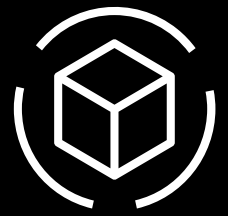
Efficient pipelined training



Automated model partitioning



Managed SageMaker training



Tight framework integration

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Benchmarks

ModelParallel training

Model	Instances	Performance with `modelparallel`	Performance without `modelparallel`
GPT-2 (13B)	16 P4d.24xlarge	53.5 samples/sec	OOM
GPT-2 (175B)	256 P4d.24xlarge	35 days	OOM

DataParallel training

Model	Instances	Performance with `dataparallel`	Performance without `dataparallel`
RoBERTa (1.3B)	16 P4d.24xlarge	894,744 words/sec	784,495 words/sec

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

LATENT SPACE

Latent Space's mission is to democratize creativity by making the world's first fully AI-rendered 3D game engine.

> Our current focus is on utilizing information retrieval to augment large-scale model training, for which we have sophisticated ML pipelines. This setup presents a challenge on top of distributed training since there are multiple data sources and models being trained at the same time. As such, we're leveraging the new distributed training capabilities in Amazon SageMaker to efficiently scale training for large generative models."
>
> **Sarah Jane Hong, Co-founder and Chief Science Officer, Latent Space**

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

turbine

Turbine is a simulation-driven drug discovery company delivering targeted cancer therapies to patients. Turbine uses machine learning to train an in silico human cell model, called Simulated Cell, based on a proprietary network architecture. By accurately predicting various interventions on the molecular level, Simulated Cell helps Turbine discover new cancer drugs and find combination partners for existing therapies.

> “Training of our simulation is something we continuously iterate on, but on a single machine each training takes days, hindering our ability to iterate on new ideas quickly. We are very excited about distributed training on Amazon SageMaker, which we are expecting to decrease our training times by 90% and to help us focus on our main task: to write a best-of-the-breed codebase for the cell model training. SageMaker ultimately allows us to become more effective in our primary mission: to identify and develop novel cancer drugs for patients.”
>
> **Kristóf Szalay, Chief Technical Officer, Turbine AI**

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Lyft, one of the largest transportation networks in the United States and Canada, launched its Level 5 autonomous vehicle division in 2017 to develop a self-driving system to help millions of riders. Lyft Level 5 aggregates over 10 terabytes of data each day to train ML models for its fleet of autonomous vehicles. Managing ML workloads on its own was becoming time-consuming and expensive.

> Using Amazon SageMaker distributed training, we reduced our model training time from days to a couple of hours. By running our ML workloads on AWS, we streamlined our development cycles and reduced costs, ultimately accelerating our mission to deliver self-driving capabilities to our customers."
>
> **Alex Bain, Lead for ML Systems, Lyft Level 5**

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Resources

Amazon SageMaker distributed training webpage

https://aws.amazon.com/sagemaker/distributed-training/

Technical documentation

https://docs.aws.amazon.com/sagemaker/latest/dg/distributed-training.html

Example notebooks

https://docs.aws.amazon.com/sagemaker/latest/dg/distributed-training-notebook-examples.html

Demo and deep dive videos

- Fast distributed training and near linear scaling with distributed DataParallel in Amazon SageMaker
  https://youtu.be/EXmz5g8F2zU
- Train billion-parameter models with model parallelism on Amazon SageMaker
  https://youtu.be/vv52RsBM8o4

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!

Aditya Bindal
bindala@amazon.com

Can Karakus
cakarak@amazon.com

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.