




Data Parallelism: How to Train Deep Learning Models on Multiple GPUs

NVIDIA Deep Learning Institute



Workshop Overview

Modern deep learning challenges leverage increasingly larger datasets and more complex models. As a result, significant computational power is required to train models effectively and efficiently. Learning to distribute data across multiple GPUs during training makes possible an incredible wealth of new applications that utilize deep learning.

Effectively using systems with multiple GPUs also reduces training time, allowing for faster application development and much faster iteration cycles. Teams who can train with multiple GPUs have an edge, building models trained on more data in shorter periods of time and with greater engineer productivity.

This workshop teaches you techniques for data-parallel deep learning training on multiple GPUs to shorten the training time required for data-intensive applications. Working with deep learning tools, frameworks, and workflows to perform neural network training, you'll learn how to decrease model training time by distributing data to multiple GPUs, while retaining the accuracy of training on a single GPU.

Learning Objectives

By participating in this workshop, you'll learn how to:

- > Perform data-parallel deep learning training with multiple GPUs
- > Achieve maximum throughput when training for the best use of multiple GPUs
- > Distribute training to multiple GPUs using Pytorch Distributed Data Parallel (DDP)
- > Utilize algorithmic considerations specific to multi-GPU training performance and accuracy

Overview

Duration	8 hours
Price	Contact us for pricing.
Prerequisites	Experience with deep learning training using Python
Tools, libraries, and frameworks	PyTorch, PyTorch Distributed Data Parallel, NVIDIA Collective Communications Library (NCCL)
Assessment type	Skills-based coding assessments evaluate learners' ability to train deep learning models on multiple GPUs.
Certificate	Upon successful completion of the assessment, participants will receive an NVIDIA DLI certificate to recognize their subject matter competency and support professional career growth.
Hardware Requirements	Desktop or laptop computer capable of running the latest version of Chrome or Firefox. Each participant will be provided with dedicated access to a fully configured, GPU-accelerated workstation in the cloud.
Language	English

Workshop Outline

Introduction (15 minutes)

Meet the instructor.

- > Create an account at courses.nvidia.com/join

Stochastic Gradient Descent and the Effects of Batch Size (120 minutes)

Learn the significance of stochastic gradient descent when training on multiple GPUs:

- > Understand the issues with sequential single-thread data processing and the theory behind speeding up applications with parallel processing.
- > Understand loss function, gradient descent, and stochastic gradient descent (SGD).
- > Understand the effect of batch size on accuracy and training time with an eye toward its use on multi-GPU systems.

Break (60 minutes)

Training on Multiple GPUs with PyTorch Distributed Data Parallel (DDP) (120 minutes)

Learn to convert single-GPU training to multiple GPUs using PyTorch Distributed Data Parallel:

- > Understand how DDP coordinates training among multiple GPUs.
- > Refactor single-GPU training programs to run on multiple GPUs with DDP.

Break (15 minutes)

Maintaining Model Accuracy when Scaling to Multiple GPUs (90 minutes)

Understand and apply key algorithmic considerations to retain accuracy when training on multiple GPUs:

- > Understand what might cause accuracy to decrease when parallelizing training on multiple GPUs.
- > Learn and understand techniques for maintaining accuracy when scaling training to multiple GPUs.

Workshop Assessment (30 minutes)

Use what you have learned during the workshop:

- > Complete the workshop assessment to earn a certificate of competency.

Final Review (15 minutes)

Review key learnings and wrap up questions.

Take the workshop survey.

Next Steps

Continue learning with these DLI trainings:

- > [Model Parallelism: Building and Deploying Large Neural Networks](#)

Why Choose NVIDIA Deep Learning Institute for Hands-On Training?

- > Access workshops from anywhere with just your desktop/laptop and an internet connection. Each participant will have access to a fully configured, GPU-accelerated server in the cloud.
- > Obtain hands-on experience with the most widely used, industry-standard software, tools, and frameworks.
- > Learn to build deep learning and accelerated computing applications for industries, such as healthcare, robotics, manufacturing, accelerated computing, and more.
- > Gain real-world experience through content designed in collaboration with industry leaders, such as the Children's Hospital of Los Angeles, Mayo Clinic, and PwC.
- > Earn an NVIDIA DLI certificate to demonstrate your subject matter competency and support your career growth.

Ready to Get Started?

For the latest DLI workshops and trainings, visit

www.nvidia.com/dli

For questions, contact us at nvdl@nvidia.com