# Fundamentals of Deep Learning for Multi-GPUs

This workshop teaches you techniques for training deep neural networks on multi-GPU technology 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 concepts for implementing Horovod multi-GPUs to reduce the complexity of writing efficient distributed software and to maintain accuracy when training a model across many GPUs.

| | |
|---|---|
| **Duration:** | 8 hours |
| **Price:** | Contact us for pricing. During the workshop, each participant will have dedicated access to a fully configured, GPU-accelerated workstation in the cloud. |
| **Assessment type:** | Code-based |
| **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. |
| **Prerequisites:** | Experience with gradient descent model training |
| **Languages:** | English |
| **Tools, libraries, and frameworks:** | TensorFlow, Keras, Horovod |

## Learning Objectives

At the conclusion of the workshop, you'll have an understanding of:

- Stochastic Gradient Descent, a crucial tool in parallelized training
- Batch size and its effect on training time and accuracy
- Transforming a single-GPU implementation to a Horovod multi-GPU implementation
- Techniques for maintaining high accuracy when training across multiple GPUs

## Why Deep Learning Institute Hands-On Training?

- Learn to build deep learning and accelerated computing applications for industries such as autonomous vehicles, finance, game development, healthcare, robotics, and more.
- Obtain hands-on experience with the most widely used, industry-standard software, tools, and frameworks.
- Gain real-world expertise 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 career growth.
- Access content anywhere, anytime with a fully configured, GPU-accelerated workstation in the cloud.

## Workshop Outline

| TOPIC | DESCRIPTION |
|---|---|
| **Introduction**<br>(15 mins) | > Meet the instructor.<br>> Create an account at **courses.nvidia.com/join** |
| **Stochastic Gradient Descent and the Effects of Batch Size**<br>(120 mins) | > Understand the issues with sequential single-thread data processing and the theory behind speeding up applications with parallel processing.<br>> Explore loss function, gradient descent, and stochastic gradient descent (SGD).<br>> Learn the effect of batch size on accuracy and training time. |
| **Break** (60 mins) | |
| **Training on Multiple GPUs with Horovod**<br>(120 mins) | > Discover the benefits of training on multiple GPUs with Horovod.<br>> Learn to transform single-GPU training on the Fashion MNIST dataset to Horovod multi-GPU implementation. |
| **Break** (15 mins) | |
| **Maintaining Model Accuracy when Scaling to Multiple GPUs**<br>(120 mins) | > Understand why accuracy can decrease when parallelizing training on multiple GPUs.<br>> Explore tools for maintaining accuracy when scaling training to multiple GPUs. |
| **Final Review**<br>(15 mins) | > Review key learnings and wrap up questions.<br>> Complete the assessment to earn a certificate.<br>> Take the workshop survey. |