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.

Duration:	8 hours
Price:	$10,000 for groups of up to 20 (price increase for larger groups). 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 stochastic gradient descent
Languages:	English
Tools, libraries, and frameworks:	TensorFlow

Learning Objectives

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

- Various approaches to multi-GPU training
- Algorithmic and engineering challenges to the large-scale training of a neural network
- The linear neuron model and the loss function and optimization logic for gradient descent
- Concepts for transforming single-GPU implementation to Horovod multi-GPU implementation to reduce the complexity of writing efficient distributed software
- Techniques that improve overall performance of the entire pipeline

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 (15 mins)	> Meet the instructor. > Create an account at **courses.nvidia.com/join**
Theory of Data Parallelism and Introduction to Multi-GPU Training (120 mins)	> 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). > Iteratively calculate the gradient of the cost function and model parameters using the SGD optimization algorithm.
Break (60 mins)	
Algorithmic Challenges to Multi-GPU training (120 mins)	> Learn to transform single-GPU implementation to Horovod multi-GPU implementation to simplify the writing of efficient distributed software. > Understand data loading, augmentation, and training logic using AlexNet.
Break (15 mins)	
Engineering Challenges to Multi-GPU training (120 mins)	> Understand the aspects of the data input pipeline, communication, and reference architecture. > Take a deeper dive into the concepts of job scheduling.
Final Review (15 mins)	> Review key learnings and wrap up questions. > Complete the assessment to earn a certificate. > Take the workshop survey.