



NVIDIA DLI HANDS-ON TRAINING

COURSE CATALOG

Valid Through July 31, 2018



DEEP
LEARNING
INSTITUTE

To start DLI hands-on training, visit
www.nvidia.com/dli

INTRODUCTION

The NVIDIA Deep Learning Institute (DLI) trains developers, data scientists, and researchers on how to use artificial intelligence and accelerated computing to solve real-world problems across a wide range of domains.

In deep learning courses, you'll learn how to train, optimize, and deploy neural networks. In accelerated computing courses, you'll learn how to assess, parallelize, optimize, and deploy GPU-accelerated computing applications.

DLI offers training in three formats:

INSTRUCTOR-LED WORKSHOPS

In-person workshops teach you how to implement and deploy an end-to-end project in eight hours. Offered at customer sites, conferences, and universities, full-day workshops include hands-on training and lectures delivered by DLI-certified instructors.

ONLINE COURSES

Online courses teach you how to implement and deploy an end-to-end project through hands-on training in eight hours. Online courses can be taken anytime, anywhere, and provide access to a fully configured GPU-accelerated workstation in the cloud.

ONLINE MINI COURSES

Online mini courses explore how to apply a specific technology or development technique through hands-on training in two hours. Like full-length online courses, online mini courses can be taken anytime, anywhere, with access to GPUs in the cloud.



CERTIFICATION

Participants can earn certification to prove subject matter competency and support professional career growth. Certification is offered for select instructor-led workshops and online courses.

INSTRUCTOR-LED WORKSHOPS

INTRODUCTION TO DEEP LEARNING

Fundamentals of Deep Learning for Computer Vision

Explore the fundamentals of deep learning by training neural networks and using results to improve performance and capabilities. Learn how to start solving problems with deep learning.

PREREQUISITES: None

LANGUAGES: English

FRAMEWORKS: Caffe

Fundamentals of Deep Learning for Multiple Data Types

Explore how convolutional and recurrent neural networks can be combined to generate effective descriptions of content within images and video clips. Learn how to train a network using TensorFlow and the Microsoft Common Objects in Context (COCO) dataset to generate captions from images and video.

PREREQUISITES: "Fundamentals of Deep Learning for Computer Vision" or similar experience

LANGUAGES: English

FRAMEWORKS: TensorFlow

INTRODUCTION TO ACCELERATED COMPUTING

Fundamentals of Accelerated Computing with CUDA C/C++

The CUDA computing platform enables the acceleration of CPU-only applications to run on the world's fastest massively parallel GPUs. Learn how to accelerate and optimize existing C/C++ CPU-only applications using the most essential CUDA tools and techniques.

PREREQUISITES: Basic experience with C/C++

LANGUAGES: English

DEEP LEARNING WORKSHOP BY INDUSTRY

AUTONOMOUS VEHICLES

Deep Learning for Autonomous Vehicles—Perception

Learn how to design, train, and deploy deep neural networks for autonomous vehicles using the NVIDIA DRIVE PX development platform. Create and optimize perception components for autonomous vehicles using NVIDIA DRIVE PX.

PREREQUISITES: Experience with CNNs

LANGUAGES: English

FRAMEWORKS: TensorFlow, DIGITS, TensorRT

FINANCE

Deep Learning for Finance Trading Strategy

Finance trading strategies can be advanced with the power of deep neural networks. Learn how to use time series financial data to make predictions and exploit arbitrage using neural networks.

PREREQUISITES: Experience with neural networks and knowledge of the financial industry

LANGUAGES: English

FRAMEWORKS: TensorFlow

GAME DEVELOPMENT AND DIGITAL CONTENT

Deep Learning for Digital Content Creation Using GANs and Autoencoders

Explore the latest techniques for designing, training, and deploying neural networks for digital content creation. Learn how to train a generative adversarial network (GAN) to generate images, explore techniques to make video style transfer, and train a denoiser for rendered images.

PREREQUISITES: Experience with CNNs

LANGUAGES: English

FRAMEWORKS: Torch, TensorFlow

HEALTHCARE

Deep Learning for Healthcare Image Analysis

Learn how to apply convolutional neural networks (CNNs) to MRI scans to perform a variety of medical tasks and calculations.

PREREQUISITES: Basic experience with CNNs and Python

LANGUAGES: English

FRAMEWORKS: Caffe, DIGITS, MXNet, TensorFlow

Deep Learning for Healthcare Genomics

Learn how convolutional neural networks (CNNs) work and how to apply deep learning to detect chromosome co-deletion and search for motifs in genomic sequences.

PREREQUISITES: Basic experience with CNNs and Python

LANGUAGES: English

FRAMEWORKS: Caffe, DIGITS, MXNet, TensorFlow

ONLINE COURSES

INTRODUCTION TO DEEP LEARNING

COURSES

Fundamentals of Deep Learning for Computer Vision

Explore how to start solving problems with deep learning by training neural networks and using results to improve performance and capabilities.

PREREQUISITES: None

LANGUAGES: English

FRAMEWORKS: Caffe

PRICE: \$90

MINI COURSES

Image Classification with DIGITS

Learn how to train a deep neural network to recognize handwritten digits by loading image data into a training environment, choosing and training a network, testing with new data, and iterating to improve performance.

PREREQUISITES: None

LANGUAGES: English, Chinese, Japanese

FRAMEWORKS: Caffe (with DIGITS interface)

PRICE: Free

Object Detection with DIGITS

Learn how to detect objects using computer vision and deep learning by identifying a purpose-built network and using end-to-end labeled data.

PREREQUISITES: Basic experience with neural networks

LANGUAGES: English, Chinese

FRAMEWORKS: Caffe (with DIGITS interface)

PRICE: Free

Neural Network Deployment with DIGITS and TensorRT

Learn to deploy deep learning to applications that recognize and detect images in real time.

PREREQUISITES: Basic experience with neural networks

LANGUAGES: English, Chinese

FRAMEWORKS: DIGITS, TensorRT

PRICE: \$30

Applications of Deep Learning with Caffe, Theano, and Torch

Explore how deep learning works and how it will change the future of computing.

PREREQUISITES: None

LANGUAGES: English

FRAMEWORKS: Caffe, Theano, Torch

PRICE: \$30

Deep Learning Workflows with TensorFlow, MXNet, and NVIDIA Docker

Learn how to use the NVIDIA Docker plug-in to containerize production-grade deep learning workflows using GPUs.

PREREQUISITES: Basic experience with a bash terminal

LANGUAGES: English, Japanese

FRAMEWORKS: TensorFlow, MXNet

PRICE: \$30

Image Segmentation with TensorFlow

Learn how to combine computer vision and natural language processing to describe scenes using deep learning.

PREREQUISITES: Basic experience with neural networks

LANGUAGES: English

FRAMEWORKS: TensorFlow

PRICE: \$30

Image Classification with Microsoft Cognitive Toolkit

Learn how to train a neural network using the Microsoft Cognitive Toolkit framework.

PREREQUISITES: None

LANGUAGES: English

FRAMEWORKS: Microsoft Cognitive Toolkit

PRICE: \$30

Linear Classification with TensorFlow

Learn how to make predictions from structured data using TensorFlow's TFLearn application programming interface (API).

PREREQUISITES: None

LANGUAGES: English

FRAMEWORKS: TensorFlow

PRICE: \$30

Signal Processing with DIGITS

Learn how to classify both image and image-like data using deep learning by converting radio frequency (RF) signals into images to detect a weak signal corrupted by noise.

PREREQUISITES: Basic experience training neural networks

LANGUAGES: English, Chinese

FRAMEWORKS: Caffe, DIGITS

PRICE: \$30

INTRODUCTION TO ACCELERATING COMPUTING

COURSES

Fundamentals of Accelerated Computing with CUDA C/C++

The CUDA computing platform enables the acceleration of CPU-only applications to run on the world's fastest massively parallel GPUs. Learn how to accelerate and optimize existing C/C++ CPU-only applications using the most essential CUDA tools and techniques.

PREREQUISITES: Basic experience with C/C++

LANGUAGES: English

PRICE: \$90

Fundamentals of Accelerated Computing with CUDA Python

Explore how to use Numba—the just-in-time, type-specializing Python function compiler—to create and launch CUDA kernels to accelerate Python programs on massively parallel NVIDIA GPUs.

PREREQUISITES: Basic experience with Python and NumPy

LANGUAGES: English

PRICE: \$90

Fundamentals of Accelerated Computing with OpenACC

Learn the basics of OpenACC, a high-level programming language for programming on GPUs. Explore how to build and optimize accelerated heterogeneous applications on multiple GPU clusters using a combination of OpenACC, CUDA-aware MPI, and NVIDIA profiling tools.

PREREQUISITES: Basic experience with C/C++

LANGUAGES: English

PRICE: \$116

MINI COURSES

Accelerating Applications with CUDA C/C++

Learn how to accelerate your C/C++ application using CUDA to harness the massively parallel power of NVIDIA GPUs.

PREREQUISITES: Basic experience with C/C++

LANGUAGES: English

PRICE: Free

OpenACC – 2X in 4 Steps

Learn how to accelerate C/C++ or Fortran applications using OpenACC to harness the massively parallel power of NVIDIA GPUs.

PREREQUISITES: Basic experience with C/C++

LANGUAGES: English

PRICE: Free

Introduction to Accelerated Computing

Explore a variety of techniques for accelerating applications, including CUDA and OpenACC.

PREREQUISITES: None

LANGUAGES: English

PRICE: \$30

GPU Memory Optimizations with CUDA C/C++

Learn useful memory optimization techniques for programming with CUDA C/C++ on an NVIDIA GPU and how to use the NVIDIA Visual Profiler (NVVP) to support these optimizations.

PREREQUISITES: Basic experience with C/C++

LANGUAGES: English

PRICE: \$30

Accelerating Applications with GPU-Accelerated Libraries in C/C++

Learn how to accelerate your C/C++ application using CUDA-optimized libraries to harness the massively parallel power of NVIDIA GPUs.

PREREQUISITES: "Accelerating Applications with CUDA C/C++" or similar experience

LANGUAGES: English

Accelerating Applications with GPU-Accelerated Libraries in Python

Learn how to accelerate your Python application using CUDA-optimized libraries to harness the massively parallel power of NVIDIA GPUs.

PREREQUISITES: Basic experience with Python

LANGUAGES: English

PRICE: \$30

Using Thrust to Accelerate C++

Discover how to build GPU-accelerated applications in C/C++ that utilize the powerful Thrust library.

PREREQUISITES: "Accelerating Applications with CUDA C/C++" or similar experience

LANGUAGES: English

PRICE: \$30

Profiling and Parallelizing with OpenACC

Get hands-on experience with the first two steps of the OpenACC programming cycle.

PREREQUISITES: "OpenACC – 2X in 4 Steps" or similar experience

LANGUAGES: English

PRICE: \$30

Expressing Data Movement and Optimizing Loops with OpenACC

Learn how to add data management and loop directives to optimize OpenACC-accelerated code.

PREREQUISITES: "Profiling and Parallelizing with OpenACC" or similar experience

LANGUAGES: English

PRICE: \$30

Introduction to Multi-GPU Programming with MPI and OpenACC

Explore how to program multi-GPU systems or GPU clusters using the Message Passing Interface (MPI) and OpenACC.

PREREQUISITES: "OpenACC – 2X in 4 Steps" or similar experience

LANGUAGES: English

PRICE: \$30

Advanced Multi-GPU Programming with MPI and OpenACC

Learn how to improve multi-GPU Message Passing Interface (MPI) and OpenACC programs by overlapping communication with computation and handling noncontiguous halo updates.

PREREQUISITES: "Introduction to Multi-GPU Programming with MPI and OpenACC" or similar experience

LANGUAGES: English

PRICE: \$30

Pipelining Work on the GPU with OpenACC

Learn how to optimize data copies in OpenACC applications to overlap with GPU computation using a simple technique called pipelining.

PREREQUISITES: "Expressing Data Movement and Optimizing Loops with OpenACC" or similar experience

LANGUAGES: English

PRICE: \$30

Profile-Driven Approach to Accelerate Seismic Applications with OpenACC

Learn how to use PGI Profiler (PGPROF), a host and GPU-profiling tool, with OpenACC to accelerate your C/C++ applications.

PREREQUISITES: None

LANGUAGES: English

PRICE: \$30

Accelerating Applications with CUDA Fortran

Learn how to accelerate your Fortran application using CUDA to harness the massively parallel power of NVIDIA GPUs.

PREREQUISITES: Basic experience with Fortran

LANGUAGES: English

PRICE: \$30

GPU Memory Optimizations with CUDA Fortran

Discover useful memory optimization techniques for programming with CUDA Fortran on an NVIDIA GPU and how to use the NVIDIA Visual Profiler (NVVP) to support these optimizations.

PREREQUISITES: "Accelerating Applications with CUDA Fortran" or similar experience

LANGUAGES: English

PRICE: \$30

Accelerating Applications with GPU-Accelerated Libraries in Fortran

Learn how to accelerate your Fortran application using CUDA-optimized libraries.

PREREQUISITES: Basic experience with Fortran

LANGUAGES: English

PRICE: \$30

DEEP LEARNING COURSES BY INDUSTRY

GAME DEVELOPMENT AND DIGITAL CONTENT

MINI COURSES

Image Creation Using Generative Adversarial Networks (GANs) with TensorFlow and DIGITS

Discover how to train a generative adversarial network (GAN) to generate image content in DIGITS.

PREREQUISITES: Experience with CNNs

LANGUAGES: English

FRAMEWORKS: TensorFlow

PRICE: \$30

Image Style Transfer with Torch

Learn how to transfer the look and feel of one image to another image by extracting distinct visual features using convolutional neural networks (CNNs).

PREREQUISITES: Experience with CNNs

LANGUAGES: English

FRAMEWORKS: Torch

PRICE: \$30

Rendered Image Denoising Using Autoencoders

Explore how a neural network with an autoencoder can be used to dramatically speed up the removal of noise in ray-traced images.

PREREQUISITES: Experience with CNNs

LANGUAGES: English

FRAMEWORKS: TensorFlow

PRICE: \$30

HEALTHCARE

COURSES

Deep Learning for Healthcare Image Analysis

Learn how to apply convolutional neural networks (CNNs) to MRI scans to perform a variety of medical tasks and calculations.

PREREQUISITES: Basic experience with CNNs and Python **LANGUAGES:** English

FRAMEWORKS: Caffe, DIGITS, MXNet, TensorFlow **PRICE:** \$90

Deep Learning for Healthcare Genomics

Learn how convolutional neural networks (CNNs) work and how to apply deep learning to detect chromosome co-deletion and search for motifs in genomic sequences.

PREREQUISITES: Basic experience with CNNs and Python **LANGUAGES:** English

FRAMEWORKS: Caffe, TensorFlow, Theano **PRICE:** \$60

MINI COURSES

Modeling Time Series Data with Recurrent Neural Networks in Keras

Explore how to classify and forecast time series data using recurrent neural networks (RNNs), such as modeling a patient's health over time.

PREREQUISITES: Basic experience with deep learning **LANGUAGES:** English

FRAMEWORKS: Keras **PRICE:** Free

For full course details or to enroll in an online course, visit
www.nvidia.com/dli

