



NVIDIA DLI HANDS-ON TRAINING

CATALOG

Published on August 1, 2019



DEEP
LEARNING
INSTITUTE

To get started with 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 deep learning and accelerated computing to solve real-world problems across a wide range of domains.

With access to GPU-accelerated workstations in the cloud, you'll learn how to train, optimize, and deploy neural networks using the latest deep learning tools, frameworks, and SDKs. You'll also learn how to assess, parallelize, optimize, and deploy GPU-accelerated computing applications.

DLI offers training in two formats:

INSTRUCTOR-LED WORKSHOPS

In-person workshops teach you how to implement and deploy an end-to-end project in one day. Offered at customer sites, conferences, and universities, full-day workshops include hands-on training and lectures delivered by DLI-certified instructors. Workshops offer a certificate of competency upon completion of the built-in assessment.

ONLINE COURSES

Online, self-paced courses teach you how to implement and deploy an end-to-end project in eight hours or how to apply a specific technology or development technique in two hours. Online courses can be taken anytime, anywhere with a laptop and internet connection. Most eight-hour courses offer a certificate of competency upon completion of the built-in assessment.



CERTIFICATE

Participants can earn a certificate to prove subject matter competency and support professional career growth. Certificates are offered for select instructor-led workshops and online courses.

INSTRUCTOR-LED WORKSHOPS

DEEP LEARNING FUNDAMENTALS

Fundamentals of Deep Learning for Computer Vision

Learn how to build, train, and deploy a neural network for common deep learning workflows to solve real world problems.

PREREQUISITES: Familiarity with basic programming fundamentals, such as functions and variables

TOOLS AND FRAMEWORKS: Caffe, DIGITS

LANGUAGES: English, Japanese, Korean, Simplified Chinese, Traditional Chinese

Fundamentals of Deep Learning for Multiple Data Types

Learn how to train convolutional neural networks (CNNs) and recurrent neural networks (RNNs) to generate captions from images and video using TensorFlow and the Microsoft Common Objects in Context (COCO) dataset.

PREREQUISITES: Familiarity with basic Python (functions and variables), prior experience training neural networks

TOOLS AND FRAMEWORKS: TensorFlow

LANGUAGES: English, Japanese, Korean, Traditional Chinese

Fundamentals of Deep Learning for Natural Language Processing

Explore how to convert text to machine understandable representation and train Machine Translators from one language to another using natural language processing (NLP).

PREREQUISITES: Basic experience with neural networks and Python, familiarity with linguistics

TOOLS AND FRAMEWORKS: TensorFlow, Keras

LANGUAGES: English

Fundamentals of Deep Learning for Multi-GPU

Learn how to use multiple GPUs to train neural networks and effectively parallelize training of deep neural networks using TensorFlow.

PREREQUISITES: Experience with stochastic gradient descent mechanics, network architecture, and parallel computing

TOOLS AND FRAMEWORKS: TensorFlow

LANGUAGES: English

DEEP LEARNING BY INDUSTRY

Deep Learning for Digital Content Creation with Autoencoders

Learn how to animate characters with phase-function neural networks, explore techniques to make arbitrary photo and video style transfer, and train your own denoiser for rendered images.

PREREQUISITES: Basic familiarity with deep learning concepts such as CNNs, experience with Python

TOOLS AND FRAMEWORKS: TensorFlow, Torch

LANGUAGES: English

Deep Learning for Healthcare Image Analysis

Learn how to apply CNNs to MRI scans to perform a variety of medical tasks and calculations.

PREREQUISITES: Basic familiarity with deep neural networks, basic coding experience in Python or similar language

TOOLS AND FRAMEWORKS: R, MXNet, TensorFlow, Caffe, DIGITS

LANGUAGES: English

Deep Learning for Industrial Inspection

Learn how to design, train, test, and deploy building blocks of a hardware-accelerated industrial inspection pipeline.

PREREQUISITES: Familiarity with deep neural networks; experience with Python and deep learning frameworks such as TensorFlow, Keras, and PyTorch

TOOLS AND FRAMEWORKS: TensorFlow, TensorRT, Keras

LANGUAGES: English, Traditional Chinese

Deep Learning for Intelligent Video Analytics

Explore how to deploy object detection and tracking networks to evaluate real-time, large-scale video streams.

PREREQUISITES: Experience with deep networks (specifically variations of CNNs), intermediate-level experience with C++ and Python

TOOLS AND FRAMEWORKS: DeepStream 3.0, TensorFlow

LANGUAGES: English, Korean

Deep Learning for Robotics

Explore how to create robotic solutions on an NVIDIA Jetson for embedded applications.

PREREQUISITES: Basic familiarity with deep neural networks, basic coding experience in Python or similar language

TOOLS AND FRAMEWORKS: ROS, DIGITS, NVIDIA Jetson

LANGUAGES: English

ACCELERATED COMPUTING FUNDAMENTALS

Fundamentals of Accelerated Computing with CUDA C/C++

Learn how to accelerate and optimize existing C/C++ CPU-only applications to run on massively parallel GPUs using essential CUDA tools and techniques.

PREREQUISITES: Basic C/C++ competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations.

LANGUAGES: English, Korean, Traditional Chinese

Fundamentals of Accelerated Computing with CUDA Python

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

PREREQUISITES: Basic Python competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. NumPy competency including the use of ndarrays and ufuncs.

LANGUAGES: English

ONLINE COURSES

DEEP LEARNING FUNDAMENTALS

Fundamentals of Deep Learning for Computer Vision

Learn how to build, train, and deploy a neural network for common deep learning workflows to solve real world problems.

PREREQUISITES: Familiarity with basic programming fundamentals, such as functions and variables

LANGUAGES: English, Japanese, Korean, Simplified Chinese, Traditional Chinese

TOOLS AND FRAMEWORKS: Caffe, DIGITS

DURATION: 8 hours

PRICE: \$90

Getting Started with Deep Learning on Jetson Nano

Explore how to build a deep learning classification project with computer vision models using your NVIDIA Jetson Nano Developer Kit.

PREREQUISITES: Basic familiarity with Python (helpful, not required)

TOOLS AND FRAMEWORKS: PyTorch, Jetson Nano

LANGUAGES: English

DURATION: 8 hours

PRICE: Free (hardware required)

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, Japanese, Simplified Chinese

TOOLS AND FRAMEWORKS: Caffe (with DIGITS interface)

DURATION: 2 hours

PRICE: \$30

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

TOOLS AND FRAMEWORKS: Caffe (with DIGITS interface)

LANGUAGES: English, Simplified Chinese

DURATION: 2 hours

PRICE: \$30

Optimization and Deployment of TensorFlow Models with TensorRT

Learn how to optimize TensorFlow models to generate fast inference engines in the deployment stage.

PREREQUISITES: Experience with TensorFlow and Python

TOOLS AND FRAMEWORKS: TensorFlow, Python, TensorRT (TF-TRT)

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

Accelerating Data Science Workflows with RAPIDS

Learn to build a GPU-accelerated, end-to-end data science workflow using RAPIDS open-source libraries for massive performance gains.

PREREQUISITES: Advanced competency in Pandas, NumPy, and scikit-learn

TOOLS AND FRAMEWORKS: None

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

Deep Learning at Scale with Horovod

Learn how to scale deep learning training to multiple GPUs with Horovod, the open-source distributed training framework originally built by Uber and hosted by the LF AI Foundation.

PREREQUISITES: Competency in Python and professional experience training deep learning models in Python

TOOLS AND FRAMEWORKS: Horovod, TensorFlow, Keras

LANGUAGES: English

DURATION: 2 hours

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

TOOLS AND FRAMEWORKS: TensorFlow

LANGUAGES: English

DURATION: 2 hours

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

TOOLS AND FRAMEWORKS: Caffe, DIGITS

LANGUAGES: English, Simplified Chinese

DURATION: 2 hours

PRICE: \$30

DEEP LEARNING BY INDUSTRY

GAME DEVELOPMENT AND DIGITAL CONTENT

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 CNNs.

PREREQUISITES: Experience with CNNs

TOOLS AND FRAMEWORKS: Torch

DURATION: 2 hours

LANGUAGES: English

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

TOOLS AND FRAMEWORKS: TensorFlow

DURATION: 2 hours

LANGUAGES: English

PRICE: \$30

Image Super Resolution Using Autoencoders

Leverage the power of a neural network with autoencoders to create high-quality images from low-quality source images.

PREREQUISITES: Experience with CNNs

TOOLS AND FRAMEWORKS: Keras

DURATION: 2 hours

LANGUAGES: English,
Simplified Chinese

PRICE: \$30

HEALTHCARE

Modeling Time Series Data with Recurrent Neural Networks in Keras

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

PREREQUISITES: Basic experience with deep learning

TOOLS AND FRAMEWORKS: Keras

DURATION: 2 hours

LANGUAGES: English

PRICE: \$30

Medical Image Classification Using the MedNIST Dataset

Explore an introduction to deep learning for radiology and medical imaging by applying CNNs to classify images in a medical imaging dataset.

PREREQUISITES: Basic experience in Python

TOOLS AND FRAMEWORKS: PyTorch

DURATION: 2 hours

LANGUAGES: English,
Simplified Chinese

PRICE: \$30

Data Science Workflows for Deep Learning in Medical Applications

Learn how to apply data augmentation and standardization techniques to a medical imaging dataset and validate your techniques by training a CNN on the dataset.

PREREQUISITES: Basic experience with Python and CNNs

TOOLS AND FRAMEWORKS: PyTorch

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

Medical Image Segmentation with DIGITS

Explore how to segment MRI images to measure parts of the heart by experimenting with TensorFlow tools, such as TensorBoard and the TensorFlow Python API.

PREREQUISITES: Basic experience with CNNs and Python

TOOLS AND FRAMEWORKS: DIGITS, Caffe

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

Image Classification with TensorFlow: Radiomics—1p19q Chromosome Status Classification

Learn how to train CNNs to detect radiomics from MRI imaging.

PREREQUISITES: Basic experience with CNNs and Python

LANGUAGES: English, Simplified Chinese

TOOLS AND FRAMEWORKS: TensorFlow

DURATION: 2 hours

PRICE: \$30

Medical Image Analysis with R and MXNet

Learn how to train a CNN to infer the volume of the left ventricle of the human heart from time-series MRI data.

PREREQUISITES: Basic experience with CNNs and Python

TOOLS AND FRAMEWORKS: MXNet

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

Data Augmentation and Segmentation with Generative Networks for Medical Imaging

Learn how to use GANs for medical imaging by applying them to the creation and segmentation of brain MRIs.

PREREQUISITES: Experience with CNNs

TOOLS AND FRAMEWORKS: TensorFlow

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

Coarse-to-Fine Contextual Memory for Medical Imaging

Learn how to use Coarse-to-Fine Context Memory (CFCM) to improve traditional architectures for medical image segmentation and classification tasks.

PREREQUISITES: Experience with CNNs and long short-term memory (LSTM)

TOOLS AND FRAMEWORKS: TensorFlow

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

INTELLIGENT VIDEO ANALYTICS

AI Workflows for Intelligent Video Analytics with DeepStream

Learn how to build hardware-accelerated applications for intelligent video analytics (IVA) with DeepStream and deploy them at scale to transform video streams into insights.

PREREQUISITES: Experience with C++ and GStreamer

TOOLS AND FRAMEWORKS: DeepStream 3.0

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

ACCELERATED COMPUTING FUNDAMENTALS

Fundamentals of Accelerated Computing with CUDA C/C++

Learn how to accelerate and optimize existing C/C++ CPU-only applications to run on massively parallel GPUs using essential CUDA tools and techniques.

PREREQUISITES: Basic C/C++ competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations.

LANGUAGES: English, Japanese, Korean, Simplified Chinese, Traditional Chinese

PRICE: \$90

DURATION: 8 hours

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 Python competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. NumPy competency including the use of ndarrays and ufuncs.

LANGUAGES: English

PRICE: \$90

DURATION: 8 hours

Fundamentals of Accelerated Computing with OpenACC

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

DURATION: 8 hours

PRICE: \$90

High-Performance Computing with Containers

Learn how to reduce complexity and improve portability and efficiency of your code by using a containerized environment for high-performance computing (HPC) application development.

PREREQUISITES: Proficiency programming in C/C++ and professional experience working on HPC applications

TOOLS AND FRAMEWORKS: Docker, Singularity, HPC Container Maker (HPCCM)

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

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, Japanese

DURATION: 2 hours

PRICE: \$30

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

DURATION: 2 hours

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 accelerating applications with CUDA C/C++

LANGUAGES: English

DURATION: 2 hours

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: Basic experience accelerating applications with CUDA C/C++

LANGUAGES: English, Japanese

DURATION: 2 hours

PRICE: \$30

Using Thrust to Accelerate C++

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

PREREQUISITES: Basic experience accelerating applications with CUDA C/C++

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

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

