



NVIDIA DLI HANDS-ON TRAINING

CATALOG

Published on May 1, 2019



DEEP
LEARNING
INSTITUTE

To start DLI hands-on training, visit
www.nvidia.co.uk/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 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. Most 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

FRAMEWORKS: Caffe, DIGITS

LANGUAGES: English, Chinese, Japanese, Korean

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

FRAMEWORKS: TensorFlow

LANGUAGES: English, Chinese, Japanese, Korean

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

FRAMEWORKS: TensorFlow, Keras

LANGUAGES: English, Chinese

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

FRAMEWORKS: TensorFlow

LANGUAGES: English

DEEP LEARNING BY INDUSTRY

Deep Learning for Autonomous Vehicles—Perception

Learn how to train a semantic segmentation neural network and deploy the neural network using TensorRT on the NVIDIA Drive AGX development platform.

PREREQUISITES: Experience with CNNs

LANGUAGES: English, Chinese, Japanese

FRAMEWORKS: TensorFlow, DIGITS, TensorRT

Deep Learning for Finance Trading Strategy

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

FRAMEWORKS: TensorFlow

LANGUAGES: English

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

FRAMEWORKS: TensorFlow, Torch

LANGUAGES: English, Chinese

Deep Learning for Digital Content Creation with GANs



Learn how to train a generative adversarial network (GAN) to generate images, convert text to images using deep learning, and create analogous images from one theme to another.

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

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

FRAMEWORKS: Caffe, DIGITS, R, MXNet, TensorFlow

LANGUAGES: English, Japanese

Deep Learning for Healthcare Genomics



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

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

FRAMEWORKS: TensorFlow, Caffe, DIGITS, Theano, deep regulatory genomics neural network (DragoNN)

LANGUAGES: English, Japanese

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

FRAMEWORKS: DeepStream 3.0, TensorFlow

LANGUAGES: English

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

FRAMEWORKS: DIGITS

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, Chinese, Japanese, Korean

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

FRAMEWORKS: Caffe, DIGITS

LANGUAGES: English

DURATION: 8 hours

PRICE: \$90

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)

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

FRAMEWORKS: Caffe (with DIGITS interface)

LANGUAGES: English, 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

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

FRAMEWORKS: None

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

FRAMEWORKS: TensorFlow

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

Image Classification with Microsoft Cognitive Toolkit

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

PREREQUISITES: None

FRAMEWORKS: Microsoft Cognitive Toolkit

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

Linear Classification with TensorFlow

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

PREREQUISITES: None

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

FRAMEWORKS: Caffe, DIGITS

LANGUAGES: English, Chinese

DURATION: 2 hours

PRICE: \$30

DEEP LEARNING BY INDUSTRY

GAME DEVELOPMENT AND DIGITAL CONTENT

Image Creation Using Generative Adversarial Networks with TensorFlow and DIGITS

Discover how to train a GAN to generate image content in DIGITS.

PREREQUISITES: Experience with CNNs

FRAMEWORKS: TensorFlow

LANGUAGES: English

DURATION: 2 hours

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 CNNs.

PREREQUISITES: Experience with CNNs

FRAMEWORKS: Torch

LANGUAGES: English

DURATION: 2 hours

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

FRAMEWORKS: TensorFlow

LANGUAGES: English

DURATION: 2 hours

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

FRAMEWORKS: Keras

LANGUAGES: English

DURATION: 2 hours

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

FRAMEWORKS: Keras

LANGUAGES: English

DURATION: 2 hours

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

FRAMEWORKS: PyTorch

LANGUAGES: English, Chinese

DURATION: 2 hours

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

FRAMEWORKS: PyTorch

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

Medical Image Segmentation Using 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

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

FRAMEWORKS: TensorFlow

LANGUAGES: English

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

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

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

FRAMEWORKS: TensorFlow

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

Deep Learning for Genomics Using DragoNN with Keras and Theano

Learn to interpret deep learning models to discover predictive genome sequence patterns using the DragoNN toolkit.

PREREQUISITES: Basic experience with CNNs and Python

FRAMEWORKS: Keras, Theano

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

INTELLIGENT VIDEO ANALYTICS

Deployment for Intelligent Video Analytics Using TensorRT

Learn how to use TensorRT to accelerate inferencing performance for neural networks.

PREREQUISITES: Basic experience with CNNs and C++

FRAMEWORKS: TensorRT

LANGUAGES: English

DURATION: 2 hours

PRICE: \$30

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

FRAMEWORKS: DeepStream3

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

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

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: "Accelerating Applications with CUDA C/C++" or similar experience **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: "Accelerating Applications with CUDA C/C++" or similar experience **LANGUAGES:** English

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: "Accelerating Applications with CUDA C/C++" or similar experience **LANGUAGES:** English

DURATION: 2 hours

PRICE: \$30

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

