NVIDIA DLI COURSE CATALOG

INTRODUCTION

The NVIDIA Deep Learning Institute (DLI) provides hands-on training in AI, accelerated computing and accelerated data science to help developers, data scientists and other professionals solve their most challenging problems. And IT professionals can learn how to design and manage infrastructure for AI, data science, and high-performance computing (HPC) workloads across their organizations.

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

DLI workshops teach you how to implement and deploy an end-to-end project in one day. These in-depth classes are taught by experts in their respective fields, delivering industry-leading technical knowledge to drive breakthrough results for individuals and organizations. Workshops are delivered remotely via a virtual classroom for customers, conferences, and universities. Participants can earn a certificate of competency to support their long-term professional growth.

ONLINE COURSES

Online, self-paced courses show you how to set up 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—as long as you have computer (desktop or laptop) and an internet connection. Most eight-hour courses offer a certificate of competency upon completion of the built-in assessment.

WHY CHOOSE THE NVIDIA DEEP LEARNING INSTITUTE FOR HANDS-ON TRAINING?

- Access instructor-led workshops and online courses from anywhere with just your computer and an internet connection. Each participant will have access to a fully configured, GPU-accelerated workstation in the cloud.
- Obtain hands-on experience with the most widely used, industry-standard software, tools, and frameworks.
- Learn to build deep learning and accelerated computing applications for industries, such as healthcare, robotics, manufacturing, and more.
- 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 Deep Learning Institute certificate to demonstrate your subject matter competency and support your career growth.

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

Learn how deep learning works through hands-on exercises in computer vision and natural language processing. You will train deep learning models from scratch, learning tools and tricks to achieve highly accurate results. You'll also learn to leverage freely available, state-of-the-art pre-trained models to save time and get your deep learning application up and running quickly.

PREREQUISITES: Understanding of fundamental programming concepts in Python such as functions, loops, dictionaries, and arrays.

TOOLS, LIBRARIES, FRAMEWORKS: Tensorflow, Keras, Pandas, Numpy

LANGUAGE: English

> **Datasheet**

Building Intelligent Recommender Systems

Explore the fundamental tools and techniques for building highly effective recommender systems, as well as how to deploy GPU-accelerated solutions for real-time recommendations.

PREREQUISITES: Intermediate knowledge of Python, including understanding of list comprehension. Data science experience using Python and familiarity with NumPy and matrix mathematics.

TOOLS, LIBRARIES, FRAMEWORKS: CuDF, CuPy, TensorFlow 2, and NVIDIA Triton™ Inference Server

LANGUAGE: English

> **Datasheet**

Building Transformer-Based Natural Language Processing

Learn how to use Transformer-based natural language processing models for text classification tasks, such as categorizing documents. You will also get insight on how to leverage Transformer-based models for named-entity recognition (NER) tasks and analyze various model features, constraints, and characteristics to determine which model is best suited for a particular use case based on metrics, domain specificity, and available resources.

PREREQUISITES: Experience with Python coding and use of library functions and parameters. Fundamental understanding of a deep learning framework such as TensorFlow, PyTorch, or Keras. And basic understanding of neural networks

TOOLS, LIBRARIES, FRAMEWORKS: PyTorch, Pandas, NVIDIA NeMo, NVIDIA Triton™ Inference Server

LANGUAGE: English

> **Datasheet**

Fundamentals of Deep Learning for Multi-GPUs

Find out 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, LIBRARIES, FRAMEWORKS: TensorFlow

LANGUAGE: English

> **Datasheet**

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) data set.

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

TOOLS, LIBRARIES, FRAMEWORKS: TensorFlow

LANGUAGES: Japanese, Korean, Traditional Chinese

> **Datasheet**

DEEP LEARNING BY INDUSTRY

Deep Learning for Autonomous Vehicles—Perception

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

PREREQUISITES: Experience with CNNs and C++

TOOLS, LIBRARIES, FRAMEWORKS: TensorFlow, NVIDIA TensorRT™, Python, NVIDIA CUDA® C++, DIGITS

LANGUAGES: English, Simplified Chinese

> **Datasheet**

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 and basic coding experience in Python or similar language

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

LANGUAGE: English

> **Datasheet**

Deep Learning for Industrial Inspection

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

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

TOOLS, LIBRARIES, FRAMEWORKS: TensorFlow, TensorRT, Keras

LANGUAGES: English, Traditional Chinese

> **Datasheet**

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) and intermediate-level experience with C++ and Python

TOOLS, LIBRARIES, FRAMEWORKS: DeepStream 3.0, TensorFlow

LANGUAGES: English, Korean

> **Datasheet**

Deep Learning for Robotics

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

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

TOOLS, LIBRARIES, FRAMEWORKS: ROS, DIGITS, NVIDIA Jetson

LANGUAGE: English

> **Datasheet**

Applications of AI for Anomaly Detection

Learn to detect anomalies in large data sets to identify network intrusions using supervised and unsupervised machine learning techniques, such as accelerated XGBoost, autoencoders, and generative adversarial networks (GANs).

PREREQUISITES: Experience with CNNs and Python

TOOLS, LIBRARIES, FRAMEWORKS: RAPIDS, Keras, GANs, XGBoost

LANGUAGE: English

> **Datasheet**

Applications of AI for Predictive Maintenance

Discover how to identify anomalies and failures in time-series data, estimate the remaining useful life of the corresponding parts, and use this information to map anomalies to failure conditions.

PREREQUISITES: Experience with Python and deep networks

TOOLS, LIBRARIES, FRAMEWORKS: TensorFlow, Keras

LANGUAGE: English

> **Datasheet**

ACCELERATED COMPUTING

Fundamentals of Accelerated Computing with CUDA C/C++

Learn how to accelerate and optimize existing C/C++ CPU-only applications to leverage the power of GPUs using the most essential CUDA techniques and the Nsight Systems profiler.

PREREQUISITES: Basic C/C++ competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. No previous knowledge of CUDA programming is assumed.

TOOLS, LIBRARIES, FRAMEWORKS: C/C++, CUDA

LANGUAGES: English, Korean, Traditional Chinese

> **Datasheet**

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. Also, must have NumPy competancy, including the use of ndarrays and ufuncs

TOOLS, LIBRARIES, FRAMEWORKS: CUDA, Python, Numba, NumPy

LANGUAGE: English

> **Datasheet**

ACCELERATED DATA SCIENCE

Fundamentals of Accelerated Data Science with RAPIDS

Learn how to perform multiple analysis tasks on large data sets using RAPIDS, a collection of data science libraries that allows end-to-end GPU acceleration for data science workflows.

PREREQUISITES: Professional data science experience with Python, including proficiency in pandas and NumPy. Also, must have familiarity with common machine learning algorithms, including XGBoost, linear regression, DBSCAN, K-Means, and SSSP

TOOLS, LIBRARIES, FRAMEWORKS: RAPIDS, NumPy, XGBoost, DBSCAN, K-Means, SSSP, Python

LANGUAGE: English

> **Datasheet**

NETWORKING

The NVIDIA Mellanox Academy offers customizable training and certification on dozens of networking topics, including InfiniBand, Cumulus-Linux, protocols configuration such as Virtual Extensible LAN (VXLAN), Multi-Chassis Link Aggregation (MLAG), Border Gateway Protocol Ethernet VPN (BGP EVPN), and much more. The training combines hands-on practice and theoretical concepts to match job requirements and prepare participants for immediate productivity.

> To explore what's available, visit **academy.mellanox.com**

ONLINE COURSES

DEEP LEARNING FUNDAMENTALS

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.

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

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

TOOLS, LIBRARIES, FRAMEWORKS: Caffe, DIGITS

DURATION: 8 hours

PRICE: $90 (excludes tax, if applicable)

> **Learn More**

Getting Started with AI on Jetson Nano

Discover how to build a deep learning classification project with computer vision models using the NVIDIA Jetson Nano™ Developer Kit.

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

TOOLS, LIBRARIES, FRAMEWORKS: PyTorch, Jetson Nano

LANGUAGE: English

DURATION: 8 hours

PRICE: Free (hardware required)

> **Learn More**

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, LIBRARIES, FRAMEWORKS: TensorFlow, Python, TensorRT (TF-TRT)

LANGUAGE: English

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

Deep Learning at Scale with Horovod

Find out how to scale deep learning training to multiple GPUs with Horovod, the open-source distributed training framework originally built by Uber.

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

TOOLS, LIBRARIES, FRAMEWORKS: Horovod, TensorFlow 2, Keras

LANGUAGE: English

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

Getting Started with Image Segmentation

Learn how to categorize segments of an image.

PREREQUISITES: Basic experience training neural networks

TOOLS, LIBRARIES, FRAMEWORKS: TensorFlow

LANGUAGE: English

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

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, LIBRARIES, FRAMEWORKS: Keras

LANGUAGE: English

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

DEEP LEARNING BY INDUSTRY

HEALTHCARE

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 data set.

PREREQUISITES: Basic experience in Python

LANGUAGES: English, Simplified Chinese

TOOLS, LIBRARIES, FRAMEWORKS: PyTorch

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

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, LIBRARIES, FRAMEWORKS: TensorFlow

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

Data Augmentation and Segmentation with Generative Networks for Medical Imaging

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

PREREQUISITES: Experience with CNNs

TOOLS, LIBRARIES, FRAMEWORKS: TensorFlow

LANGUAGE: English

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

Coarse-to-Fine Contextual Memory for Medical Imaging

Find out 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, LIBRARIES, FRAMEWORKS: TensorFlow

LANGUAGE: English

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

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, LIBRARIES, FRAMEWORKS: DeepStream 3.0

LANGUAGE: English

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

Getting Started with DeepStream for Video Analytics on Jetson Nano

Explore how to build DeepStream applications to annotate video streams using object detection and classification networks.

PREREQUISITES: Basic familiarity with C

TOOLS, LIBRARIES, FRAMEWORKS: DeepStream, TensorRT, Jetson Nano

LANGUAGES: English

DURATION: 8 hours; Self-paced

PRICE: Free

> **Learn More**

ACCELERATED COMPUTING FUNDAMENTALS

Fundamentals of Accelerated Computing with CUDA C/C++

Discover how to accelerate and optimize existing C/C++ CPU-only applications to leverage the power of GPUs using the most essential CUDA techniques and the Nsight Systems profiler.

PREREQUISITES: Basic C/C++ competency, including familiarity with variable types, loops, conditional statements, functions, and array manipulations. No previous knowledge of CUDA programming is assumed.

TOOLS, LIBRARIES, FRAMEWORKS: C/C++, CUDA

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

DURATION: 8 hours

PRICE: $90 (excludes tax, if applicable)

> **Learn More**

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. Also, must have NumPy competency, including the use of ndarrays and ufuncs

TOOLS, LIBRARIES, FRAMEWORKS: CUDA, Python, Numba, NumPy

LANGUAGE: English

DURATION: 8 hours

PRICE: $90 (excludes tax, if applicable)

> **Learn More**

Fundamentals of Accelerated Computing with OpenACC

Find out 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++

TOOLS, LIBRARIES, FRAMEWORKS: OpenACC, C/C++

LANGUAGE: English

DURATION: 8 hours

PRICE: $90 (excludes tax, if applicable)

> **Learn More**

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, LIBRARIES, FRAMEWORKS: Docker, Singularity, HPC Container Maker (HPCCM), C/C++

LANGUAGE: English

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

OpenACC—2X in 4 Steps

Discover 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++

TOOLS, LIBRARIES, FRAMEWORKS: C/C++, OpenACC

LANGUAGE: English

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

ACCELERATED DATA SCIENCE

Fundamentals of Accelerated Data Science with RAPIDS

Find out how to perform multiple analysis tasks on large data sets using RAPIDS, a collection of data science libraries that allows end-to-end GPU acceleration for data science workflows.

PREREQUISITES: Professional data science experience with Python, including proficiency in pandas and NumPy. Also, must have familiarity with common machine learning algorithms, including XGBoost, linear regression, DBSCAN, K-Means, and SSSP

TOOLS, LIBRARIES, FRAMEWORKS: RAPIDS, NumPy, XGBoost, DBSCAN, K-Means, SSSP, Python

LANGUAGE: English

DURATION: 8 hours

PRICE: $90 (excludes tax, if applicable)

> **Learn More**

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, LIBRARIES, FRAMEWORKS: RAPIDS, cuDF, cuML, XGBoost

LANGUAGE: English

DURATION: 2 hours

PRICE: $30 (excludes tax, if applicable)

> **Learn More**

GPU COMPUTING IN THE DATA CENTER

Introduction to AI in the Data Center

Explore AI, GPU computing, NVIDIA AI software architecture, and how to implement and scale AI workloads in the enterprise data center.

PREREQUISITES: Basic knowledge of enterprise networking, storage, and data center operations

TOOLS, LIBRARIES, FRAMEWORKS: Artificial intelligence, machine learning, deep learning, GPU hardware and software

LANGUAGE: English

DURATION: 4 hours

> **Learn More**

NETWORKING

The NVIDIA Mellanox Academy offers dozens of online, self-paced courses and certifications on networking topics such as InfiniBand, remote direct memory access (RDMA) programming, Cumulus-Linux, data center protocols configuration, network automation tools, and much more.

> To explore the offerings, visit **academy.mellanox.com**

To get started with DLI hands-on training, visit
www.nvidia.com/dli

For questions, contact us at
nvdli@nvidia.com

NVIDIA | DEEP LEARNING INSTITUTE