



NVIDIA DLI COURSE CATALOGUE

BOSTON
TRAINING ACADEMY



DEEP
LEARNING
INSTITUTE

BOSTON TRAINING ACADEMY

EXPERT AI AND DEEP LEARNING TRAINING FOR YOUR STAFF AND CUSTOMERS

The mission of the Boston Training Academy (BTA) is to become a renowned developmental ground for talent engagement, education and solutions across a variety of disciplines. The BTA offers structured, face-to-face, labs and training delivered by world-class trainers that are tailored to the knowledge-base of the attendees.

TEACHING YOU TO SOLVE PROBLEMS WITH DEEP LEARNING

NVIDIA Deep Learning Institute (DLI) workshops, hosted by Boston, offer hands-on training for developers, data scientists, and researchers looking to solve challenging problems with deep learning.

Through self-paced labs and instructor-led workshops, the Deep Learning Institute teaches the latest techniques for designing, architecting, and deploying neural network-powered machine learning across a variety of application domains.

Students of the DLI will explore widely used open-source frameworks and NVIDIA's latest GPU-accelerated deep learning platforms. Boston are pleased to be DLI Delivery Partners providing training globally. Whilst we organise several public courses through the year, and our trainers are TA's at GTC's across the world, Boston recommends Private Workshops for organisations, so that they may benefit from our trainers' expertise to tailor workshops to their requirements.

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](#)

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 competency, including the use of ndarrays and ufuncs

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

LANGUAGE: English

> [Datasheet](#)

Accelerating CUDA C++ Applications with Multiple GPUs

Learn how to write CUDA C++ applications that efficiently and correctly utilize all available GPUs in a single node, dramatically improving the performance of applications and making the most cost-effective use of systems with multiple GPUs

PREREQUISITES: Professional experience programming CUDA C/C++ applications, including the use of the NVIDIA CUDA Compiler (NVCC), kernel launches, grid-stride loops, host-to-device and device-to-host memory transfers, and CUDA error handling. Familiarity with the Linux command line and experience using makefiles to compile C/C++ code.

TOOLS, LIBRARIES, FRAMEWORKS: CUDA C++, NVCC, Nsight Systems

LANGUAGES: English, Simplified Chinese

> [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



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

For questions, contact us at

UK (Headquarters)

+44 (0) 1727 876 100
sales@boston.co.uk
www.boston.co.uk

India

+91 22 5002 3262 (Mumbai office)
+91 80 4308 4000 (Bangalore office)
sales@bostonindia.in
www.bostonindia.in

Germany

+49 (0) 89 9090 1993
sales@boston-it.de
www.boston-it.de