NVIDIA DEEP LEARNING INSTITUTE

TRAINING CATALOG

Valid Through March 25, 2018

NVIDIA | DEEP LEARNING INSTITUTE

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. These include autonomous vehicles, robotics, healthcare image analysis, healthcare genomics, video analytics, financial trading, and more.

You'll learn how to train, optimize, and deploy neural networks using the latest tools, frameworks, and techniques for deep learning. DLI training for accelerated computing teaches you how to assess, parallelize, optimize, and deploy GPU-accelerated computing applications across a wide range of application domains.

DLI offers training in three formats:

- **Courses**
  Courses include 1+ days of hands-on training, during which you'll work through the implementation and deployment of an end-to-end project. Courses are available in both self-paced online format and instructor-led sessions.
- **Labs**
  Labs teach a specific technology or development technique through hands-on practice in a couple of hours. Labs are available in both self-paced online format and instructor-led sessions delivered at conferences.
- **Lectures**
  Expert-led talks are designed to cover core concepts and industry-specific use cases, and to prepare participants for hands-on training.

You can access DLI training in two ways:

- **Self-Paced Online Training**
  This training can be taken anytime from anywhere with access to a fully-configured GPU-accelerated workstation in the cloud. Get started with online training at **www.nvidia.com/dlilabs**.
- **Instructor-Led Workshops**
  Live, in-person workshops led by DLI-certified instructors are available around the world. To view upcoming workshops near you or request a workshop for your organization, visit **www.nvidia.com/dli**.

CONTENT BY APPLICATION DOMAIN

FUNDAMENTALS OF DEEP LEARNING

Courses

Labs

Lectures

AUTONOMOUS VEHICLES

Courses

GAME DEVELOPMENT AND DIGITAL CONTENT

Courses

HEALTHCARE

Courses

FINANCE

Courses

INTELLIGENT VIDEO ANALYTICS

Courses

ACCELERATED COMPUTING

Courses

Labs

FUNDAMENTALS OF DEEP LEARNING

COURSES

FUNDAMENTALS OF DEEP LEARNING FOR COMPUTER VISION

PREREQUISITES: None
FRAMEWORK: Caffe, DIGITS
LANGUAGES: English, Chinese

DURATION: 8 hours
FORMAT: Self-paced online or instructor-led

Explore the fundamentals of deep learning by training neural networks and using results to improve performance and capabilities.

In this hands-on course, you'll learn the basics of deep learning by training and deploying neural networks. You'll learn how to:

- Implement common deep learning workflows, such as image classification and object detection.
- Experiment with data, training parameters, network structure, and other strategies to increase performance and capability.
- Deploy your neural networks to start solving real-world problems.

Upon completion, you'll be able to start solving problems on your own with deep learning.

FUNDAMENTALS OF DEEP LEARNING FOR MULTIPLE DATA TYPES

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

DURATION: 8 hours

FRAMEWORK: TensorFlow

FORMAT: Instructor-led

This hands-on course explores 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 MSCOCO dataset to generate captions from images and video by:

> Implementing deep learning workflows like image segmentation and text generation

> Comparing and contrasting data types, workflows, and frameworks

> Combining computer vision and natural language processing

Upon completion, you'll be able to solve deep learning problems that require multiple types of data inputs.

FUNDAMENTALS OF DEEP LEARNING FOR NATURAL LANGUAGE PROCESSING

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

DURATION: 8 hours

FRAMEWORK: TensorFlow, Keras

FORMAT: Instructor-led

In this course, you will receive hands-on training on the latest techniques for understanding textual input using Natural Language Processing. You'll learn how to:

> Classify words to accurately understand their meaning

> Handle factual queries and their semantic meaning

> Train Machine Translators from one language to another

Upon completion of this course, you'll be proficient in Natural Language Processing using neural networks in any application.

LABS

APPLICATIONS OF DEEP LEARNING WITH CAFFE, THEANO, AND TORCH

PREREQUISITES: None
FRAMEWORK: Caffe, Theano, Torch
LANGUAGES: English
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led

Learn how deep learning will change the future of computing. In this hands-on session (no technical background required), you'll:

> Compare deep learning to traditional methods

> Run training and inference with three different deep learning frameworks

> Learn how deep learning works and why the GPU is integral

Upon completion, you'll be better equipped to decide how you or your organization can get started with deep learning.

DEEP LEARNING WORKFLOWS WITH TENSORFLOW, MXNET, AND NVIDIA DOCKER

PREREQUISITES: Bash terminal familiarity
FRAMEWORK: TensorFlow and MXNet
LANGUAGES: English
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led

The NVIDIA Docker plugin makes it possible to containerize production-grade deep learning workflows using GPUs. Learn to reduce host configuration and administration by:

> Learning to work with Docker images and manage the container lifestyle

> Accessing images on the public Docker image registry—DockerHub—for maximum reuse in creating composable lightweight containers

> Training neural networks using both TensorFlow and MXNet frameworks

Upon completion, you'll be able to containerize and distribute pre-configured images for deep learning.

IMAGE CLASSIFICATION WITH MICROSOFT COGNITIVE TOOLKIT

PREREQUISITES: None
FRAMEWORK: Microsoft Cognitive Toolkit
LANGUAGES: English
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led

Learn to train a neural network using the Microsoft Cognitive Toolkit framework. You'll build and train increasingly complex networks to:

> Compare the expression of a neural network using BrainScript's "Simple Network Builder" vs. the more generalizable "Network Builder"

> Visualize neural network graphs

> Train and test neural network to classify handwritten digits

Upon completion, you'll have basic knowledge of convolutional neural networks and be prepared to move to the more advanced usage of Microsoft Cognitive Toolkit.

LINEAR CLASSIFICATION WITH TENSORFLOW

PREREQUISITES: None
FRAMEWORK: TensorFlow
LANGUAGES: English
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led

Learn to make predictions from structured data using TensorFlow's TFLearn API. Through the challenge of predicting personal income when given census data, you'll:

> Load, view, and organize data from a CSV for machine learning

> Split an existing dataset into features and labels (input, output) of a neural network

> Build from linear to deep models and assess the difference in performance

Upon completion, you'll be able to make predictions from your own structured data.

SIGNAL PROCESSING WITH DIGITS

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

DURATION: 2 hours

FRAMEWORK: Caffe, DIGITS

FORMAT: Self-paced online or instructor-led

LANGUAGES: English

The fact that deep neural networks are better at classifying images than humans has implications beyond what we typically think of computer vision.

In this hands-on lab, you'll convert radio frequency (RF) signals into images to detect a weak signal corrupted by noise. You'll be trained how to:

> Treat non-image data as image data

> Implement a deep learning workflow (load, train, test, adjust) in DIGITS

> Test performance programmatically and guide performance improvements

Upon completion, you'll be able to classify both image and image-like data using deep learning.

LECTURES

DEEP LEARNING DEMYSTIFIED

LANGUAGES: English, Chinese

This lecture provides an overview on deep learning that includes key terminology, industry use cases, how deep learning differs from the previous algorithmic approach, and how a deep neural network gets trained, optimized, and deployed. You'll learn how deep learning can be applied to challenging problems, what types of problems benefit most from deep learning, what skills and knowledge are needed to use deep learning, and the characteristics of successful deep learning projects.

BEST PRACTICES FOR STARTING A DEEP LEARNING PROJECT

LANGUAGES: English

This lecture examines common characteristics that successful deep learning applications share. You'll explore the types of problems, dataset/model characteristics, workflows, and mindsets that contribute to successful problem solving with deep learning.

AUTONOMOUS VEHICLES

COURSES

PERCEPTION FOR AUTONOMOUS VEHICLES

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

DURATION: 8 hours

FRAMEWORK: TensorFlow, DIGITS, TensorRT

FORMAT: Instructor-led

LANGUAGES: English

In this course, you'll learn how to design, train, and deploy deep neural networks for autonomous vehicles using the NVIDIA DRIVE™ PX2 development platform. Learn how to:

> Integrate sensor input using the DriveWorks software stack

> Train a semantic segmentation neural network

> Optimize, validate, and deploy a trained neural network using TensorRT

Upon completion of this course, students will be able to create and optimize perception components for autonomous vehicles using DRIVE PX2.

GAME DEVELOPMENT AND DIGITAL CONTENT

COURSES

DIGITAL CONTENT CREATION USING GANS AND AUTOENCODERS

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

DURATION: 8 hours

FRAMEWORK: TensorFlow, Theanos, DIGITS

FORMAT: Instructor-led

LANGUAGES: English

Get hands-on training on the latest techniques for designing, training, and deploying neural networks for digital content creation. You'll learn how to:

- Train a Generative Adversarial Network (GAN) to generate images
- Explore the architectural innovations and training techniques used to make arbitrary video style transfer
- Train your own denoiser for rendered images

Upon completion of this course, you'll be able to start creating digital assets using deep learning approaches.

HEALTHCARE

COURSES

DEEP LEARNING FOR HEALTHCARE IMAGE ANALYSIS

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

DURATION: 8 hours

FRAMEWORK: Caffe, MXNet, TensorFlow

FORMAT: Instructor-led

LANGUAGES: English

This hands-on course explores how to apply Convolutional Neural Networks (CNNs) to MRI scans to perform a variety of medical tasks and calculations. You'll learn how to:

- Perform image segmentation on MRI images to determine the location of the left ventricle.
- Calculate ejection fractions by measuring differences between diastole and systole using CNNs applied to MRI scans to detect heart disease.
- Apply CNNs to MRI scans of LGGs to determine 1p/19q chromosome co-deletion status.

Upon completion of this course, you'll be able to apply CNNs to MRI scans to conduct a variety of medical tasks.

DEEP LEARNING FOR HEALTHCARE GENOMICS

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

DURATION: 8 hours

FRAMEWORK: Caffe, TensorFlow, Theano

FORMAT: Instructor-led

LANGUAGES: English

This hands-on course teaches you how to apply deep learning to detect chromosome co-deletion and search for motifs in genomic sequences. You'll learn how to:

> Understand the basics of Convolutional Neural Networks (CNNs) and how they work.

> Apply CNNs to MRI scans of LGGs to determine 1p/19q chromosome co-deletion status.

> Use the DragoNN toolkit to simulate genomic data and to search for motifs.

Upon completion of this course, you'll be able to: understand how CNNs work, evaluate MRI images using CNNs, and use real regulatory genomic data to research new motifs.

FINANCE

COURSES

DEEP LEARNING FOR FINANCE TRADING STRATEGY

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

DURATION: 8 hours

FORMAT: Instructor-led

LANGUAGES: English

Linear techniques like principal component analysis (PCA) are the workhorses of creating 'eigenportfolios' for use in statistical arbitrage strategies. Other techniques using time series financial data are also prevalent. But now, trading strategies can be advanced with the power of deep neural networks.

In this course, you'll learn how to:

- Prepare time series data and test network performance using training and test datasets
- Structure and train a LSTM network to accept vector inputs and make predictions
- Use the Autoencoder as anomaly detector to create an arbitrage strategy

Upon completion, you'll be able to use time series financial data to make predictions and exploit arbitrage using neural networks.

INTELLIGENT VIDEO ANALYTICS

COURSES

DEEP LEARNING FOR FULL MOTION VIDEO ANALYTICS

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

DURATION: 8 hours

FORMAT: Instructor-led

LANGUAGES: English

Traffic cameras, drones, and aerial sensor platforms are collecting huge amounts of video footage, which requires advanced deep learning techniques to transform data into actionable insights. The first step in more complex deep learning workflows is detecting specific types of objects, which involves identification, classification, segmentation, prediction, and recommendation.

In this course, you'll learn how to:

- Train and evaluate deep learning models using the TensorFlow Object Detection API
- Explore the strategies and trade-offs involved in developing high-quality neural network models for track moving objects in large-scale video datasets
- Optimize inference times using TensorRT for real-time applications

Upon completion, you'll be able to deploy object detection and tracking networks to work on real-time, large-scale video streams.

ACCELERATED COMPUTING

COURSES

FUNDAMENTALS OF ACCELERATED COMPUTING WITH CUDA C/C++

PREREQUISITES: None
DURATION: 8 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

The CUDA computing platform enables the acceleration of CPU-only applications to run on the world's fastest massively parallel GPUs. Experience C/C++ application acceleration by:

- Accelerating CPU-only applications to run their latent parallelism on GPUs
- Utilizing essential CUDA memory management techniques to optimize accelerated applications
- Exposing accelerated application potential for concurrency and exploiting it with CUDA streams
- Leveraging command line and visual profiling to guide and check your work

Upon completion, you'll be able to accelerate and optimize existing C/C++ CPU-only applications using the most essential CUDA tools and techniques. You'll understand an iterative style of CUDA development that will allow you to ship accelerated applications fast.

FUNDAMENTALS OF ACCELERATED COMPUTING WITH OPENACC

PREREQUISITES: None
DURATION: 8 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

Learn the basics of OpenACC, a high-level programming language for programming on GPUs. This course is for anyone with some C/C++ experience who is interested in accelerating the performance of their applications beyond the limits of CPU-only programming. In this course, you'll learn:

> Four simple steps to accelerating your already existing application with OpenACC

> How to profile and optimize your OpenACC codebase

> How to program on multi-GPU systems by combining OpenACC with MPI

Upon completion, you'll be able to build and optimize accelerated heterogeneous applications on multiple GPU clusters using a combination of OpenACC, CUDA-aware MPI, and NVIDIA profiling tools.

LABS

GPU MEMORY OPTIMIZATIONS WITH C/C++

PREREQUISITES: Basic CUDA C/C++ competency
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

In this lab, you'll learn about a number of useful memory optimization techniques to use when programming with CUDA C/C++ on an NVIDIA GPU, and how to use the NVIDIA Visual Profiler (NVVP) to support these optimizations. You'll:

> Implement a naive matrix transposing algorithm

> Perform several cycles of profiling the algorithm with NVVP and then optimizing its performance

Upon completion, you'll have learned how to analyze and improve both global and shared memory access patterns, and be able to optimize your accelerated C/C++ applications.

ACCELERATING APPLICATIONS WITH GPU-ACCELERATED LIBRARIES IN C/C++

PREREQUISITES: Basic CUDA C/C++ competency
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

Learn how to accelerate your C/C++ application using drop-in libraries to harness the massively parallel power of NVIDIA GPUs. In about two hours, you'll work through three exercises, including:

> Using cuBLAS to accelerate a matrix multiplication algorithm

> Combining libraries by adding cuRAND API calls to the previous cuBLAS calls

> Using nvprof to profile code and optimize your accelerated code

Upon completion, you'll be able to use CUDA-optimized libraries to accelerate your C/C++ applications.

USING THRUST TO ACCELERATE C++

PREREQUISITES: Basic CUDA C/C++ competency
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

Thrust is a parallel algorithms library loosely based on the C++ Standard Template Library, which allows developers to quickly embrace the power of parallel computing. Thrust code can be compiled to run on the massively parallel NVIDIA GPUs, as well as OpenMP and Intel's Threading Building Blocks.

In this lab, you'll learn the following Thrust features, and incorporate them all into a case study:

> Iterators, containers, and functions

> Porting to CPU processing

> Exception and error handling

Upon completion, you'll be able to build GPU-accelerated applications in C/C++ that use the powerful Thrust library.

ACCELERATING APPLICATIONS WITH GPU-ACCELERATED LIBRARIES IN PYTHON

PREREQUISITES: None
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

Learn how to accelerate your Python application using GPU drop-in libraries to harness the massively parallel power of NVIDIA GPUs. You'll work through three exercises, including:

> Using a Python profiler to determine which part of the would benefit most from acceleration

> Using a cuRAND API call to optimize the application

> Profile and optimize again using the CUDA Runtime API to optimize data movement

Upon completion, you'll be ready to start accelerating your Python applications using CUDA and CUDA-optimized libraries.

ACCELERATING APPLICATIONS WITH CUDA FORTRAN

PREREQUISITES: None
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

Learn how to accelerate your Fortran application using CUDA to harness the massively parallel power of NVIDIA GPUs. You'll work through three exercises, including:

> Accelerating SAXPY algorithms

> Accelerating Matrix Multiply algorithms

> Accelerating heat conduction algorithms

Upon completion, you'll be able to use the CUDA platform to accelerate Fortran applications.

GPU MEMORY OPTIMIZATIONS WITH FORTRAN

PREREQUISITES: Basic CUDA C/C++ competency
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

In this lab, you'll learn about a number of useful memory optimization techniques to use when programming with CUDA Fortran on an NVIDIA GPU, and how to use the NVIDIA Visual Profiler (NVVP) to support these optimizations. You'll:

> Implement a naive matrix transposing algorithm

> Perform several cycles of profiling the algorithm with NVVP and then optimize its performance

Upon completion, you'll have learned how to analyze and improve both global and shared memory access patterns, and be able to optimize your accelerated Fortran applications.

ACCELERATING APPLICATIONS WITH GPU-ACCELERATED LIBRARIES IN FORTRAN

PREREQUISITES: Basic CUDA Fortran competency
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

Learn how to accelerate your Fortran application using GPU Libraries to harness the massively parallel power of NVIDIA GPUs. You'll work through three exercises, including:

> Using cuBLAS to accelerate a matrix multiplication algorithm

> Combining libraries by adding cuRAND API calls to the previous cuBLAS calls

> Using nvprof to profile code and optimize your accelerated code

Upon completion, you'll be able to use CUDA-optimized libraries to accelerate your Fortran applications.

INTRODUCTION TO ACCELERATED COMPUTING

PREREQUISITES: None
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

This lab will expose you to a collection of techniques for accelerating applications. You'll:

> Use CUDA-accelerated libraries to accelerate application code

> Use compiler directives like OpenACC to accelerate application code

> Use the CUDA platform to accelerate application code

Upon completion, you'll be ready to accelerate your applications using a number of acceleration techniques.

PROFILE-DRIVEN APPROACH TO ACCELERATE SEISMIC APPLICATION WITH OPENACC

PREREQUISITES: None
DURATION: 2 hours
FORMAT: Self-paced online or instructor-led
LANGUAGES: English

In this lab, you'll use PGPROF, a host and GPU profiling tool, to accelerate an open source seismic processing application. This lab follows the 4-stage APOD development cycle and will show you how to:

> Assess critical regions of the application, profile baseline CPU code, and decorate key loops with directives in order to parallelize them

> Use profile and verbose compiler output to decorate data directives, optimize, and measure best performance

> Use the compiler "multicore" option with OpenACC directives for portable performance

Upon completion, you'll be able to use PGPROF with OpenACC to accelerate your C/C++ applications.

Visit **www.nvidia.com/dli** to get started.

NVIDIA