

NVIDIA DEEP LEARNING INSTITUTE

TRAINING CATALOG

Valid Through July 31, 2018

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, digital content creation, finance, healthcare, intelligent video analytics, and more.

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

DLI offers training in three formats:

1. Instructor-Led Workshops

Full-day workshops include hands-on training and lectures delivered onsite by DLI certified instructors. To view upcoming workshops near you or request a workshop for your organization, visit www.nvidia.com/dli.

2. Online Courses

Online courses teach you how to implement and deploy an end-to-end project through hands-on training in 6 hours.

3. Online Mini Courses

Online mini courses teach a specific technology or development technique through hands-on training in 2 hours.

DLI online training is designed as a self-paced learning experience for independent learners. Online 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.

Certification

Through built-in assessments, participants can earn certification to prove subject matter competency and support professional career growth. Certification is available for a handful of workshops and online courses today.

Enterprise Training Solutions

For organizations interested in transforming their workforce with deep learning and accelerated computing, DLI offers enterprise solutions that include hands-on online and onsite training for employees, executive briefings, and enterprise-level reporting. DLI enterprise solutions help organizations solve challenging problems, improve staff productivity, and become leaders in AI. Contact dli-request@nvidia.com for more information.

WORKSHOPS BY INDUSTRY AT A GLANCE

FUNDAMENTALS OF DEEP LEARNING

- >>Fundamentals of Deep Learning for Computer Vision
- >>Fundamentals of Deep Learning for Multiple Data Types

AUTONOMOUS VEHICLES

- >>Deep Learning for Autonomous Vehicles – Perception

GAME DEVELOPMENT AND DIGITAL CONTENT

- >>Deep Learning for Digital Content Creation using GANs and Autoencoders

HEALTHCARE

- >>Deep Learning for Healthcare Image Analysis
- >>Deep Learning for Healthcare Genomics

FINANCE

- >>Deep Learning for Finance Trading Strategy

INTELLIGENT VIDEO ANALYTICS

- >>Deep Learning for Full Motion Video Analytics

ACCELERATED COMPUTING

- >>Fundamentals of Accelerated Computing with CUDA C/C++

ONLINE TRAINING BY INDUSTRY AT A GLANCE

FUNDAMENTALS OF DEEP LEARNING

Courses

>> Fundamentals of Deep Learning for Computer Vision

Mini Courses

>> Image Classification with DIGITS

>> Object Detection with DIGITS

>> Neural Network Deployment with DIGITS and TensorRT

>> Applications of Deep Learning with Caffe, Theano, and Torch

>> Deep Learning Workflows with TensorFlow, MXNet, and NVIDIA-Docker

>> Image Segmentation with TensorFlow

>> Image Classification with Microsoft Cognitive Toolkit

>> Linear Classification with TensorFlow

>> Signal Processing with DIGITS

GAME DEVELOPMENT AND DIGITAL CONTENT

Mini Courses

>> Image Creation using GANs with TensorFlow and DIGITS

>> Image Style Transfer with Torch

>> Rendered Image Denoising using Autoencoders

HEALTHCARE

Courses

>> Deep Learning for Healthcare Image Analysis

>> Deep Learning for Healthcare Genomics

Mini Courses

>> Modeling Time Series Data with Recurrent Neural Networks in Keras

INTELLIGENT VIDEO ANALYTICS

Mini Courses

>> Object Detection for Full Motion Video

>> Object Tracking for Large Scale Full Motion Video

>> Deployment for Intelligent Video Analytics using TensorRT

ACCELERATED COMPUTING

Courses

- >> Fundamentals of Accelerated Computing with CUDA C/C++
- >> Fundamentals of Accelerated Computing with CUDA Python
- >> Fundamentals of Accelerated Computing with OpenACC

Mini Courses

- >> Accelerated Applications with CUDA C/C++
- >> GPU Memory Optimizations with C/C++
- >> Accelerating Applications with GPU-Accelerated Libraries in C/C++
- >> Using Thrust to Accelerate C++
- >> Accelerating Applications with GPU-Accelerated Libraries in Python
- >> Accelerating Applications with CUDA Fortran
- >> GPU Memory Optimizations with Fortran
- >> Accelerating Applications with GPU-Accelerated Libraries in Fortran
- >> Introduction to Accelerated Computing
- >> Profile-Driven Approach to Accelerate Seismic Applications with OpenACC
- >> OpenACC – 2X in 4 Steps
- >> Introduction to Multi-GPU Programming with MPI and OpenACC
- >> Advanced Multi-GPU Programming with MPI and OpenACC
- >> Profiling and Parallelizing with OpenACC
- >> Expressing Data Movement and Optimizing Loops with OpenACC
- >> Pipelining Work on the GPU with OpenACC

WORKSHOPS

FUNDAMENTALS OF DEEP LEARNING FOR COMPUTER VISION

PREREQUISITES: Technical background

FRAMEWORKS: Caffe, DIGITS

LANGUAGES: English, Chinese

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

In this 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 for Computer Vision" or similar experience

FRAMEWORKS: TensorFlow

LANGUAGES: English

This 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.

DEEP LEARNING FOR AUTONOMOUS VEHICLES – PERCEPTION

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar experience

FRAMEWORKS: TensorFlow, DIGITS, TensorRT

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, students will be able to create and optimize perception components for autonomous vehicles using DRIVE PX2.

DEEP LEARNING FOR DIGITAL CONTENT CREATION USING GANS AND AUTOENCODERS

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar experience

FRAMEWORKS: TensorFlow, Theano, DIGITS

LANGUAGES: English

Explore 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, you’ll be able to start creating digital assets using deep learning approaches.

DEEP LEARNING FOR HEALTHCARE IMAGE ANALYSIS

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar experience

FRAMEWORKS: Caffe, MXNet, TensorFlow

LANGUAGES: English

This 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, 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 for Computer Vision” or similar experience

FRAMEWORKS: Caffe, TensorFlow, Theano

LANGUAGES: English

This 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, you’ll be able to: understand how CNNs work, evaluate MRI images using CNNs, and use real regulatory genomic data to research new motifs.

DEEP LEARNING FOR FINANCE TRADING STRATEGY

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar experience

FRAMEWORKS: TensorFlow

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.

DEEP LEARNING FOR FULL MOTION VIDEO ANALYTICS

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar experience

FRAMEWORKS: TensorFlow

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.

FUNDAMENTALS OF ACCELERATED COMPUTING WITH CUDA C/C++

PREREQUISITES: Basic C/C++ competency

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.

ONLINE TRAINING

FUNDAMENTALS OF DEEP LEARNING

COURSES

FUNDAMENTALS OF DEEP LEARNING FOR COMPUTER VISION

PREREQUISITES: Technical background

FRAMEWORKS: Caffe, DIGITS

LANGUAGES: English, Chinese

PRICE: \$90

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

In this 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.

MINI COURSES

IMAGE CLASSIFICATION WITH DIGITS

PREREQUISITES: Technical background

FRAMEWORKS: DIGITS

LANGUAGES: English, Chinese, Japanese

PRICE: Free

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.

OBJECT DETECTION WITH DIGITS

PREREQUISITES: Technical background

FRAMEWORKS: DIGITS

LANGUAGES: English, Chinese

PRICE: Free

Learn how to detect objects using computer vision and deep learning by identifying a purpose-built network and using end-to-end labeled data.

NEURAL NETWORK DEPLOYMENT WITH DIGITS AND TENSORRT

PREREQUISITES: Technical background

FRAMEWORKS: DIGITS, TensorRT

LANGUAGES: English, Chinese

PRICE: \$30

Learn to deploy deep learning to applications that recognize and detect images in real-time.

APPLICATIONS OF DEEP LEARNING WITH CAFFE, THEANO, AND TORCH

PREREQUISITES: None

FRAMEWORKS: Caffe, Theano, Torch

LANGUAGES: English

PRICE: \$30

Explore how deep learning works and how it will change the future of computing.

DEEP LEARNING WORKFLOWS WITH TENSORFLOW, MXNET, AND NVIDIA DOCKER

PREREQUISITES: Bash terminal familiarity

FRAMEWORKS: TensorFlow and MXNet

LANGUAGES: English

PRICE: \$30

Learn how to use the NVIDIA Docker plug-in to containerize production-grade deep learning workflows using GPUs.

IMAGE SEGMENTATION WITH TENSORFLOW

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

FRAMEWORKS: TensorFlow

LANGUAGES: English

PRICE: \$30

Learn to combine computer vision and natural language processing to describe scenes using deep learning.

IMAGE CLASSIFICATION WITH MICROSOFT COGNITIVE TOOLKIT

PREREQUISITES: None

FRAMEWORKS: Microsoft Cognitive Toolkit

LANGUAGES: English

PRICE: \$30

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

LINEAR CLASSIFICATION WITH TENSORFLOW

PREREQUISITES: None

FRAMEWORKS: TensorFlow

LANGUAGES: English

PRICE: \$30

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

SIGNAL PROCESSING WITH DIGITS

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

LANGUAGES: English

PRICE: \$30

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.

GAME DEVELOPMENT AND DIGITAL CONTENT

MINI COURSES

IMAGE CREATION USING GANS WITH TENSORFLOW AND DIGITS

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar deep learning experience

FRAMEWORKS: TensorFlow, DIGITS

LANGUAGES: English

PRICE: \$30

Discover how to train a generative adversarial network (GAN) to generate image content in DIGITS.

IMAGE STYLE TRANSFER WITH TORCH

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar deep learning experience

FRAMEWORKS: Torch

LANGUAGES: English

PRICE: \$30

Learn how to transfer the look and feel of one image to another image by extracting distinct visual features using convolutional neural networks (CNNs).

RENDERED IMAGE DENOISING USING AUTOENCODERS

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar deep learning experience

FRAMEWORKS: TensorFlow

LANGUAGES: English

PRICE: \$30

Explore how a neural network with an autoencoder can be used to dramatically speed up the removal of noise in ray-traced images.

HEALTHCARE

COURSES

DEEP LEARNING FOR HEALTHCARE IMAGE ANALYSIS

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar experience

FRAMEWORKS: Caffe, MXNet, TensorFlow

LANGUAGES: English

PRICE: \$90

This 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 CNN s to MRI scans of LGGs to determine 1p/19q chromosome co-deletion status.

Upon completion, 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 for Computer Vision” or similar experience

FRAMEWORKS: Caffe, TensorFlow, Theano

LANGUAGES: English

PRICE: \$60

This 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, you’ll be able to: understand how CNNs work, evaluate MRI images using CNNs, and use real regulatory genomic data to research new motifs.

MINI COURSES

MODELING TIME SERIES DATA WITH RECURRENT NEURAL NETWORKS IN KERAS

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar deep learning experience

FRAMEWORKS: Keras

LANGUAGES: English

PRICE: Free

Explore how to classify and forecast time-series data using recurrent neural networks (RNNs), such as modeling a patient’s health over time.

INTELLIGENT VIDEO ANALYTICS

MINI COURSES

OBJECT DETECTION FOR FULL MOTION VIDEO

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar deep learning experience

FRAMEWORKS: TensorFlow

LANGUAGES: English

PRICE: \$30

Discover how to analyze video data by implementing object detection methods using deep learning.

OBJECT TRACKING FOR LARGE SCALE FULL MOTION VIDEO

PREREQUISITES: “Object Detection for Full-Motion Video” or similar deep learning experience

FRAMEWORKS: TensorFlow

LANGUAGES: English

PRICE: \$30

Explore how to track moving objects in large-scale video datasets using a deep neural network model.

DEPLOYMENT FOR INTELLIGENT VIDEO ANALYTICS USING TENSORRT

PREREQUISITES: “Fundamentals of Deep Learning for Computer Vision” or similar deep learning experience

FRAMEWORKS: TensorRT

LANGUAGES: English

PRICE: \$30

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

ACCELERATED COMPUTING

COURSES

FUNDAMENTALS OF ACCELERATED COMPUTING WITH CUDA C/C++

PREREQUISITES: Basic C/C++ competency

LANGUAGES: English

PRICE: \$90

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.

FUNDAMENTALS OF ACCELERATED COMPUTING WITH CUDA PYTHON

PREREQUISITES: Basic Python and NumPy competency

LANGUAGES: English

PRICE: \$90

This course explores how to use Numba—the just-in-time, type-specializing, Python function compiler—to accelerate Python programs to run on massively parallel NVIDIA GPUs. You'll learn how to:

- >>Use Numba to compile CUDA kernels from NumPy universal functions (ufuncs)
- >>Use Numba to create and launch custom CUDA kernels
- >>Apply key GPU memory management techniques

Upon completion, you'll be able to use Numba to compile and launch CUDA kernels to accelerate your Python applications on NVIDIA GPUs.

FUNDAMENTALS OF ACCELERATED COMPUTING WITH OPENACC

PREREQUISITES: Basic C/C++ competency

LANGUAGES: English

PRICE: \$116

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.

MINI COURSES

ACCELERATING APPLICATIONS WITH CUDA C/C++

PREREQUISITES: Basic C/C++ competency

LANGUAGES: English

PRICE: Free

Learn how to accelerate your C/C++ application using CUDA to harness the massively parallel power of NVIDIA GPUs.

OPENACC – 2X IN 4 STEPS

PREREQUISITES: Basic C/C++ competency

LANGUAGES: English

PRICE: \$30

Learn how to accelerate C/C++ or Fortran applications using OpenACC to harness the massively parallel power of NVIDIA GPUs.

INTRODUCTION TO ACCELERATED COMPUTING

PREREQUISITES: None

LANGUAGES: English

PRICE: \$30

Explore a variety of techniques for accelerating applications, including CUDA and OpenACC.

GPU MEMORY OPTIMIZATIONS WITH C/C++

PREREQUISITES: Basic CUDA C/C++ competency

LANGUAGES: English

PRICE: \$30

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.

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

PREREQUISITES: Basic CUDA C/C++ competency

LANGUAGES: English

PRICE: \$30

Learn how to accelerate your C/C++ application using CUDA-optimized libraries to harness the massively parallel power of NVIDIA GPUs.

ACCELERATING APPLICATIONS WITH GPU-ACCELERATED LIBRARIES IN PYTHON

PREREQUISITES: None

LANGUAGES: English

PRICE: \$30

Learn how to accelerate your Python application using CUDA-optimized libraries to harness the massively parallel power of NVIDIA GPUs.

USING THRUST TO ACCELERATE C++

PREREQUISITES: Basic CUDA C/C++ competency

LANGUAGES: English

PRICE: \$30

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

PROFILING AND PARALLELIZING WITH OPENACC

PREREQUISITES: “OpenACC – 2X in 4 Steps”

LANGUAGES: English

PRICE: \$10

Get hands-on experience with the first two steps of the OpenACC programming cycle.

EXPRESSING DATA MOVEMENT AND OPTIMIZING LOOPS WITH OPENACC

PREREQUISITES: “OpenACC - 2X in 4 Steps”, “Profiling and Parallelizing with OpenACC”

LANGUAGES: English

PRICE: \$15

Learn how to add data management and loop directives to optimize OpenACC accelerated code.

INTRODUCTION TO MULTI-GPU PROGRAMMING WITH MPI AND OPENACC

PREREQUISITES: “OpenACC – 2X in 4 Steps”

LANGUAGES: English

PRICE: \$30

Explore how to program multi-GPU systems or GPU clusters using the Message Passing Interface (MPI) and OpenACC.

ADVANCED MULTI-GPU PROGRAMMING WITH MPI AND OPENACC

PREREQUISITES: “Introduction to Multi-GPU Programming with MPI and OpenACC”

LANGUAGES: English

PRICE: \$30

Learn how to improve multi-GPU MPI+OpenACC programs by overlapping communication with computation and handling noncontiguous halo updates.

PIPELINING WORK ON THE GPU WITH OPENACC

PREREQUISITES: “OpenACC - 2X in 4 Steps”, “Profiling and Parallelizing with OpenACC”, and “Expressing Data Movement and Optimizing Loops with OpenACC”

LANGUAGES: English

PRICE: \$30

Learn how to optimize data copies in OpenACC applications to overlap with GPU computation using a simple technique called pipelining.

PROFILE-DRIVEN APPROACH TO ACCELERATE SEISMIC APPLICATIONS WITH OPENACC

PREREQUISITES: None

LANGUAGES: English

PRICE: \$30

Learn how to use PGI Profiler (PGPROF), a host and GPU profiling tool, with OpenACC to accelerate your C/C++ applications.

ACCELERATING APPLICATIONS WITH CUDA FORTRAN

PREREQUISITES: None

LANGUAGES: English

PRICE: \$30

Learn how to accelerate your Fortran application using CUDA to harness the massively parallel power of NVIDIA GPUs.

GPU MEMORY OPTIMIZATIONS WITH CUDA FORTRAN

PREREQUISITES: “Accelerating Applications with CUDA Fortran”

LANGUAGES: English

PRICE: \$30

Discover useful memory optimization techniques for programming with CUDA Fortran on an NVIDIA GPU and how to use the NVIDIA Visual Profiler (NVVP) to support these optimizations.

ACCELERATING APPLICATIONS WITH GPU-ACCELERATED LIBRARIES IN FORTRAN

PREREQUISITES: Basic CUDA Fortran competency

LANGUAGES: English

PRICE: \$30

Learn how to accelerate your Fortran application using CUDA-optimized libraries.

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